Approximating Minimum-Weight Triangulations in Three Dimensions

Boris Aronov[1]
Computer and Information Science
Polytechnic University
Brooklyn, NY 11201
aronov@ziggy.poly.edu

Steven Fortune
Bell Laboratories
Murray Hill, New Jersey 07974
sjf@research.bell-labs.com

November 18, 1998

Abstract

Let S be a set of non-crossing triangular *obstacles* in $\mathbb{R}^3$ with convex hull H. A triangulation $\mathcal{T}$ of H is *compatible* with S if every triangle of S is the union of a subset of the faces of $\mathcal{T}$. The *weight* of $\mathcal{T}$ is the sum of the areas of the triangles of $\mathcal{T}$. We give a polynomial-time algorithm that computes a triangulation compatible with S whose weight is at most a constant times the weight of any compatible triangulation.

One motivation for studying minimum-weight triangulations is a connection with ray shooting. A particularly simple way to answer a ray-shooting query ("Report the first obstacle hit by a query ray") is to walk through a triangulation along the ray, stopping at the first obstacle. Under a reasonably natural distribution of query rays, the average cost of a ray-shooting query is proportional to triangulation weight. A similar connection exists for line-stabbing queries ("Report all obstacles hit by a query line").

1 Introduction

Let S be a finite set of non-crossing *obstacles* (line segments in $\mathbb{R}^2$, triangles in $\mathbb{R}^3$) with convex hull H. A triangulation $\mathcal{T}$ of H is *compatible* with S if each obstacle is the union of a subset of the faces of $\mathcal{T}$. $\mathcal{T}$ may have *Steiner vertices*, i.e., vertices that are not vertices of S. The *weight* of a facet f, $|f|$, is edge length in $\mathbb{R}^2$ and triangle area in $\mathbb{R}^3$); the *weight* of $\mathcal{T}$, $|\mathcal{T}|$, is the sum of the weights of its facets.

We give a polynomial-time algorithm that computes a triangulation $\mathcal{T}$ compatible with a three-dimensional obstacle set S. The weight of $\mathcal{T}$ is within a constant factor of the smallest possible compatible triangulation. The algorithm is a generalization of Eppstein's algorithm [9] to compute a constant-factor approximation to the minimum-length Steiner triangulation of a set of points in two dimensions. As with Eppstein's analysis, the approximation ratio is large, though constant.

The algorithm has two steps. The first step produces a depth-bounded octtree from the obstacles, where the rule is that an octtree cube is split if it meets an obstacle vertex or edge and is not

[1]Work of the first author was performed while he was visiting Bell Laboratories, Murray Hill, NJ.

too small. The second step triangulates each leaf cube in a manner compatible with the obstacles and neighboring leaf cubes. For the worst-case set of n obstacles, the algorithm runs in time $O(n^6)$ and produces as many tetrahedra. An improvement of an order of magnitude in both running time and number of tetrahedra is possible if the obstacle set is well-shaped (see section 3.3). A further improvement is possible if obstacles are just points (though this is not discussed in this paper).

The proof of the approximation ratio has two parts. The first part is to show that the total surface area of the octtree is at most a constant factor times the area of an arbitrary triangulation $\hat{\mathcal{T}}$. To do this we charge the surface area of each leaf cube to some local feature of $\hat{\mathcal{T}}$. There are essentially two cases: if the central subcube of the leaf cube meets a vertex, edge, or face of $\hat{\mathcal{T}}$, then the area of $\hat{\mathcal{T}}$ within the leaf cube must be at least proportional to its surface area. Otherwise the central subcube must be contained in a tetrahedron of $\hat{\mathcal{T}}$, some face of which must have area at least proportional to the surface area of the leaf cube. (The actual argument is more complex, to guarantee that a single tetrahedral face is not charged by too many leaf cubes.) The second part of the approximation-ratio bound is to triangulate all leaf cubes with total area at most a constant times the tree surface area plus the obstacle area.

We also briefly consider the two-dimensional case of compatible triangulations. As in three dimensions, there is a polynomial-time algorithm that approximates the minimum-weight compatible triangulation. In section 2 we show that the minimal Steiner triangulation weight is approximately the length of S plus the length of the MST (within a logarithmic factor). Here MST is the minimum Steiner spanning tree of S, i.e., the minimum-length set of line segments so that $MST \cup S$ is connected.

Other work. A long-standing open problem is the question of whether there is a polynomial-time algorithm that finds the minimum-length triangulation of a point set in two dimensions, without using additional Steiner points. Beirouti and Snoeyink [5] report recent progress and give many references. Eppstein [9] describes a polymial-time algorithm that gives a constant-factor approximation to the minimum-weight triangulation of a two-dimensional point set, allowing Steiner points.

Average-case line stabbing and ray shooting. We now describe a connection between triangulation weight and the average cost of simple algorithms for ray shooting ("Report the first obstacle hit by a ray") and line stabbing ("Report all obstacle hit by a line").

Line-stabbing queries can be answered in a particularly simple way given a triangulation of space compatible with the obstacles. It suffices to walk through the triangulation along the line, reporting each encountered obstacle. The walk takes constant time per visited triangle, so the total cost of the walk is proportional to the number of triangle faces crossed. In the worst case, the walk can be long. For example, Agarwal *et al* [2] describe a configuration of n obstacle triangles in $\mathbb{R}^3$ so that there is a line missing all the obstacles yet hitting $\Omega(n)$ faces of any triangulation compatible with the obstacles.

We consider instead the average-case cost of line stabbing, using the standard rigid-motion invariant measure μ on lines [16]. Let $L(U)$ be the set of lines that meet a set U. A basic fact from integral geometry is that for a facet f, $\mu(L(f))$ is the weight of f (length in two dimensions, area in three dimensions) times a constant depending upon dimension.

Let S be a set of obstacles so that any facet of its convex hull is the union of obstacles. Let $\mathcal{T}$ be a Steiner triangulation compatible with S. For line ℓ, let $s(\ell)$ and $t(\ell)$ be the number of obstacle facets and triangulation facets intersected by ℓ, respectively. Then

$$\frac{\int t(\ell)\ d\mu}{\mu(L(H))}$$

is the average walk length to answer a line-stabbing query that meets the convex hull H of S. The ratio

$$c(\mathcal{T}) = \frac{\int t(\ell)\ d\mu}{\int s(\ell)\ d\mu}$$

is the average walk length per reported obstacle facet. Letting $\delta_f(l)$ be 1 if line l hits facet f and 0 otherwise, we have that

$$c(\mathcal{T}) = \frac{\int t(\ell)\ d\mu}{\int s(\ell)\ d\mu} = \frac{\sum_{f\in\mathcal{T}} \int \delta_f(l)\ d\mu}{\sum_{f\in\mathcal{S}} \int \delta_f(l)\ d\mu} = \frac{\sum_{f\in\mathcal{T}} |f|}{\sum_{f\in\mathcal{S}} |f|} = \frac{|\mathcal{T}|}{|S|}.$$

A ray-shooting query can be answered by an algorithm similar to the line-stabbing algorithm: locate the endpoint of the ray in the triangulation, and walk along the ray through the triangulation until an obstacle facet is encountered. For a ray r, let $t(r)$ be the number of triangulation facets encountered in a walk along r up to an obstacle. Formally, $t(r)$ should include the facet containing the ray endpoint (if any) and should count all facets up to but not including the first obstacle facet. Clearly, $t(r)$ is proportional to the cost of answering a ray-shooting query, ignoring the cost of locating the triangle containing the endpoint of t.

The ratio $c(\mathcal{T})$ is also the average value of $t(r)$, for a particular distribution of query rays. Let ℓ be a directed line, and assume that the intersection of ℓ with each obstacle is either empty or a point. (Note that $c(\mathcal{T})$ is changed neither by assuming that lines are directed, nor by ignoring the measure-zero set of lines that overlap an obstacle.) Associate with ℓ the set of rays in the same direction as ℓ and with endpoints at an intersection of ℓ with an obstacle. Clearly, $s(\ell)$ is the number of such rays and $t(\ell) = \sum t(r)$, r in the associated set of rays. Hence

$$c(\mathcal{T}) = \frac{\int t(\ell)\ d\mu}{\int s(\ell)\ d\mu}$$

is the average of $t(r)$, under the distribution on rays induced by the distribution on lines. Using integral geometry [16, Section 12.7, eq. (12.60)], the induced ray distribution is $\sin\theta\ dA \wedge du$, where dA is the uniform area distribution, du is the uniform solid angle distribution, and θ is the angle between the ray and the surface of the obstacle. Informally, a ray is chosen with endpoint uniformly at random from an obstacle and with direction proportional to the sine of the angle θ between the ray and the obstacle.

Let

$$c(S) = \inf_{\hat{\mathcal{T}}}\ c(\hat{\mathcal{T}}),$$

where $\hat{\mathcal{T}}$ ranges over all triangulations compatible with S. Clearly, to answer line-stabbing or ray-shooting queries it is desirable to choose a triangulation $\mathcal{T}$ with $c(\mathcal{T})$ as close as possible to $c(S)$.

It is not obvious that the lower bound $c(S)$ can be attained [6]; for example, it is conceivable that it is always possible to decrease weight and hence $c(\mathcal{T})$ by adding Steiner points.

In two dimensions, $c(S) \approx (|MST| + |S|)/|S|$ (see section 2). Hence in cases where the MST is short, for example if the obstacle set is connected or nearly connected, the average cost of ray-shooting by walking through a triangulation should be small. This good behavior has been observed experimentally [10] (even without explicitly minimizing the weight of the triangulation).

Other work. The ray-shooting problem has been studied extensively in computational geometry (see [1] or [15] for a survey of theoretical results). Assuming roughly linear data-structure storage, the best theoretical algorithms for ray shooting have worst-case query time $O(\log n)$ for a simple polygon [11], roughly $O(\sqrt{n})$ for a set of planar line-segments [4], and roughly $O(n^{3/4})$ for a set of triangles in three dimensions [3]. The last two query times can be improved to $O(\log n)$ with polynomial storage. Agarwal *et al* [2] consider the line-stabbing number of triangulations consistent with a set of obstacles. Mitchell *et al* [12] consider segment shooting, a variant of ray shooting. They show that the cost of a segment-shooting query in an octtree can be bounded up to a constant factor by the "cover complexity" of the segment.

2 The two-dimensional case

Throughout this section S is a set of n planar obstacle segments that meet only at endpoints; S must include segments partitioning the boundary of its convex hull. A triangulation $\mathcal{T}$ is *compatible* with S if any edge in S is the union of closed edges of $\mathcal{T}$. Vertical bars $|\cdot|$ denote length, thus $|S|$ is the sum of the lengths of the segments in S. Let M be $\inf_{\mathcal{T}} |\mathcal{T}|$, where $\mathcal{T}$ varies over triangulations compatible with S.

Lemma 2.1 *In polynomial time it is possible to compute a polynomial-size Steiner triangulation $\mathcal{T}$ compatible with S so that $|\mathcal{T}| = O(M)$.*

We omit a detailed proof of lemma 2.1; it can be proven using techniques of Eppstein [9] or of the proof in section 3. The basic strategy is to build a depth-bounded balanced quadtree using only the vertices of S, and then triangulate each square in a fashion compatible with the edges of S meeting the square. We remark that lemma 2.1 depends upon including the length of the obstacle set S in the length of the triangulation; if the length is not included, a constant-factor polynomial-time approximation algorithm is not known [9].

Let MST be a *minimum (Steiner) spanning tree* of S, i.e., a set of segments of minimum total length so that $S \cup \mathrm{MST}$ is connected. An easy compactness argument shows that the minimum length can actually be attained.

Lemma 2.2 $|S| + |MST| \leq M \leq O((|S| + |MST|) \log n)$.

Proof: The first inequality is obvious, since any triangulation compatible with S must be connected. For the second, notice that $\mathrm{MST} \cup S$ partitions the convex hull of S into simple polygons (recall that S includes the convex hull boundary). The total number of vertices is $O(n)$, since the MST is a forest with no vertices of degree 2 and the number of leaves is the number of connected components of S. By a result of Clarkson [7], each simple polygon can be triangulated with weight proportional to $\log n$ times the perimeter of the polygon. $\square$

Let D be the diameter of S, i.e., the length of the longest segment contained in the convex hull of S. The minimum spanning tree of the vertices of S has length at most $O(D\sqrt{n})$, hence $|\mathrm{MST}| \le O(D\sqrt{n})$. The following lemma improves the worst-case bound that can be obtained from lemma 2.2 and this estimate by a factor of $\log n$. Again we omit a detailed proof of the lemma (see section 3.7 for a similar proof in three dimensions).

Lemma 2.3 *$M \le O(|S| + D\sqrt{n})$.*

Corollary 2.4 *$|MST|/|S| + 1 \le c(S) \le \min(O((1 + |MST|/|S|)\log n), O(1 + D\sqrt{n}/|S|))$.*

3 The three-dimensional case

This section describes an algorithm that produces a Steiner triangulation compatible with a set of polyhedral obstacles in three dimensions. The triangulation has area within a constant factor of the smallest possible. Section 3.1 reviews some basic definitions. Section 3.2 states the main theorem in the case that the obstacles are "wide," a condition on the aspect ratio of their convex hull. The main theorem is proved in sections 3.3 through 3.5. Section 3.3 gives an algorithm that constructs an octtree from a set of polyhedral obstacles. In section 3.4 we prove that the surface area of the octtree is at most a constant factor times the area of any Steiner triangulation compatible with the obstacles, while section 3.5 gives an algorithm to triangulate the octtree with total area proportional to the area of the octtree plus the area of the obstacles. The wideness condition on obstacles is removed in section 3.6. Finally in section 3.7 we show some worst-case bounds on the ratio of triangulation area to obstacle area.

3.1 Definitions

We use terminology from the theory of convex polyhedra [8]. For example, a *polyhedral set* is any set obtained from open and closed halfspaces by a finite number of unions and intersections. A *face* of a convex polyhedron C is the relative interior of the intersection of the closure of C with a hyperplane supporting C; ∂C is the relative boundary of C.

A *polyhedral subdivision* $\mathcal{Q}$ of a polyhedral set Q is a finite partition of Q into relatively open convex polyhedral *cells* so that every face of every cell is a union of cells in $\mathcal{Q}$. A subdivision is *proper* if every face of every cell is itself a cell in the subdivision. Not all subdivisions are proper, since some face of a cell may be subdivided into more than one cell. If f is a face of a cell C, we will occasionally refer to f as a *polyhedral face* of C to distinguish it from a cell contained in ∂C that is not a face of C. The k-skeleton $\mathcal{Q}^{(k)}$ of $\mathcal{Q}$ is the subdivision consisting of all cells of $\mathcal{Q}$ of dimension at most k. If C is a convex polytope that is the union of cells of $\mathcal{Q}$, then $\mathrm{bdry}(C, \mathcal{Q})$ is the subdivision consisting of the cells in $\mathcal{Q}$ whose union is ∂C. For convenience we define $\mathrm{bdry}(\mathcal{Q})$ to be $\mathrm{bdry}(Q, \mathcal{Q})$. The *area* of a subdivision $\mathcal{Q}$, $\mathrm{area}(\mathcal{Q})$, is the sum of the areas of the 2-cells in $\mathcal{Q}$;

similarly the *length* of $\mathcal{Q}$, length($\mathcal{Q}$), is the sum of the lengths of the 1-cells in $\mathcal{Q}$. A *triangulation* of a polyhedral set is a proper polyhedral subdivision so that all cells are simplices. A triangulation $\mathcal{T}$ is *compatible* with a subdivision $\mathcal{S}$ if every cell in $\mathcal{S}$ is a union of cells in $\mathcal{T}$.

If C is a square or cube and k a positive real, then kC is the square or cube with the same center and orientation as C and side length k times the side length of C. We write for example $C/2$ for $(1/2)C$.

3.2 The main theorem for wide obstacles

The *obstacle set* S is a subdivision in $\mathbb{R}^3$ consisting only of 0-, 1-, and 2-simplices. For simplicity we assume that there are no isolated edges and vertices in S, i.e., every edge or vertex lies on the boundary of a 2-simplex; with minor modifications the algorithm can be extended to handle isolated edges and vertices. We let H be the convex hull of S; we assume that S contains a triangulation of ∂H. Throughout n is the total number of cells in S.

Let B be the smallest cube containing S, and let its side length be b. Obstacle set S is *wide* if the area of ∂H is $\Omega(b^2)$. Informally, H can look like a ball or pancake, but not like a pencil. The following theorem holds for wide obstacles; it follows immediately from theorem 3.2 in section 3.4 and theorem 3.9 in section 3.5, using the fact that the area of ∂H and hence S is $\Omega(b^2)$. The wideness assumption is removed in section 3.6.

Theorem 3.1 *Let S be a wide set of obstacles with a total of n simplices. In time $O(n^5)$ it is possible to compute a triangulation compatible with S whose area is within a constant factor of the smallest possible. The triangulation has $O(n^5)$ tetrahedra and partitions the convex hull of S.*

3.3 The octtree algorithm

We now describe how to build an octtree T from the obstacle set S. If necessary, perturb the minimal enclosing cube B slightly so that no obstacle face is parallel to a face of B. Recall that b is the side length of B. Let m be the first power of two greater than or equal to n. Define $s_0 = b/m$. Then b/s_0 is a power of two and $b/(2n) < s_0 \leq b/n$.

We first build an octtree T_0. Every node of the tree is a cube. B is the root node. A cube is *skewered* if it meets the 1-skeleton of S. A cube C is subdivided (i.e., is not a leaf) if it is skewered and has side length greater than s_0; its children are the eight subcubes obtained by cutting it by the three planes through its center parallel to its facets. This process is repeated until no cube can be further subdivided. Clearly, each leaf cube in the octtree has side length $2^i s_0$, for some integral $i \geq 0$, and each skewered leaf cube has side length s_0.

Octtree T_0 must then be made balanced, resulting in the octtree T. Two cubes are *adjacent* if they have overlapping 2-faces; an octtree is *balanced* if the side lengths of any two adjacent leaf cubes differ by at most two. If an octtree is unbalanced, it can be made balanced by repeatedly choosing a pair of adjacent cubes that violate the balance property and subdividing the larger cube. Standard results [13] imply that balancing T increases the number of cubes by at most a constant factor.

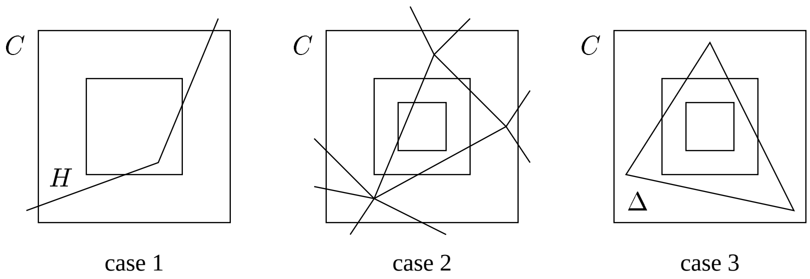

Figure 1: Cases of theorem 3.2. In case 1, ∂H meets $C/2$; in case 2, the 2-skeleton of $\hat{\mathcal{T}}$ meets $C/4$; in case 3, the 2-skeleton of $\hat{\mathcal{T}}$ avoids $C/4$.

An obstacle triangle Δ *fully cuts* a cube C if Δ intersects C and $\partial\Delta$ avoids C. Notice that a cube may be fully cut by an obstacle triangle independently of whether or not it is skewered.

3.4 Area of the octtree

The *area* of octtree T, $\mathrm{area}(T)$, is $\sum_C \mathrm{area}(\partial C \cap H)$, where C varies over the leaf cubes of T. Let $M = \inf_{\hat{\mathcal{T}}} \mathrm{area}(\hat{\mathcal{T}})$, where $\hat{\mathcal{T}}$ ranges over triangulations compatible with S.

Theorem 3.2 $\mathrm{area}(T) \leq c \cdot M$*, for some absolute constant* c*.*

Proof: Let $\hat{\mathcal{T}}$ be an arbitrary triangulation compatible with S and let C be a leaf cube of T. We can assume $C \cap H$ is not empty. We charge $\mathrm{area}(\partial C \cap H)$ to features of $\hat{\mathcal{T}}$, with cases as follows (see figure 1):

1. $C/2 \not\subseteq H$. Necessarily ∂H meets $C/2$ (and C). By lemma 3.3 below, $\mathrm{area}(\partial C \cap H)$ is $O(\mathrm{area}(\partial H \cap C))$. We charge $\mathrm{area}(\partial C \cap H)$ to $\partial H \cap C$.

2. $C/2 \subseteq H$ and $C/4$ meets the 2-skeleton of $\hat{\mathcal{T}}$. By lemma 3.4 below, $\mathrm{area}(\partial C \cap H)$ is $O(\mathrm{area}(\hat{\mathcal{T}} \cap C))$. We charge $\mathrm{area}(\partial C \cap H)$ to $\hat{\mathcal{T}} \cap C$.

3. $C/2 \subseteq H$ and $C/4$ avoids the 2-skeleton of $\hat{\mathcal{T}}$. This is the most complex case. Let α, β, and γ be constants (chosen in the proof of lemma 3.6 below, with $0 < \alpha, \beta < 1$ and $\gamma > 1$). Let the side length of C be s. A *charging pair* (Δ, e) for C is a triangle Δ of $\hat{\mathcal{T}}$ and an edge e of Δ so that $\mathrm{area}(\Delta \cap \gamma C) \geq \alpha s^2$ and $\mathrm{length}(e \cap \gamma C) \geq \beta s$. We charge the area of $\partial C \cap H$ (which is $O(s^2)$) to the triangle of a charging pair, whose existence is guaranteed by lemma 3.6 below. By lemma 3.8 the total charge to any triangle Δ of $\hat{\mathcal{T}}$, over all leaf cubes of $\mathcal{T}$, is at most $O(\mathrm{area}(\Delta))$.

In cases 1 and 2, the charges are to disjoint portions of ∂H and $\hat{\mathcal{T}}$, respectively, so the total charge is $O(\mathrm{area}(\partial H) + \mathrm{area}(\hat{\mathcal{T}}))$ which is $O(\mathrm{area}(\hat{\mathcal{T}}))$. In case 3, the total charge to any triangle Δ of $\hat{\mathcal{T}}$ is $O(\mathrm{area}(\Delta))$, so the total charge over all triangles is $O(\mathrm{area}(\hat{\mathcal{T}}))$. $\square$

Lemma 3.3 *If* C *is a leaf cube and* $C/2 \not\subseteq H$*, then* $\mathrm{area}(\partial C \cap H)$ *is* $O(\mathrm{area}(\partial H \cap C))$*.*

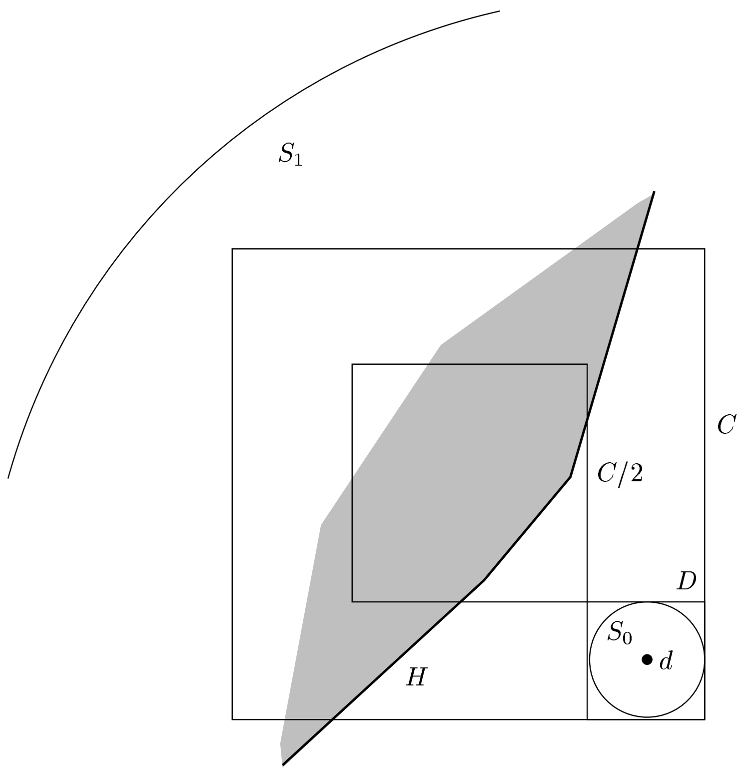

Figure 2: Proof of Lemma 3.3.

Proof: Let s be the side length of C. Pair each corner of $C/2$ with the corresponding corner of C, and consider the 8 subcubes whose opposite corners are formed by the pairs. At least one such subcube D must be entirely outside H, else H would contain $C/2$. Let d be the center of D. See figure 2.

Let S_0 and S_1 be the spheres of radius $s/8$ and $\sqrt{3}s$, respectively, with center d. S_0 is contained in D and is outside H, and S_1 contains C. Let mappings π_0 and π_1 be central projections from d into S_0 and S_1, respectively. Any ray from d to a point of $\partial C \cap H$ must first hit a point of $\partial H \cap C$, since d is outside H and inside C. Hence $\pi_1(\partial C \cap H) \subseteq \pi_1(\partial H \cap C)$. For some constant c, we have

$$c \cdot \mathrm{area}(\partial C \cap H) \leq \mathrm{area}(\pi_1(\partial C \cap H))$$

since for any point $p \in \partial C \cap H$, the angle that a ray $\vec{dp}$ makes with $\partial C \cap H$ is bounded away from zero, as d is bounded away from ∂C. We also have

$$\mathrm{area}(\pi_1(\partial C \cap H)) \leq \mathrm{area}(\pi_1(\partial H \cap C)) = (8/\sqrt{3})^2 \cdot \mathrm{area}(\pi_0(\partial H \cap C)) \leq 64/3 \cdot \mathrm{area}(\partial H \cap C),$$

proving the lemma. □

Lemma 3.4 *If C is a leaf cube and $B = C/2 \subseteq H$ and $B/2 = C/4$ meets $\hat{\mathcal{T}}^{(2)}$ (the 2-skeleton of $\hat{\mathcal{T}}$), then* $\mathrm{area}(\partial C \cap H)$ *is* $O(\mathrm{area}(\hat{\mathcal{T}} \cap C))$.

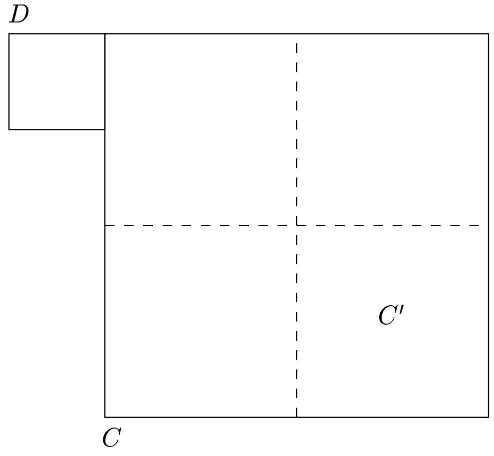

Figure 3: Proof of lemma 3.5.

Proof: Let s be the edge length of B. Clearly area$(\partial C \cap H)$ is $O(s^2)$; we show area$(\hat{\mathcal{T}} \cap B)$ is $\Omega(s^2)$. Let D be the open cube of side length $s/2$ centered at a point x of the 2-skeleton of $\hat{\mathcal{T}}$ in $B/2$, with the same orientation as $B/2$; plainly $D \subset B$. Choose three orthogonal edges e_1, e_2, e_3 of ∂D, and for any $p \in D$ let $l_i = l_i(p)$, $i = 1, 2, 3$, be the segment through p parallel to e_i connecting opposite faces of D. We claim at least one of l_1, l_2, l_3 must meet $\hat{\mathcal{T}}^{(2)}$. Suppose not, then they must all lie in the same tetrahedron Δ of $\hat{\mathcal{T}}$, since they share point p. Hence the convex hull of l_1, l_2, l_3 is contained in Δ, a contradiction since the convex hull contains x, which is in $\hat{\mathcal{T}}^{(2)}$.

Let D_i, $i = 1, 2, 3$, be the set of points $p \in D$ for which $l_i(p)$ meets $\hat{\mathcal{T}}^{(2)}$. Since $D = D_1 \cup D_2 \cup D_3$, we can assume, say, that the volume of D_1 is at least $(s/2)^3/3 = s^3/24$. Now D_1 must be the Cartesian product of a segment of length $s/2$ parallel to e_1 with the projection of $\hat{\mathcal{T}}^{(2)} \cap D$ onto a facet of D perpendicular to e_1. Hence the projection of $\hat{\mathcal{T}}^{(2)}$ has area at least $s^2/12$, and $\hat{\mathcal{T}}^{(2)} \cap D$ itself has at least the same area. $\Box$

Lemma 3.5 *For any leaf cube C of T, $7C$ meets the 1-skeleton of S.*

Proof: Let a *refinement step* during balancing be the replacement of a leaf node by an interior node with 8 children. We show that the lemma holds for T_0 and is maintained by every refinement step.

By construction, every skewered leaf in T_0 meets the 1-skeleton of S. Every unskewered leaf C in T_0 has a skewered parent, so certainly $3C$ meets the 1-skeleton of S.

A refinement step maintains the condition of the lemma for every unchanged cube. So suppose C is refined into 8 subcubes, with s the original side length of C and C' an arbitrary child of C. See figure 3. C must share a common 2-face with a cube D whose side length is at most $s/4$. By inductive hypothesis, $7D$ meets the 1-skeleton of S. Easy calculations show that the L_∞ distance between the centers of C' and D is at most $7s/8$ and that $7D \subset 7C'$. Hence $7C'$ meets the 1-skeleton of S. $\Box$

Lemma 3.6 *If C is a leaf cube, $B = C/2 \subseteq H$, and $B/2$ avoids the 2-skeleton of $\hat{\mathcal{T}}$, then there is*

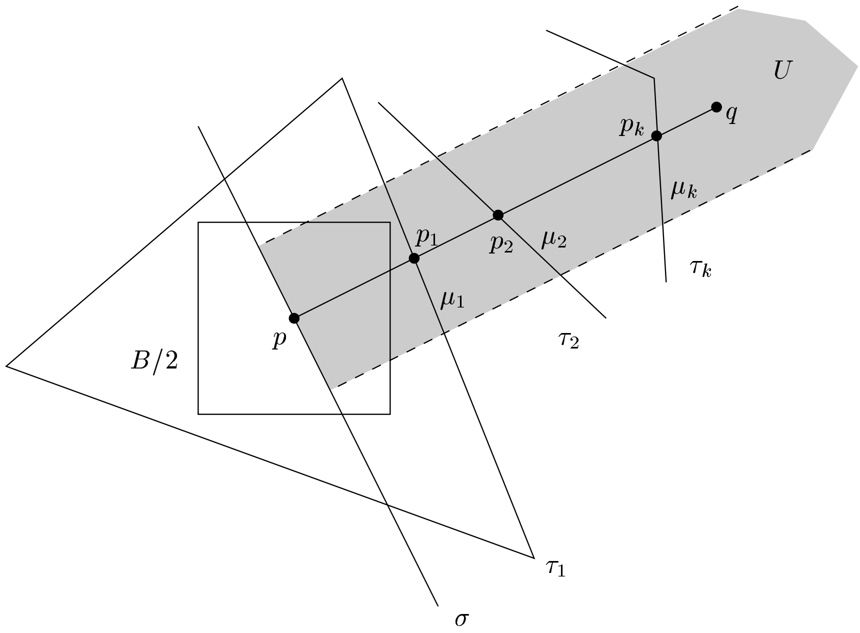

Figure 4: Proof of Lemma 3.6.

a charging pair for C.

Proof: Let s be the side length of B, so $B/2$ has side length $s/2$. By lemma 3.5 there is a point q on the 1-skeleton of $\hat{\mathcal{T}}$ within $7C = 14B$, hence at distance $7\sqrt{3}s < 14s$ from the center p of B. We assume that open segment pq meets the 2-skeleton of $\hat{\mathcal{T}}$ only at interior points of 2-cells; otherwise a slightly modified argument with a perturbed segment $p'q$ is necessary.

Let σ be the plane orthogonal to pq through p and let π_σ be orthogonal projection onto σ. Let U be the infinite cylinder with axis pq and radius $s/4$; then the disc $D = U \cap \sigma$ is contained in $B/2$. Let $p_1, \ldots, p_{l-1}$ be the intersections of 2-cells of $\hat{\mathcal{T}}$ with open segment pq in order from p to q, and set $p_0 = p$ and $p_l = q$. For $i = 1, \ldots, l$, let τ_i be the 3-simplex of $\hat{\mathcal{T}}$ containing the open segment $p_{i-1}p_i$. See figure 4. We have that $B/2 \subseteq \tau_1$, since $p_0 = p \in B/2$ and the 2-skeleton of $\hat{\mathcal{T}}$ avoids $B/2$.

Let μ_i be the connected component of $\partial\tau_i \cap U$ containing p_i. Choose k minimal so that μ_k contains a point of the 1-skeleton of $\hat{\mathcal{T}}$; such k must exist since $p_l = q$ is on the 1-skeleton of $\hat{\mathcal{T}}$.

We claim that μ_i *fully cuts* U, i.e., $\pi_\sigma(\mu_i) = D$, for $i = 1, \ldots, k$. Now μ_1 fully cuts U, since $D \subseteq B/2 \subseteq \tau_1$. For $1 < i \leq k$, μ_{i-1} must be a portion of a single 2-face of τ_i, since it contains no point on the 1-skeleton of $\hat{\mathcal{T}}$. Hence there must be two connected components to $U \cap \partial\tau_i$, specifically μ_{i-1} and μ_i, and μ_i must fully cut U.

We now show that μ_k lies at distance at most $28s$ from σ. Let f_{k-1} and f_k be the faces of τ_k containing p_{k-1} and p_k, and let P_{k-1} and P_k be the planes containing f_{k-1} and f_k, respectively. Now $P_{k-1} \cap D = P_{k-1} \cap U \cap \sigma$ must be empty, since $P_{k-1} \cap U = \mu_{k-1}$ lies entirely in face f_{k-1} but $D = U \cap \sigma$ lies entirely in the interior of simplex τ_1. Also $P_k \cap D$ must be empty: if not, choose $u \in P_k \cap D$; then p_k and u lie on opposite sides of P_{k-1} within U, so segment up_k meets an interior point of face f_{k-1} of τ_k; however up_k also lies entirely in the plane P_k of face f_k of τ_k, a contradiction. The intersection of P_k with ∂U is an ellipse, whose center p_k is at distance at most $14s$ from σ. The ellipse avoids D, since P_k avoids D. Hence the farthest point of the ellipse is at most $28s$ from σ. Since the closed half-space of P_k containing D also contains τ_k and μ_k, μ_k is at most distance $28s$ from σ.

We set the charging-pair constants $\alpha = \pi/48$, $\beta = 1/8$, and $\gamma = 43$. To finish the proof, notice that $\text{area}(\mu_k) \geq \text{area}(\pi_\sigma(\mu_k)) = \pi(s/4)^2$. Since μ_k consists of portions of at most three triangular faces of τ_k, at least one such triangle Δ satisfies $\text{area}(\Delta \cap U) \geq \pi s^2/48 = \alpha s^2$. Δ is the triangle of $\hat{\mathcal{T}}$ promised by the lemma statement. We have $\Delta \cap U$ within distance $28s$ from σ, in fact within distance $28s$ from $\sigma \cap U \subset B/2 \subset B$. Since the side length of B is s, $\Delta \cap U \subset (2 \cdot 28 + 1)B \subset \gamma C$.

It remains to find a piece of an edge of Δ of length at least βs within γC. Let U' be the cylinder coaxial with U but with double the radius. By analogous reasoning, we have $\Delta \cap U'$ within distance $42s$ from $\sigma \cap U' \subset B$, so $\Delta \cap U' \subset (2 \cdot 42 + 1)B \subset \gamma C$. If Δ is completely contained within U', then by the isoperimetric inequality [16] the perimeter L of Δ satisfies $L \geq \sqrt{4\pi \cdot \text{area}(\Delta)} \geq \pi s / 2\sqrt{3}$, so one edge of Δ has length at least $\pi s / 6\sqrt{3} \geq \beta s$. Otherwise $\partial\Delta$ intersects both $\partial U'$ and the interior of U. Since ∂U and $\partial U'$ are separated by $s/4$, there is a portion of an edge of $\partial\Delta$ lying within U' of length at least $s/8 = \beta s$. $\Box$

The *width* of a compact planar set S, $\text{width}(S)$, is the smallest distance w so that S is contained in the closed region between two parallel lines w apart. The *diameter* of S, $\text{diam}(S)$, is the length of the longest segment contained in the convex hull of S. It is easy to see that the area of S is at most $\text{width}(S) \cdot \text{diam}(S)$.

Lemma 3.7 *Let Δ be an obstacle triangle and r the radius of its inscribed circle. If leaf cube C with side length s has charging pair (Δ, e), then s is $O(r)$.*

Proof: The width of Δ is at most $3r$, by some elementary geometry. Since (Δ, e) is a charging pair for C, we have

$$\begin{aligned} \alpha s^2 &\leq \text{area}(\Delta \cap \gamma C) \\ &\leq \text{width}(\Delta \cap \gamma C) \cdot \text{diam}(\Delta \cap \gamma C) \\ &\leq 3r \cdot \sqrt{3}\gamma s. \end{aligned}$$

Hence $s \leq 3\sqrt{3}\gamma r/\alpha$, which is $O(r)$. $\Box$

Lemma 3.8 *The total charge to triangle Δ of $\hat{\mathcal{T}}$ from all leaf cubes in case 3 above is $O(\text{area}(\Delta))$.*

Proof: Let p and r be the perimeter and inscribed-circle radius of Δ, respectively. Some elementary geometry gives $\text{area}(\Delta) = pr/2$.

Consider the set C_s of all leaf cubes with a fixed side length s. Since no two cubes in C_s overlap, no point in space is covered by more than γ^3 open supercubes γC, for $C \in C_s$. Since the intersection of γC with $\partial\Delta$ has length at least βs, no more than $\gamma^3 p/(\beta s) = O(p/s)$ cubes in C_s

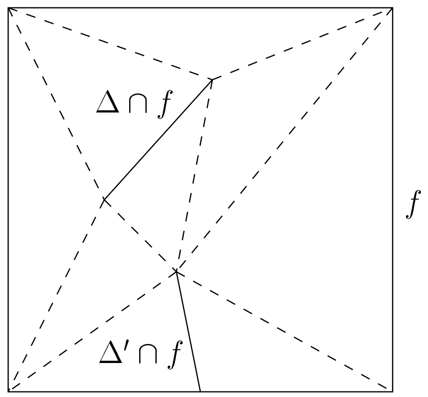

Figure 5: Δ and Δ' are obstacle triangles that meet minimal cube 2-face f. In $\mathcal{P}$, f is replaced by a triangulation of f compatible with $\Delta \cap f$ and $\Delta' \cap f$.

can be charged to Δ. The charge for each cube is $O(s^2)$, so the total charge to Δ for cubes in C_s is $O(ps)$.

Let $\hat{s}$ be the largest cube side length for which a cube is charged to Δ; by lemma 3.7, $\hat{s}$ is $O(r)$. The total charge over all cube sizes is $O(p(s_0 + 2s_0 + \ldots + \hat{s})) = O(p\hat{s}) = O(pr)$, which is $O(\text{area}(\Delta))$. $\square$

3.5 Triangulating the tree

Theorem 3.9 *In time $O(n^5)$ it is possible to construct a triangulation $\mathcal{T}$ from the octtree T so that* $\text{area}(\mathcal{T}) \leq c' \cdot (\text{area}(T) + \text{area}(S) + b^2)$, *where c' is an absolute constant and b is the side-length of the root cube of T. $\mathcal{T}$ has $O(n^5)$ tetrahedra and partitions the convex hull H of S.*

For C a leaf cube of T, the *clipped cube* C_H is $C \cap H$. We choose below a subdivision $\mathcal{P}$ of H whose 3-cells are clipped-cube interiors and whose 2-skeleton is a triangulation. In sections 3.5.2 and 3.5.3 we show how to triangulate a clipped cube so that its boundary triangulation matches $\mathcal{P}$. The desired triangulation $\mathcal{T}$ is then obtained simply by pasting the clipped-cube triangulations into $\mathcal{P}$.

In octtree T, two adjacent cubes may have overlapping but distinct k-faces, $k = 1, 2$, if the cubes are of different sizes. A clipped-cube k-face is *minimal* if no other clipped-cube face of the same dimension is properly contained within it; a minimal k-face is a *tree-partitioning* face if it is properly contained in another clipped-cube face (possibly of higher dimension). It is easy to see that any clipped-cube k-face is the union of minimal faces of dimension at most k. In fact, since T is balanced, any 1-face is the union of at most four minimal 1-faces and three vertices, and any 2-face is the union of at most four minimal 2-faces, four minimal 1-faces, and a vertex (forming a '+'-shape in the middle of the 2-face).

Let $\mathcal{P}'$ consist of clipped-cube interiors plus all minimal 0-, 1-, and 2-faces. Subdivision $\mathcal{P}$ is obtained from $\mathcal{P}'$ by replacing 1- and 2-faces. A 1-face e of $\mathcal{P}'$ is replaced with the chain of 0- and 1-cells formed by subdividing e at each point of intersection with an obstacle triangle. Let f be a

2-face of $\mathcal{P}'$. Form the set of line segments obtained by intersecting f with all obstacle triangles, then triangulate this set together with the set of 0- and 1-cells partitioning ∂f. The resulting triangulation replaces f in $\mathcal{P}$. See figure 5. It is easy to check that $\mathcal{P}$ is a subdivision of H whose 2-skeleton has been triangulated.

The *combinatorial complexity* of leaf cube C, n_C, is the number of edges of C plus the number of obstacle triangles meeting C. It is easy to check that $\mathrm{bdry}(C_H, \mathcal{P})$ has $O(n_C)$ edges and vertices.

3.5.1 Central triangulations

Let Q be a 3-cell in a polyhedral subdivision $\mathcal{Q}$ with $\mathrm{bdry}(Q, \mathcal{Q})$ a triangulation and let q be an interior point of Q. Recall that by the definition of polyhedral subdivision, Q is convex. The *central triangulation* of Q from q consists of the tetrahedra formed by q and the triangles in $\mathrm{bdry}(Q, \mathcal{Q})$, and all tetrahedral faces.

Proposition 3.10 *If* $\mathrm{bdry}(Q, \mathcal{Q})$ *has* k *vertices, then the area of any central triangulation of* Q *is at most* $3k/2$ *times the area of* ∂Q*.*

Proof: Let q be the central triangulation vertex. There are at most $3k$ edges of $\mathrm{bdry}(Q, \mathcal{Q})$. Each new tetrahedral 2-face is formed by q and such an edge, and has area at most half the area of ∂Q. $\square$

3.5.2 Triangulating an unskewered cube

For this section, C is an unskewered leaf cube (i.e., C avoids the 1-skeleton of S). A triangulation $\mathcal{T}_C$ of C_H is obtained as follows. Start with the subdivision consisting of C_H and $\mathrm{bdry}(C_H, \mathcal{P})$. Subdivide C_H by all obstacle triangles that meet it. Notice that all such obstacle triangles must cut C fully and cannot meet within C; furthermore $\mathrm{bdry}(C_H, \mathcal{P})$ already contains all edges of intersection between obstacle triangles and ∂C_H. Now triangulate any new 2-cells, and centrally triangulate each 3-cell.

Lemma 3.11 $\mathcal{T}_C$ *has area* $O(\mathrm{area}(S \cap C) + \mathrm{area}(\partial C \cap H))$*.*

Proof: Let $\mathcal{T}_0$ be the subdivision before central triangulation, and let C_0 be a 3-cell of $\mathcal{T}_0$. We show that the number of vertices in $\mathrm{bdry}(C_0, \mathcal{T}_0)$ is constant. The lemma then follows using Proposition 3.10 and the fact that $\mathrm{area}(\mathcal{T}_0)$ is $\mathrm{area}(S \cap C) + \mathrm{area}(\partial C \cap H)$.

First notice that C_0 has at most 14 polyhedral 2-faces, 6 faces that are subsets of a 2-face of C and at most 8 faces separating a vertex of C from C_0. Hence C_0 has a constant number of polyhedral vertices. Any vertex of $\mathrm{bdry}(C_0, \mathcal{T}_0)$ that is not a polyhedral vertex must be a cube partitioning vertex, of which there are at most a constant number. $\square$

3.5.3 Triangulating a skewered cube

For this section, C is a skewered leaf cube (i.e., C meets the 1-skeleton of S); recall C has side length s_0. A triangulation $\mathcal{T}_C$ of C_H is obtained in two steps. We first compute a triangulation

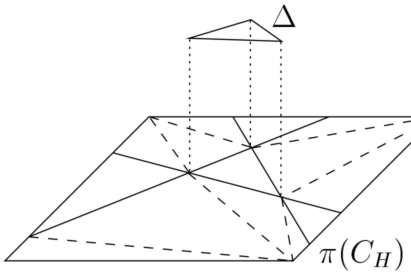

Figure 6: The edges of obstacle Δ are projected onto a horizontal plane, extended to lines, and clipped to the projection of C_H (labeled $\pi(C_H)$). The result is triangulated and each triangle is lifted to an infinite vertical prism.

$\mathcal{T}'_C$ so that $\mathrm{bdry}(\mathcal{T}'_C) = \mathrm{bdry}(C_H, \mathcal{T}'_C)$ is a refinement of $\mathrm{bdry}(C_H, \mathcal{P})$. Using lemma 3.15, we then compute a triangulation $\mathcal{T}_C$ so that $\mathrm{bdry}(\mathcal{T}_C) = \mathrm{bdry}(C_H, \mathcal{P})$.

The first step has four substeps. In the following, the "vertical" direction can be chosen to be any direction not parallel to a face of C_H or a face of S and so that no vertical line meets three obstacle edges that do not already meet at a common vertex. See figure 6.

1. Orthogonally project each edge of $\mathrm{bdry}(C_H, \mathcal{P})$ onto a horizontal plane (i.e., a plane orthogonal to the vertical direction), and extend the projection to a line. Similarly project and extend each obstacle edge meeting C.

2. Let $\mathcal{T}_2$ be a (two-dimensional) triangulation of the resulting arrangement, truncated to the projection of C_H. Lift $\mathcal{T}_2$ to a set of infinite vertical triangular prisms.

3. Subdivide C_H using both the vertical prisms and any obstacle triangles meeting C, forming a subdivision $\mathcal{T}_3$.

4. Extend the triangulation of $\mathrm{bdry}(C_H, \mathcal{P})$ to form a triangulation of the 2-skeleton of $\mathcal{T}_3$. Then centrally triangulate each 3-cell, forming the triangulation $\mathcal{T}'_C$.

Lemma 3.12 *The total length of $\mathcal{T}_2$ is $O(n_C s_0)$.*

Proof: Let ℓ be a line in the plane and let $t(\ell)$ be the number of edges of $\mathcal{T}_2$ met by ℓ. Consider the arrangement (in substep 2 above) before triangulation and truncation. By the zone theorem for lines [8], the total combinatorial complexity of the cells of the arrangement intersected by a line ℓ is $O(n_C)$, so $t(\ell)$ is $O(n_C)$.

Let μ be the rigid-motion invariant measure on sets of lines in the plane [16]; up to a constant multiple, the length of $\mathcal{T}_2$ is $\int t(\ell)\, d\mu$. The measure of the set of line intersecting $\mathcal{T}_2$ is $O(s_0)$, since the perimeter of $\mathcal{T}_2$ is $O(s_0)$ [16]. Hence $\int t(\ell)\, d\mu$ is bounded by $O(s_0 n_C)$. $\square$

Lemma 3.13 *Triangulation $\mathcal{T}'_C$ has area $O(\mathrm{area}(S \cap C) + \mathrm{area}(\partial C \cap H) + n_C s_0^2)$.*

Proof: The total area of the vertical 2-cells in the subdivision $\mathcal{T}_3$ is $O(n_C s_0^2)$, since the total length of $\mathcal{T}_2$ is $O(n_C s_0)$ and the height of C_H at most $O(s_0)$. Any nonvertical 2-cell is a portion of either $S \cap C$ or $\partial C \cap H$, hence the total area of all nonvertical 2-cells is $\text{area}(S \cap C) + \text{area}(\partial C \cap H)$. Let P be a 3-cell in $\mathcal{T}_3$; we show that the number of vertices in $\text{bdry}(P, \mathcal{T}_3)$ is constant. The lemma then follows using proposition 3.10 and the bound on the sum of the areas of all 2-cells.

Cell P is a section of a triangular prism and has at most 5 polyhedral facets, hence a constant number of polyhedral vertices. We claim $\text{bdry}(P, \mathcal{T}_3)$ has at most two additional vertices per vertical polyhedral edge of P. To see this, note that by construction of $\mathcal{T}_3$, any vertex of $\text{bdry}(P, \mathcal{T}_3)$ must either be a vertex of $\text{bdry}(C_H, \mathcal{T}_3)$, and hence a polyhedral vertex of P, or must lie in the interior of a collinear vertical chain of edges of $\mathcal{T}_3$. In the latter case, the vertex must either lie in the interior of an obstacle triangle, and hence must be a polyhedral vertex of P, or on the closure of an obstacle edge. However, there can be at most two such vertices of this last type per vertical chain, since the vertical direction was chosen so that no three obstacle edges lie on a common vertical line. $\square$

It is easy to see that $\mathcal{T}'_C$ has $O(n_C^3)$ tetrahedra, has $O(n_C^2)$ vertices on $\text{bdry}(\mathcal{T}'_C)$, and can be computed in time $O(n_C^3)$.

The second step is to transform $\mathcal{T}'_C$ into $\mathcal{T}_C$ by removing the vertices of $\text{bdry}(C_H, \mathcal{T}'_C)$ not in $\text{bdry}(C_H, \mathcal{P})$. As will be seen in the proof of lemma 3.15, removing a vertex that lies in the interior of a 2-face of $\text{bdry}(C_H, \mathcal{P})$ is slightly different from removing a vertex that lies in the interior of an edge of $\text{bdry}(C_H, \mathcal{P})$.

Let $G = (V, E)$ be the 1-skeleton of $\text{bdry}(C_H, \mathcal{T}'_C)$ as a graph, and similarly let $G_P = (V_P, E_P)$ be the 1-skeleton of $\text{bdry}(C_H, \mathcal{P})$. A subset of V is *independent* if no two vertices are connected by an edge of E.

Lemma 3.14 *There is an independent subset of $V \setminus V_P$ whose size is a constant fraction of $V \setminus V_P$ so that each vertex has constant degree in G.*

Proof: The average degree of vertices in the subgraph of G induced by $V \setminus V_P$ is at most 6, by planarity. Any vertex v in $V \setminus V_P$ is incident to at most 4 vertices of V_P: either v lies in the interior of a triangle of $\text{bdry}(C_H, \mathcal{P})$, in which case it can be incident to at most 3 vertices of V_P, or v lies on an edge of $\text{bdry}(C_H, \mathcal{P})$, in which case it can be incident to at most 4 vertices of V_P. Hence the average degree of vertices $V \setminus V_P$ in G is at most ten. By standard techniques [14], it follows that there is a bounded-degree independent set whose size is a constant fraction of $V \setminus V_P$. $\square$

Lemma 3.15 *There is a triangulation $\mathcal{T}_C$ of C with $\text{bdry}(\mathcal{T}_C) = \text{bdry}(C_H, \mathcal{P})$. $\mathcal{T}_C$ can be computed from $\mathcal{T}'_C$ in time $O(n_C^3)$, adding $O(n_C^3)$ tetrahedra and area $O(s_0^2 \log n_C + \text{area}(S \cap C))$.*

Proof: $\mathcal{T}'_C$ is updated in stages, with the final stage yielding $\mathcal{T}_C$. Each stage removes the independent set of vertices guaranteed by lemma 3.14. Clearly, the number of stages is logarithmic in the number of vertices of $\mathcal{T}'_C$, i.e., $O(\log n_C)$.

For simplicity, we first assume that every vertex to be removed lies in the interior of some 2-face of $\text{bdry}(C_H, \mathcal{P})$ and not on an edge of $\text{bdry}(C_H, \mathcal{P})$. Consider the triangles incident to v; all such triangles are coplanar. Let S_v be the polygon formed by the triangle edges opposite v. It is possible to push v slightly inside C_H maintaining the combinatorial structure of $\mathcal{T}'_C$ (in particular, v must not be pushed through the plane of the face opposite v in any tetrahedron incident to v). Pushing v leaves a dimple in the triangulation of C_H. The dimple can be filled by triangulating S_v and then adding to $\mathcal{T}'_C$ the tetrahedra formed by v and the new triangles.

The area added to $\mathcal{T}'_C$ at each stage is at most $O(s_C^2)$, where s_C is the side length of cube C. To see this, first note that when a vertex v is pushed, the change in area of the tetrahedra incident to v is negligible, since the perturbation of v can be made arbitrarily small. The area of the new tetrahedral faces is at most a constant times the area of S_v, since S_v has at most a constant number of edges. Since the vertices at each stage form an independent set, the polygons $\{S_v\}$ have disjoint interiors and their total area is at most the surface area of C, i.e., $O(s_C^2)$. Hence the total area added at all stages is $O(s_C^2 \log n_C)$.

Now suppose some vertex v lies on an edge e of $\mathrm{bdry}(C_H, \mathcal{P})$. Edge e may be a portion of a polyhedral edge of ∂C_H or may be interior to a polyhedral face of ∂C_H. If v does not lie on an obstacle triangle, then the approach is similar, except that the triangulation of S_v must use a piece of edge e.

Now suppose v lies on an edge e of $\mathrm{bdry}(C_H, \mathcal{P})$ and also on an obstacle triangle. It does not suffice to simply push v into the interior of C_H as above, since the resulting triangulation would not be compatible with the obstacles. Instead, vertex v is removed as follows. Necessarily v lies in the interior of a segment s of the intersection of the obstacle triangle with ∂C_H. Let p and q be the vertices on either side of v along s. Let $\Delta_1, \ldots, \Delta_k$ be the obstacle triangles incident to s that enter the interior of C_H, in cyclic order around s. We can assume $k>0$, i.e., not all triangles incident to s lie on ∂C_H, otherwise v can be perturbed as before. Let P_i be the plane through Δ_i, and choose the positive and negative open halfspaces of P_i so that the positive halfspace contains $\Delta_1, \ldots, \Delta_{i-1}$ and the negative halfspace contains $\Delta_{i+1}, \ldots, \Delta_k$. Let S_v be defined as above, and let S_v^+ be the polygon bounded by pq and the portion of S_v in the positive halfspace of P_1; similarly let S_v^- be the polygon bounded by pq and the portion of S_v in the negative halfspace of P_k. For each triangle Δ_i, $i = 1, \ldots, k-1$ in turn, split v into two vertices v and v_i connected by an edge, and perturb v_i slightly into the interior of ∂C_H while staying on Δ_i. See figure 7. Incidences to v are adjusted as follows. Within the positive halfspace of P_{i+1}, any edge, triangle, or tetrahedron previously incident to v should be made incident to v_i; in the negative halfspace of P_{i+1}, incidences remain with v. Incidences to v on P_{i+1} expand by a dimension: an edge uv on P_{i+1} becomes a triangle uvv_i; a triangle tuv on P_{i+1} becomes a tetrahedron $tuvv_i$. Finally, for the last triangle Δ_k, simply perturb $v_k = v$ to the interior of ∂C_H while staying on Δ_k. After all splitting and perturbation there is again a dimple on ∂C_H. Specifically, the perturbation of v_1 formed a dimple bounded by S_v^+; each subsequent perturbation of v_i, $i>1$, increased the dimple by the tetrahedron pqv_iv_{i-1}; in addition, the final perturbation of $v = v_k$ increased the dimple by S_v^-. To fill the dimple, choose triangulations of S_v^- and S_v^+; then add the tetrahedra formed by v_1 and the triangles of S_v^+, the tetrahedra pqv_iv_{i-1}, $i = 2, \ldots, k$, and the tetrahedra formed by v_k and the triangles of S_v^-.

To check that the resulting triangulation is compatible with all obstacles, observe that the perturbation of v_i maintains a two-dimensional triangulation of obstacle Δ_i, with the exception of the triangle pqv_i. However triangle pqv_i is a face of tetrahedron pqv_iv_{i-1} (and tetrahedron $pqv_{i+1}v_i$) and hence is added when the dimple is filled.

The perturbation increases area by adding tetrahedra pqv_iv_{i+1}, by replacing an edge uv with a triangle uvv_i, and by replacing a triangle tuv with a tetrahedron $tuvv_i$. In the first two cases, the

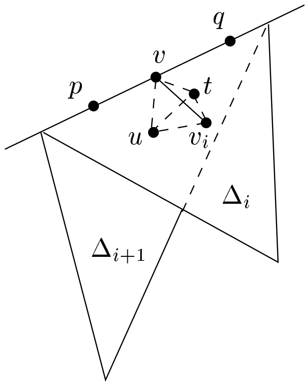

Figure 7: Vertex v lies on ∂C_H (not shown); both Δ_i and Δ_{i+1} are interior to C_H; v_i lies on Δ_i; u and t lie on Δ_{i+1} though they need not be close to v.

additional area is negligible since the distances vv_i can be made arbitrarily small. In the last case, the triangle tuv is essentially duplicated, adding the triangle tuv_i. However, notice that after the entire perturbation, neither v in its new position at v_k nor any vertex split from v lies on ∂C_H. Hence over all stages each triangle tuv can be duplicated at most three times, once per vertex. Since tuv is part of an obstacle lying in C_H, the total additional area is $O(\text{area}(S \cap C_H))$.

The number of tetrahedra required to fill the dimple resulting from perturbing v is constant if v is not on an obstacle and $O(n_C)$ if v is on an obstacle. Since $O(n_C^2)$ vertices have to be removed, the total number of tetrahedra over all stages is $O(n_C^3)$. New tetrahedra are also created by expanding a triangle to a tetrahedron. However, this happens at most once per vertex of each triangle, hence at most $O(n_C^3)$ times altogether.

The running time of the algorithm $\mathrm{O}(n_C^3)$, since it takes time $O(n_C^2)$ to find independent sets and $O(n_C^3)$ to add tetrahedra. □

3.5.4 Accounting

Lemma 3.16 $\sum_C n_C s_0^2 = O(\text{area}(S) + b^2)$, *where the sum runs over all skewered leaf cubes C and b is the side length of the root bounding cube.*

Proof: Let m_C be the total number of obstacle triangles that intersect cube C. Since every skewered cube meets some obstacle triangle, n_C is $O(m_C)$. We can express $\sum_C m_C s_0^2$ as

$$\sum_{\Delta \in S} |S_\Delta| s_0^2$$

where S_Δ is the set of skewered leaf cubes intersected by obstacle triangle Δ. We can write

$$|S_\Delta| = a_\Delta + m_\Delta$$

where a_Δ and m_Δ are respectively the number of cubes C in S_Δ so that the doubled cube $2C$ avoids or meets the boundary of Δ.

Consider first cubes C in S_Δ so that $2C$ avoids the boundary of Δ. Plainly the area of $2C \cap \Delta$ is $\Theta(s_0^2)$. Since no point in space is covered by more than 8 open doubled skewered cubes, $a_\Delta \leq 8 \cdot \text{area}(\Delta)/\Omega(s_0^2) = O(\text{area}(\Delta)/s_0^2)$. Hence $a_\Delta s_0^2$ is $O(\text{area}(\Delta))$, and $\sum_{\Delta \in S} a_\Delta s_0^2$ is $O(\text{area}(S))$.

Now consider cubes C in S_Δ so that $2C$ meets the boundary of Δ. The boundaries of the doubled cubes lie on the planes of a cubic grid of step size s_0. An obstacle triangle edge e meets at most $3\lceil \text{length}(e)/s_0 \rceil$ planes, which is $O(b/s_0)$ (recall that b is the side length of the box bounding the obstacle set S). Hence each edge meets $O(b/s_0)$ doubled cubes, so m_Δ is $O(b/s_0)$. We have

$$\sum_{\Delta \in S} m_\Delta s_0^2 = O(n \cdot b/s_0 \cdot s_0^2) = O(nbs_0) = O(b^2)$$

using $b/(2n) < s_0 \leq b/n$. $\square$

Proof of Theorem 3.9: By lemmas 3.11, 3.13, and 3.15, the total area of $\mathcal{T}$ is

$$\sum_C \text{area}(\partial C \cap H) + \sum_C \text{area}(S \cap C) + \sum_C n_C s_0^2$$

where the first two summations run over all leaf cubes and the last over skewered leaf cubes. The first two summations add to $O(\text{area}(T) + \text{area}(S))$. Using lemma 3.16, the last summation is $O(\text{area}(S) + b^2)$.

Since the minimum cube size is about b/n, the total number of leaf cubes is $O(n^3)$. Both the running time to compute $\mathcal{T}$ and the number of tetrahedra are bounded by $O(\sum_C n_C^3) = O(\sum_C (m_C+1)^3) = O(\sum_C m_C^3) + O(n^3)$, where m_C is the number of triangles meeting cube C. All leaf cubes C have faces lying on a cubic grid of planes with step size s_0; a triangle can hit at most $O(n^2)$ such grid cubes and thus at most $O(n^2)$ leaf cubes of T. Hence the number of incidences between leaf cubes and triangles, $\sum_C m_C$ is $O(n^3)$. The sum $\sum_C m_C^3$ is maximized if m_C is as large as possible, i.e., n, for as many cubes C as possible, i.e., $O(n^2)$, yielding an upper bound of $O(n^3 \cdot n^2) = O(n^5)$. $\square$

3.6 General obstacle sets

Theorem 3.17 *Let S be an arbitrary obstacle set with a total of n simplices. In time $O(n^6)$ it is possible to compute a triangulation compatible with S whose area is within a constant factor of the smallest possible. The triangulation has $O(n^6)$ tetrahedra and partitions the convex hull of S.*

We begin the proof by finding a bounding box of S. Choose a diametrical segment of the convex hull H of S. Project H onto a plane orthogonal to the segment, and choose a smallest square containing the projection. S fits into a box B_S which is a translation of the Cartesian product of the diametrical segment with the square. As before, perturb B_S slightly to guarantee that no obstacle face is parallel to a face of B_S. Let B_S have size $h \times h \times \ell h$. We can assume that ℓ is an integer. It is easy to see that the area of ∂H is $\Omega(\ell h^2)$.

We choose a family $\mathcal{B}$ of *boxes* that partition B_S. Conceptually split B_S into ℓ consecutive $h \times h \times h$ cubes. Any cube that contains a vertex of S is a box in $\mathcal{B}$. Any maximal union of consecutive cubes not containing a vertex is also a box in $\mathcal{B}$. Clearly, there are at most $2v - 1$ boxes in $\mathcal{B}$, where v is the number of vertices in S.

Refine the obstacle set S using the planes separating adjacent boxes in $\mathcal{B}$, as follows. For each plane P, if P hits an edge e of S, split the edge into two subedges and the vertex $e \cap P$; if P hits a 2-simplex Δ of S, split the simplex into two 2-faces and the edge $e \cap \Delta$; finally, triangulate the region $P \cap H$ compatibly with all intersection vertices $e \cap P$ and intersection edges $\Delta \cap P$, and add all resulting 2-faces to S. After examining all planes, triangulate any remaining 2-face that is not a triangle. Let S' be the refined obstacle set and S_B be the subset of S' lying inside box B. Recall $|\mathcal{T}|$ is the area of triangulation $\mathcal{T}$. Let $M = \inf_{\mathcal{T}} |\mathcal{T}|$, where $\mathcal{T}$ varies over all triangulations compatible with S; for $B \in \mathcal{B}$, let $M_B = \inf_{\mathcal{T}} |\mathcal{T}|$, where $\mathcal{T}$ varies over all triangulations compatible with S_B.

Lemma 3.18 $\sum_{B \in \mathcal{B}} M_B = O(M)$.

Proof: Let $\hat{\mathcal{T}}$ be an arbitrary triangulation compatible with S. Use each plane P that separates adjacent boxes in $\mathcal{B}$ to refine $\hat{\mathcal{T}}$, as follows: if P hits 0-, 1-, or 2-cells, then they are refined as S was refined to S'; if P hits a 3-cell Δ, then Δ is split into two 3-cells and the 2-cell $P \cap \Delta$. The area of $P \cap \Delta$ is $O(h^2)$, for a total over all planes of $O(|\mathcal{B}| \cdot h^2) = O(\text{area}(\partial H)) = O(\text{area}(\hat{\mathcal{T}}))$. Each 3-cell in the resulting subdivision is either a tetrahedron or has a constant number of polyhedral faces; in the latter case the cell can be centrally triangulated (after triangulating any nontriangular faces), increasing area by at most a constant factor. The resulting triangulation can be split into a triangulation of each box $B \in \mathcal{B}$, with triangulation areas summing to $O(\text{area}(\hat{\mathcal{T}}))$. $\square$

The obstacle set S' is triangulated by triangulating each set S_B in turn, and then pasting the resulting triangulations together. A detail is that adjacent triangulations must be compatible along their common boundary; however this is easily guaranteed using the technique of section 3.5. Choose a box $B \in \mathcal{B}$ of size $h \times h \times \ell_B \cdot h$, $\ell_B \geq 1$ an integer. Let L_B be the affine transformation that fixes one of the square sides of B and shrinks the orthogonal direction by a factor of ℓ_B; then $L_B(B)$ is a cube. Triangulate $L_B(B)$ using $L_B(S_B)$ as before: build an octtree with root cube $L_B(B)$ and subdivide it using the obstacles $L_B(S_B)$, as described in section 3.3, and then triangulate each octtree leaf cube as in section 3.5, again using $L_B(S_B)$. Apply the inverse transformation L_B^{-1} to obtain a triangulation $\mathcal{T}_B$ compatible with S_B.

Let T_B be the image under L_B^{-1} of the octtree, then $L_B(T_B)$ is the octtree with root cube $L(B)$ and T_B is an octtree-like structure formed from blocks with side-length ratio $1 \times 1 \times \ell_B$. The *area* of T_B, $\text{area}(T_B)$, is $\sum_C \text{area}(C \cap H)$, where the sum runs over all leaf boxes C in T_B.

Lemma 3.19 $\text{area}(T_B) = O(M_B)$.

Proof: We can assume that $\ell_B > 1$, otherwise the lemma follows from theorem 3.2. Choose a leaf box C in T_B and conceptually partition it into ℓ_B consecutive cubes. We show below that for each such cube D, $21D$ meets the 1-skeleton of S_B. Let $\hat{\mathcal{T}}$ be an arbitrary triangulation of S_B. Using the same argument as the proof of 3.2, we charge $\text{area}(D \cap H)$ to features of $\hat{\mathcal{T}}$ (with an appropriate modification to the charging-pair constants, since $21D$ rather than $7D$ meets the 1-skeleton). Since $\text{area}(C \cap H)$ is bounded by $\sum_D \text{area}(D \cap H)$, D in the partition of C, the lemma follows.

By lemma 3.5, $7L_B(C)$ meets the 1-skeleton of $L_B(S_B)$, so $7C$ meets the 1-skeleton of S_B. Since S has no vertices within B, $7C$ must meet an edge e of S_B with endpoints on opposite square faces of B. Consider the subsegment e' of e lying between the planes through the square faces of C. Clearly there is a translate of $7C$ that contains e' and overlaps $7C$. Hence for any of the cubes D partitioning C, $21D$ meets e'. $\square$

Lemma 3.20 $\ell_B \cdot \text{area}(L_B(T_B)) = O(\text{area}(T_B))$ *and* $\ell_B \cdot \text{area}(L_B(S_B)) = O(\text{area}(S_B))$.

Proof: We show the second statement; the first is easier. Assume $\ell_B > 1$, otherwise the lemma is trivial. Choose triangle $\Delta \in S_B$. Since B contains no vertices of S, Δ must result from refining a triangle of S by the planes through the two square sides of B. Hence Δ is a triangle, with an edge e on one square side of B and a vertex on the opposite square side of B. Similarly, $L_B(\Delta)$ has an edge e' of the same length as e on one side of $L_B(B)$, and a vertex on the opposite side. The height of Δ opposite e is at least $h\ell_B$, and the height of $L_B(\Delta)$ opposite e' is at most $h\sqrt{3}$. Hence $\ell_B \cdot \text{area}(L_B(\Delta)) = O(\text{area}(\Delta))$. $\square$

Lemma 3.21 $\text{area}(\mathcal{T}_B) = O(\text{area}(T_B) + \text{area}(S_B) + \ell_B \cdot h^2)$.

Proof: By theorem 3.9, we have $\text{area}(L_B(\mathcal{T}_B)) = O(\text{area}(L_B(T_B)) + \text{area}(L_B(S_B)) + h^2)$. Hence

$$\begin{aligned}
\text{area}(\mathcal{T}_B) &= O(\ell_B \cdot \text{area}(L_B(\mathcal{T}_B))) \\
&= O(\ell_B \cdot \text{area}(L_B(T_B)) + \ell_B \cdot \text{area}(L_B(S_B)) + \ell_B h^2) \\
&= O(\text{area}(T_B) + \text{area}(S_B) + \ell_B h^2)
\end{aligned}$$

using lemma 3.20. $\square$

Proof of theorem 3.17: By lemmas 3.19 and 3.21, the triangulation $\mathcal{T}_B$ of each box $B \in \mathcal{B}$ has area $O(M_B + \ell_B \cdot h^2)$. Hence the whole triangulation has area

$$\sum_{B\in\mathcal{B}} O(M_B) + \sum_{B\in\mathcal{B}} O(\ell_B \cdot h^2) = O(M) + O(\ell h^2) = O(M),$$

using lemma 3.18 and $\ell h^2 = O(\text{area}(\partial H)) = O(M)$. The bounds on running time and tetrahedra follow from theorem 3.9 since there are $O(n)$ boxes in $\mathcal{B}$. $\square$

3.7 Worst-case bounds

For obstacle set S in $\mathbb{R}^3$, recall that $c(S) = \inf_{\mathcal{T}} \text{area}(\mathcal{T})/\text{area}(S)$, where $\mathcal{T}$ varies over all triangulations compatible with S. The following lemma gives worst-case bounds on $c(S)$. As before, S must contain faces partitioning the boundary of its convex hull.

Lemma 3.22 *For any wide obstacle set S in $\mathbb{R}^3$, $c(S) = O(\sqrt{n})$, where n is the number of simplices in S. There is a wide set $\hat{S}$ of $O(n)$ obstacles with $c(\hat{S}) = \Omega(\sqrt{n})$.*

Proof: Let B be the minimum-size bounding cube of S, perturbed slightly so that no obstacle face is parallel to a face of B. Let B have side length b. Split B into a grid of identical cubes, where each cube has side length $s \approx b/\sqrt{n}$, so there are about $n^{3/2}$ cubes altogether. Using the algorithms of section 3.5, triangulate each clipped cube and paste the triangulations together, yielding a triangulation $\mathcal{T}$ compatible with S. $\mathcal{T}$ has area $O(n^{3/2} \cdot s^2) = O(b^2\sqrt{n})$ (the surface area of the cubes) plus $O(\text{area}(S))$ plus $O(\sum_C n_C s^2)$, where n_C is the combinatorial complexity of cube C. An analysis similar to lemma 3.16 shows that $\sum_C n_C s^2$ is $O(\text{area}(S) + n \cdot (b/s) \cdot s^2) = O(\text{area}(S) + b^2\sqrt{n})$. Hence the ratio $\text{area}(\mathcal{T})/\text{area}(S)$ is $O(\sqrt{n})$, as area(S) is $\Omega(b^2)$.

For the second statement, choose an axis-aligned unit cube B. Subdivide B into a cubical grid of about $\sqrt{n} \times \sqrt{n} \times \sqrt{n}$ identical subcubes. For each one-dimensional row of subcubes parallel to the x-, y-, or z-axis, choose a very thin obstacle triangle that covers all the subcube centers in the row. Slightly perturb the resulting set of $\Theta(n)$ triangles so that no two intersect. Add to the obstacle triangles a triangulation of their convex hull, forming the obstacle set $\hat{S}$. Since the 1-skeleton of $\hat{S}$ passes very near the center of each subcube, Lemma 3.4 implies that any triangulation $\mathcal{T}$ compatible with $\hat{S}$ must have area $\Omega((1/\sqrt{n})^2) = \Omega(1/n)$ within each subcube. Since there are $n^{3/2}$ subcubes, the total area of $\mathcal{T}$ is $\Omega(\sqrt{n})$. Since the area of the convex hull is $O(1)$ and the area of the remaining obstacles can be made arbitrarily small, $c(\hat{S}) = \Omega(\sqrt{n})$. $\square$

An argument similar to this proof shows that if S contains n points (and the faces of the convex hull) then $c(S) = O(n^{1/3})$. Furthermore there is a set $\hat{S}$ of n points with $c(\hat{S}) = \Omega(n^{1/3})$. These results contrast with the results of Agarwal *et al* [2]. They show there is a set S of n points in $\mathbb{R}^3$ so that in any triangulation of S, some line meets $\sqrt{n}$ triangulation faces. Similarly there is a set S of n obstacle triangles in $\mathbb{R}^3$ so that in any triangulation compatible with S', some line meets $\Omega(n)$ triangulation triangles, even though it misses all obstacles.

4 Discussion

We have not tried to estimate the approximation ratio for the construction in section 3. Our algorithm is based on Eppstein's algorithm, which approximates the minimum-length Steiner triangulation of a set of points in two dimensions. Eppstein is able to prove an approximation ratio of 316 (though he suspects the true ratio is much smaller, perhaps around 20); our proof is much less careful about constants than his. A challenging open problem is to construct triangulations of approximately minimum weight in two or three dimensions with reasonable constants and with a reasonable number of vertices.

In two dimensions, the minimum spanning tree provides an intrinsic measure of the minimum weight triangulation, in the sense that their weights differ by at most a logarithmic factor (lemma 2.2). In three dimensions, the surface area of the octtree constructed in section 3 is an intrinsic measure of the minimum weight triangulation (theorem 3.2). It would be of interest to obtain a more natural intrinsic measure.

The analysis of ray-shooting-by-walking can be extended to other subdivisions besides triangulations. For example, consider the leaf cubes of the octtree constructed in section 3, with the modification that each unskewered cube is partitioned by all obstacles that cut it fully. Label each skewered leaf cube with the number of obstacle triangles that meet it, and label all the other 3-cells 1. Then it is possible to walk through the partitioned octtree along a line ℓ with total cost proportional to $w(\ell)$, where $w(\ell)$ is the sum of the labels of the 3-cells intersected by ℓ. The analysis in section 3 shows that $\int w(\ell) d\mu$ is the area of the minimum weight triangulation, to within a constant factor. Two problems arise naturally when considering alternative subdivisions for ray-shooting queries: first, to determine if the area on the minimum weight triangulation is always a relevant bound; second, to provide an analytic comparison of the constants that arise from different subdivisions.

References

[1] P. Agarwal, Range searching, *CRC Handbook of Discrete and Computational Geometry*, J. E. Goodman and J. O'Rourke, eds., CRC Press, Boca Raton, 1997.

[2] P. Agarwal, B. Aronov, and S. Suri, Stabbing triangulations by lines in 3D, *Proc. Eleventh Ann. Symp. Comp. Geometry*, 1995, pp. 267–276.

[3] P. Agarwal and J. Matoušek, On range searching with semialgebraic sets, *Discrete and Comput. Geom.* **11**:393–418, 1994.

[4] P. Agarwal and M. Sharir, Applications of a new partition scheme, *Discrete and Comput. Geom.* **9**:11–38, 1993.

[5] R. Berirouti, J. Snoeyink, Implementations of the LMT Heuristic for Minimum Weight Triangulation, *Proc. Fourteenth Ann. Symp. Comp. Geom*, pp. 96–105, 1998.

[6] M. Bern and D. Eppstein, Approximation algorithms for geometric problems, in *Approximation Algorithms for NP-hard Problems*, D. Hochbaum, ed., PWS Publishing Co., 1997, pp. 296–345.

[7] K. Clarkson, Approximation algorithms for planar traveling salesman tours and minimum-length triangulations, *Proc. 2nd ACM-SIAM Symp. on Discrete Algorithms*, 1991, pp. 17–23.

[8] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Berlin, 1987.

[9] D. Eppstein, Approximating the minimum weight Steiner triangulation, *Discrete and Computational Geometry*, **11**:163–191, 1994.

[10] S. Fortune, A beam-tracing algorithm for prediction of indoor radio propagation, *Applied Computational Geometry: Towards Geoemtric Engineering*, M. Lin, D. Manocha, eds, *LNCS* 1148, Springer-Verlag, 1996.

[11] J. Hershberger and S. Suri, A pedestrian approach to ray shooting: shoot a ray, take a walk, *J. Algorithms* **18**, pp. 403–431, 1995.

[12] J.S.B. Mitchell, D.M. Mount, and S. Suri, Query-sensitive ray shooting, *Int. J. Comp. Geom. Appl* **7**:4, pp. 317–347, 1997.

[13] D.W. Moore, *Simplicial mesh generation with applications.* Ph.D. Thesis, Dept. of Computer Science, Cornell Univ., 1992. Report 92-1322.

[14] J. O'Rourke, *Computational Geometry in C*, Cambridge University Press, Cambridge, 1994.

[15] M. Pellegrini, *Ray-shooting and lines in space*, *CRC Handbook of Discrete and Computational Geometry*, J. E. Goodman and J. O'Rourke, eds., CRC Press, Boca Raton, 1997.

[16] L. Santaló, *Integral Probability and Geometric Probability*, Encyclopedia of Mathematics and its Applications, Volume 1, Addison-Wesley, 1979.