Optimistic and Pessimistic Shortest Paths on Uncertain Terrains

Yury Kholondyrev* William Evans*

Abstract

We consider the problem of finding shortest paths on the surface of uncertain terrains. In this paper, a terrain is a triangulated 2D surface in 3D such that every vertical line intersects the surface at most once. Terrains of this type are used to represent, for example, a piece of the earth's surface, and are typically inexact. We model their uncertainty by allowing the terrain vertices to have a range of possible heights (Z-coordinates), while fixing the triangulation (i.e. the adjacency of vertices). This defines a set of feasible (certain) terrains. We are looking for a "shortest" path between two vertices s and t (defined by its projection to the XY-plane) but the length of any particular path may depend on the actual feasible terrain. We consider both pessimistic (a path's length is its maximum length over all feasible terrains) and optimistic (a path's length is its minimum feasible length) scenarios.

If we are allowed to walk on the faces of the terrain, the problem is NP-hard in both pessimistic and optimistic scenarios [5, 4]. In this paper, we prove that if we can walk only on terrain edges, the pessimistic problem is still NP-hard (and we give a fully-polynomial time approximation scheme for it) while the optimistic problem is solvable in polynomial time.

1 Introduction

We model an uncertain terrain as an undirected graph $G = (V, E)$ with an *uncertainty interval* for each vertex $v \in V$ that is specified by its two *extreme points* v^- and v^+ in 3D. The X and Y coordinates of v^- and v^+ are identical but their Z-coordinates may differ (v^- has smaller Z-coordinate). The graph G embedded in the XY-plane using the XY-coordinates of its vertices forms a triangulation, possibly with holes. An uncertain terrain defines a set of feasible (certain) terrains; those whose projections to the XY-plane match the projection of the uncertain terrain and whose vertices lie within the corresponding Z-ranges.

Within this model, a number of shortest path problems can be formulated. The problem can be *unrestricted*, when you are allowed to traverse the faces of the terrain, or *edge-restricted*, when you are only allowed to travel along the edges. The problem can be solved with the *optimistic* assumption, when we assume that the actual terrain is the one that minimizes the length of the path that we choose, or the *pessimistic* assumption, when we assume that the terrain will have the worst shape possible. We mainly consider the edge-restricted versions of the optimistic and the pessimistic shortest path problems.

*Department of Computer Science, University of British Columbia, yurykh@gmail.com will@cs.ubc.ca

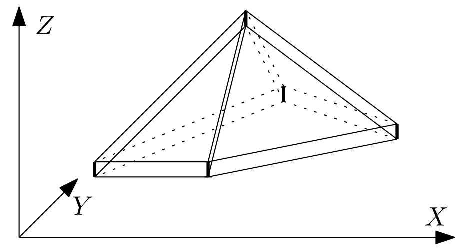

Figure 1: An uncertain terrain.

The model only considers uncertainty of Z-coordinates (and not X, Y), first, because altitude error in elevation models derived from satellite images can be significantly greater than planimetric (horizontal) error especially in mountainous areas [6] (though the reverse may be true for flatter areas and other surveying techniques [1, 8]) and, second, because planimetric error can be seen as producing altitude error [3].

Chris Gray [5, 4] considered the unrestricted version of the shortest path on an uncertain terrain [5]. He proved that finding either optimistic or pessimistic shortest paths on an uncertain terrain is NP-hard using techniques similar to those Canny and Reif [2] used to prove the NP-hardness of finding the Euclidian shortest path among polyhedral obstacles in 3D.

The problem of finding shortest paths on (certain) terrains is well studied and there are a number of algorithms that solve it in polynomial time. Mitchell *et al.* [7] showed how to solve the more general problem of finding the shortest path on an arbitrary 2D polyhedral surface in 3D in $O(n^2 \log n)$ time, where n is the number of edges in the polyhedra.

2 Pessimistic Edge-Restricted Shortest Path

In the pessimistic case, we want the *guaranteed* shortest path, that is, the path, among all possible edge-restricted paths from s to t, that has the shortest length when measured on the terrain that maximizes the path's length (i.e. the path's worst-case terrain). Such a path has the minimum (over paths from s to t), maximum (over feasible terrains) length. The worst-case terrain for any path places each vertex of the path at one of the vertex's extreme points [5].

Let us define the *measure* of a path from s to t as the pair (a, b) where a (resp. b) is the worst-case length of the path from s to t^+ (resp. t^-). (The worst-case path length from s to t is $\max\{a, b\}$.) We say that a measure (a, b) *dominates* a measure (c, d) if $a \leq c$ and $b < d$, or $a < c$ and $b \leq d$. We also say that the path with measure (a, b) dominates the path with measure (c, d) and that measure (a, b) is *better* than measure (c, d). It is possible that any non-dominated path may be the prefix of a guaranteed shortest path.

Consider the terrain with five vertices, called the *gadget* shown in Figure 2. Vertex z has such a large Z-

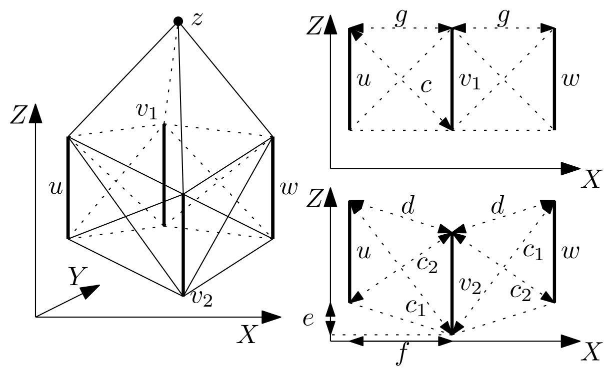

Figure 2: The terrain at the left, in which vertex z is very high, allows two potential shortest paths from u to w, which, to permit annotation, are shown separately at the right.

coordinate that no path containing z could be a shortest path, thus a path with measure (a, b) ending at u can be extended in two ways, via v_1 or v_2, to obtain two paths ending at w with measures: $(\max\{a + 2c, a + 2g, b + c + g\}, \max\{b + 2c, b + 2g, a + c + g\})$ and $(\max\{a+2c_1, a+2d, b+c_2+d, b+d+c_1\}, \max\{b+2c_2, b+2d, a + d + c_2, a + c_1 + d\})$. If the difference between a and b is negligible compared to $c - g$ and $c_2 - d$ then these measures are $(a + 2c, b + 2c)$ and $(a + 2c_1, b + 2c_2)$. We construct the gadget so that $c_1 + c_2 = 2c$. Thus if we set $\phi = c_1 + c_2 = 2c$ and $\alpha = c_1 - c_2$, the gadget takes measure (a, b) at u and produces two non-dominated measures $(a + \phi, b + \phi)$ and $(a + \phi + \alpha, b + \phi - \alpha)$ at vertex w. Given α, the precise coordinates of the vertices for a gadget with parameter α are: $u = (0, 0, [0, f])$, $v_1 = (f, y_1, [0, f])$, $v_2 = (f, 0, [-e, f - e])$, and $w = (2f, 0, [0, f])$, where, in order to satisfy $c_1 - c_2 = \alpha$, $e = \alpha\sqrt{(8f^2 - \alpha^2)/(4f^2 - \alpha^2)}/2$, and, in order to satisfy $2c = c_1 + c_2$, $y_1 = \sqrt{((c_1 + c_2)/2)^2 - 2f^2}$. Note that y_1 is a positive real number since $c_1 + c_2$ is greater than $2\sqrt{2}f$ (by triangle inequality). We must also set f to be large enough so that the worst-case terrain alternates top and bottom extreme points for either path. It can be shown that, if for measure (a, b) both $|a - b|$ and α are at most some constant β, setting $f = 10\beta$ will suffice.

Theorem 1 *The pessimistic edge-constrained shortest path problem on uncertain terrains is NP-hard.*

Proof. Given a set $S = \{\alpha_1, \alpha_2, \ldots, \alpha_N\}$, of positive integers and a target sum, T, construct a shortest path problem instance as follows (see Figure 3): Set the parameter f of all gadgets to $f = 20\sum_{i=1}^{N} \alpha_i$. Construct a chain of N gadgets from a vertex s to a vertex w, such that the parameter α for the i-th gadget equals α_i (let ϕ_i be the resulting ϕ). Note that our construction guarantees that $|a - b|$ for every path measure (a, b) will never become greater than $2\sum_{i=1}^{N} \alpha_i$. Create a vertex t, put it at distance $f + T$ from w^- and at distance $f - T$ from w^+, set $t^+ = t^-$, and connect t and w with an edge.

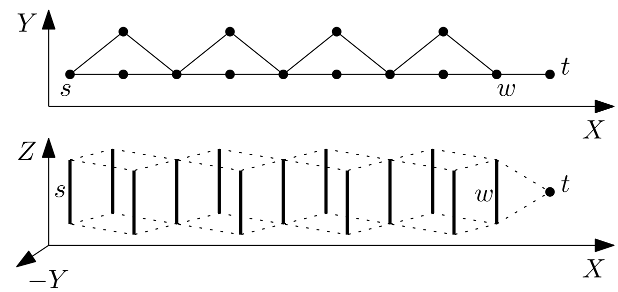

Figure 3: Reduction from SUBSET-SUM.

Let $\Phi = \sum_{i=1}^{N} \phi_i$ be the sum of the ϕ's for all gadgets. If the pessimistic shortest path length from s to t equals $\Phi + f$, there is a way to make up the target sum using the numbers in S, otherwise (you can only get a bigger answer) it is impossible.

Because we start with the measure $(0, 0)$ at s and only apply transformations of the form $(a, b) \to (a + \phi, b + \phi)$ and $(a, b) \to (a + \phi + \alpha, b + \phi - \alpha)$, where α's are the numbers from S, every potential shortest path measure at the last vertex, w, of the chain has the form $(\Phi + \tau, \Phi - \tau)$ where τ is the sum of a subset of the numbers in S. Thus, a shortest path of length $\Phi + f$ implies that there is a way to get to w via a path with a measure $(\Phi + T, \Phi - T)$, and that T is a subset sum of S.

As stated, the reduction produces terrain vertices whose coordinates may require an infinite number of bits. However, if these coordinates are accurate to within $\pm 1/(8(2N + 1))$, all path measures will be accurate to within $\pm 1/2$ in each component (since every path has at most $2N+1$ edges). That accuracy suffices for the reduction. □

2.1 Approximation Algorithm

Using path measures, we can generalize the Bellman-Ford algorithm to solve our problem. We associate a set of non-dominated path measures with every vertex. Initially, all sets are empty except s's set which contains measure $(0,0)$. An edge relaxation step using the edge (u,v) takes every measure (a,b) from the set at vertex u and adds (c,d) to the set of measures at v, where $c = \max\{a + \|u^+ - v^+\|, b + \|u^- - v^+\|\}$ and $d = \max\{a + \|u^+ - v^-\|, b + \|u^- - v^-\|\}$, as long as (c,d) is not dominated by an existing measure at v. If we relax all the edges $|V|-1$ times (the number of edges in the pessimistic shortest path can be at most $|V|-1$), the sets of non-dominated path measures at every vertex stop changing. At that point, we can compute the pessimistic shortest path distance to a vertex by looking at all the path measures in its set and choosing the measure that guarantees the shortest distance in the worst case. Unfortunately, the number of non-dominated path measures associated with a vertex can be exponential.

To construct a $(1+\epsilon)$-approximation algorithm, we store only $2\lceil 1/\epsilon \rceil$ measures at any vertex. Note that any path measure (a,b) at a vertex v will have $|a-b| < \|v^+ - v^-\|$. Measure (a,b) goes into bucket $\left\lceil \frac{a-b}{\epsilon\|v^+ - v^-\|} \right\rceil$ at vertex v if it is the measure with smallest sum in that bucket.

Theorem 2 *The approximation scheme that stores $2\lceil 1/\epsilon \rceil$ measures at every vertex (as described above) will find a path with guaranteed length at most $(1+\epsilon)$ times the optimal path length and run in $O(|E||V|/\epsilon)$ time.*

Proof. (Omitted) □

3 Optimistic Edge-Restricted Shortest Path

Having realized that the pessimistic version is NP-hard, we will look at the optimistic version of the problem. Now the problem is to find the shortest path from s to t where path length is measured on the best-case terrain (i.e. the terrain that minimizes the path's length). Within this formulation, the problem seems to become even more difficult; the best-case terrain does not have to force the path to traverse only the extreme points of a vertex's Z-range. A *traversal* of a a simple path $\phi = v_1, v_2, \ldots, v_k$ (a sequence of uncertain terrain vertices) is a sequence of points (in 3D) $p_1, p_2, \ldots, p_k$ such that, for all i, p_i lies in v_i's uncertainty interval. If some optimistic shortest path traversal does not go through an extreme point, it should have the same slope where it comes to the vertex and where it leaves it. Otherwise it would be possible to find a feasible terrain on which the path is shorter.

A *pseudo-straight traversal* is a traversal in which the line segments $\overline{p_{i-1}p_i}$ and $\overline{p_i p_{i+1}}$ have the same slope $i = 2, \ldots, k-1$. We call a pseudo-straight traversal that starts at an extreme point a *pseudo-straight ray*. If it also ends at an extreme point, we call it a *pseudo-straight path*.

Any shortest path (assuming both source and destination points have a zero Z-range) will be composed of one or more pseudo-straight paths connected at the extreme points of some vertices. (See Figure 4.) With this

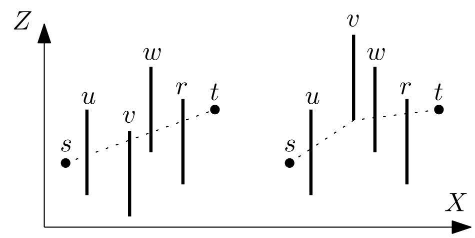

Figure 4: Any optimal path must be piecewise pseudo-straight.

in mind, the most intuitive attempt to solve the problem is to find all the pseudo-straight paths and treat them as the only edges that can be used to traverse the uncertain terrain. All that we need to do (provided we have precomputed all the pseudo-straight paths) is to run any shortest path algorithm on the resulting graph (with $2n$ vertices). Unfortunately, not only can the number of pseudo-straight paths be exponential, but also the problem of deciding if there is a pseudo-straight path between any two given points is NP-hard.

Theorem 3 *The problem of deciding if there is a pseudo-straight path between two certain (zero Z-range) vertices on an uncertain terrain is NP-complete.*

Proof. (sketch) If we construct an uncertain terrain in such a way that there is an exponential number of pseudo-straight paths from a (zero Z-range) vertex s to a (zero Z-range) vertex g then the slopes of those paths will be determined by the paths' lengths only. We can then add a vertex t (also zero Z-range) and connect it to the vertex g in such a way that only the pseudo-straight path with a certain slope (length) can make it through while remaining pseudo-straight. If we use the idea of the earlier NP-hardness proof of creating a chain such that every pseudo-straight path from s to g corresponds to a sub-set sum of a given set of numbers, we can decide the SUBSET-SUM problem by reducing it to the pseudo-straight path existence problem. □

Even though determining the existence of a pseudo-straight path is NP-hard, we can still find the shortest optimistic path between s and t in polynomial time. The key is to notice that some pseudo-straight paths cannot be pieces of a shortest optimistic path. A pseudo-straight traversal $p_1, p_2 \ldots p_k$ of a path $\phi = v_1, v_2 \ldots v_k$ is *dominated* if there exists a path ϕ' from v_1 to v_k with a shorter traversal from p_1 to p_k. Note that the traversal from p_1 to p_k via ϕ' might not be realized by a single pseudo-straight traversal.

The algorithm builds a graph G_{ps} whose vertices are the extreme points of the terrain vertices, with an edge of length ℓ between two extreme points if there is a non-dominated pseudo-straight path of length ℓ between them. Assuming for the moment that s and t are certain vertices (i.e. $s^+ = s^-$ and $t^+ = t^-$), the shortest path from s to t in G_{ps} is the optimistic shortest path we desire. Even if G_{ps} contains some edges corresponding to dominated pseudo-straight paths, the algorithm will still work. It will take polynomial time (in the size of the original terrain) if the number of edges we add is polynomial (and we can find them in polynomial time). If s and/or t are uncertain then the shortest path from s to t may start and/or end with a horizontal (0-slope) pseudo-straight ray. The algorithm adds edges corresponding to these rays to G_{ps} as well.

The *upper (resp. lower) cone* $C(\phi)$ for a path $\phi = v_1, v_2, \ldots, v_k$ is the set of all non-negative (resp. non-positive) slope pseudo-straight rays through ϕ starting at a specified extreme point e of v_1. Each cone has an origin e and two bounding extreme points u and l that limit the maximum and minimum slopes of rays within the cone. (Either u, for lower cones, or l, for upper cones, may be $\emptyset$ if the cone's bounding ray has slope 0.) Let (e, u, l, v_k) be the *label* of the cone. We call a cone *dominated* if every pseudo-straight ray belonging to the cone is dominated.

Lemma 4 *If two upper (or two lower) cones are identically labeled then either one of them is dominated or they are equivalent (their pseudo-straight rays are the same length).*

Proof. (sketch) For two upper cones C_1 and C_2 both with label (e, u, l, f), let $d_i(x)$ be the pseudo-straight ray distance in cone C_i to x. If $(d_1(u), d_1(l), d_1(f)) < (d_2(u), d_2(l), d_2(f))$ then C_2 is dominated, where "<" is lexicographic order. If the distance triples are equal then the cones are equivalent. □

Theorem 5 *We can find the edge-restricted optimistic shortest path on an uncertain terrain $G = (V, E)$ in $O(|V|^4|E|)$ time.*

Proof. For every extreme point e, we calculate the non-dominated upper and lower cones from e, and add the corresponding edges to G_{ps}. This takes $O(|V|^3|E|)$ time per extreme point (assuming $|V| < |E|$). The final application of Dijkstra's algorithm takes $O(|V|^2)$ time because G_{ps} contains $O(|V|)$ vertices (the extreme points). That adds up to $O(|V|^4|E|)$ total running time, dominated by the time required to construct G_{ps}. □

4 Conclusions and Future Work

We have shown that the pessimistic edge-constrained problem is NP-hard and have given an approximation algorithm for it. It is important to notice that the construction used for the NP-hardness proof is not too artificial. It seems that similar (hard) problem instances can arise even if the input is an almost flat surface that was built by placing vertices (with the same Z-ranges) at grid points and then moving the vertices slightly in random directions. In other words, it seems that even the problem of finding the shortest path in a uniformly triangulated plane with noise is hard.

The polynomial time algorithm for the optimistic edge-constrained problem that we presented was not designed for performance, but rather to show that the problem can be solved in polynomial time. Almost certainly, a faster algorithm exists.

References

[1] N. Al-Rousan and G. Petrie. System calibration, geometric accuracy testing and validation of DEM & orthoimage data extracted from SPOT stereopairs using commercially available image processing systems. *International Archives of Photogrammetry and Remote Sensing*, 34(4):8–15, 1998.

[2] John Canny and John Reif. New lower bound techniques for robot motion planning problems. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*, pages 49–60, 1987.

[3] P. Fisher and N. Tate. Causes and consequences of error in digital elevation models. *Progress in Physical Geography*, 30(4):467–489, 2006.

[4] C. Gray and W. Evans. Optimistic shortest paths on uncertain terrains. In *Proceedings of the 16th Canadian Conference on Computational Geometry (CCCG)*, pages 68–71, August 2004.

[5] Chris Gray. Shortest paths on uncertain terrains. Master's thesis, Department of Computer Science, University of British Columbia, August 2004.

[6] Andreas Kääb. Monitoring high-mountain terrain deformation from repeated air- and spaceborne optical data: examples using digital aerial imagery and ASTER data. *ISPRS Journal of Photogrammetry & Remote Sensing*, 57:39–52, 2002.

[7] Joseph S. B. Mitchell, David M. Mount, and Christos H. Papadimitriou. The discrete geodesic problem. *SIAM Journal on Computing*, 16(4):647–668, 1987.

[8] T. Schenk. Modeling and analyzing systematic errors in airborne laser scanners. Technical Report Photogrammetry No. 19, Department of Civil and Environmental Engineering and Geodetic Science, OSU, 2000.