# The Complexity of Bisectors and Voronoi Diagrams on Realistic Terrains

Boris Aronov[1], Mark de Berg[2], and Shripad Thite[3]

[1] Department of Computer and Information Science, Polytechnic University, USA
aronov@poly.edu
[2] Department of Computing Science, TU-Eindhoven, the Netherlands
mdberg@win.tue.nl
[3] California Institute of Technology, Center for the Mathematics of Information, USA
shripad@caltech.edu

**Abstract.** We prove tight bounds on the complexity of bisectors and Voronoi diagrams on so-called realistic terrains, under the geodesic distance. In particular, if $n$ denotes the number of triangles in the terrain, we show the following two results.

(i) If the triangles of the terrain have bounded slope and the projection of the set of triangles onto the $xy$-plane has low density, then the worst-case complexity of a bisector is $\Theta(n)$.

(ii) If, in addition, the triangles have similar sizes and the domain of the terrain is a rectangle of bounded aspect ratio, then the worst-case complexity of the Voronoi diagram of $m$ point sites is $\Theta(n + m\sqrt{n})$.

## 1 Introduction

*Motivation.* The *Voronoi diagram* of a set $S$ of $m$ sites in a metric space is the decomposition of the space into $m$ cells, one per site, such that the cell corresponding to a site $p \in S$ contains exactly those points for which $p$ is the closest site. Often the sites are points and the ambient space is a Euclidean space, but there are many other interesting settings. Voronoi diagrams play a role in numerous applications and they have been studied extensively—see for example the book by Okabe *et al.* [9] or one of the several surveys [1, 2, 6] dedicated to Voronoi diagrams. One of the areas where Voronoi diagrams are frequently used is geographic information systems. A natural setting in this application is where the sites are points in a mountainous terrain, and the distance between any two points on the terrain is the geodesic distance. (The *geodesic distance* between two points is the length of a shortest path on the terrain connecting them.) This is the setting of our paper.

A standard way to model a terrain is using a *triangulated irregular network*, or *TIN* for short: a triangulation of a convex polygonal domain in the $xy$-plane—usually the domain is simply a rectangle—where each vertex is given an elevation. In computational geometry, a TIN is called a *polyhedral terrain*. Hereafter, the term *terrain* refers to a polyhedral terrain.

A fundamental issue in the study of Voronoi diagrams is their combinatorial complexity. It is well known that the complexity—that is, the number of vertices, edges, and cells—of the Voronoi diagram of $m$ sites in the plane under the Euclidean distance is $\Theta(m)$. This follows easily from the fact that the Voronoi diagram is a planar subdivision with $m$ faces, whose vertices have degree at least three and whose edges are line segments or half lines. For Voronoi diagrams on a terrain, things are more complicated. Here, the complexity depends not only on the number of sites but also on the number of triangles in the terrain. Indeed, a single bisector—the *bisector* of two sites is the set of points equidistant from both sites—on a terrain consisting of $n$ triangles can already have complexity $\Omega(n^2)$ [8] and the Voronoi diagram of $m$ sites can have complexity $\Omega((n+m)n)$. Fortunately, these high-complexity bisectors and Voronoi diagrams seem to arise only on carefully constructed, artificial terrains—terrains in practical applications probably behave much better. Thus the question arises: how can we formalize the notion of a "well-behaved terrain" and what is the worst-case complexity of bisectors and Voronoi diagrams on such terrains?

*Previous results.* These considerations lead Moet *et al.* [8] to study terrains with the following properties: (i) the triangles in the terrain have bounded slope; (ii) the set of terrain triangles has low density; (iii) the domain of the terrain has bounded aspect ratio; (iv) all terrain triangles have roughly the same size. (A more formal definition of these properties is given in Section 2.) They call a terrain with these four properties a *realistic terrain.* Moet investigates [7] whether the assumption of bounded slope, density etc. is pragmatic by measuring these parameters of terrain models of various mountainous regions in the US, which she concludes indeed have the properties listed above.

Moet *et al.* prove that the complexity of a single bisector on a realistic terrain with $n$ triangles is $O(n\sqrt{n})$ and can sometimes be $\Omega(n)$. Moreover, they show that the complexity of the Voronoi diagram on a realistic terrain is $O((n+m)\sqrt{n})$, and can sometimes be $\Omega(n+m\sqrt{n})$.

Recently, Schreiber [11] studied the computation of shortest paths on realistic terrains (or, more generally, realistic polyhedra). Schreiber computes an *implicit* representation of the Voronoi diagram on a realistic terrain in $O((n+m)\log(n+m))$ time, so that the site closest to a query point can be reported in $O(\log(n+m))$ time. For some applications it will be sufficient to have such an implicit representation; for others one needs an explicit representation. The explicit Voronoi diagram can be constructed in $O((n+m)\log(n+m)+k)$ time, where $k$ is the combinatorial complexity of the Voronoi diagram, by an extension of Schreiber's algorithm [12]. The question now arises: what is the maximum combinatorial complexity of the Voronoi diagram on a realistic terrain? This is the question studied by Moet *et al.* [7, 8] and explored further in this paper.

*Our results.* We improve on the results by Moet *et al.* [8] and give tight bounds on the complexity of bisectors and Voronoi diagrams on realistic terrains.

First, we prove that the worst-case complexity of a single bisector on a realistic terrain is $\Theta(n)$. We obtain our improved bound by studying the global shape of the bisector and showing essentially that it cannot "wiggle" too wildly. More precisely, we prove that the set of pieces forming the bisector has low density. We believe that this result is of independent interest. Interestingly, our proof only requires the terrain to have properties (i) and (ii) listed above; thus it yields not only a significantly better bound than what was known, but it also applies to a wider class of terrains.

Second, we show that the worst-case complexity of the Voronoi diagram on a realistic terrain is $\Theta(n + m\sqrt{n})$. This result is based partially on our improved bound on the complexity of a bisector and partially on a careful investigation of the structure of the Voronoi diagram.

## 2 Preliminaries

Let $\mathcal{T}$ be a terrain with $n$ triangles. In this section, we denote the vertical projection of any subset $o \subset \mathcal{T}$ to the $xy$-plane by $\overline{o}$. We will use $D$ to denote the domain of $\mathcal{T}$, which is a subset of the $xy$-plane. For simplicity, and because this is mostly the case in practice, we assume that $D$ is a rectangle; our results can easily be extended to the case where $D$ is an arbitrary convex region. Notice that $\overline{\mathcal{T}}$ is a triangulation of $D$.

Next we formally define the parameters that measure how well-behaved a terrain is.

- The *slope* of a triangle $\Delta$ in $\mathbb{R}^3$ is the maximum slope of any line segment contained in $\Delta$. For example, a triangle parallel to the $xy$-plane has slope 0, while a vertical triangle—a triangle parallel to the $z$-axis—has infinite slope. The slope $\xi$ of the terrain $\mathcal{T}$ is the maximum slope of any of its triangles. Note that a terrain does not contain vertical triangles by definition, so it has finite slope.
- The *density* [4] of a set $S$ of objects in the plane is defined as the smallest number $\lambda$ such that any disk $B$ intersects at most $\lambda$ objects $o \in S$ such that $\text{diam}(o) \geqslant \text{diam}(B)$, where $\text{diam}(\cdot)$ denotes the diameter. The density $\lambda$ of the terrain $\mathcal{T}$ is the density of the set of edges of $\overline{\mathcal{T}}$. In other words, the density refers to the edges of the triangulation of the domain $D$ that corresponds to $\mathcal{T}$.
- The *aspect ratio* of a rectangle with width $w$ and height $h$ is defined as $\max(w/h, h/w)$. The aspect ratio $\rho$ of $\mathcal{T}$ is the aspect ratio of its domain $D$.
- The *scale factor* $\sigma$ of $\mathcal{T}$ is the ratio between the maximum and the minimum length of any edge of $\overline{\mathcal{T}}$.

Moet *et al.* [8] define a realistic terrain as a terrain whose slope $\xi$, density $\lambda$, aspect ratio $\rho$, and scale factor $\sigma$ are constants independent of $n$, and then prove bounds on the complexity of bisectors and Voronoi diagrams as a function of $n$ only, with the dependence on $\xi$, $\lambda$, $\rho$, and $\sigma$ hidden in the asymptotic notation. We make this dependence explicit in all our bounds.

For two points $p, q \in \mathcal{T}$, we use $\text{dist}(p, q)$ to denote the geodesic distance between $p$ and $q$. In other words, $\text{dist}(p, q)$ is the length of a shortest path from

$p$ to $q$ on $\mathcal{T}$. Furthermore, we use $|\overline{pq}|$ to denote the Euclidean distance between $\overline{p}$ and $\overline{q}$. As already observed by Moet *et al.* [8], the geodesic distance between two points on a terrain with bounded slope is closely related to the Euclidean distance between their projections:

**Lemma 1.** [**8**] *For any two points $p, q$ on a terrain $\mathcal{T}$ with slope $\xi$, we have* $\mathrm{dist}(p, q) \leqslant \sqrt{\xi^2 + 1} \cdot |\overline{pq}|$.

A second basic fact that we will use is that shortest paths on a realistic terrain cross $O(\sqrt{n})$ triangles.

**Lemma 2.** [**8**] *Let $\mathcal{T}$ be a terrain with slope $\xi$, density $\lambda$, aspect ratio $\rho$, and scale factor $\sigma$. Then any shortest path on $\mathcal{T}$ crosses $O(c\sqrt{n})$ terrain edges, where $c = \xi\lambda\sigma\sqrt{\rho}$.*

Finally, we will use the following result, which follows easily from the definition of density (just charge every intersecting pair $(o_1, o_2)$ to the object with smaller diameter). Similar results have been used in previous papers [3,13] dealing with low-density scenes.

**Lemma 3.** *Let $S_1$ be a set of $n_1$ objects and density $\lambda_1$, and let $S_2$ be a set of $n_2$ objects and density $\lambda_2$. Then the number of pairs $(o_1, o_2) \in S_1 \times S_2$ such that $o_1$ intersects $o_2$ is $O(\lambda_2 n_1 + \lambda_1 n_2)$.*

There is a natural one-to-one correspondence (obtained by vertical projection) between points on the terrain $\mathcal{T}$ and points in the domain $D$. Hence, we can view Voronoi diagrams and bisectors as subsets of $\mathcal{T}$, or as subsets of $D$. From now on, we will take the latter view and consider these structures to be subsets of $D$. It is then also convenient to no longer make an explicit distinction between geometric entities—points, shortest paths, bisectors, etc.—on the terrain $\mathcal{T}$ and their projections to the domain $D$, and drop the notation $\overline{o}$ for the projection of an object $o$. Thus, for example, when we speak of a shortest path $\pi$ between two points $s$ and $t$ on the terrain, we actually refer to the path $\overline{\pi}$ that connects $\overline{s}$ to $\overline{t}$. (When it is important to make the distinction between an object and its projection, we will explicitly do so.) Moreover, $|xy|$ refers to the Euclidean distance between points $x$ and $y$ on $D$, while $\mathrm{dist}(x, y)$ refers to the length of a shortest path between the corresponding points on $\mathcal{T}$.

*The structure of shortest paths on a terrain.* A shortest path $\pi(x, y)$ between $x, y \in \mathcal{T}$ is a polygonal path that stays straight within individual terrain triangles and unfolds to a straight line segment whenever it crosses a terrain edge away from a vertex. A shortest path may pass through a terrain vertex (the vertex has to be non-convex in a technical sense that is not important in this paper). Two shortest paths $\pi(x, y)$ and $\pi(x, z)$ emanating from the same point $x$ do not properly cross, nor overlap and then diverge, except possibly at (non-convex) vertices of $\mathcal{T}$.

If two sites are equidistant from a terrain vertex, their bisector need not be a curve; it may contain entire two-dimensional regions. So, in order for bisectors and Voronoi diagrams to be properly defined, following previous work [8], we therefore make the *general position assumption* that no two sites are equidistant from a terrain vertex. This assumption guarantees that bisectors are 1-dimensional and that Voronoi cells are regions that cover $\mathcal{T}$ without overlap, except at their common boundaries.

Moreover, in this version of the paper, we also add another *non-degeneracy assumption* that is not needed for the results presented to hold but that simplifies the presentation; the assumption is removed in the full version of the paper. Namely, we assume that each site $s$ is connected to every vertex $v$ of the terrain by a *unique* shortest path. Non-degeneracy implies that, for every point $x \in \mathcal{T}$, all shortest paths between $s$ and $x$ are pairwise non-crossing; they may overlap but they cannot cross. Now, for any two shortest paths $\pi(s,x)$ and $\pi(s,y)$, for $x \neq y$, there must exist a point $z$ (which might coincide with $s$) so that $\pi(s,x) \cap \pi(s,y) = \pi(s,z)$. Moreover, for any two distinct shortest paths $\pi_1(s,x)$ and $\pi_2(s,x)$ from $s$ to the same point $x$ there must exist a point $z \neq x$ with a unique shortest path $\pi(s,z)$ from $s$ so that $\pi_1(s,x) \cap \pi_2(s,x) = \pi(s,z) \cup \{x\}$.

## 3 The bisector

Let $s$ and $t$ be two point sites (not necessarily vertices) on a terrain $\mathcal{T}$. In this section we study the complexity of the bisector $b = b(s,t)$ of $s$ and $t$ on $\mathcal{T}$. We will do the analysis in terms of $n$, the number of triangles of the terrain, and its slope $\xi$ and density $\lambda$.

The bisector $b$, by definition, consists of all points $p \in \mathcal{T}$ such that $\text{dist}(p,s) = \text{dist}(p,t)$. It partitions $\mathcal{T}$ into two regions: $\mathcal{V}(s)$, the *Voronoi cell* of $s$, which contains the points closer to $s$, and $\mathcal{V}(t)$, the *Voronoi cell* of $t$, which contains the points closer to $t$. Since Voronoi cells are connected, $b$ is a simple curve that is either closed—this can happen, for instance, when $s$ is the peak of a mountain and $t$ is at the foot of the mountain—or connects two points on the boundary of the terrain.

For most points on $b$, there is a unique shortest path to $s$ and a unique shortest path to $t$. For some points, however, there are multiple shortest paths to $s$ and/or to $t$. We call such points *breakpoints*. The number of breakpoints on $b$ is at most $n$, because each of them can be attributed to a terrain vertex [10]. The breakpoints partition $b$ into *pieces*; the intersection of a piece with a terrain triangle is a line segment or hyperbolic arc [10]. The *complexity* of $b$ is now defined as the number of breakpoints plus the number of times that $b$ crosses a terrain edge.

We denote the set of all bisector pieces by $\Gamma$. Moet *et al.* [8] prove that on a realistic terrain any piece $\gamma \in \Gamma$ can cross only $O(\sqrt{n})$ triangles. Since $|\Gamma| \leqslant n+1$, this implies that the total complexity of the bisector is $O(n\sqrt{n})$. To improve upon this, we take a more global look at the bisector and show that the set $\Gamma$ has low density. (Here it is important to recall that we view $\Gamma$ as a set of curves in the $xy$-plane, that is, as a collection of subsets of $D$.) The result will then readily follow from Lemma 3 and the fact that $\mathcal{T}$ has low density.

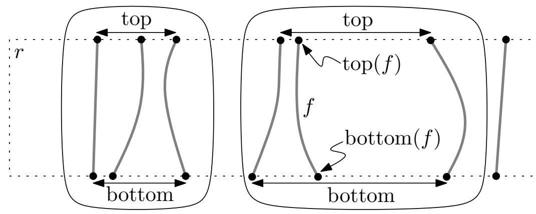


**Fig. 1.** Seven fragments and the two blocks defined by them. The rightmost fragment does not belong to any block.

Let $r$ be a rectangle in the $xy$-plane, and assume without loss of generality that $r$ is axis-parallel. We say that a piece $\gamma$ *crosses* $r$ if $\gamma \cap r$ has a connected component with one endpoint on the top edge of $r$ and one endpoint on the bottom edge of $r$. We call such a component a *fragment* of $\gamma$. To bound the density of $\Gamma$, we first show that $r$ cannot be crossed too many times.

We denote the top endpoint of a fragment $f$ by $\mathrm{top}(f)$ and its bottom endpoint by $\mathrm{bottom}(f)$. For each piece $\gamma$ that crosses $r$, we pick one of its fragments, and we let $F$ denote the set of all such fragments. Since each fragment $f \in F$ connects the top edge of $r$ to the bottom edge of $r$, we can order the fragments from left to right. We group the fragments from left to right in triples, and we call such an (ordered) triple a *block*—see Fig. 1. We start by proving a lemma on the structure of the shortest paths from the endpoints of the fragments in a block to $s$ and $t$.

Define the *top* of a block to be the line segment connecting the top of the leftmost fragment of the triple to that of its rightmost fragment. Define the *bottom* of a block analogously.

**Lemma 4.** *Let $(f_1, f_2, f_3)$ be a block. Then at least one of the three top endpoints has a shortest path to $s$ or $t$ that intersects the bottom of the block. Similarly, at least one of the three bottom endpoints has a shortest path to $s$ or $t$ that intersects the top of the block.*

*Proof.* We will prove the lemma for the top endpoints; the proof for the bottom endpoints is symmetric.

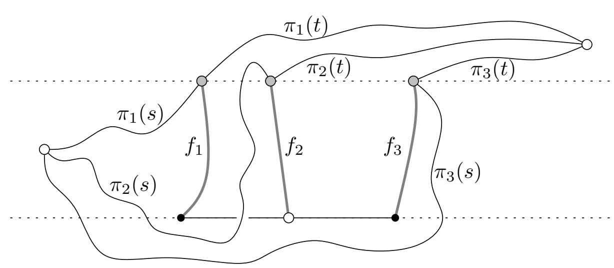


**Fig. 2.** One of the paths from the top endpoints must intersect the bottom of the block.

For $i \in \{1,2,3\}$ and $q \in \{s,t\}$, let $\pi_i(q)$ denote the shortest path from $\mathrm{top}(f_i)$ to $q$. Let $V_1 := \{\mathrm{top}(f_1), \mathrm{top}(f_2), \mathrm{top}(f_3)\}$ and $V_2 := \{s, t, \mathrm{bottom}(f_2)\}$.

Consider a geometric realization of the graph $K_{3,3}$, where $V_1 \cup V_2$ is the set of nodes and the arcs are realized as follows—see also Fig. 2. The arcs from nodes in $V_1$ to $s$ and $t$ are given by the shortest paths $\pi_i(s)$ and $\pi_i(t)$; the arc from $\mathrm{top}(f_2)$ to $\mathrm{bottom}(f_2)$ is given by $f_2$; and for $i = 1,3$ the arc from $\mathrm{top}(f_i)$ to $\mathrm{bottom}(f_2)$ is given by the concatenation of $f_i$ and the segment connecting $\mathrm{bottom}(f_i)$ to $\mathrm{bottom}(f_2)$.

Since $K_{3,3}$ is non-planar, there must be an intersection between some pair of arcs. Recall that the fragments $f_i$ are all part of the bisector $b$, which also means that the points $\mathrm{top}(f_i)$ lie on $b$. Hence, the paths $\pi_i(s)$ and $\pi_i(t)$ lie inside $\mathcal{V}(s)$ and $\mathcal{V}(t)$, respectively. This implies that the paths $\pi_i(s)$ and $\pi_i(t)$ do not intersect any of the fragments $f_j$, and also that a path $\pi_i(s)$ does not intersect any path $\pi_j(t)$ (except possibly at common endpoints). Furthermore, after two shortest paths $\pi_i(s)$ and $\pi_j(s)$ meet for the first time, they follow the same subpath, by our non-degeneracy assumption. Hence, a small perturbation yields paths that are disjoint (except at $s$). Similarly, we can enforce that the paths $\pi_i(t)$ and $\pi_j(t)$ are disjoint except at $t$. Finally, two fragments $f_i, f_j$ do not intersect each other, by construction. The only remaining possibility is that one of the paths $\pi_i(s)$ or $\pi_i(t)$ intersects one of the segments connecting $\mathrm{bottom}(f_i)$ to $\mathrm{bottom}(f_2)$ for $i = 1,3$. In other words, one of these paths intersects the bottom of the block. □

Now define the *top width* of a block $(f_1, f_2, f_3)$ as the length of the top of the block, define its *bottom width* analogously, and define the *width* of a block as the maximum of its top and bottom widths. The previous lemma allows us to prove a lower bound on the width of a block.

**Lemma 5.** *The width of any block* $(f_1, f_2, f_3)$ *is at least* $h/\sqrt{\xi^2+1}$*, where* $h$ *is the height of the rectangle* $r$*.*

*Proof.* By Lemma 4 one of the top endpoints, say $\mathrm{top}(f_i)$, has a shortest path to $s$ or $t$ that intersects the bottom of the block. Similarly, one of the bottom endpoints, say $\mathrm{bottom}(f_j)$, has a shortest path to $s$ or $t$ that intersects the top of the block. We can assume, without loss of generality, that $\mathrm{dist}(\mathrm{top}(f_i), s) \leqslant \mathrm{dist}(\mathrm{bottom}(f_j), s)$. Since $\mathrm{top}(f_i)$ and $\mathrm{bottom}(f_j)$ lie on the bisector $b$, we get

$$\mathrm{dist}(\mathrm{top}(f_i), t) = \mathrm{dist}(\mathrm{top}(f_i), s) \leqslant \mathrm{dist}(\mathrm{bottom}(f_j), s) = \mathrm{dist}(\mathrm{bottom}(f_j), t).$$

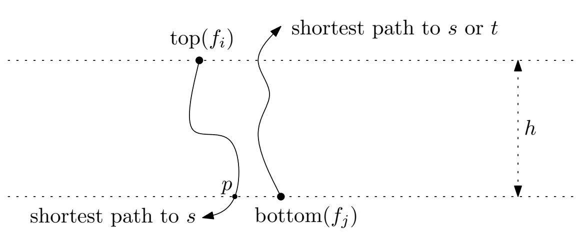


**Fig. 3.** The width of a block cannot be too small.

Now assume, also without loss of generality, that the shortest path from $\text{top}(f_i)$ that crosses the bottom of the block is the shortest path to $s$, and let the point where it crosses the bottom be denoted by $p$—see Fig. 3. Then we have

$$\begin{aligned}\text{dist}(\text{top}(f_i), s) &\leqslant \text{dist}(\text{bottom}(f_j), s) \\ &\leqslant \text{dist}(\text{bottom}(f_j), p) + \text{dist}(p, s) \\ &\leqslant \sqrt{\xi^2+1} \cdot |\text{bottom}(f_j)p| + \text{dist}(p, s) && \text{by Lemma 1} \\ &\leqslant \sqrt{\xi^2+1} \cdot |\text{bottom}(f_j)p| + \text{dist}(\text{top}(f_i), s) - h\end{aligned}$$

It follows that $|\text{bottom}(f_j)p| \geqslant h/\sqrt{\xi^2+1}$. □

The previous lemma implies that a rectangle $r$ of small aspect ratio cannot be crossed by too many bisector pieces.

**Lemma 6.** *The rectangle $r$ is crossed by at most $2 + 6w\sqrt{\xi^2+1}/h$ bisector pieces, where $h$ is the height of the rectangle $r$ and $w$ is its width.*

*Proof.* Consider the set of fragments induced by the bisector pieces crossing $r$, as defined above. These fragments are grouped into blocks of three, with at most two fragments not belonging to any block. Hence, we must show that the number of blocks is at most $2w\sqrt{\xi^2+1}/h$. By Lemma 5, either the top width or the bottom width of any block is at least $h/\sqrt{\xi^2+1}$. Since the width of $r$ is $w$, there can be at most $w/\sqrt{\xi^2+1}/h$ blocks whose top width is at least at least $h/\sqrt{\xi^2+1}$. Similarly, there can be at most $w\sqrt{\xi^2+1}/h$ blocks whose bottom width is at least $h/\sqrt{\xi^2+1}$, and so the total number of blocks is as claimed. □

Using Lemma 6 we obtain the following theorem. Its proof is very similar to a low-density proof by De Berg [3], and therefore omitted.

**Theorem 1.** *Let $s$ and $t$ be any two points on a terrain $\mathcal{T}$, let $b(s,t)$ be their bisector, and let $\Gamma$ be the collection of (projected) bisector pieces obtained by splitting $b(s,t)$ at breakpoints as defined above. Then $\Gamma$ has density $O(\xi)$, where $\xi$ is the slope of $\mathcal{T}$.*

Combining Theorem 1 with Lemma 3 immediately leads to the following result.

**Corollary 1.** *The bisector of two points on a terrain $\mathcal{T}$ with $n$ triangles has complexity $O((\xi+\lambda)n)$, where $\xi$ is the slope of $\mathcal{T}$ and $\lambda$ is its density.*

## 4 The Voronoi diagram

Let $S := \{s_1, \ldots, s_m\}$ be a set of $m$ point sites on a terrain $\mathcal{T}$ with $n$ triangles, and let $\mathcal{VD}(S)$ denote the Voronoi diagram of $S$. Each Voronoi edge is a portion of some bisector $b(s_i, s_j)$, and the number of Voronoi edges is $O(m)$. (Note that a Voronoi edge is not necessarily incident to two Voronoi vertices; it can also be a closed curve or have one or both endpoints on the boundary of $D$.) Bounding the complexity of the Voronoi diagram amounts to bounding the total complexity of all Voronoi edges. In the previous section we bounded the complexity of a single bisector as a function of $n$, and the slope $\xi$ and density $\lambda$ of $\mathcal{T}$. The bound on the complexity of the Voronoi diagram that we will prove in this section also depends on the aspect ratio $\rho$ and the scale factor $\sigma$ of $\mathcal{T}$.

Recall that a bisector $b(s_i, s_j)$ is partitioned into pieces at breakpoints—points where the shortest path to $s_i$ and/or to $s_j$ is not unique—and that the intersection of such a piece with a terrain triangle is a hyperbolic arc or line segment. Because a breakpoint on a Voronoi edge that is part of $b(s_i, s_j)$ can be uniquely attributed to a vertex lying inside one of the Voronoi cells $\mathcal{V}(s_i)$ or $\mathcal{V}(s_j)$, the total number of breakpoints over all Voronoi edges is $O(n)$ [8]. Hence, the total complexity of the Voronoi edges is proportional to $m+n$ plus the total number of intersections between Voronoi edges and terrain edges. The rest of this section is devoted to bounding the number of intersections between Voronoi edges and terrain edges.

Consider a breakpoint $p$ on the Voronoi edge generated by $b(s_i, s_j)$. We call $p$ a *special breakpoint* if it has two shortest paths to $s_i$ that enclose at least one hole in $\mathcal{V}(s_i)$, or two shortest paths to $s_j$ that enclose at least one hole in $\mathcal{V}(s_j)$. (A *hole* in $\mathcal{V}(s_i)$ is formed by one or more other Voronoi cells $\mathcal{V}(s_k)$ enclosed by $\mathcal{V}(s_i)$. As remarked earlier, this can happen for instance if $s_i$ is at the foot of a steep mountain and $s_k$ is at the peak.) Let $B$ be the set of all special breakpoints. The special breakpoints subdivide the Voronoi edges into *subedges*. To simplify the presentation, we augment $B$ with $O(m)$ additional points to ensure that $B$ contains, for each site $s_i$, at least two points on every component of $\partial\mathcal{V}(s_i)$.

**Lemma 7.** *Let $p$ and $q$ be the endpoints of a subedge $\gamma$ on $\partial\mathcal{V}(s_i)$. Then there exists a shortest path $\pi(p)$ from $p$ to $s_i$ and a shortest path $\pi(q)$ from $q$ to $s_i$, such that the region $\mathcal{V}(\gamma, s_i) \subset \mathcal{V}(s_i)$ enclosed by $\gamma$, $\pi(p)$ and $\pi(q)$ is simply connected. Moreover, there is a choice of shortest paths $\pi(p), \pi(q)$ for all subedges $\gamma$ that guarantees that $\mathcal{V}(\gamma, s_i)$ and $\mathcal{V}(\gamma', s_j)$ do not overlap for $(\gamma, s_i) \neq (\gamma', s_j)$.*

*Proof.* Since we augmented $B$ with extra points, $p$ and $q$ cannot coincide. Take any point $r$ in the interior of $\gamma$, and draw a shortest path $\pi(r)$ from $r$ to $s_i$. Imagine moving $r$ towards $p$. As we move $r$ continuously, we can also transform $\pi(r)$ continuously such that it stays shortest, except when $r$ moves over a breakpoint: at that point $\pi(r)$ jumps. (More precisely, when $r$ reaches a breakpoint which, by definition, has more than one shortest path to $s_i$, $\pi(r)$ coincides with one of these paths. To be able to continuously deform the path further while remaining shortest, we have to switch $\pi(r)$ to one of the other shortest paths. This is what we refer to as "jumping" of a shortest path.) Since $\gamma$ is a subedge, however, a breakpoint in its interior cannot be a special breakpoint. Hence, $\pi(r)$ does not jump over a hole of $\mathcal{V}(s_i)$, i.e., the region bounded by the two shortest paths to the breakpoint is fully contained in $\mathcal{V}(s_i)$ and is simply connected. The same argument shows that $\pi(r)$ will not jump over a hole when we move $r$ to $q$.

Hence, we can find shortest paths from $p$ and $q$ to $s_i$ such that the region $\mathcal{V}(\gamma, s_i)$ enclosed by them and $\gamma$ is simply connected.

Since $\mathcal{V}(\gamma, s_i) \subset \mathcal{V}(s_i)$, regions belonging to different sites cannot overlap. We claim that $\mathcal{V}(\gamma, s_i)$ and $\mathcal{V}(\gamma', s_i)$, for different subedges $\gamma, \gamma'$, do not overlap either. Recall from Section 2 that the non-degeneracy assumption implies that shortest paths from $s_i$ to points $x, y$, after diverging for the first time (as seen from $s_i$), do not meet again (except at $x$ when $x = y$). Hence, the shortest paths that bound the regions $\mathcal{V}(\gamma, s_i)$ form a tree-like structure, which implies the regions $\mathcal{V}(\gamma, s_i)$ and $\mathcal{V}(\gamma', s_i)$ for different subedges do not overlap. □

Next we bound the total number of subedges.

**Lemma 8.** *The total number of subedges is $O(m)$.*

*Proof.* The number of edges of the Voronoi diagram is $O(m)$, so to prove the lemma we must show that the total number of special breakpoints is $O(m)$. For each special breakpoint $p$, there is a Voronoi cell $\mathcal{V}(s_i)$ such that $p$ has two shortest paths $\alpha, \beta$ to $s_i$ so that the loop $\alpha \cup \beta$ encloses a hole of $\mathcal{V}(s_i)$. Because shortest paths do not cross (by the non-degeneracy assumption), two such cycles can never cross, though they may partially overlap. Thus we obtain a collection of non-crossing loops, each containing one or more holes (and thus sites) in its interior. These loops may nest, but each loop contains a different subset of sites. We can conclude that the overall number of loops—and, hence, the overall number of special breakpoints—is $O(m)$. □

We are now ready to prove a bound on the number of intersections between Voronoi edges and terrain edges and, hence, obtain our main result.

**Theorem 2.** *The complexity of the Voronoi diagram of $m$ point sites on a terrain $\mathcal{T}$ with $n$ triangles is $O(c_1 n + c_2 m\sqrt{n})$, where $c_1 = \xi + \lambda$ and $c_2 = \xi^2 \lambda\sigma\sqrt{\rho}$. Here $\xi$, $\lambda$, $\rho$, and $\sigma$ denote the slope, density, aspect ratio, and scale factor of $\mathcal{T}$, respectively.*

*Proof.* We already observed that, up to an additive $O(n + m)$ term, the complexity of the Voronoi diagram is bounded by the total number of intersections between Voronoi edges and terrain edges.

Let $\Gamma$ denote the set of subedges as defined above. Let $\gamma \in \Gamma$ be a subedge lying on the common boundary of Voronoi cells $\mathcal{V}(s_i)$ and $\mathcal{V}(s_j)$. For any terrain edge intersecting $\gamma$, at least one of the following two conditions holds: (i) it has an endpoint inside $\mathcal{V}(\gamma, s_i)$ or $\mathcal{V}(\gamma, s_j)$, or (ii) it intersects a shortest path from an endpoint of $\gamma$ to $s_i$ or $s_j$, i.e., the boundary of $\mathcal{V}(\gamma, s_i) \cup \mathcal{V}(\gamma, s_j)$. In Fig. 4, the edge $e$ illustrates case (i) and $e'$ case (ii). Let $E_{\mathcal{T}}(\gamma)$ denote the set of all terrain edges for which one of the two cases (i) and (ii) hold. From Lemmas 8 and 2 we can conclude that

$$\sum_{\gamma \in \Gamma} |E_{\mathcal{T}}(\gamma)| = O(n + cm\sqrt{n}),$$

where $c = \xi\lambda\sigma\sqrt{\rho}$, because every edge of $\mathcal{T}$ can contribute at most two to the count in case (i) and each shortest path contributes at most $O(c\sqrt{n})$ edges in

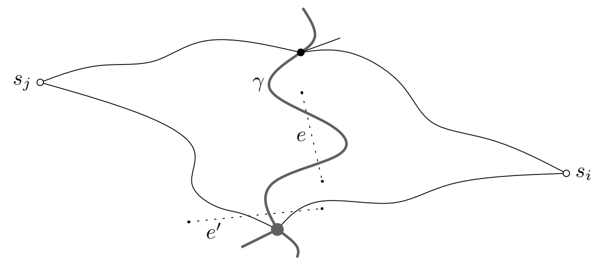


**Fig. 4.** Two terrain edges, $e$ and $e'$, intersecting a subedge $\gamma$.

case (ii). Now partition each subedge $\gamma$ into pieces, by adding all the breakpoints on $\gamma$; let $E_{\mathcal{VD}}(\gamma)$ denote the set of these pieces. Since the overall number of breakpoints is $O(n)$, we have $\sum_{\gamma \in \Gamma} |E_{\mathcal{VD}}(\gamma)| = O(n+m)$. Trivially, the total number of intersections between Voronoi edges and terrain edges is equal to the sum, over all $\gamma \in \Gamma$, of the number of intersections between pieces in $E_{\mathcal{VD}}(\gamma)$ and terrain edges in $E_{\mathcal{T}}(\gamma)$. Moreover, $E_{\mathcal{VD}}(\gamma)$ has density $O(\xi)$ by Theorem 1, and $E_{\mathcal{T}}(\gamma)$ has density $\lambda$ by definition. Hence, by Lemma 3, the number of intersections between pieces in $E_{\mathcal{VD}}(\gamma)$ and terrain edges in $E_{\mathcal{T}}(\gamma)$ is

$$O(\xi \cdot |E_{\mathcal{T}}(\gamma)| + \lambda \cdot |E_{\mathcal{VD}}(\gamma)|).$$

It follows that the total number of intersections is bounded by

$$\sum_{\gamma \in \Gamma} O(\xi \cdot |E_{\mathcal{T}}(\gamma)| + \lambda \cdot |E_{\mathcal{VD}}(\gamma)|) = O(\xi(n + cm\sqrt{n}) + \lambda(n+m)),$$

which gives the claimed bound. □

## 5 Conclusion

We proved tight bounds on the complexity of bisectors and Voronoi diagrams on the *realistic terrains* introduced by Moet *et al.* [8]. Even though our bounds are tight, there are still some interesting open questions. In particular, we have shown that the total number of intersections between bisector pieces and terrain edges is $O(n)$, but we suspect that this bound is not tight. Improving this bound will not lead to a better bound on the complexity of the bisector, since the number of bisector pieces can already be $\Omega(n)$. Nevertheless, a tight bound on the number of intersections between a bisector and terrain edges would give more insight into the global shape of bisectors. Moreover, an $O(\sqrt{n})$ bound on the number of triangle edges crossed by a bisector would immediately imply an $O(n + m\sqrt{n})$ bound on the complexity of the Voronoi diagram—something that requires a more involved argument in our current paper.

Another direction for further research is to see under what conditions one can prove an $O(n+m)$ bound on the complexity of the Voronoi diagram. For this one would probably also need to make assumptions on how the sites are distributed, and not only on the properties of the terrain.

As remarked in the introduction, Moet [7] studied some terrain models for mountainous regions in the US and found that the values of the four parameters defined in Section 2 are indeed bounded by a constant independent of the terrain size. Some of these values, however, are still fairly high. Usually this is caused by only a few triangles in the terrain. Hence, it would be interesting to obtain bounds that depend on the average slope of the triangles rather than the maximum slope.

*Acknowledgments.* Work by Boris Aronov has been partially supported by a grant from the U.S.-Israel Binational Science Foundation and by NSA MSP Grant H98230-06-1-0016. Research by Mark de Berg has been supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 639.023.301. Part of the research was carried out by Shripad Thite at TU/e, supported by NWO project no. 639.023.301, and by Boris Aronov while visiting TU/e in January 2007 and January 2008.

## References

1. F. Aurenhammer. Voronoi diagrams: A survey of a fundamental geometric data structure. *ACM Comput. Surv.* 23:345–405 (1991).
2. F. Aurenhammer and R. Klein. Voronoi diagrams. Chapter 5 in: J.-R. Sack and J. Urrutia (eds.), *Handbook of Computational Geometry*, Elsevier, 1999.
3. M. de Berg. Improved bounds for the union complexity of fat objects. *Discr. Comput. Geom*, 2008. In print.
4. M. de Berg, A. F. van der Stappen, J. Vleugels, and M. J. Katz. Realistic input models for geometric algorithms. *Algorithmica* 34:81–97 (2002).
5. J. Chen and Y. Han. Shortest paths on a polyhedron. *Int. J. Comput. Geom. Appl.* 6:127–144 (1996).
6. S. Fortune. Voronoi diagrams and Delaunay triangulations. Chapter 23 in: J.E. Goodman and J. O'Rourke (eds.), *Handbook of Discrete and Computational Geometry*, CRC Press, 2004.
7. E. Moet. *Computation and complexity of visibility in geometric environments.* PhD thesis, Utrecht University, 2008.
8. E. Moet, M. van Kreveld, and A.F. van der Stappen. On realistic terrains. In *Proc. 22nd ACM Sympos. Comput. Geom*, pages 177–186, 2006.
9. A. Okabe, B. Boots, and K. Sugihara. *Spatial tesselations: Concepts and applications of Voronoi diagrams.* John Wiley & Sons, 1992.
10. J.S.B. Mitchell, D.M. Mount, and C.H. Papadimitriou. The discrete geodesic problem. *SIAM J. Comput.* 16:647–668 (1987).
11. Y. Schreiber. Shortest paths on realistic polyhedra. In *Proc. 23rd ACM Sympos. Comput. Geom*, pages 74–83, 2007.
12. Y. Schreiber. Personal communication. April 2008.
13. A. F. van der Stappen. *Motion planning amidst fat obstacles.* Ph.D. thesis, Utrecht University, 1994.