On Realistic Terrains

Esther Moet

Marc van Kreveld

A. Frank van der Stappen

Department of Information and Computing Sciences, Utrecht University

Technical Report UU-CS-2006-029

www.cs.uu.nl

ISSN: 0924-3275

On Realistic Terrains

Esther Moet, Marc van Kreveld[*] and A. Frank van der Stappen[†]

June 2006

Abstract

We study worst-case complexities of visibility and distance structures on terrains under realistic assumptions on edge length ratios and the angles of the triangles, and a more general low-density assumption. We show that the visibility map of a point for a realistic terrain with n triangles has complexity $\Theta(n\sqrt{n})$. We also prove that the shortest path between two points p and q on a realistic terrain passes through $\Theta(\sqrt{n})$ triangles, and that the bisector of p and q has complexity $O(n\sqrt{n})$. We use these results to show that the shortest path map for any point on a realistic terrain has complexity $\Theta(n\sqrt{n})$, and that the Voronoi diagram for any set of m points on a realistic terrain has complexity $\Omega(n + m\sqrt{n})$ and $O((n + m)\sqrt{n})$. Our results immediately imply more efficient algorithms for computing the various structures on realistic terrains.

Keywords: Realistic input models, polyhedral terrains, visibility maps, shortest paths, Voronoi diagrams.

1 Introduction

One of the main objectives of computational geometry is to uncover the computational complexity of geometric problems. It provides a theory that explains how efficiently geometric problems can be solved that arise in applications. However, in many cases a discrepancy exists between the provable worst-case computational complexity of an algorithm and the actual running time behavior of that algorithm on inputs that arise in applications. This has led to the study of *fatness* and *realistic input models*. By making assumptions on the input, certain hypothetical worst-case scenarios can no longer occur, and a more efficient solution to a problem can be shown for all inputs that satisfy the assumptions. Either an existing algorithm can be shown to be more efficient for realistic inputs, or a special algorithm is designed that assumes that it will run on realistic inputs only. Our algorithmic results are of the first type.

Among the first applications of realistic input models in computational geometry, Alt et al. [1] consider motion planning for a rectangular robot. The efficiency depends on the aspect ratio of this rectangle. Matoušek et al. [13] show that if all triangles of a set of n triangles have their angles bounded away from zero (at least α, for some constant $\alpha > 0$), then the union of these triangles has $O(n)$ holes rather than $O(n^2)$ for the general case, and the boundary complexity of the union is $O(n \log\log n)$. Such triangles are called *fat*. Since then, various definitions of fatness [2, 9, 18, 21] have been proposed. Other realistic models—such as low density [20], unclutteredness [22], and simple-cover complexity [16]—consider the spatial distribution of objects or their features. An overview of reduced combinatorial complexities and improved algorithmic efficiencies for inputs satisfying these models is given in [8, 22] along with a model hierarchy. Most papers concern union complexities of fat objects (e.g. [2, 9]) or motion planning in realistic environments (e.g. [7, 20]), but many other results exist as well. For example, de Berg studies linear-size binary space partitions for uncluttered scenes [5] and vertical ray shooting in fat objects [6], Mitchell et al. [16] consider

[*]The second author was partially supported by BRICKS through the project GADGET.
[†]The third author was partially supported by the Dutch BSIK/BRICKS project.

ray shooting in scenes with small simple cover complexity, Overmars and van der Stappen [18] and Vleugels and Schwarzkopf [19] consider point location and range searching in sets of fat objects and low density scenes respectively, and Erickson [10] studies Boolean combinations and Minkowski sums of polygons and polyhedra whose features satisfy certain realistic locality assumptions.

Realistic assumptions have not yet been studied for polyhedral terrains. However, several geometric structures on terrains have complexities much higher than typical in applications. For example, the visibility map of a polyhedral terrain of n triangles has worst-case complexity $\Theta(n^2)$ [14], a shortest path passes through $\Theta(n)$ triangles in the worst case, and even a bisector of two points on a terrain can have quadratic size. Hence, the discrepancy between theoretical complexity bounds and typical complexity bounds exists on terrains as well. In this paper, we analyze this discrepancy by studying realistic assumptions on terrains.

For visibility maps, we give three assumptions whose combination provides an $\Theta(n\sqrt{n})$ bound on the worst-case complexity visibility map of a terrain for both views from infinity and perspective views. In fact, our results apply to the transparent visibility map, that is, our bounds hold even when occlusion is not taken into account. Dropping any of the three assumptions immediately makes an $\Omega(n^2)$ lower bound construction possible. It is interesting to note that the assumptions all refer to the xy-projection of the terrain.

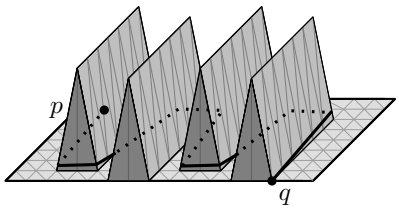

Figure 1: The shortest path between p and q crosses $\Omega(n)$ triangles.

Next, we consider distance structures on terrains. Using only the three assumptions for visibility, we can still have shortest paths that have linear complexity; the terrain in Figure 1 can be constructed such that in the projection, all its vertices lie on a regular grid. Therefore, we introduce a fourth assumption that relates to the steepness of the terrain, and show that any shortest path between two points passes through only $\Theta(\sqrt{n})$ triangles. For a bisector of two points, we show that the same set of four assumptions gives an $O(n\sqrt{n})$ complexity bound rather than quadratic. We give an $\Omega(n)$ size lower bound. The shortest path map for a source point s is the subdivision of the terrain into regions where the combinatorial structure of shortest paths from s is the same. In general, it has worst-case complexity $\Theta(n^2)$, but we show that under our assumptions it is $\Theta(n\sqrt{n})$. Finally, we study Voronoi diagrams for m point sites on terrains. In general, if m is $O(n)$, these also have quadratic worst-case complexity; this is even the case for only two sites. Our assumptions allow us to prove an upper bound of $O((n+m)\sqrt{n})$, and we give a lower bound of $\Omega(n + m\sqrt{n})$.

Our results immediately imply faster computation of visibility maps on realistic terrains. The output-sensitive construction of the visibility map of a terrain by Katz et al. [12] implies that for realistic terrains, it can be computed in $O(n\sqrt{n}\log n)$ time.

Shortest paths on general terrains can be computed in $O(n\log^2 n)$ time [11]. This result improved earlier algorithms [4, 15]. The description of Kapoor [11] does not provide sufficient detail to analyze whether a slightly faster algorithm can be obtained for realistic terrains. For shortest path maps and Voronoi diagrams, the algorithm of Mitchell et al. [15] leads to $O(n\sqrt{n}\log n)$ time bounds for their construction on realistic terrains.

This paper is structured as follows. In Section 2, we list three assumptions on realistic terrains and argue that no subset of these three is sufficient to prove a subquadratic upper bound on the complexity of the visibility map. We also prove several auxiliary results using these assumptions. In Section 3, we obtain a $\Theta(n\sqrt{n})$ bound on the size of the visibility map for both views from infinity and perspective views. In Section 4, we add a fourth assumption to study shortest paths,

bisectors, and the shortest path map; we give a $\Theta(n\sqrt{n})$ size bound for the latter structure in Section 4.2. In Section 4.3, we show that the Voronoi diagram of m sites has worst-case complexity $\Omega(n + m\sqrt{n})$ and $O((n + m)\sqrt{n})$ on realistic terrains. In Section 5, we give the conclusions and open problems.

2 Input model

Let $\mathcal{T}$ be a polyhedral terrain, comprising a set T of n triangles, a set E of n_e edges, and a set V of n_v vertices, where n_e and n_v are $O(n)$. Throughout this section, with the exception of Section 2.3, we refer to the projection of $\mathcal{T}$. We give two sets of assumptions on the properties of the projection of $\mathcal{T}$ that allow us to prove a subquadratic upper bound on the complexity of the visibility map. Below, we give the set of assumptions that we used in an earlier version of this paper [17]; we prove the results from [17] in this paper with a weaker set of assumptions, which we introduce further on in this section.

1. the minimum angle of any triangle in T in the projection to the xy-plane is at least α,
2. the boundary of this projection is a rectangle with side lengths 1 and c, and
3. the longest xy-projection over all edges in E is at most d times as long as the shortest one.

The values α, c, and d in the above assumptions are positive constants. A projected triangle that satisfies the first assumption is *fat* [13]. We used fatness of the triangles in the projection; alternatively, we could assume that the angles of the triangles in 3D are all larger than some $\alpha > 0$. However, as will become clear from Figure 5 in the following section, this assumption does not help to obtain a subquadratic complexity bound for the visibility map.

Since fatness probably is an overly restrictive assumption—it is unlikely that a constant number of non-fat triangles will induce a quadratic-size visibility map—we now consider a weaker assumption on the input, which we use in the remainder of this paper to prove subquadratic upper bounds on the combinatorial complexity of various structures. Moreover, we show that any fat triangulation satisfies this weaker assumption; therefore, all our results hold for triangulations with fat triangles as well.

Because our input is a planar triangulation instead of a general set of objects, we have adapted the original definition of low-density in [22] to the following tailored version.

Definition 1 *Let k be a positive integer, let T be a planar triangulation with n triangles, and let E be the set of $O(n)$ edges of T. We call T a* k-low-density *triangulation if for any square S, the number of edges $e \in E$ with $|e| \geq s$ that intersect S is at most k, where s is the side length of S. The smallest k for which T is a k-low-density triangulation is the* density *of T.*

Lemma 2 *Let T be a triangulation in the plane in which each angle of each triangle is at least $\alpha > 0$. Then T is a k-low-density triangulation for some $k = O(\frac{1}{\alpha})$.*

Proof. We have to show an upper bound on the number of edges of T of length at least ℓ that intersect an arbitrary square B_ℓ with side length ℓ. Each edge that intersects B_ℓ belongs to some triangle t of T, where all three angles in t are at least α. Let B^* be the box that is B_ℓ scaled by a factor 3 with respect to its center; see Figure 2.

Let t be a triangle that has an edge e of length at least ℓ which intersects B_ℓ. By the fatness of t, the area of $B^* \cap t$ is at least $\frac{1}{2}\ell^2 \sin\alpha$. This implies that there are at most

$$\frac{18\ell^2}{\ell^2 \sin\alpha}$$

such triangles, and thus the density of T is at most

$$\frac{54}{\sin\alpha} = O\left(\frac{1}{\alpha}\right).$$

This result holds for all $0 < \alpha < \frac{\pi}{2}$. □

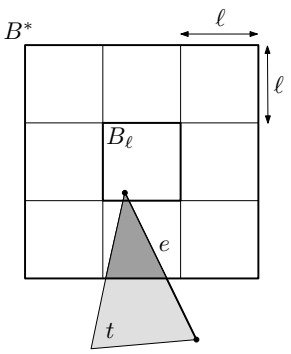

Figure 2: Illustration for the proof of Lemma 2.

From now on, we assume that $\mathcal{T}$ satisfies the following three properties:

1a. the projection of $\mathcal{T}$ onto the xy-plane is a k-low-density triangulation,
2. the boundary of this projection is a rectangle R with side lengths 1 and c, and
3. the longest xy-projection over all edges in E is at most d times as long as the shortest one.

The values k, c, and d in the above assumptions are positive constants.

Definition 3 *We call a polyhedral terrain $\mathcal{T}$ a* realistic terrain *if $\mathcal{T}$ satisfies Assumptions 1a, 2, and 3.*

Although we have not yet verified whether most terrains that arise in applications in fact fit our model, we do employ the term *realistic terrain* to associate our assumptions with the widely accepted term *realistic input model*. Our first assumption (low-density) is very common to assume and has been studied extensively. Experiments show that the value k in scenes that arise in practice varies between approximately fifteen and twenty-five [8].

The second assumption (constant aspect ratio of the projection) guarantees that the terrain cannot be unnaturally long and/or skinny, which is quite reasonable to assume. In fact, our results also hold for terrains that have a polygonal projected boundary which is contained in a rectangle of constant area.

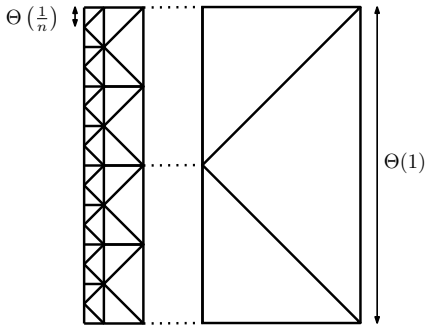

Figure 3: A low-density triangulation does not necessarily induce a bounded edge length ratio.

Assumption 3 seems quite strong, but it turns out to be crucial to obtain a non-trivial lower bound on the minimum edge length, an auxiliary result that we use extensively in this paper. The triangulation in Figure 3 is a 7-low-density triangulation, in which all triangles are α-fat for $\alpha = \frac{\pi}{4}$, its boundary side lengths are 1 and 2, and it consists of at most $3n$ triangles. The edge length ratio is not bounded by a constant, which makes it possible that edges of the triangulation have length $O(\frac{1}{n})$ or $\Omega(1)$.

2.1 No subset of assumptions is sufficient

In this section, we show that if we drop any of the three assumptions above, it is possible to construct a visibility map of quadratic size. Thus, no subset of our assumptions is sufficient to give a subquadratic upper bound.

A vertex of the visibility map of a point p directly corresponds to a line through p that is a common tangent of two terrain edges. The terrain that is displayed in Figure 4 produces an $\Omega(n^2)$ size visibility map; the $\Omega(n)$ hills in the front induce a linear size subdivision on each of the $\Omega(n)$ triangles in the back.

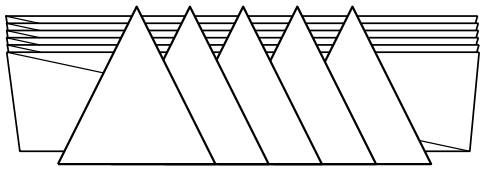

Figure 4: A visibility map of a terrain, for a viewpoint at infinity, which has complexity $\Omega(n^2)$.

First, if we allow the projected triangulation to have non-constant density, we can place two sets of $\Omega(n)$ triangles in a rectangular region of constant area in such a way that *every* triangle in the first set produces a constant number of visibility map vertices with *every* triangle of the second set, which results in a visibility map of complexity $\Omega(n^2)$. A schematic illustration of such a construction is shown in Figure 5(a); the Θ's refer to distances in the xy-projection. Note that this constructuin satisfies Assumptions 2 and 3, that all triangles in this construction are not α-fat for any constant $\alpha > 0$, and that the same construction can be achieved with a terrain in which the triangles are fat in 3D, but not in the projection.

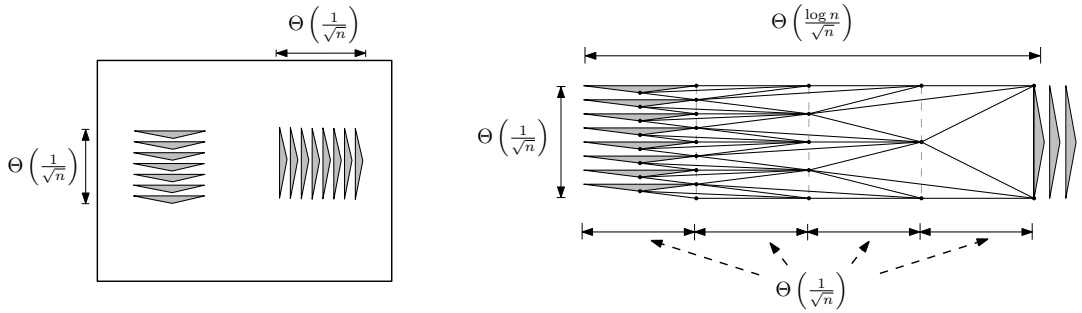

Figure 5: (a) A terrain that produces a quadratic-size visibility map if the viewpoint is located infinitely far to the left. The terrain satisfies Assumptions 2 and 3. (b) The subconstruction used to obtain a valid terrain from (a).

To obtain a valid terrain from this construction, we need to triangulate the white parts in Figure 5(a) such that the resulting triangles do not violate Assumption 3. The subconstruction that achieves this is shown in Figure 5(b); we need a rectangular region with side lengths $O\left(\frac{\log n}{\sqrt{n}}\right)$ and $O\left(\frac{1}{\sqrt{n}}\right)$ to adjoin the two subconstructions shown in Figure 5(a). This construction uses $O(n)$ triangles. Similar constructions in which the width of the triangles grows, can be used to triangulate the other white parts of Figure 5(a) and thus connect the two subconstructions with the boundary.

We can also construct a visibility map of complexity $\Omega(n^2)$ in a terrain that obeys the first and third assumption, but not the second one. Since the boundary aspect ratio is no longer a constant, we can again place two sets of $\Omega(n)$ triangles that mutually interact in the visibility map; see Figure 6. Because the area that we are able to use is not bounded by a constant, it is straightforward to triangulate the rest of the terrain with $O(n)$ triangles, while not violating Assumptions 1a and 3.

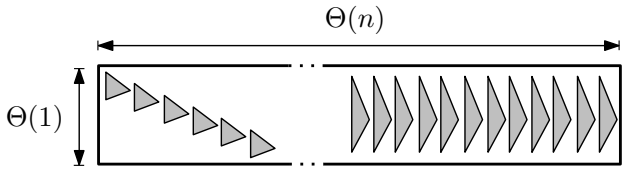

Figure 6: A terrain that produces a quadratic-size visibility map if the viewpoint is located infinitely far to the left. The terrain satisfies Assumptions 1a and 3.

Finally, if we allow an unbounded edge length ratio, it is again possible to construct a terrain with a visibility map of complexity $\Omega(n^2)$. A schematic display of this construction is shown in Figure 7(a).

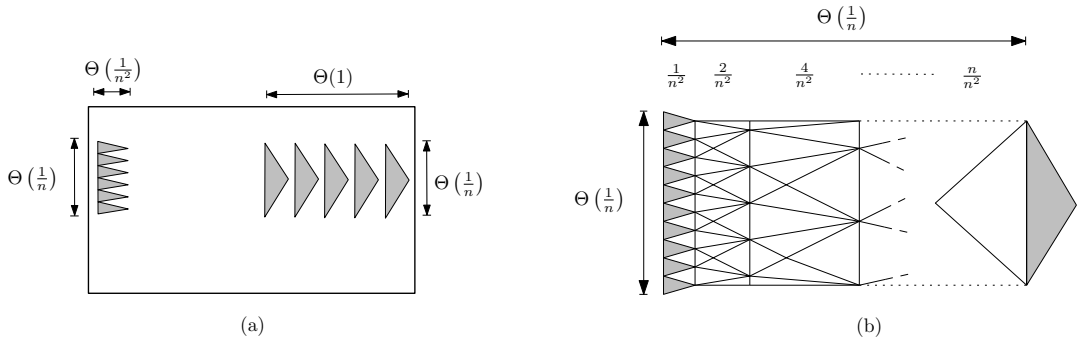

Figure 7: (a) A terrain that produces a quadratic-size visibility map if the viewpoint is located infinitely far to the left. The terrain satisfies Assumptions 1a and 2. (b) The subconstruction used to obtain a valid terrain from (a).

Again, we need to triangulate the white parts of Figure 7(a) to complete the terrain, but now we can only use $O(n)$ triangles, while not violating Assumption 1a. In order to adjoin the two sets of triangles shown in Figure 7(a), we need to grow the lengths of the shortest edges by a factor n. We use a distance of $O\left(\frac{1}{n}\right)$ and $O(n)$ triangles to achieve this. This subconstruction is shown in Figure 7(b). The other white parts in Figure 7(a) can also be triangulated with $O(n)$ triangles, such that we do not violate Assumption 1a, by letting the edge lengths grow in a similar way.

Finally, we observe that in Figure 7, the ratio of edge lengths is $O(n)$, and hence the assumption of Erickson [10] of polynomially bounded edge lengths is not strong enough in our case.

2.2 Auxiliary results

We now give several auxiliary lemmas on properties of the projection of $\mathcal{T}$, which we use to obtain our main results in Sections 3 and 4. Throughout this section, R is the rectangular region that contains the projection of $\mathcal{T}$, and x is the length of the shortest edge of the projection of $\mathcal{T}$. We let B_ℓ denote an axis-aligned square (box) with side lengths ℓ, $B_\ell(p)$ denotes the square B_ℓ translated such that the point p is its center, and $\mathcal{E}(B_\ell)$ denotes the set of edges of length at least ℓ that intersect the square B_ℓ.

Lemma 4 *Each vertex v of V has degree at most k.*

PROOF. By Assumption 1a, $\mathcal{E}(B_x(v))$ is a set of at most k edges from E. The lemma follows from the fact that this set includes all edges that are incident to v. □

Lemma 5 *All edges in E have length $\Theta\left(\frac{1}{\sqrt{n}}\right)$.*

PROOF. First, we observe that all edges in E have length at most dx. Any triangle in T is at most as large as the equilateral triangle with edge lengths dx. The maximum area of any triangle is therefore

$$\frac{\sqrt{3}(dx)^2}{4}.$$

The sum of the areas of the triangles in T is at most

$$c \leq \frac{n\sqrt{3}(dx)^2}{4},$$

and therefore

$$x \geq \frac{2\sqrt{c}}{\sqrt[4]{3}\ d\sqrt{n}} = \Omega\left(\frac{1}{\sqrt{n}}\right).$$

For the upper bound, it is sufficient to show that x is $O\left(\frac{1}{\sqrt{n}}\right)$, since the longest edge has length dx. To this purpose, we cover R by a set of copies of the square B_x. By Assumption 2, we need $\lceil \frac{c}{x^2} \rceil$ such squares to cover R completely. Because of the low-density assumption, every square B_x intersects at most k edges of E. Because x is the minimum edge length, every single edge that intersects a given B_x is indeed included in $\mathcal{E}(B_x)$. Since every edge intersects a constant number of squares, and since there are $\Omega(n)$ edges in E, we have that $k\lceil \frac{c}{x^2} \rceil$ is $\Omega(n)$, which implies that x is $O\left(\frac{1}{\sqrt{n}}\right)$. □

Lemma 6 *Let $\mathcal{R}$ be a rectangle that intersects the projection of $\mathcal{T}$, of which both side lengths are $\Omega\left(\frac{1}{\sqrt{n}}\right)$, and that has total area A. Then $\mathcal{R}$ intersects $O(knA)$ triangles of T.*

PROOF. We cover $\mathcal{R}$ by a set of (possibly rotated) copies of the square B_x; see Figure 8. Since x is $\Omega\left(\frac{1}{\sqrt{n}}\right)$, this is possible with $O(An)$ such squares. Because x is the minimum edge length, every single edge that intersects a given square B_x is indeed included in $\mathcal{E}(B_x)$. By Assumption 1a, we have that $|\mathcal{E}(B_x)| \leq k$ for each B_x. Therefore, the sum of all values $|\mathcal{E}(B_x)|$, and thus the number of edges (and triangles) that intersect $\mathcal{R}$, is $O(knA)$. □

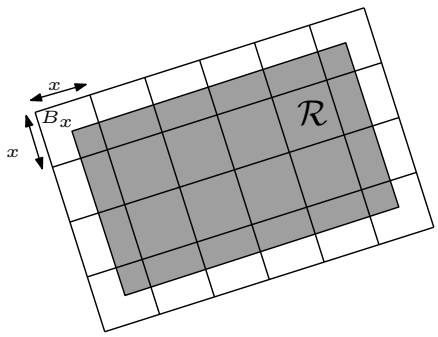

Figure 8: Covering a rectangle $\mathcal{R}$ (shaded) with squares of side length x.

Lemma 7 *Let s be a straight line segment s that intersects the projection of $\mathcal{T}$. Then s intersects $O(k\sqrt{n})$ triangles of T.*

PROOF. By Assumption 2, the length of s within R is $O(1)$. Let L be a slab of width $\Theta\left(\frac{1}{\sqrt{n}}\right)$, enclosed by two lines parallel to s, and which contains s strictly in its interior. All edges that intersect s must intersect L as well. By Assumption 2, the area of L is $O\left(\frac{1}{\sqrt{n}}\right)$, and so, by the previous lemma, L intersects $O(k\sqrt{n})$ triangles. □

2.3 Dihedral angles

So far, we have only considered assumptions that refer to the projected triangulation induced by a terrain. All complexity bounds for the visibility map of a realistic terrain in Section 3 are achieved with only the first three assumptions, which only deal with properties of the terrain in the projection. In particular, we do not need a bound on the dihedral angles to prove the upper bounds, and the lower bound constructions are all possible with bounded dihedral angle.

The terrain in Figure 1 does satisfy Assumptions 1a, 2, and 3, but a shortest path between two points on the terrain still passes through $\Theta(n)$ triangles in the worst case. In Section 4, we discuss distance structures on terrains, and to bound their complexity, we introduce an additional assumption:

4. The dihedral angle of the supporting plane of any triangle in T with the xy-plane is at most β, where $\beta < \frac{\pi}{2}$ is some constant.

Assumption 4 implies that the maximum slope of a line segment on any triangle of $\mathcal{T}$ is $\tan\beta = O(1)$. In this section and from Section 4 onwards, we call a terrain *a realistic terrain* if it satisfies all four assumptions. It is easily verified that no subset of our four assumptions is sufficient to obtain a sublinear bound on the complexity of a shortest path; for each subset of three assumptions, we can create an $\Omega(n)$ lower bound.

Lemma 8 *The shortest path $P(p,q)$ between two points p and q on a realistic terrain $\mathcal{T}$ has $O(1)$ length.*

PROOF. Let $P'(p,q)$ be the path from p to q over $\mathcal{T}$ whose projection is the straight line segment $\overline{pq}$. Obviously, $P(p,q)$ is at most as long as $P'(p,q)$. By Assumption 4, the length of $P'(p,q)$ is at most $\tan\beta \cdot |\overline{pq}|$. By the second assumption, $\overline{pq}$ has length at most $\sqrt{1+c^2}$, and thus the length of $P(p,q)$ is bounded from above by $\tan\beta \cdot \sqrt{1+c^2}$, which is $O(1)$. □

Lemma 9 *The shortest path $P(p,q)$ between two points p and q on a realistic terrain $\mathcal{T}$ crosses $O(k\sqrt{n})$ triangles.*

PROOF. First we observe that $P(p,q)$ intersects any triangle at most once, and within each triangle, $P(p,q)$ is a straight line segment. We cover the projection of $P(p,q)$ completely by copies of the square B_x. The first square that we place is $B_x(p)$. The next square is $B_x(p')$, where p' is the point on the projected path that is closest to p and which has not yet been covered. We keep placing squares in this way until the entire projection of $P(p,q)$ is covered. See Figure 9 for an example of such a covered path.

Every square covers at least the length $\frac{1}{2}x$ of the projection of $P(p,q)$. By the previous lemma, and because x is $\Omega\left(\frac{1}{\sqrt{n}}\right)$, we can do this with $O(\sqrt{n})$ such squares. By Assumption 1a, each square B_x intersects at most k edges of E, and thus $P(p,q)$ crosses $O(k\sqrt{n})$ triangles. □

It is easy to see that on any realistic terrain $\mathcal{T}$, two points exist whose shortest path passes through $\Omega(\sqrt{n})$ triangles, and thus we conclude this section with the following lemma.

Lemma 10 *Let $\mathcal{T}$ be a realistic terrain with n triangles. The shortest path over $\mathcal{T}$ between two points p and q on $\mathcal{T}$ passes through $\Theta(\sqrt{n})$ triangles in the worst case.*

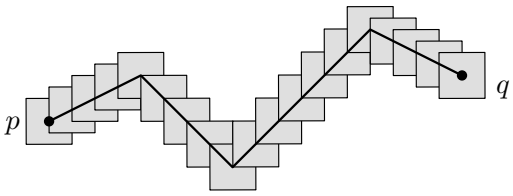

Figure 9: Covering the path $P(p, q)$ by squares with side length x.

3 The visibility map

3.1 Upper bounds

Let p be the viewpoint of the visibility map $\mathrm{VM}(p, \mathcal{T})$ that we consider. For a given edge e, let $\#I_e$ be the number of terrain edges with which e can *interact*, i.e., with which it can create features of $\mathrm{VM}(p, \mathcal{T})$, and which lie closer to the viewpoint than e. We bound the combinatorial complexity of $\mathrm{VM}(p, \mathcal{T})$ by giving an upper bound on $\#I_e$. Recall that a vertex of the visibility map of p directly corresponds to a line through p that is a common tangent to a pair of terrain edges. Since any two edges create $\Theta(1)$ vertices of $\mathrm{VM}(p, \mathcal{T})$ in the worst case, we get the following expression.

$$\text{Complexity of } \mathrm{VM}(p, \mathcal{T}) = O\left(\sum_{e \in E} \#I_e\right) \tag{1}$$

For an edge e, we define the *influence region* of e as the locus of the edges that (i) lie in between p and e, and (ii) with which e interacts in the (transparent) visibility map of p. This definition implies that we charge a visibility map vertex to the *last* edge that the corresponding tangent touches. We denote the influence region of an edge e by $\mathcal{R}_e$, and we give an upper bound on its area. Using this bound, we can bound $\#I_e$ from above by the number of edges in E whose projection intersects $\mathcal{R}_e$. We distinguish two cases based on the location of the viewpoint: (i) p is located infinitely far away from $\mathcal{T}$ (parallel projection), and (ii) p does not lie at infinity (perspective projection).

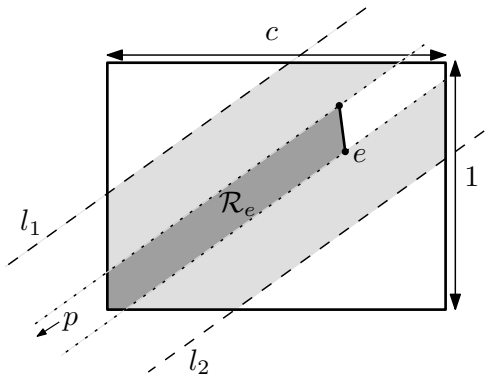

Figure 10: The influence region for an edge e and a viewpoint p at infinity.

If the viewpoint p is located at infinity, then for any edge e, in the projection to the xy-plane, all triangles that intersect the influence region $\mathcal{R}_e$ lie either entirely within the influence region, or partly in the influence region and partly outside this region, but in a buffer region that is enclosed between two parallel lines. Figure 10 illustrates this; please note that the proportions in this figure are distorted for illustration purposes. Let l_1 and l_2 be the two lines that pass through

the two endpoints of e and are parallel to the viewing direction induced by the viewpoint p. Now we translate l_1 and l_2 in the perpendicular direction, away from each other, over a distance of the length of the longest edge. Let L_e be the slab that is bounded by l_1 and l_2; it is easy to verify that L_e completely contains all the triangles that intersect $\mathcal{R}_e$. By Lemma 5, for any edge e from E, $L_e \cap R$ is completely contained in a rectangle of area $O\left(\frac{1}{\sqrt{n}}\right)$, and thus, by Lemma 6, the number of triangles (and edges) in L_e is $O(\sqrt{n})$. Using Equation (1), we obtain the following lemma.

Lemma 11 *Let $\mathcal{T}$ be a realistic terrain with n triangles, and let p be a viewpoint at infinity. Then* $\mathrm{VM}(p, \mathcal{T})$ *has complexity $O(n\sqrt{n})$ in the worst case.*

The influence region for an edge e and a perspective view is very similar to the influence region of e and some viewpoint at infinity; see Figure 11. In fact, for some viewing direction, $\mathcal{R}_e$ is contained in the influence region from infinity, and hence $\mathcal{R}_e$ has area $O\left(\frac{1}{\sqrt{n}}\right)$. The two lightly shaded buffer regions around $\mathcal{R}_e$ are both enclosed by a pair of parallel lines at a distance of the longest (projected) edge of E of each other. By Lemma 5, this distance is $\Theta\left(\frac{1}{\sqrt{n}}\right)$, which together with Assumption 2 gives that both buffers have area $O\left(\frac{1}{\sqrt{n}}\right)$ as well. Thus, by Lemma 6, the number of triangles (and edges) in that intersect $\mathcal{R}_e$ in this case is also $O(\sqrt{n})$, which, together with Equation (1), yields the lemma below.

Lemma 12 *Let $\mathcal{T}$ be a realistic terrain with n triangles, and let p be a point located on or above $\mathcal{T}$. Then* $\mathrm{VM}(p, \mathcal{T})$ *has complexity $O(n\sqrt{n})$ in the worst case.*

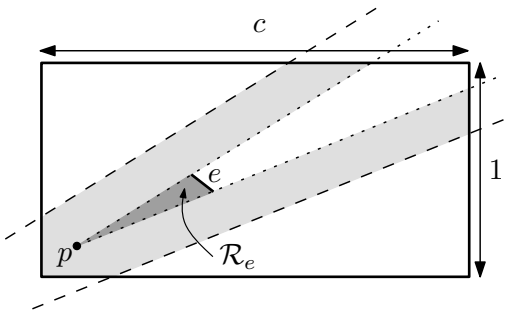

Figure 11: The influence region for an edge e and a viewpoint p close to or above the terrain. (Proportions are distorted for illustration purposes.)

Note that in the figure above, we displayed the projection of the viewpoint inside the projection of the terrain, but this is not necessarily the case. If we move the viewpoint further and further away from the terrain, the perspective influence region converges into the parallel influence region.

3.2 Lower bounds

In this section, we show that the upper bounds from the previous section are tight.

We start by describing a construction of a visibility map for parallel views that has $\Omega(n\sqrt{n})$ vertices. Under the assumptions of Section 2, we can place $\Theta(\sqrt{n})$ triangles in a rectangle of constant length and of width $\Theta(\frac{1}{\sqrt{n}})$. This can be done in such a way that these $\Theta(\sqrt{n})$ triangles together form a smaller version of the construction in Figure 4; in the front, we place $\Omega(\sqrt{n})$ triangles, each of which interacts with all $\Omega(\sqrt{n})$ triangles in the back. For such a rectangle, the visibility map for a point at infinity has complexity $\Omega(n)$.

Since the projection of $\mathcal{T}$ is a rectangle with side lengths 1 and c, we can replicate this construction $\Omega(\sqrt{n})$ times; see Figure 12 for a schematic illustration. If the viewpoint is located infinitely far down, this terrain has a visibility map of complexity $\Omega(n\sqrt{n})$.

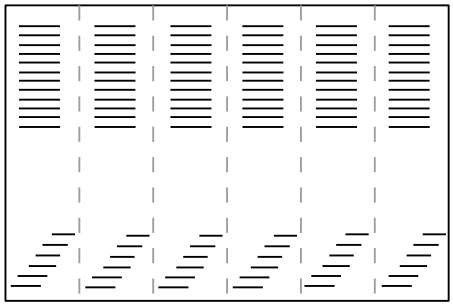

Figure 12: The visibility map of a viewpoint at infinity can have $\Omega(n\sqrt{n})$ vertices.

Lemma 13 *The visibility map of a realistic terrain for a viewpoint at infinity can have $\Omega(n\sqrt{n})$ vertices.*

For perspective views, we show a very similar lower bound. By a transformation of the construction in Figure 4, we can create a subconstruction with $\Omega(\sqrt{n})$ triangles that are located in a truncated wedge of area $\Theta(\frac{1}{\sqrt{n}})$, instead of in a rectangle of the same area, as we did in Figure 12. Note that this is possible with such a transformation that the edge lengths and angles of the transformed triangles are all within a constant factor of the original values. If we place the viewpoint at the apex of the (non-truncated) wedge, then this subconstruction produces a visibility map with $\Omega(n)$ vertices. We can replicate this construction $\Omega(\sqrt{n})$ times in a rectangle of $\Theta(1)$ area; Figure 13 displays the total construction schematically.

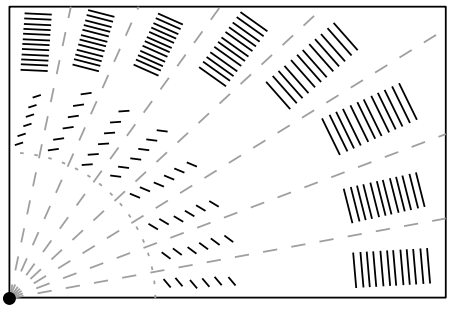

Figure 13: The perspective visibility map can have $\Omega(n\sqrt{n})$ vertices.

Lemma 14 *The visibility map of a realistic terrain for a perspective view can have $\Omega(n\sqrt{n})$ vertices.*

We conclude with a theorem that summarizes the results of this section.

Theorem 15 *Let $\mathcal{T}$ be a realistic terrain with n triangles. Then both the parallel and the perspective visibility map have complexity $\Theta(n\sqrt{n})$ in the worst case.*

4 The shortest path map and the Voronoi diagram

Recall from Section 2 that from now on, we have an additional, fourth assumption, on the dihedral angle of the triangles of $\mathcal{T}$. In order to study the complexity of shortest path maps and Voronoi diagrams on realistic terrains, we first need some definitions on shortest paths and bisectors, and results on their complexity.

Shortest paths on terrains were studied extensively in [4, 11, 15]. A geodesic between two points is a path that is locally shortest. An actual shortest path is a geodesic that has minimum length. A geodesic that crosses an edge between two triangles is such that if the triangles are made coplanar by rotation about their common edge, then the geodesic becomes a straight line segment on the union of the two triangles. The rotation is referred to as *unfolding*. Following [15], a geodesic path "is an alternating sequence of vertices and (possibly empty) edge sequences such that the unfolded image of the path along any edge sequence is a straight line segment, and the angle passing through a vertex is greater than or equal to π." A shortest path has the additional property that no edge can occur more than once. The *root* of a path from p to p' is the last vertex we encounter on a shortest path from p to p'; if there is no such vertex, then p is the root. Recall from Lemma 10 in Section 2.3 that a shortest path between two points on a realistic terrain has complexity $\Theta(\sqrt{n})$.

4.1 Bisectors

The bisector $B(p, q)$ of a point p and a point q on $\mathcal{T}$ is the set of points on $\mathcal{T}$ with equal distance to p and q. It is a simple curve (open or closed) that consists of line segments and hyperbolic arcs [15]. Imagine that we walk along $B(p, q)$, while we 'sweep' the shortest path of the current point to p across the terrain. At certain points on $B(p, q)$, we encounter a discontinuity in the sweep; the edge sequence of the shortest path from either p or q to the current point on $B(p, q)$ changes. We call such points on $B(p, q)$ *breakpoints*. Note that breakpoints on the bisector $B(p, q)$ are points that have two shortest paths from p or two shortest paths from q.

The worst-case complexity of a bisector is $\Theta(n^2)$ on general terrains. In the terrain in Figure 14(a), a set of $\Theta(n)$ skinny peaks form two fans, which are directed away from both p and q. Along the northwest–southeast diagonal, p and q are alternatingly closest. The bisector $B(p, q)$ has $\Omega(n)$ breakpoints, and between each two consecutive breakpoints, it crosses $\Omega(n)$ edges of the terrain.

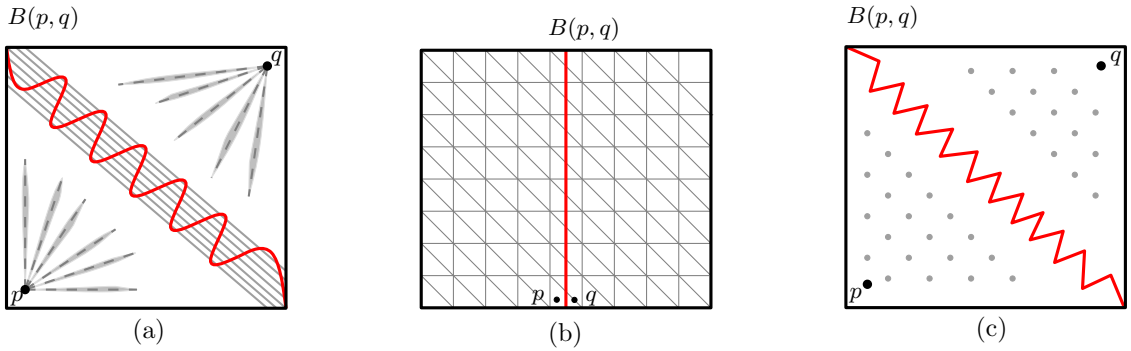

Figure 14: (a) In general terrains, the bisector of two points has complexity $\Omega(n^2)$. (b) In a realistic terrain, a bisector can cross $\Omega(\sqrt{n})$ edges between two breakpoints. (c) A bisector on a realistic terrain can have $\Omega(n)$ breakpoints.

In a realistic terrain, we cannot create such a construction. However, it is possible that a bisector in a realistic terrain crosses $\Omega(\sqrt{n})$ edges before it encounters a breakpoint; see Figure 14(b). the terrain is horizontal, i.e., it is parallel to the xy-plane, and its vertices lie on a regular $\sqrt{n} \times \sqrt{n}$ grid. The bisector of p and q, which does not have any breakpoints, is a straight line segment which intersects $\Omega(\sqrt{n})$ triangles.

Figure 14(c) shows that a bisector on a realistic terrain has $\Omega(n)$ breakpoints. In fact, it has $\Theta(n)$ breakpoints, since it has at most one breakpoint for every vertex.

Lemma 16 *A bisector of two points on a realistic terrain intersects $\Theta(\sqrt{n})$ triangles between two breakpoints in the worst case.*

PROOF. Figure 14 shows a lower bound construction; it remains to prove the upper bound. Let b_1 and b_2 be two consecutive breakpoints on $B(p,q)$, and let $P(p,b_1)$ and $P(p,b_2)$ be the shortest paths from p to these two breakpoints, respectively. If there is more than one shortest path to b_1 and/or b_2, choose them so that the area enclosed by $P(p,b_1)$, $P(p,b_2)$, and $B(p,q)$ is smallest. Let r_p be the last common vertex of the shortest paths $P(p,b_1)$ and $P(p,b_2)$, and let r_q be defined analogously; see Figure 15.

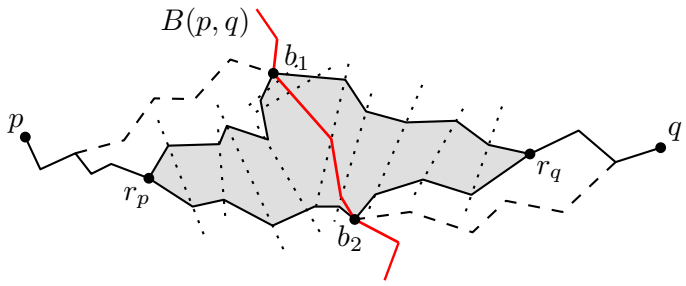

Figure 15: Illustration of the proof of Lemma 16.

By Lemma 10, all four paths intersect $O(\sqrt{n})$ triangles/edges. Because b_1 and b_2 are consecutive breakpoints on $B(p,q)$, the region on $\mathcal{T}$ that is enclosed by $P(r_p,b_1)$, $P(r_p,b_2)$, $P(r_q,b_1)$, and $P(r_q,b_2)$ does not contain any vertices of $\mathcal{T}$; in Figure 15 this region is shaded. Consequently, the edges that intersect $B(p,q)$ between b_1 and b_2 have to intersect, or end on, at least two of the four paths. As we observed above, there are only $O(\sqrt{n})$ such edges, which completes the proof. □

We conclude this section with the following, summarizing lemma.

Lemma 17 *For two points p and q on a realistic terrain $\mathcal{T}$, the bisector $B(p,q)$ has complexity $\Omega(n)$ and $O(n\sqrt{n})$.*

4.2 Shortest path maps

The *shortest path map* of a source point s is the subdivision of $\mathcal{T}$ into cells such that the vertex and edge sequence of the shortest path to any point in that cell from the source is the same. For simplicity of argument, we assume that for all vertices of $\mathcal{T}$ the shortest path from s is unique. Obviously, other points on $\mathcal{T}$ can have more than one shortest path.

The analysis of [15] shows that in general terrains, the shortest path map has complexity $\Theta(n)$ on a single triangle, and $\Theta(n^2)$ overall. Furthermore, from the description in [15], it is known that in the shortest path map, every terrain edge is partitioned into $O(n)$ intervals such that the vertex and edge sequence of the shortest paths from the source s to all points within an interval is the same. Any triangle is partitioned into $O(n)$ cells that arise from the additively weighted Voronoi diagram of some set of points; these points are obtained by planar unfoldings of terrain vertices along different edge sequences. The weights are given by the shortest path distances from the source s to these roots.

It is probably true that in a realistic terrain the shortest path map can still have linear complexity on a single triangle. Therefore, it is unlikely that the analysis in [15] will help to obtain a subquadratic complexity bound for realistic terrains. We give a more global analysis that describes where the boundaries of the shortest path map cells originate from to obtain an $O(n\sqrt{n})$ size bound.

Triangle edges necessarily give rise to boundaries of cells in the shortest path map. The cell boundaries that cut through triangles capture whether a shortest path passes to the one side or the other side of a vertex of $\mathcal{T}$, or through the vertex itself. In a terrain, there are two types of vertices: (1) vertices of which the spatial angle, i.e., the sum of the angles of the triangles incident to the vertex, is at least 2π, and (2) vertices with spatial angle less than 2π. From earlier papers

on shortest paths on terrains [4, 11, 15], it is known that a shortest path from s to any point on the terrain can only pass through vertices of type (1). These two types of vertices give different boundary types in the shortest path map.

Definition 18 *The* vertex trace *of a vertex v with total spatial angle $< 2\pi$ is the set of points on the terrain that each have two different shortest paths from s, and such that for any point p on the vertex trace, the two shortest paths from s have the same edge sequences except for edges incident to v.*

Every vertex with spatial angle $< 2\pi$ has one vertex trace. The vertex trace is a single straight line segment on each triangle on which it occurs. It is also the bisector of two different realizations of the same point (some vertex, or s) in the unfolding of the triangles that the shortest paths cross.

Definition 19 *The* geodesic region *of a vertex v with total spatial angle $\geq 2\pi$ is the region of points on $\mathcal{T}$ such that for any point p in that region, the shortest path from the source s to p passes through v as the last vertex. A* geodesic trace *of a vertex v with total spatial angle $\geq 2\pi$ is the maximal connected set of points on the boundary of the geodesic region of v that is a subset of a shortest path from s through v.*

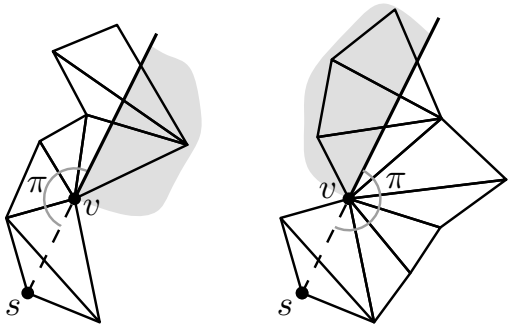

Figure 16: Two different unfoldings around v give rise to two geodesic traces (in bold) starting at v, which enclose the geodesic region of v (shaded).

Intuitively, the geodesic traces of a vertex v are the outermost shortest paths that still pass through v. More formally, if the shortest path from s to v leaves v in a certain direction, then the two geodesic traces continue beyond v at angles exactly π from this direction in the unfolding at v. Every vertex with spatial angle $> 2\pi$ has exactly two geodesic traces. If the spatial angle at v is exactly 2π, then the geodesic region degenerates to a linear feature and the geodesic trace is the same as the geodesic region. A geodesic trace is a single straight line segment on each triangle on which it occurs; see Figure 16 for an illustration.

We consider traces to be directed away from the source, which means that for a point that moves over the trace in the given direction, it gets further from the source s. Terrain vertices are always starting points of vertex or geodesic traces. The only way in which a trace can end is when it reaches the boundary of the terrain, or when it reaches another trace. The latter case gives rise to junction traces.

Definition 20 *The* junction trace *of a point q on the terrain, where q is a point where two traces end (any combination of the three types), is the set of points with two different shortest paths from s, and such that the edge sequences of these two shortest paths are the same starting at one edge of the triangle that contains q.*

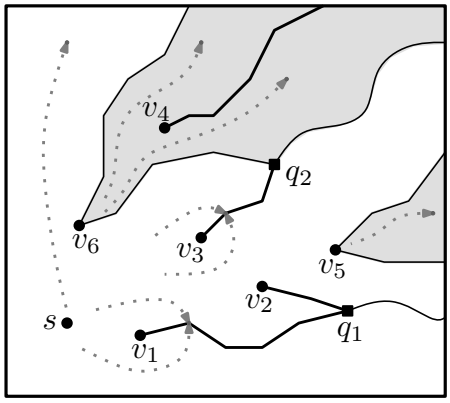

Figure 17: Different traces on a terrain: vertices v_1, v_2, v_3, and v_4 give rise to vertex traces, v_5 and v_6 to geodesic regions and corresponding traces, and q_1 and q_2 are starting points of junction traces. The dotted arrows show the directions of shortest paths from s.

A junction trace cannot start at a vertex, but starts in the interior of a triangle or on an edge. A junction trace is a single straight line segment or a hyperbolic arc on each triangle on which it occurs (or two or three hyperbolic arcs that are part of the same hyperbola). It is also the bisector of two vertices of the terrain (possibly the same but in different unfoldings), with additive weights. The three different types of traces are illustrated in Figure 17.

On a convex polytope, Aronov and O'Rourke [3] define the ridge tree as the set of points that have two or more distinct paths from the source. A convex polytope or terrain has no geodesic traces, and the vertex traces and junction traces together are the same as the ridge tree.

Lemma 21 *No pair of traces intersects. No trace properly intersects a shortest path from s.*

Proof. The first part is true by construction. Geodesic traces are parts of a shortest path from s, therefore another shortest path from s cannot intersect a geodesic trace properly.

Vertex traces and junction traces are bisectors of two vertices v and v' (after unfolding), possibly with additive weights. A proper intersection of a shortest path σ with such a trace τ implies that one point on σ has its shortest path from s via v, whereas a point just across τ on σ has its shortest path from s via v'. This contradicts the fact that σ is (part of) a shortest path from s. ◻

Lemma 22 *There are $O(n)$ traces on a terrain with n triangles.*

Proof. There are $O(n)$ vertex and geodesic traces. At a junction trace, two other traces always end and one begins. Hence there are at most $O(n)$ junction traces as well. ◻

Lemma 23 *Any trace has complexity $O(\sqrt{n})$.*

Proof. A geodesic trace is part of a shortest path from s, so by Lemma 9 it has complexity $O(\sqrt{n})$.

Vertex and junction traces are bisectors of two roots r_1 and r_2 in different unfoldings. If we consider the two endpoints p and q of some trace τ, then there are two shortest paths from s to p, one via root r_1 and one via root r_2. The same is true for q. This implies that the trace τ is a straight line segment or hyperbolic arc in the unfolding of the triangles on which it occurs, and hence it has no breakpoints. Therefore, by the proof of Lemma 16, each vertex or junction trace has complexity $O(\sqrt{n})$. ◻

Lemma 24 *The subdivision induced by the traces and the triangles is the shortest path map.*

PROOF. The proof consists of two parts: we start by showing that two points p and q in the same cell of the subdivision induced by the traces and triangles have the same edge sequence from s, then we show the converse.

Let p and q lie in the same cell of the subdivision induced by the edges of T and the traces. Clearly p and q must be in the same triangle of T. Consider the shortest paths from s to p and q and the corresponding edge sequences. Assume for a contradiction that the edge sequences differ, and consider the triangle t closest to p and q where this occurs. The shortest paths from s to p and q leave t through the same edge, but do not enter it through the same edge.

First, assume that the shortest paths enter t through different edges, and let the shared endpoint of these edges be v. This case is illustrated by the left part of Figure 18. Then v has a trace, or possibly two. By Lemma 21, traces and shortest paths do not intersect, so the trace of v must be between the shortest paths of p and q on all triangles intersected by the shortest paths between t and p and q. Hence, p and q cannot be in the same cell, a contradiction with the assumption that they have a different edge sequence.

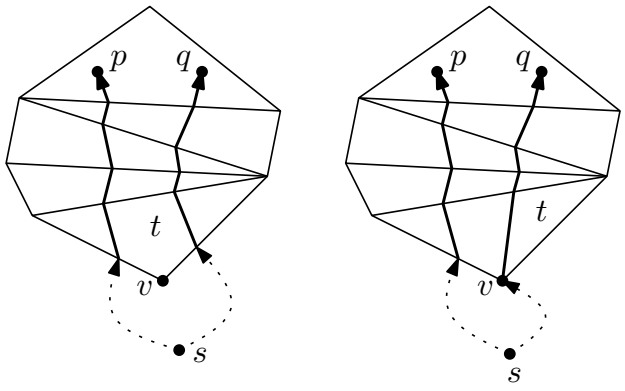

Figure 18: Illustration of the proof of Lemma 24.

Second, assume that one shortest path (say, to p) enters t through an edge and the other shortest path (to q) enters t through a vertex; this vertex v is opposite to the edge through which both shortest paths leave t; see the right of Figure 18 for an illustration. The total spatial angle at v is at least 2π, so v has one or two geodesic traces. By the angle properties at v, we have two cases: (i) the shortest path to q lies between the two geodesic traces of v, which is a contradiction to the fact that p and q lie in the same cell of the subdivison, or (ii) the shortest path to q coincides partially with (one of the) geodesic trace(s) of v. In the latter case, we distinguish three subcases: (a) the other geodesic trace lies strictly between the two shortest paths to p and q, in which case p and q cannot lie in the same cell of the subdivision, which is a contradiction again, (b) q itself lies on a geodesic trace of v, again a contradiction to the fact that p and q are in the same cell, or (c) a junction trace starts at some point z on the shortest path between v and q. Assuming z lies in the interior of a triangle, the shortest path of q goes straight through z, whereas the three traces have angles less than π between them (for hyperbolic arcs we refer to the angle of the tangent at z). This is true because on a triangle, the traces are the additively weighted Voronoi diagram of some set of sites. Hence, locally at z, either the other trace arriving at z or the junction trace leaving z is between the shortest paths of p and q. Since traces do not intersect shortest paths, both cases yield that p and q do not lie in the same cell. Note that the proof arguments also apply to the case where t is the triangle that contains p and q.

In the second part of this proof, we show that the converse is also true, i.e., two points p and q that have the same edge sequence from s lie in the same cell of the subdivision induced by the traces and triangles.

Assume for a contradiction that p and q do not lie in the same cell of the subdivision. If p and q lie in different triangles of $\mathcal{T}$, then they clearly have different edge sequences from s and we are done. Otherwise, they lie in the same triangle t of $\mathcal{T}$ and there is a trace τ that separates p and q on t. We only consider the case that p and q lie in adjacents cells that share τ as a common boundary; the proof in the other cases proceeds similarly. Since the shortest path from s to p does not intersect any trace, it must stay on the other side of τ than the shortest path to q. If τ is a vertex or geodesic trace induced by vertex v, then the shortest paths from s to p and q pass on different sides of v, and thus these paths have different edge sequences. If τ is a junction trace, then there exists at least two vertices v and w, which (indirectly) induce τ, such that the shortest paths from s to p and q pass on opposite sides of both v and w, and thus they have different edge sequences, which completes the proof. □

Theorem 25 *The shortest path map of a realistic terrain has complexity $\Theta(n\sqrt{n})$ in the worst case.*

Proof. The upper bound directly follows from the lemmas above. For the lower bound, take a convex terrain so that there are no geodesic traces. If we place the vertices on a nearly flat surface, then vertex traces never merge into junction traces. At least $\Omega(n)$ vertex traces will intersect $\Omega(\sqrt{n})$ edges of $\mathcal{T}$. □

Although we could not use the combinatorial analysis of Mitchell et al. [15] to obtain better complexity bounds, it is easy to verify that the total number of points for which their algorithm computes the additively weighted Voronoi diagram is $O(n\sqrt{n})$ on a realistic terrain, and therefore the shortest path map for a point on a realistic terrain can be computed in $O(n\sqrt{n}\log n)$ time.

4.3 Voronoi diagram

Let S be a set of m sites on a realistic terrain $\mathcal{T}$. We are interested in the maximum complexity of the Voronoi diagram of S on $\mathcal{T}$, denoted by $\mathrm{VD}(S,\mathcal{T})$, where distances are shortest path distances on $\mathcal{T}$. Each of the m Voronoi cells, for the sites $s_i \in S$, on $\mathcal{T}$ is connected, but not necessarily simply-connected. By Euler's formula this implies that there are $O(m)$ Voronoi vertices of degree 3 or higher. Furthermore, only $O(m)$ bisectors appear in the Voronoi diagram of S on $\mathcal{T}$. Together with Lemma 17, this immediately leads to an $O(mn\sqrt{n})$ bound on the complexity, but we will show that it is $O((n+m)\sqrt{n})$. First we give a lower bound of $\Omega(n+m\sqrt{n})$.

Lemma 26 *The Voronoi diagram of m sites on a realistic terrain has complexity $\Omega(n+m\sqrt{n})$.*

Proof. If the $\Omega(n)$ term dominates the $\Omega(m\sqrt{n})$ term, we simply use two sites that give a bisector of complexity $\Omega(n)$. Otherwise, we place m points on a flat, horizontal terrain in a row, so that they have the same y-coordinates. If we place the terrain vertices on a regular grid, then each of the $m-1$ bisectors intersects $\Omega(\sqrt{n})$ triangles between the boundaries with the terrain. □

As we mentioned before, the Voronoi diagram $\mathrm{VD}(S,\mathcal{T})$ has $O(m)$ vertices of degree 3 (or higher), and consists of parts of bisectors that intersect edges and contain breakpoints. Since any vertex can only give a breakpoint in a bisector of a site in whose cell that vertex lies, we can prove the following lemma.

Lemma 27 *The total number of breakpoints in $\mathrm{VD}(S,\mathcal{T})$ is $O(n+m)$.*

Proof. Every vertex of $\mathcal{T}$ lies in exactly one Voronoi cell; let v be some vertex and assume it lies in the cell c_i of site s_i. We consider only those traces that start at v and are created by source s_i. Vertex v may have one vertex trace or two geodesic traces incident to it. By Lemma 22, there are $O(n_i)$ such traces inside c_i, where n_i is the number of vertices inside c_i. We will show that the number of intersections of traces with the boundary of c_i is $O(n_i+h_i)$, where h_i is the number of Voronoi cells that lie in the interior of c_i, and share at least one Voronoi edge with c_i. These intersections are exactly the breakpoints of the bisectors associated with s_i that occur in

$\mathrm{VD}(S, \mathcal{T})$. We have two cases: (i) the Voronoi cell c_i does not contain any other Voronoi cells in its interior, and (ii) c_i does contain other cells as 'holes'. We analyse these two cases separately. We only consider vertex traces; the proof is analogous for the case that v gives rise to a geodesic trace, or the trace starting at v has merged with another trace into a junction trace before reaching the boundary of c_i.

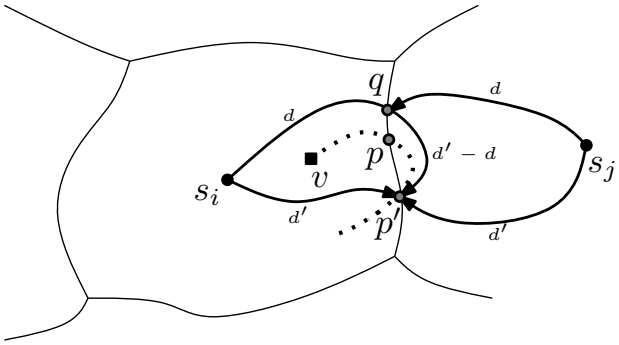

Figure 19: The vertex trace of v cannot reenter the Voronoi cell c_i.

For the first case, assume the vertex trace of v reaches the boundary of cell c_i at a point p, disappears when it leaves c_i, and reappears in c_i at another point p'; see Figure 19. Then at p', there are two equally long shortest paths from s_i and one also equally long shortest path from another site s_j. The two shortest paths from s_i enclose the vertex v and they do not intersect the trace of v. Therefore, one of the shortest paths from s_i must contain points strictly in the interior of the cell c_j. Assume one shortest path to p' enters c_i at a point q. If the distance from s_i and s_j to q is d and the distance from to s_i and s_j to p' is d', then the distance from p' to q is $d' - d$. Consequently, there is also a second shortest path from s_j to p', and the whole shortest path between p' and q has equal distance from s_i and s_j. But then it is part of the bisector, which contradicts the fact that the path from q to p' lies inside c_j. This concludes the proof for case (i).

In the second case, let c_j be the Voronoi cell of a site s_j which lies in the interior of c_i. Again, let v be a vertex in c_i, and assume that the vertex trace τ_v of v reaches the boundary between cell c_i and c_j at a point p, disappears, and reappears in c_i at another point p'. This time, we cannot argue that one of the shortest paths from p' to s_i must contain points strictly in the interior of the cell c_j, since these paths can enclose c_j, and indeed, it is possible that the trace τ_v reappears in c_i. However, we can show that this can only happen once for every cell c_j inside c_i. For a contradiction, assume that another trace τ_u also reappears in c_i, at a point q; see Figure 20.

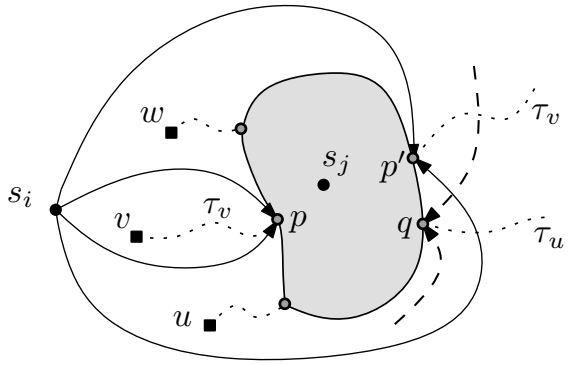

Figure 20: The Voronoi cell c_j is contained in the cell c_i.

By the argumentation for the first case, the two shortest paths from s to q cannot go through c_j, which implies that one of these paths must intersect some trace, which is not possible by Lemma 21. Thus, there is only one 'reappearing trace' per Voronoi cell that lies within c_i. Therefore, in case (ii), the number of intersections of traces with the cell boundary is $O(n_i + h_i)$.

Finally, we observe that the total number of breakpoints is $\sum_{i=1}^{m}(n_i + h_i)$, which is $O(n + m)$. □

Between the $O(m)$ Voronoi vertices and the $O(n+m)$ breakpoints, the bisectors have complexity $O(\sqrt{n})$ as shown in Lemma 17, and hence we conclude with the following theorem.

Theorem 28 *The Voronoi diagram of a set of m sites on a realistic terrain with n triangles has complexity $\Omega(n + m\sqrt{n})$ and $O((n + m)\sqrt{n})$ in the worst case.*

5 Concluding Remarks

This paper studies realistic input models for polyhedral terrains, a topic that has not been considered so far. We have made three input assumptions that together are sufficient to show a subquadratic upper bound on the complexity of the visibility map. Furthermore, we have shown that no subset of our assumptions is sufficient to achieve the same. For shortest paths, bisectors, shortest path maps, and Voronoi diagrams on terrains, we have used a fourth input assumption and proved upper and lower bounds on their complexities. Our research helps to explain the discrepancy between the worst-case performance of algorithms on polyhedral terrains and their efficiency in practice.

The upper and lower bounds for visibility maps, shortest paths, and the shortest path map are tight, but there is a considerable gap for bisectors and Voronoi diagrams. This is the most important open problem that arises from this paper.

Most of our bounds are $O(n\sqrt{n})$ for realistic terrains, in contrast with quadratic for general terrains. However, in practice we may expect even lower bounds for the visibility map, like close to linear. The question is whether a (slightly) stronger set of realistic assumptions exists that leads to such a bound. On the other hand, this study is a first attempt at defining a realistic input model for polyhedral terrains, and therefore, it would be interesting to replace one (or more) of our assumptions with a weaker version, while still obtaining the same bounds.

Acknowledgements

The authors thank Mark de Berg and several anonymous reviewers for useful comments.

References

[1] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra, and C. Uhrig. Approximate motion planning and the complexity of the boundary of the union of simple geometric figures. *Algorithmica*, 8:391–406, 1992.

[2] B. Aronov, A. Efrat, V. Koltun, and M. Sharir. On the union of κ-round objects in three and four dimensions. In *Proc. 20st Annu. ACM Sympos. Comput. Geom.*, pages 383–390, 2004.

[3] B. Aronov and J. O'Rourke. Nonoverlap of the star unfolding. *Discrete Comput. Geom.*, 8:219–250, 1992.

[4] J. Chen and Y. Han. Shortest paths on a polyhedron. *Internat. J. Comput. Geom. Appl.*, 6(2):127–144, 1996.

[5] M. de Berg. Linear size binary space partitions for uncluttered scenes. *Algorithmica*, 28:353–366, 2000.

[6] M. de Berg. Vertical ray shooting for fat objects. In *Proc. 21st Annu. ACM Sympos. Comput. Geom.*, pages 288–295, 2005.

[7] M. de Berg, M. J. Katz, M. Overmars, A. F. van der Stappen, and J. Vleugels. Models and motion planning. *Comput. Geom. Theory Appl.*, 23:53–68, 2002.

[8] M. de Berg, M. J. Katz, A. F. van der Stappen, and J. Vleugels. Realistic input models for geometric algorithms. *Algorithmica*, 34:81–97, 2002.

[9] A. Efrat. The complexity of the union of (α, β)-covered objects. *SIAM J. Comput.*, 34:775–787, 2005.

[10] J. Erickson. Local polyhedra and geometric graphs. *Comput. Geom. Theory Appl.*, 31:101–125, 2005.

[11] S. Kapoor. Efficient computation of geodesic shortest paths. In *Proc. 31st Annu. ACM Sympos. Theory of Computing*, pages 770–779, 1999.

[12] M. J. Katz, M. H. Overmars, and M. Sharir. Efficient hidden surface removal for objects with small union size. *Comput. Geom. Theory Appl.*, 2:223–234, 1992.

[13] J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl. Fat triangles determine linearly many holes. *SIAM J. Comput.*, 23:154–169, 1994.

[14] M. McKenna. Worst-case optimal hidden-surface removal. *ACM Trans. Graph.*, 6:19–28, 1987.

[15] J. S. B. Mitchell, D. M. Mount, and C. H. Papadimitriou. The discrete geodesic problem. *SIAM J. Comput.*, 16:647–668, 1987.

[16] J. S. B. Mitchell, D. M. Mount, and S. Suri. Query-sensitive ray shooting. *Internat. J. Comput. Geom. Appl.*, 7(4):317–347, Aug. 1997.

[17] E. Moet, M. van Kreveld, and A. F. van der Stappen. On realistic terrains. To appear in *Proc. 22nd Annu. ACM Sympos. Comput. Geom.*, 2006.

[18] M. H. Overmars and A. F. van der Stappen. Range searching and point location among fat objects. *J. Algorithms*, 21:629–656, 1996.

[19] O. Schwarzkopf and J. Vleugels. Range searching in low-density environments. *Inform. Process. Lett.*, 60:121–127, 1996.

[20] A. F. van der Stappen, M. H. Overmars, M. de Berg, and J. Vleugels. Motion planning in environments with low obstacle density. *Discrete Comput. Geom.*, 20(4):561–587, 1998.

[21] M. van Kreveld. On fat partitioning, fat covering, and the union size of polygons. *Comput. Geom. Theory Appl.*, 9(4):197–210, 1998.

[22] J. Vleugels. *On Fatness and Fitness — Realistic Input Models for Geometric Algorithms*. Ph.D. thesis, Dept. Comput. Sci., Univ. Utrecht, Utrecht, The Netherlands, 1997.