Computational Geometry 42 (2009) 464–470

Contents lists available at ScienceDirect

Computational Geometry: Theory and Applications

www.elsevier.com/locate/comgeo

Shortest descending paths through given faces

Mustaq Ahmed *, Anna Lubiw

David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, N2L 3G1, Canada

ARTICLE INFO

Article history:
Received 11 December 2006
Received in revised form 17 October 2007
Accepted 17 October 2007
Available online 24 December 2008
Communicated by H. Meijer and D. Rappaport

Keywords:
Shortest descending paths
Shortest paths
Terrain
Convex optimization

ABSTRACT

A path from s to t on a polyhedral terrain is *descending* if the height of a point p never increases while we move p along the path from s to t. No efficient algorithm is known to find a shortest descending path from s to t in a polyhedral terrain. We give some properties of such paths. In the case where the face sequence is specified, we show that the shortest descending path is unique, and use convex optimization to give an ϵ-approximation algorithm that computes the path in $O(n^{3.5}\log(\frac{1}{\epsilon}))$ time.

1. Introduction

The problem of determining a shortest path in a polyhedral terrain has many applications in robotics, industrial automation, Geographic Information Systems and wire routing. In certain applications, the feasibility of a path is determined by the height of the points. For example, for laying a canal of minimum length from the source of water at the top of a mountain to fields for irrigation purpose [11], and for skiing down a mountain along a shortest route, we need to compute a shortest path whose height never increases as we move from source to destination. The problem of finding *descending paths in a polyhedral terrain* was first studied by de Berg and van Kreveld [4], who gave an $O(n\log n)$ time algorithm to decide if there is a descending path between two points. They stated as open the problem of finding a shortest descending path (SDP). In a subsequent paper, Roy, Das and Nandy [11] consider some special cases; in particular, they give an $O(n^2\log n)$ time algorithm to compute an SDP in a convex or concave terrain, and an $O(n\log n)$ time algorithm to compute an SDP through a sequence of parallel edges.

In this paper we first establish some properties of locally shortest descending paths, and show that they are much more complicated than geodesic paths. Then we turn to the case where the path must go through a given sequence of faces. We prove that the SDP is unique by proving that the length function is strictly convex. This approach was used by Mitchell and Papadimitriou [8] and Mitchell and Sharir [9] for respectively the weighted region problem and the shortest paths over parallel walls problem. Finally, we formulate the problem as a convex optimization problem and show that a $(1+\epsilon)$-approximation of the path can be computed in $O(n^{3.5}\log(\frac{1}{\epsilon}))$ time.

Our main long-term goal for the shortest descending path problem is to give an approximation algorithm that will find a shortest path even if the face sequence is not known. We will use an approach like that of Chen and Han [3] to decide which face sequence is best. We will use the uniqueness result for a given face sequence from the current paper, but we need significantly more. To be precise, our algorithm requires the ability to extend a locally shortest path. This is easy

* Corresponding author.
E-mail addresses: m6ahmed@uwaterloo.ca (M. Ahmed), alubiw@uwaterloo.ca (A. Lubiw).

0925-7721/$ – see front matter
doi:10.1016/j.comgeo.2007.10.011

for geodesic paths because they unfold to straight lines, but for descending paths we need a significantly more detailed analysis of the bend angles. Although the convex optimization technique offered in the current paper is not part of our general solution, we think it is interesting in its own right.

2. Preliminaries

A terrain is a 2-dimensional surface in 3-dimensional space with the property that every line parallel to the z-axis intersects it in a point [5]. We assume that the terrain is triangulated. For any point p in the terrain, $h(p)$ denotes the height of p, i.e., the z-coordinate of p.

A path P from s to t on the terrain is *descending* if the z-coordinate of a point p never increases while we move p along the path from s to t. A *shortest descending path* (*SDP*) from s to t is a descending path that is not longer than any other descending path from s to t in the terrain. A line segment of a descending path in face f is called a *free segment* if moving either of its endpoints by an arbitrarily small amount to a new position in f keeps the segment descending. Otherwise, the segment is called a *constrained segment*. All the points in a constrained segment are at the same height, though not all constant height segments are constrained. For example, a segment in a horizontal face is free, although all its points are at the same height. A path consisting solely of constrained segments is called a *constrained path*.

We will now define a *locally shortest descending path* (*LSDP*), which is analogous to a *geodesic path* (i.e., a *locally shortest path*) [7]. An LSDP between two nodes is a descending path that cannot be shortened by slight perturbation of the intermediate nodes. Note that perturbing a single node in a descending path may make the path infeasible (i.e., not descending), and hence, we allow more than one node to be perturbed simultaneously. For example, if we increase the height of a node p to H, all the points before p on the path must be moved to height at least H to keep the path descending. Also note that a constrained path is an LSDP.

Our uniqueness result relies on a certain function being strictly convex, which is defined as follows [2, Section 3.1.1]. A function $f : \mathbb{R}^n \to \mathbb{R}$ is called *convex* if the domain of f is a convex set, and for all x, y in that domain and all $c \in [0, 1]$,

$$f\big(cx + (1-c)y\big) \leqslant cf(x) + (1-c)f(y). \tag{1}$$

The function f is called *strictly convex* if strict inequality holds in Eq. (1) whenever $x \neq y$ and $0 < c < 1$.

For ease of discussion, we will use the term "edge" to denote a line segment of the terrain, and the term "segment" to denote a line segment of a path. Similarly, an endpoint of an edge is called a "vertex", while an endpoint of a segment is called a "node". We assume that all paths in our discussion are directed. In our figures, a dashed line, which may contain an arrow to indicate the direction of ascent, denotes an edge, and a heavy arrow denotes a constrained segment.

3. Characteristics of an LSDP

An LSDP and a geodesic path over a terrain are similar in many respects. The following lemmas establish two properties of an LSDP that make an LSDP analogous to a geodesic path [7].

Lemma 1. *Any subpath of an LSDP is an LSDP.*

Proof. If a subpath P_1 of an LSDP P is not locally shortest, there exists an LSDP P_1' corresponding to P_1 such that the length of P_1' is less than that of P_1. In that case, we can modify P by replacing the subpath P_1 with P_1' and get an LSDP of length less than that of P. This leads to a contradiction. □

Lemma 2. *An LSDP consists of straight line segments, and bends only at the edges of the terrain.*

Proof. Suppose that an LSDP P bends at point p that is an interior point of some face f of the terrain. Let P_1 be the connected subpath of $P \cap f$ that contains p, and p_1 and p_2 be respectively the starting and ending points of P_1. Since P is a descending path, $h(p_1) \geqslant h(p_2)$. Therefore, the line segment from p_1 to p_2 is a descending path, and its length is less

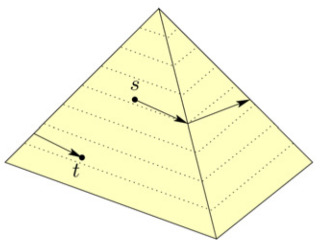

Fig. 1. An LSDP visiting a face twice.

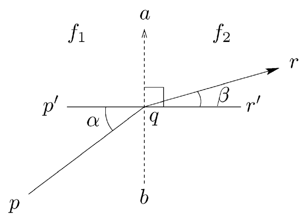

Fig. 2. Entering and exiting angles.

than that of P_1 because P_1 bends at p. That means, P_1 is not an LSDP. Then, by Lemma 1, P is not an LSDP, which is a contradiction. □

As in the case of a geodesic path [7], an LSDP may visit a single face more than once. For example, a string tightly wrapped around a pyramid as shown in Fig. 1 is an LSDP from s to t, and it visits a face twice. However, like a shortest path, an SDP visits a face at most once:

Lemma 3. *(See [11].) The intersection of an SDP P with a face of the terrain is either empty or a line segment.*

One important difference between an LSDP and a geodesic path is that unlike a geodesic path [7], two consecutive segments of an LSDP through an edge ab do not always become a straight line segment when the two faces of the terrain adjacent to ab are unfolded onto a plane. Before proving this claim, we define two angles at every edge intersected by an LSDP to quantify the amount of deflection at that edge. Let $P = (p, q, r)$ be a descending path from an interior point p in face f_1 to an interior point r in face f_2 adjacent to f_1 such that P crosses edge $ab = f_1 \cap f_2$ at q where $h(a) \geqslant h(b)$ (Fig. 2). Unfold the two faces onto a common plane, and then let $p'r'$ be a line segment perpendicular to ab at q such that $p' \in f_1$ and $r' \in f_2$. The angle $\angle pqp'$ is called the *entering angle of P at ab*, and is considered positive if and only if p and b are on the same side of $p'r'$. The angle $\angle rqr'$ is called the *exiting angle of P at ab*, and is considered positive if and only if r and a are on the same side of $p'r'$. In Fig. 2, α and β are respectively the entering angle and the exiting angle of P at ab. When $h(a) > h(b)$, we say that P deflects downward at q if $\alpha > \beta$, and that P deflects upward at q if $\alpha < \beta$. Note that if $h(a) = h(b)$, entering and leaving angles can be defined in two ways. Our discussion is valid for any of these definitions.

Lemma 4. *The descending path $P = (p, q, r)$ is an LSDP if and only if one of the following holds*:

(i) $\alpha = \beta$;
(ii) $\alpha > \beta$, *and qr is constrained*; *or*
(iii) $\alpha < \beta$, *and pq is constrained.*

Proof. ($\Leftarrow$)

(i) If $\alpha = \beta$, then for any point $q' \in ab$ such that $q' \neq q$, length of the path (p, q', r) is more than that of P. Therefore, P is an LSDP.
(ii) If $\alpha > \beta$, and qr is constrained, then for any point $q' \in ab$ such that $h(q') > h(q)$, q lies inside the smaller angle made by pq' and $q'r$ at q' (Fig. 3(a)). Clearly, the length of the path (p, q', r) is more than that of P. On the other hand, for any point $q' \in ab$ such that $h(q') < h(q)$, (p, q', r) is not a descending path because $h(r) = h(q) > h(q')$. Therefore, P is an LSDP.
(iii) If $\alpha < \beta$, and pq is constrained, the proof is similar to the one for Case (ii), except that the path (p, q', r) is longer than P when $h(q') < h(q)$, and is infeasible when $h(q') > h(q)$.

($\Rightarrow$) It is sufficient to show that if either $\alpha < \beta$ and pq is free, or $\alpha > \beta$ and qr is free, P is not an LSDP.

If $\alpha < \beta$ and pq is free, let q' be a point on ab slightly above q, but with $h(q') \leqslant h(p)$ (Fig. 3(b)). Such a point q' always exists because pq is a free segment and we can make q' arbitrarily close to q. We form a new path $P' = (p, q', r)$. Clearly, q' lies inside the smaller angle made by pq and qr at q, and hence, the length of P' is smaller than that of P. Because qr is descending, the segment $q'r$ is also descending. The segment pq' is descending since $h(q') \leqslant h(p)$. Therefore, P' is a descending path, and is shorter than P. So, P is not an LSDP. We can similarly show that P is not an LSDP if $\alpha > \beta$ and qr is free. □

Roy, Das and Nandy [11, Lemma 1] proved part of this lemma—in particular, for the case when the terrain is convex, in which case an SDP may bend upwards but not downwards. (Our preliminary version [1] misrepresented this aspect of their work, due to our misunderstanding.)

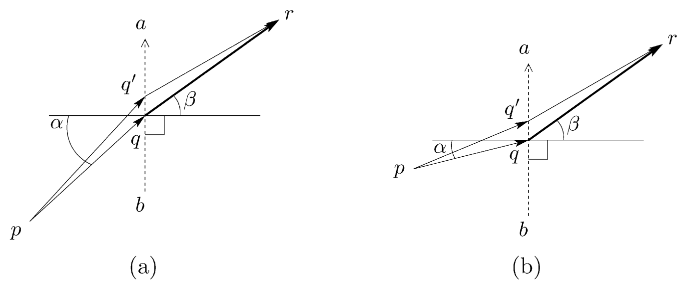

Fig. 3. Proof of Lemma 4.

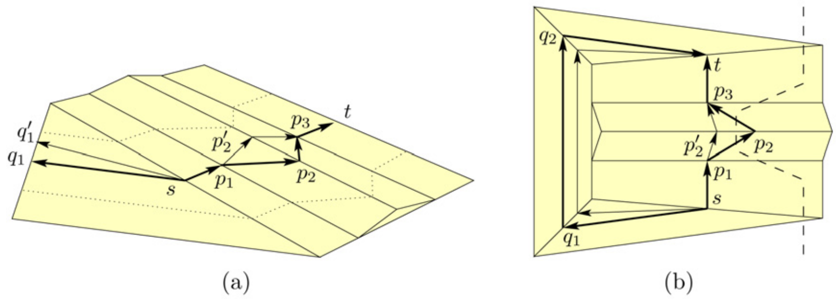

Fig. 4. Proof of Lemma 5.

In spite of all the similarities between an LSDP and a geodesic path, an SDP and a shortest path can be very different from each other. The following lemma proves this claim.

Lemma 5. *Let P_T and P'_T denote respectively an SDP and a shortest path from s to t in terrain T. There exists a terrain T for which one* (*or more*) *of the following holds*: (i) *the ratio of the lengths of P_T and P'_T is arbitrarily large even when the paths pass through the same sequence of faces*; (ii) *P_T and P'_T pass through two different face sequences; and* (iii) *there is no descending path through the face sequence crossed by P'_T.*

Proof. The idea is that a shortest path may climb over a ridge while the SDP may need to travel a long way around. We use a slightly more elaborate example to capture all the given situations. Consider a polyhedron that has a perspective view and a top view as in Fig. 4. The dotted lines in the perspective view are horizontal lines. Let s and t be two points of equal heights, and P be the constrained path (s, p_1, p_2, p_3, t), as shown in the figure.

Let T_1 be the terrain consisting of the faces crossed by P, and (s, p_1, p'_2, p_3, t) be the shortest path in T_1. Clearly, (s, p_1, p_2, p_3, t) is an SDP in T_1. Moreover, $p_1 p'_2 \perp p_2 p'_2$ (and $p_3 p'_2 \perp p_2 p'_2$ by symmetry). Now imagine rotating T_1 around the axis defined by the line through s and t. This rotation keeps the length of (s, p_1, p'_2, p_3, t) unchanged, but changes the length of (s, p_1, p_2, p_3, t). If we rotate T_1 until the face adjacent to s becomes almost horizontal, the length of (s, p_1, p_2, p_3, t) becomes arbitrarily large. This proves the first part of the lemma.

Let T_2 be the terrain consisting of the faces visible in the top view in Fig. 4. It is not hard to see that from s to t there are exactly two LSDPs (s, p_1, p_2, p_3, t) and (s, q_1, q_2, t), and exactly two geodesic paths (s, p_1, p'_2, p_3, t) and (s, q'_1, q'_2, t) in T_2. In the figure, the path (s, p_1, p'_2, p_3, t) is shorter than the path (s, q'_1, q'_2, t). So, (s, p_1, p'_2, p_3, t) is the shortest path from s to t. We can make the length of (s, p_1, p_2, p_3, t) greater than that of (s, q_1, q_2, t) by rotating the faces crossed by (s, p_1, p_2, p_3, t) as in the first part of the proof, while keeping the slopes of other faces unchanged. This makes (s, q_1, q_2, t) an SDP in T_2. Clearly, the SDP and the shortest path in T_2 pass through disjoint sets of faces, which proves the second part.

If we modify T_2 by removing the part of it to the right of the dashed lines in Fig. 4(b), it is no longer possible to construct any descending path through the face sequence crossed by the shortest path (s, p_1, p'_2, p_3, t). □

Roy, Das and Nandy [11, Section 5] suggest a heuristic for tracing an approximate SDP P from s to t, which follows a shortest path P' until reaching a point where P' is not descending, and follows constrained segments from that point until the traced path P either reaches t, or reunites with P' in which case P starts following P' again. They point out in their conclusion that the efficacy of this method depends on whether there is any relationship between the length of an SDP and the length of a descending path through the face sequence of a shortest path. Lemma 5 answers this negatively. They also state as open the problem of finding an SDP through a given face sequence (for which they resort to a heuristic). We address this problem in the following sections.

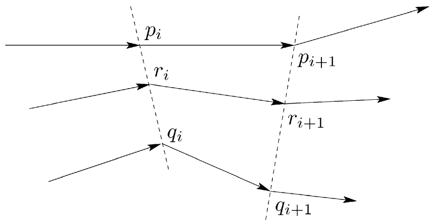

Fig. 5. Proving strict convexity of $L(x_1, x_2, \dots, x_k)$.

4. Uniqueness of an LSDP

In this section, we show that LSDPs are unique by formulating the problem of computing an LSDP as a convex optimization problem. The uniqueness of a geodesic path is evident from the fact that an unfolded geodesic path is a straight line segment. Since an unfolded LSDP is not a straight line segment, the uniqueness of an LSDP is not obvious. In our proof below, we use $\sigma_k = (f_0, f_1, \dots, f_k)$ to denote the given face sequence. We assume without loss of generality that source s is an interior point of f_0, destination t is an interior point of f_k, and for all $i \in [1,k]$, the edge between f_{i-1} and f_i is a_ib_i with $h(a_i) \geqslant h(b_i)$.

Let $F(x_1, x_2, \dots, x_k)$ denote the general path consisting of the line segments sp_1, p_1p_2, $p_2p_3, \dots, p_{k-1}p_k$ and p_kt in this order, where for all $i \in [1,k]$, p_i is any point on the line through a_ib_i, and x_i is a parameter to denote the position of p_i on line a_ib_i. For all $i \in [1,k]$ such that a_ib_i is non-horizontal, the height of p_i uniquely determines its position. So, in these cases, we use the height of p_i as parameter x_i. For each horizontal edge a_ib_i, we use as parameter x_i the signed distance of p_i from b_i. More precisely, $x_i = \overrightarrow{b_ip_i} \cdot \overrightarrow{b_ia_i}/|a_ib_i|$ in this case.

Note that parameter x_i can be defined in many ways, and our results below remain unchanged when x_i varies *linearly* with the position of p_i. However, choosing the height of p_i as our parameter x_i for non-horizontal edges makes the proof of Lemma 7 very simple and intuitive, because the constraint that the path is descending is simply expressed as $x_i \geqslant x_{i+1}$. This is why we have chosen to define x_i in this manner.

Let $L(x_1, x_2, \dots, x_k)$ denote the length of the path $F(x_1, x_2, \dots, x_k)$. In other words, $L(x_1, x_2, \dots, x_k) = \sum_{i=0}^{k} |p_ip_{i+1}|$, where $p_0 = s$ and $p_{k+1} = t$. The core of our uniqueness proof is the following property of $L(x_1, x_2, \dots, x_k)$:

Lemma 6. *$L(x_1, x_2, \dots, x_k)$ is a strictly convex function.*

A similar result was proved by Mitchell and Papadimitriou [8, Lemma 3.6] for the case of the weighted region problem. We include a proof for completeness and in order to be meticulous about faces where the function is convex but not strictly so.

Proof. Let $(u_1, u_2, \dots, u_k)$ and $(v_1, v_2, \dots, v_k)$ be two lists of constants such that $u_i \neq v_i$ for at least one $i \in [1,k]$. For any real constant $\kappa \in (0,1)$, let $w_i = \kappa u_i + (1-\kappa)v_i$ for all $i \in [1,k]$. We show that

$$L(w_1, w_2, \dots, w_k) < \kappa L(u_1, u_2, \dots, u_k) + (1-\kappa)L(v_1, v_2, \dots, v_k),$$

which proves the lemma by the definition of strict convexity.

Clearly, $F(u_1, u_2, \dots, u_k)$ and $F(v_1, v_2, \dots, v_k)$ denote two different paths. Let the ith segments of the three paths $F(u_1, u_2, \dots, u_k)$, $F(v_1, v_2, \dots, v_k)$ and $F(w_1, w_2, \dots, w_k)$ be p_ip_{i+1}, q_iq_{i+1} and r_ir_{i+1} respectively (Fig. 5). We can express the coordinates of r_i in terms of those of p_i and q_i as follows:

$$r_i = \kappa p_i + (1-\kappa)q_i.$$

Therefore,

$$\overrightarrow{r_ip_i} = p_i - r_i = p_i - \kappa p_i - (1-\kappa)q_i = (1-\kappa)(p_i - q_i) = (1-\kappa)\,\overrightarrow{q_ip_i},$$

and similarly, $\overrightarrow{r_iq_i} = \kappa\,\overrightarrow{p_iq_i}$. So, $\kappa\,\overrightarrow{r_ip_i} + (1-\kappa)\,\overrightarrow{r_iq_i} = 0$. For the same reason, $\kappa\,\overrightarrow{p_{i+1}r_{i+1}} + (1-\kappa)\,\overrightarrow{q_{i+1}r_{i+1}} = 0$. Using these two equations, we can express $\overrightarrow{r_ir_{i+1}}$ in terms of $\overrightarrow{p_ip_{i+1}}$ and $\overrightarrow{q_iq_{i+1}}$ as follows:

$$\begin{aligned}
\overrightarrow{r_ir_{i+1}} &= \kappa\overrightarrow{r_ir_{i+1}} + (1-\kappa)\,\overrightarrow{r_ir_{i+1}} \\
&= \kappa(\overrightarrow{r_ip_i} + \overrightarrow{p_ip_{i+1}} + \overrightarrow{p_{i+1}r_{i+1}}) + (1-\kappa)(\overrightarrow{r_iq_i} + \overrightarrow{q_iq_{i+1}} + \overrightarrow{q_{i+1}r_{i+1}}) \\
&= \kappa\,\overrightarrow{p_ip_{i+1}} + (1-\kappa)\,\overrightarrow{q_iq_{i+1}} + \kappa\overrightarrow{r_ip_i} + (1-\kappa)\overrightarrow{r_iq_i} + \kappa\,\overrightarrow{p_{i+1}r_{i+1}} + (1-\kappa)\overrightarrow{q_{i+1}r_{i+1}} \\
&= \kappa\,\overrightarrow{p_ip_{i+1}} + (1-\kappa)\overrightarrow{q_iq_{i+1}}.
\end{aligned}$$

Taking the lengths of $\overrightarrow{r_ir_{i+1}}$, $\overrightarrow{p_ip_{i+1}}$ and $\overrightarrow{q_iq_{i+1}}$, we get

$$|r_ir_{i+1}| \leqslant \kappa|p_ip_{i+1}| + (1-\kappa)|q_iq_{i+1}|. \tag{2}$$

In Eq. (2), the equality holds only when $p_i p_{i+1}$ and $q_i q_{i+1}$ are parallel to each other. Because $F(u_1, u_2, \ldots, u_k)$ and $F(v_1, v_2, \ldots, v_k)$ are different paths, and they both start at s and end at t, there are at least two $i \in [0, k]$ for which $p_i p_{i+1}$ and $q_i q_{i+1}$ are not parallel to each other. Considering this fact, and adding the lengths in Eq. (2) over all $i \in [0, k]$, we get:

$$\sum_{i=0}^{k} |r_i r_{i+1}| < \kappa \sum_{i=0}^{k} |p_i p_{i+1}| + (1-\kappa) \sum_{i=0}^{k} |q_i q_{i+1}|$$
$$\Rightarrow L(w_1, w_2, \ldots, w_k) < \kappa L(u_1, u_2, \ldots, u_k) + (1-\kappa) L(v_1, v_2, \ldots, v_k),$$

which completes the proof. □

We now determine the constraints on the variables x_i, $1 \leqslant i \leqslant k$, that ensure that $F(x_1, x_2, \ldots, x_k)$ is a descending path through σ_k. For all $i \in [1, k]$, the following constraints ensure that the intermediate nodes of $F(x_1, x_2, \ldots, x_k)$ are not outside the corresponding edges:

$$h(b_i) \leqslant x_i \leqslant h(a_i), \quad \text{when } h(a_i) \neq h(b_i), \tag{3}$$

$$\text{and} \quad 0 \leqslant x_i \leqslant |a_i b_i|, \quad \text{when } h(a_i) = h(b_i). \tag{4}$$

The constraints that ensure that $F(x_1, x_2, \ldots, x_k)$ is a descending path are: $h(p_i) \geqslant h(p_{i+1})$ for all $i \in [0, k]$. For each i such that $a_i b_i$ is horizontal, $h(p_i)$ is a constant of value $H_i = h(a_i)$. Moreover, $h(p_0)$ and $h(p_k)$ are also constants of values $H_0 = h(s)$ and $H_{k+1} = h(t)$ respectively. For all other $i \in [1, k]$, $h(p_i) = x_i$. Therefore, for every $i \in [1, k]$, the corresponding height constraint expressed in terms of variables x_i's has one of the following forms:

$$x_i \geqslant x_{i+1}, \tag{5}$$

$$H_i \geqslant x_{i+1}, \tag{6}$$

$$x_i \geqslant H_{i+1}, \tag{7}$$

$$H_i \geqslant H_{i+1}. \tag{8}$$

Note that the constraint in Eq. (8) is either always satisfied, or never satisfied. Clearly, the constraint is redundant in the former case, and there is no descending path through σ_k from s to t in the latter case.

Lemma 7. *There is at most one LSDP through σ_k from s to t.*

Proof. Any LSDP through σ_k from s to t is an instance of $F(x_1, x_2, \ldots, x_k)$ because an LSDP bends only on the edges of the terrain (Lemma 2). Moreover, the length of an LSDP through σ_k from s to t corresponds to a local minimum of the length of $F(x_1, x_2, \ldots, x_k)$ subject to the constraints in Eqs. (3) to (8), i.e., a local minimum of $L(x_1, x_2, \ldots, x_k)$ subject to those constraints. Now observe that the constraints in Eqs. (3) to (8) are linear, and therefore, the domain defined by them is convex. Since $L(x_1, x_2, \ldots, x_k)$ is a strictly convex function (Lemma 6), it has at most one local minimum in the convex domain defined by those constraints [2, Section 4.2.1]. Therefore, there is at most one LSDP through σ_k from s to t. □

5. Algorithm

It follows from Lemma 7 that we can determine an SDP through σ_k from s to t by solving the following convex optimization problem:

$$\begin{array}{ll} \text{minimize} & L(x_1, x_2, \ldots, x_k) = \sum_{i=0}^{k} |p_i p_{i+1}| \\ \text{subject to} & \text{the constraints in Eqs. (3) to (8).} \end{array}$$

We can use any method of solving a convex optimization problem to compute the SDP. Using the method used by Polishchuk and Mitchell [10], we get the following running-time:

Theorem 8. *Determining a $(1+\epsilon)$-approximate SDP through a sequence of k faces from s to t takes $O(k^{3.5} \log(\frac{1}{\epsilon}))$ time.*

Proof. We first convert the above convex optimization problem into the following equivalent problem on variables $x_1, x_2, \ldots, x_k$, $t_0, t_1, t_2, \ldots, t_k$:

$$\begin{aligned} \text{minimize} \quad & \sum_{i=0}^{k} t_i \\ \text{subject to} \quad & |p_i p_{i+1}| \leqslant t_i, \quad \text{for } i \in [0,k], \\ & \text{and} \\ & \text{the constraints in Eqs. (3) to (8).} \end{aligned} \tag{9}$$

It is easy to show that the coordinates of p_i vary linearly with x_i for all $i \in [1,k]$. As a result, the constraint in Eq. (9) can be written in the form $|A_i x_i + B_i x_{i+1} + C_i| \leqslant t_i$ for some scalar constants A_i, B_i and C_i for all $i \in [0,k]$. This makes the above optimization problem a Second-order Cone Program, for which finding an $(1+\epsilon)$-approximate solution takes $O(k^{3.5}\log(\frac{1}{\epsilon}))$ time [6]. $\square$

References

[1] Mustaq Ahmed, Anna Lubiw, Shortest descending paths through given faces, in: Proceedings of the 18th Canadian Conference on Computational Geometry, August 2006, pp. 35–38.
[2] Stephen Boyd, Lieven Vandenberghe, Convex Optimization, Cambridge University Press, Cambridge, UK, 2004.
[3] Jindong Chen, Yijie Han, Shortest paths on a polyhedron. I. Computing shortest paths, Internat. J. Comput. Geom. Appl. 6 (2) (1996) 127–144.
[4] Mark de Berg, Marc J. van Kreveld, Trekking in the Alps without freezing or getting tired, Algorithmica 18 (3) (1997) 306–323.
[5] Mark de Berg, Marc J. van Kreveld, Mark Overmars, Otfried Schwarzkopf, Computational Geometry: Algorithms and Applications, 2nd edition, Springer-Verlag, Berlin, 2000.
[6] Miguel S. Lobo, Lieven Vandenberghe, Stephen Boyd, Hervé Lebret, Applications of second-order cone programming, Linear Algebra Appl. 284 (1998) 193–228.
[7] Joseph S.B. Mitchell, David M. Mount, Christos H. Papadimitriou, The discrete geodesic problem, SIAM J. Comput. 16 (4) (1987) 647–668.
[8] Joseph S.B. Mitchell, Christos H. Papadimitriou, The weighted region problem: finding shortest paths through a weighted planar subdivision, J. ACM 38 (1) (1991) 18–73.
[9] Joseph S.B. Mitchell, Micha Sharir, New results on shortest paths in three dimensions, in: Proceedings of the 20th Annual Symposium on Computational Geometry, ACM Press, New York, 2004, pp. 124–133.
[10] Valentin Polishchuk, Joseph S.B. Mitchell, Touring convex bodies—a conic programming solution, in: Proceedings of the 17th Canadian Conference on Computational Geometry, 2005, pp. 290–293.
[11] Sasanka Roy, Sandip Das, Subhas C. Nandy, Shortest monotone descent path problem in polyhedral terrain, in: Proceedings of the 22nd Annual Symposium on Theoretical Aspects of Computer Science, in: Lecture Notes in Computer Science, vol. 3404, Springer-Verlag, Berlin, 2005, pp. 281–292.