Quality Mesh Generation in Three Dimensions

Scott A. Mitchell* Stephen A. Vavasis[†]

February 5, 1992

Abstract

We show how to triangulate a three dimensional polyhedral region with holes. Our triangulation is optimal in the following two senses. First, our triangulation achieves the best possible aspect ratio up to a constant. Second, for any other triangulation of the same region into m triangles with bounded aspect ratio, our triangulation has size $n = O(m)$. Such a triangulation is desired as an initial mesh for a finite element mesh refinement algorithm. Previous three dimensional triangulation schemes either worked only on a restricted class of input, or did not guarantee well-shaped tetrahedra, or were not able to bound the output size. We build on some of the ideas presented in previous work by Bern, Eppstein, and Gilbert, who have shown how to triangulate a two dimensional polyhedral region with holes, with similar quality and optimality bounds.

1 Introduction

Triangulation of polyhedral regions is a fundamental geometric problem for numerical analysis. In particular, if one wishes to solve an elliptic boundary

*Center for Applied Mathematics, Cornell University, Ithaca, NY 14853. Supported in part by Hughes Research Laboratories, Malibu, CA, and by DARPA under ONR contract N00014-88K-0591, ONR contract N00014-89J-1946, and NSF grant IRI-9006137, by a Xerox summer internship and an NSF PYI award.

[†]Department of Computer Science, Cornell University, Ithaca, NY 14853. Supported by an NSF Presidential Young Investigator award.

value problem on a three-dimensional domain, it is necessary to discretize the domain with a mesh. If the domain is sufficiently complicated, then the method of *finite elements* is commonly used. In this method the domain is divided into small convex polyhedral regions with a fixed number of faces called *elements*. A common choice for an element in three dimensions is the tetrahedron. Thus, a *triangulation* of the domain is required.

For numerical stability in the finite element method, it is necessary that the tetrahedra have bounded *aspect ratio*. This means that the angle between any adjacent pair of edges of the tetrahedron, or between any edge and a 2-dimensional face not containing it, is bounded below by a constant. For information about aspect ratio bounds in numerical analysis, see Babuška and Aziz [1976].

Our algorithm generates a triangulation for a nonconvex bounded polyhedral domain with holes. In addition, the triangulation is optimal in two respects. In particular, the best possible aspect ratio is achieved for the tetrahedra, and the number of tetrahedra is within a constant factor of the best possible for any triangulation with bounded aspect ratios.

Our work is closely based on earlier work by Bern, Eppstein and Gilbert [1990], who solved the corresponding problem for two-dimensional polyhedral domains. The main complication in extending the work to three dimensions concerns vertices of the domain. In particular, for a two-dimensional domain a vertex is adjacent to only two edges, and there are essentially three cases for vertices (the two edges make an acute angle, an obtuse angle, or an angle greater than 180°). In three dimensions, however, a large number of faces can meet at a single vertex, forming a complicated cross-section with some convex and some concave angles. Handling vertices properly necessitates a host of changes in the algorithm. Our running time bounds are probably not the best possible; this is discussed further at the end of the paper. Our optimality condition is slightly stronger than Bern et. al.'s condition. In particular, they show that overall the number of triangles they generate is no more than a constant above optimal. The triangles used in a particular region of the domain, however, could be arbitrarily smaller than what is needed for the triangulation to achieve good aspect ratio. In our triangulation, no tetrahedron is ever more than a constant factor smaller than the tetrahedron at the same location in the best possible triangulation.

Other authors have considered three-dimensional regions, but, to our knowledge, no previous work has simultaneously addressed the problems of

optimality, aspect ratio and of complicated nonconvex regions. Indeed, as far as we know, there is no previous algorithm to triangulate a nonconvex three-dimensional region with guaranteed aspect ratio (regardless of the optimality of the triangulation). Our technique is based on an octree partitioning of the domain. This technique has been used before, for example, by Carey, Sharna and Wang [1988]. Chazelle and Palios [1989] consider optimality but are not interested in angle bounds. Chew [1989b] presents an algorithm for two-dimensional polyhedral regions, but does not address the problem of optimal number of triangles. Moore and Warren [1990] address the problem of adaptive mesh generation for box-shaped regions. Mitchell [1987] and [1988] consider the problem of adaptive mesh generation of complex regions in regard to finite element error bounds. Because of the importance of mesh generation, the literature on this problem is very extensive; see for example, the conference proceedings edited by Hauser and Taylor [1986].

Recently, Bern and Eppstein [1990] have surveyed the literature on triangulations, giving more details about earlier work.

Our optimality proof is based on characteristic length functions. Similar functions have been used before by Miller and Thurston [1990], Miller and Vavasis [1990], and Miller, Teng and Vavasis [1991]. We have not seen them used, however, to analyze the construction of a triangulation.

The remainder of this paper is organized as follows. In Section 2 we give more details on the aspect ratio of a tetrahedron. In Section 3 we discuss the formation of the octree. In Section 4 we demonstrate a relation between the boxes of the octree. In Section 5 to Section 7 we discuss the steps necessary to produce a triangulation given an octree. In Section 8 to Section 10 we provide the optimality proof. For the remainder of this introduction, we describe our assumptions.

Our triangulation is denoted Δ^{OCT} . We assume we are given a three-dimensional polyhedral region as follows. There is a list of vertices with coordinates specified. There is also a list of edges and faces. The three lists are mutually linked. The list of edges for each vertex is ordered as they occur around the vertex. Similar orderings are assumed for the other lists. We assume that P is connected; if not, each component could be triangulated separately.

A *face* of P , or another polyhedral region, can have either zero dimensions, called a *vertex*, one dimension, called an *edge*, or two dimensions, called a *facet*. In some circumstances we want to regard P itself as the three-

dimensional face. Note that the containment relation induces a partial order on faces.

We assume that the polyhedral region is nondegenerate in the following senses. Every two-dimensional face of P has the interior of the polytope on exactly one side. Every one-dimensional edge is incident on exactly two two-dimensional faces. For every zero-dimensional face (vertex) v , for every small enough open neighborhood N of v , the set $(N \cap P) - \{v\}$ has exactly one connected component. These assumptions may be dropped in future work.

2 Aspect ratio definition

To argue about the aspect ratio of triangulations of P , we wish to define the angle between any two faces of P , or faces in a triangulation of P , that intersect each other non-trivially. Various concepts such as dihedral angle and solid angle have been used in the past, but we adopt a different approach.

Interior Angle We define the *interior angle* or *angle* between two faces f and g in the cases where f and g are a facet and an edge, two facets, or two edges that have a common intersection v . We additionally require that one face is not a subset of the other. We say that two rays r_1, r_2 with a common endpoint v can *see* each other if there are points v_1, v_2 on r_1, r_2 each at distance $t > 0$ from v such that the sector vv_1v_2 lies in P .

If f and g are edges, then we measure angles in the usual way for rays, except we require that the rays can see each other. We always measure angles with respect to the interior of P , so that f and g could make an angle close to 360° .

If f is an edge and g is a facet, then we consider all rays r in the plane of g , such that r lies between the two edges of g that meet at v , and such that f and r can see each other. Then the angle between f and g is defined to be the minimum angle between the ray covering f and all rays r meeting the preceding conditions. We have minimum rather than infimum because faces are assumed to be closed.

If f and g are both facets whose intersection is a vertex, then we let both the ray for f and the ray for g vary as in the last paragraph, and take the minimum angle between the rays to be the angle between f and g that can see each other. If the facets share a common edge, then the interior angle between the facets is the angle between the rays in the half planes of f and g

perpendicular to the common edge. Once again, we only consider rays that see each other.

We use the symbol “ α ” throughout the paper to refer to the smallest such interior angle between any two faces of P . We note that if the interior angle between two facets is β , then one of them has a subface edge that makes an angle of no more than β with the other. Thus there are always either two edges, or an edge and a facet, that have interior angle α between them.

Aspect ratio. A useful measure of the shape of a tetrahedron is its *aspect ratio*, which is the ratio of the radius R of the smallest containing sphere, to the radius r of the largest inscribed sphere. In general the smallest containing sphere is not the sphere through the vertices of the simplex. A simplex with bounded aspect ratio has bounds on the angles made between any pair of faces that intersect, and not just on the angles between edges.

It can be proved that the aspect ratio of a tetrahedron is bounded above and below by constant multiples of the reciprocal of the sharpest interior angle (as defined above) in the tetrahedron. We do not prove this claim because it is implicit in theorems and lemmas in upcoming sections.

3 Subdivision of P into cubes, the octree

3.1 Data structures and definitions

The main data structure we use for our algorithm is an octree. We commonly refer to each node of the tree as a *box*. We associate with each box, b , a polyhedral region of $\mathbb{R}^3$ called the embedding of the box and denoted $I(b)$. During the generation of the octree, $I(b)$ is exactly a three dimensional cube. Later boxes are warped and triangulated, changing their geometric structure.

An octree node is either a leaf, or has eight children. The embedding of the eight children of a node are the eight cubes obtained by dividing the embedding of the box in half in each of the three dimensions. We say a node is *split* if it is not a leaf. The process of creating the eight children of a node is called *splitting*.

For clarity, we use the term box face (“box facet,” “box edge,” etc.) to denote a face of a box and P face to denote a face of P . Similarly for ∂P and ∂b .

Duplicate. Some boxes in the octree will be *duplicated* into the original

node and several new nodes, called duplicates or duplicate boxes. We create duplicates whenever the intersection of the box with P has more than one connected component. Each duplicate represents the same geometric cube in $\mathbb{R}^3$, but is associated with one connected component of $P \cap I(b)$. We use the notation $P \wedge b$ to denote the component of $P \cap I(b)$ assigned to a particular box b .

Whenever we split a box b , if $P \wedge b$ is nonconvex a child box b' may have more than one component (i.e., $(P \wedge b) \cap b'$ might have more than one component even though $P \wedge b$ consists of one component). Whenever a box is split, we immediately determine whether any of its children contains more than one component of P , and if so, we make duplicates of the child box.

If a P face is incident upon $P \wedge b$ for a box b , we say that a box *contains* the face. We maintain pointers between each box and the faces it contains.

Extended Box. For a given box b , we define the extended box of b , $\text{ex}(b)$, such that $I(\text{ex}(b))$ is the cube concentric with $I(b)$ and with each dimension expanded by a factor of 5. We use the notation $P \wedge \text{ex}(b)$ to denote the component of $P \cap I(\text{ex}(b))$ that contains $P \wedge b$. The extended box contains only the P faces of $P \wedge \text{ex}(b)$. Note that $\text{ex}(b)$ is not a box of the octree, but may be constructed from boxes of the octree.

Adjacent. Two boxes are called *neighbors* if their embeddings intersect non-trivially, and one is not a duplicate of the other. The *size* of a box is the length of an edge. Suppose the size of box b_1 is larger than the size of b_2 , then we say b_1, b_2 are *adjacent* if they are neighbors and $(P \wedge \text{ex}(b_1)) \cap (P \wedge b_2) \neq \emptyset$. In certain cases such as when a box contains faces meeting at a reflex angle, a box may be adjacent to two duplicates of a second box.

Balance Condition. If we split a box, we immediately split other boxes to maintain the following invariant called the *balance condition*: No box is adjacent to a box more than twice its size. Certain boxes containing vertices or edges of P are called *protected* boxes, and are exempt from being split by the balance condition later in the algorithm. Nonetheless, the ratio of the size of a protected box to a box adjacent to it is bounded above by a constant that depends linearly on $1/\alpha$, as we prove in the next section.

3.2 Creating the octree

We generate the octree by selectively splitting and duplicating nodes. The goal of splitting and duplicating boxes is to make the boxes small enough

so that the intersection of P with the embedding of any box is as simple as possible. However, boxes should not be made too small, as this would lead to an excessive number of tetrahedra in the final triangulation.

The octree generation algorithm is divided into three phases, the *vertex phase*, the *edge phase*, and the *facet phase*.

It is easy to state the conditions under which we duplicate b , namely, whenever $P \cap I(b)$ has more than one component. Nonetheless, the box duplication process is actually the most computationally complicated part of the octree generation algorithm, because determining components of $P \cap I(b)$ is a nontrivial computational task. The details of the duplication process are described below.

On the other hand, for splitting, the conditions under which a box is split are more complicated, but the computational tasks are straightforward. We now describe the various steps of the splitting algorithm.

3.3 The vertex phase

How finely we split boxes depends on the following definition.

Vertex Cone. A *vertex cone* of a box b is a set of P faces, $F_1, F_2, \dots, F_k$, that satisfy the following:

- (a) The set is precisely one vertex and all of its superfaces. This vertex is called the *apex* of the vertex cone.
- (b) The apex is in b .
- (c) The faces $F_1, F_2, \dots, F_k$ are exactly the P faces incident upon $P \wedge \text{ex}(b)$.

Vertex Crowded. We say that a box b is *vertex crowded* if the following are satisfied.

- (a) There is a P vertex v in b .
- (b) The superfaces of v are not the only faces of P incident upon $P \wedge \text{ex}(b)$.

Equivalently, a box is vertex crowded if $P \wedge b$ contains a vertex that is not the apex of a vertex cone. We say that a face G of $P \wedge \text{ex}(b)$ *interferes* with $v \in P \wedge b$ if v is not a subface of G . Thus b is crowded if it contains a P vertex v , and $\text{ex}(b)$ contains a face that interferes with v .

Figure 1: Here a box is duplicated for two components, and one duplicate box is split because it is crowded.

We recursively split and duplicate boxes, maintaining the balance condition at each step, until every box containing a vertex contains exactly one vertex cone. Clearly this procedure will terminate once the box sizes become a constant factor smaller than the minimum path length in P between a vertex and a face that does not contain that vertex.

The details of the procedure for determining whether a box b should be split are straightforward given an enumeration of the P faces bounding $P \wedge \text{ex}(b)$. Such an enumeration is obtained from the procedure for determining whether a box should be duplicated described below.

3.3.1 The vertex centering step

The conclusion of the vertex phase is a special one-time reorganization of the boxes so that each vertex of P is far from the boundary of the box that contains it. The value of this property will become clear in Section 5 and later sections.

We first make the following observation concerning the boxes at the conclusion of the vertex phase. Let b be a box with a P vertex v . Let b_1 be adjacent to b . We claim that the size of b_1 is either the same or twice as large as the size of b . We denote the size of b with the notation $h(b)$. Note that the balance condition ensures that $h(b_1)$ is either equal to $h(b)$, $2h(b)$, or $h(b)/2$. Thus the claim made in this paragraph is that if b contains a vertex, then any adjacent box has size at least that of b .

To prove this, suppose that $h(b_1) = h(b)/2$. Let b'_1 be the parent box of b_1 . Then either b'_1 is vertex crowded, or it is split because of the balance condition. We may assume that b'_1 is split by the balance condition, since

if it is vertex crowded we simply take $k = 1$ in the following discussion and the proof holds. There is a chain of boxes $b'_1, \dots, b'_k$ such that b'_k is vertex crowded, b'_i is adjacent to b'_{i+1} (for $i = 1, \dots, k-1$), and such that $b'_i \leq 2b'_{i+1}$. These are just the conditions necessary for the balance condition to propagate from b'_k back to b'_1 . Note that $I(b'_k)$ lies interior to $I(\text{ex}(b))$, no matter how large k is, because $\text{ex}(b)$ is 5 times larger than b in every direction, and the sizes of the b'_i 's form a decreasing geometric series. Thus, the vertex causing b'_k to be crowded is in $P \wedge \text{ex}(b)$, contradicting the fact that b is not vertex crowded.

Thus, every box b with a vertex is surrounded by boxes that are either twice the size or the same size as b . The first part of the vertex centering step is as follows. For every box b with a vertex, we split the boxes adjacent to b that are twice as large, and then we propagate the balance condition (which causes every other box in the octree to be split at most one time).

We claim that after this procedure, every box adjacent to a box b with a vertex is the same size as b . This claim is clearly true before the balance condition gets propagated. Thus, we have to show that when the balance condition is propagated, this does not cause any box adjacent to a box b with a vertex to be split.

Consider a setup of boxes that might cause a box b_1 adjacent to a box b to be split by propagation of the balance condition. Suppose b_1 is a box adjacent to b , where b contains a vertex v , such that b_1 is the same size as b . Suppose elsewhere there is a box $\hat{b}$ with adjacent b_k such that b_k is the same size as $\hat{b}$, and b_k is adjacent to a box b_{k-1} four times as large, which is adjacent to a box b_{k-2} twice as large, and so on, up to b_1 , which is twice as large as b_2 . Under these conditions (and only these conditions) propagating the balance condition from b_k would cause b_1 to split.

However, we claim that the conditions described in the previous paragraph cannot occur. This is because such conditions would contradict the fact that every facet contained in $\text{ex}(b)$ has vertex v as a subface (same argument as above).

Thus, after the first part of the vertex centering step, every box has been split at most once, and every box with a vertex is adjacent to boxes of the same size. Consider such a box b . Without loss of generality, b has 26 adjacent boxes surrounding it of the same size, arranged in a $3 \times 3 \times 3$ group. These boxes may have been duplicated. If necessary, we adjoin empty boxes in the spaces of this group where no box is adjacent to b .

Figure 2: No matter how small we split a box b' , if it is outside of $\text{ex}(b)$ then the balance condition will not split b .

Now, the second part of the vertex centering step is as follows. We reorganize the octree by uniting b with some of the boxes of the $3 \times 3 \times 3$ group, and perhaps some duplicates of these as well. We form a new box B by uniting b with seven boxes in a $2 \times 2 \times 2$ group, and perhaps some of their duplicates, to obtain the result that $I(B)$ is a cube of side length $2h(b)$, and B contains the component of $P \wedge B$ containing the P vertex in b . See Figure 3 for an illustration of this procedure in two dimensions. Note that we can always pick the correct 7 neighbors of b , so that v is distance at least $h(b)/2$ from any facet of B . Since we have effectively doubled the size of b , B may be adjacent to boxes four times smaller than B . We do not take any steps to repair this violation of the balance condition, as the size ratio of adjacent boxes is still bounded.

From now on, B is *protected*, meaning that it is never split again during the algorithm.

3.4 The edge phase

The next phase of the octree generation algorithm focuses on edges of P , and is analogous to the vertex phase.

We define “edge cones” and “edge crowded” and split and duplicate boxes as for vertices. The main difference is that we take an approach different from centering to ensure that no box face is close to an edge. The splitting and

Figure 3: Uniting boxes around a vertex box in two dimensions.

duplicating algorithm proceeds as if the protected vertex boxes do not exist. In particular, the extended box of any box in this phase does not extend into a protected box, and the vertices of P are ignored.

Edge Cone. An *edge cone* of a box b is a set of P faces, F_1, F_2, F_3 , that satisfy the following:

- (a) The set is precisely one edge and its two superfaces. This edge is called the *apex* of the edge cone.
- (b) The apex is contained in b , and
- (c) The faces F_1, F_2, F_3 are exactly the P faces incident upon $P \wedge \text{ex}(b)$.

Edge Crowded. We say that a box b is *edge crowded* if it contains an edge of P , but it and its superfacets are not the only P faces incident on $P \wedge \text{ex}(b)$. Equivalently, b is edge crowded if any edge contained in b is not the apex of an edge cone.

Just as in the vertex case, we recursively duplicate an unprotected box for each edge cone it contains, and split it if it is edge crowded. We maintain the balance condition immediately after a box is split. Recall that the protected vertex cone boxes are not changed in this step.

Identifying edge cones is straightforward if we have a list of facets bounding $P \wedge \text{ex}(b)$, which is obtained from the algorithm for duplication described below.

3.4.1 Protecting edge cone boxes

We seek the same goal as the centering vertices step but for edges. The fact that an edge cone box may be adjacent to a protected vertex box prompts us to take an approach somewhat different from the one used for vertex cones. Here if a P edge is near a box face G , we surround the face with protected boxes, as in the vertex case, but do not require them to be the same size. Unlike the vertex centering phase, we do not eliminate G at this time, because G may be a face of a protected vertex box. Instead, after the entire octree is generated but before we triangulate, we warp the cube (see Section 5) into a more general shape that increases the distance between G and the P edge, and eliminates any small angles as well.

We begin by splitting every unprotected box containing a P edge and its adjacent boxes and enforcing the balance condition. Every box is split at most once. This ensures by the “crowded” definition that if a box b (other than a protected vertex box) is adjacent to a box b' containing an edge E , then b cannot be adjacent to a box containing any other edge E' .

We always protect every box containing a P edge, and every box adjacent to a box containing a P edge, regardless of whether any face of the box is close to a P edge. These boxes are never split again.

3.5 The facet phase

The third phase of the octree generation algorithm focuses on facets of P . We define “facet cones” and “facet crowded” as in the previous phases. We split and duplicate boxes as we did for vertices. As for edges, the main complications arise when centering facets, especially when a facet cone box is adjacent to a vertex or edge cone boxes. The splitting and duplicating algorithm proceeds as if the protected vertex and edge boxes do not exist. In particular, the extended box of any box does not extend into a protected box, and the edges and vertices of P are ignored.

Facet Cone. A *facet cone* of an unprotected box b is a single P facet, F_1 , that satisfies the following:

- (a) Facet F_1 is contained in b .
- (b) Facet F_1 is the one and only P face incident upon $P \wedge \text{ex}(b)$.

For the sake of consistency, we will call F_1 the *apex* of the facet cone.

Facet crowded. We say that a box b is *facet crowded* if it contains a facet that is not the only P facet incident upon $P \wedge \text{ex}(b)$. Equivalently, b is facet crowded if it contains a facet that is not the apex of a facet cone.

Just as in the vertex case, we recursively split a box if it is facet crowded. We maintain the balance condition immediately after a box is split. Recall that the protected vertex and edge cone boxes are not changed in this step.

3.5.1 Protecting facet cone boxes

We take the same approach as for edges. We begin as in the edge case by splitting every unprotected box containing a P facet and its adjacent boxes and enforcing the balance condition. This ensures that a box is adjacent to at most one facet cone apex. Thus, in the warping and triangulation step to follow in Section 5, we may consider the boxes containing a facet without reference to any other facets, except for boxes adjacent to a protected edge or vertex box.

Technically, there is no need to protect boxes because there are no more splitting phases to consider. However, for the sake of consistent terminology, we protect facet cone boxes and the boxes adjacent to them.

3.6 The duplication process

Recall that we duplicate b whenever it is determined that $P \cap I(b)$ has more than one component. In this section we explain how to identify the components of $P \cap I(b)$. This same techniques allow us to determine the components of $P \cap I(\text{ex}(b))$, which is necessary for determining whether boxes are crowded.

Because we believe that the algorithm proposed here is not the most efficient algorithm possible for the task, we skip over some of the details involved.

The first part of the duplication algorithm is a preprocessing algorithm, which is run before the octree generation begins. In the preprocessing algorithm we first identify the separate components of ∂P . If P is a convex domain (or some other fairly simple shape) then ∂P has a single component

and no further preprocessing is necessary. Otherwise, suppose ∂P has several components. One component is the *exterior* component; we call all the other components *floaters*. The component for each facet, edge, and vertex can be determined in $O(n)$ steps with a standard graph search, where n is the total number of facets edges and vertices of P .

Once the components are identified, we now perform a second procedure in which we identify a *tether* for each floater. First, for each floater C , we identify the point v (necessarily a vertex) with the largest x coordinate on C . This point is called the *base* of the tether. From the base we shoot a ray in the positive x coordinate direction until we encounter the first point v' on ∂P . This point cannot be on the same floater. Such a point exists, since the ray must eventually pass through the exterior component. The directed segment from v to v' is called the *tether* for C , and v' is called the *head* of the tether. See Figure 4 for a two dimensional example of the tethering structure of a polygon. We can carry out the search for v' in a fairly naive fashion: Simply loop over all facets, and for each facet F carry out a point-in-polygon test for the point where the ray v passes through the plane of F . This requires time proportional to the number of edges of F . Thus, the total time to locate v' for a particular tether (looping over all facets) is $O(n)$. Therefore, the time to identify all tethers is $O(n^2)$.

Note the following two properties of tethers: (1) Each tether is contained completely in P . (2) When regarded as a digraph on the components of ∂P , the tethers form an in-tree rooted at the exterior component.

Now we discuss the procedure for determining when a box should be duplicated. We first make a list of all P facets, edges and vertices contained in the box. A particular P facet F in the box may intersect the box in more than one component if the facet is nonconvex. We therefore must first determine all the components of each facet passing through the box. This may be done by sweeping a line segment across the facet, requiring $O(v \log v)$ steps if v is the number of vertices in the facet. The total time for identifying all components of all facets for a box is $O(n \log n)$.

Then we perform a combinatorial graph search to identify the various components of $\partial P \cap I(b)$. Let these components be called *sheets*. There are two kinds of sheets: those that intersect the boundary of $I(b)$, and those that are completely internal to $I(b)$. Note that a sheet entirely inside $I(b)$ must be a floater; such a floater is called an *internal floater* for b . The other sheets are called *external* sheets. Note that a floater that intersects $I(b)$ is

Figure 4: Floaters and tethers for a two dimensional polygon.

an internal floater if and only if it lies entirely in $I(b)$, otherwise it is one or more external sheets.

Clearly any sheet is incident upon a single connected component of $P \cap I(b)$ since the sheet itself is connected. Thus, the remaining task is to determine which sheets are connected to each other in $P \cap I(b)$.

For external sheets, the following algorithm exactly determines how they are connected. We consider the intersection of each sheet with the boundary of $I(b)$ (the surface of a cube). This intersection is a collection of disjoint simple polygons. We form the list of polygons for all sheets. We call these polygons the “rims.” Note that these polygons break across edges of the cube. We then sweep over this list of rims, forming a tree indicating which polygons are contained in others. If these polygons were in the plane, we could sweep with a vertical line. This would require sorting x coordinates of the polygons, hence requiring $O(n \log n)$ steps. On the surface of the cube we can sort by the sum of x, y, z coordinates, which means that we sweep over the surface of the cube with a closed polygonal segment (initially the boundary of a triangle).

In the previous paragraphs we made reference to point-in-polygon tests and plane sweeps. These are well-known tools in the computational geometry

Figure 5: Determining the components of a box. Sweeping the boundary determines two components, and following tethers determines that all internal floaters belong to the same component, and there are no internal components.

literature. The reader unfamiliar with these tools should consult Edelbrunner [1987] or Preparata and Shamos [1985].

From the sweep we can build a rooted tree indicating the containment hierarchy among the polygons. If we assume that the first point swept is outside P , then the following simple rule determines connectivity in P : Every polygon at an even level i (counting 0 as the top level) is connected in P to its children at odd level $i + 1$. The parities are interchanged if the first point swept is inside P .

Thus, from an $O(n \log n)$ procedure we can determine exactly which external sheets are in the same components of $P \cap I(b)$. This leaves the problem of determining the component of $P \cap I(b)$ for an internal floater C . This is determined from the tether (v, v') of C . If the head v' is inside $I(b)$, then the floater lies in the same component as the facet containing v' (because the entire tether is in P). This facet might be a facet of another internal floater, but because the tethers form an acyclic graph, eventually we will reach a point v' on an external sheet. See Figure 5.

The other case is that the tether's head v' lies outside $I(b)$. In this case we find the point v'' where the tether intersects the surface of $I(b)$. If we knew where in the polygon tree v'' lay, then we could determine the component

containing v'' (which is the same component containing the floater). The polygon regions containing v'' are easily determined if, before starting the polygon containment sweep, we first insert v'' into the list of polygons as a singleton polygon. This should be done for all tethers. This does not change the running time bound of $O(n \log n)$.

Thus, in time $O(n \log n)$ we can determine all the components of $P \cap I(b)$. Moreover, it is easy to see that the rims allow us to determine which neighboring boxes are adjacent to a given duplicate box.

Finally, we need to determine the facets adjacent to $P \wedge \text{ex}(b)$ for testing crowdedness. This is done by running the component-determination algorithm on $\text{ex}(b)$, and then saving the component of $P \cap I(\text{ex}(b))$ that contains a point in $P \wedge b$.

3.7 Summary of the octree algorithm

The octree algorithm has now been completely specified. The octree structure, including adjacency, duplication, and P face containment properties remain unchanged by the algorithms that follow.

1. We initialize by setting the octree equal to a single node representing a large box with all of the faces of P in its interior.
2. *Vertex Phase.* We split and duplicate boxes until it is determined that no box is vertex crowded. The balance condition is enforced each time a box is split.
3. Every box b' adjacent to a box b with a P vertex is split, in the case that $h(b') = 2h(b)$. The balance condition is enforced.
4. Every vertex is centered in its containing box by merging eight boxes around a vertex. Every box containing a vertex is now protected, so is never split again.
5. *Edge Phase.* We split and duplicate unprotected boxes until it is determined that no box is edge crowded. The balance condition is enforced each time a box is split.
6. Every box containing a P edge and its adjacent boxes are split to further separate edges from each other. The balance condition is enforced.

7. Every box containing an edge and its adjacent boxes are protected, and never split again.
8. *Facet Phase.* We split and duplicate boxes until it is determined that no box is facet crowded. The balance condition is enforced each time a box is split.
9. Every box containing a P facet and its adjacent boxes are split to further separate facets from each other. The balance condition is enforced.

An important feature of the algorithm is that when it terminates, every leaf box is either empty, lying entirely inside P , or contains a vertex, edge, or facet cone. Also recall that the running time per box is proportional to $n \log n$, where n is the number of faces of P .

We now consider the running time of constructing the octree. We believe that our current running time analysis is suboptimal, so we omit some of the details. First, there is a slight addition to the algorithm as described so far that gives us a better bound. If a box is crowded, with $O(n)$ additional time per box we determine the closest face G interfering with some F , where n is the number of faces of all dimension of P . Here for closest, distance is distance within the box. We immediately split the box nonuniformly down to a small level around the closest point of F to G , so that G does not interfere with F in one of the descendent boxes containing F . We thus get at least one uncrowded box every time we test if a box is crowded. The running time of this split is dependent on the distance in the box between F and G . This time may be amortized against the boxes created by the balance condition when balancing the octree after this split.

This allows us to claim that the total number of boxes constructed by the octree algorithm is bounded by a constant multiple of the number of leaf boxes that contain a point of P . From the warping and triangulating rules to follow, each such leaf box leads to at least one tetrahedron. In particular, the number of such leaf boxes is bounded by γ , the size of the output. Finally, the total amount of time spent on each box is at most $n \log n$ where n is the number of faces of the polyhedral region P . Thus, a bound on the running time is $O(\gamma n \log n)$. Note that this subsumes the $O(n^2)$ preprocessing to find tethers, since $\gamma = \Omega(n)$. See Section 11 for more remarks on this bound.

4 Relative sizes of adjacent octree boxes

In the last section, a box may have been protected in a vertex or edge phase of the octree generation algorithm. A protected box is not subject to the balance condition in subsequent phases, so boxes adjacent to it may be smaller than half its size. In this section we prove that the ratio of the sizes of two adjacent boxes cannot be smaller than a constant multiple of α . Here α is the smallest interior angle of P as defined in Section 2. The various kinds of boxes must be considered as separate cases.

Observe that from the centering algorithm, if the vertex was close to any face of its containing box, that face was deleted. Thus we have the following lemma.

Lemma 1 *The distance between a vertex of P and any face of a box is at least one eighth the size of the box containing the P vertex.*

Immediately after boxes are split in the vertex phase, each P vertex is a vertex cone apex in the box that contains it. That is, the only P faces in the extended box of a box containing a P vertex are superfaces of that P vertex. The extra splitting to further separate vertices decreases the size of boxes (and hence extended boxes). The merging of boxes to center the P vertices increases the size of boxes, but boxes are no bigger than they were before the extra splitting step. Hence, even after the separating and merging, the superfaces of a P vertex are the only P faces contained in the extended box of the box containing that P vertex. So each P vertex still satisfies the definition of a vertex cone apex.

4.1 Vertex cone boxes

Theorem 1 *Let B be a box with vertex cone apex v . For any box b adjacent to B , $h(b) \geq k \cdot \alpha \cdot h(B)$, where k is a constant.*

Proof. The proof divides into a series of lemmas, depending on the phase in which the parent of b was split. If the parent of b was split in a separating phase, we note that any box is split at most once in a centering phase, so that the size of b is bounded by the size of its first ancestor that was split in a crowded box phase. Let b^1 be this ancestor. So, we can ignore the separating phases, and the proof of the theorem divides into three lemmas depending on when b^1 was split, and one corollary.

Lemma 2 *If b^1 was split during the vertex box stage, then $h(b) \geq h(B)/2$.*

Proof. The balance condition ensures that $h(b^1)$ is at least one half of the size of the box containing v before the centering phase begins. ■

Lemma 3 *If b^1 was split during the edge box stage, then $h(b) \geq k \cdot \alpha \cdot h(B)$, where k is a constant and α is as in Section 2.*

Proof. The proof breaks up into two cases that consider why b^1 was split. We first consider the more difficult case, that when b^1 was split because of the balance condition. Let d be the edge crowded box that was split causing the balance condition to split b^1 . We now divide into two subcases depending on the position of d relative to B .

Subcase 1. If $I(d)$ lies entirely outside $I(\text{ex}(B))$ or intersects $\partial I(\text{ex}(B))$ the balance condition gives the result: Now matter how small a box is that intersects $\partial I(\text{ex}(B))$, a box that intersects $\partial I(B)$ will be split to a size no smaller than half the distance from $\partial I(B)$ to $\partial I(\text{ex}(B))$. This is very similar to Figure 2. If $I(d)$ actually lies outside of $I(\text{ex}(B))$, then the same bound holds.

Subcase 2. Otherwise $I(d)$ lies strictly interior to $I(\text{ex}(B))$. Observe then that for two boxes, one either contains the other or their intersection is a subset of their boundaries. So, there is a box at least half the size of d between $I(d)$ and $\partial I(\text{ex}(B))$. If d is split down to a size one half that of this buffer box, for each of the children d_1 of d , $I(\text{ex}(d_1))$ is a subset of $I(\text{ex}(B))$. Note that d may have to be split at most twice to achieve this size, so it is sufficient to bound the size of d_1 , by showing that a box d that is actually half the size of the buffer box has a lower bound on its size. That is, we may assume that d is actually half the size of the buffer box, if it was larger then this implies that b is larger.

The P vertex in B can not interfere with the edge in d , so the only way that a face can interfere with the edge contained in d is for that face to be a facet or edge of the vertex cone for B . Let f be the edge in d and g the closest face that interferes with f . Note that g may not be the face that makes the sharpest angle with f at v because of alignment odities, but of course g makes an angle β with f at least this sharpest angle.

The distance from $f \cap \text{ex}(d)$ to v is at least the distance from a face of $I(B)$ on the shortest straightline path to v . So by Lemma 1 and trigonometry, we

have that $\text{dist}(f \cap \text{ex}(d), g \cap \text{ex}(d)) \geq k \cdot \beta \cdot h(B)$. The dependence is actually on $\sin(\beta/2)$. If β is large, then the balance condition between d and B in the vertex centering step will be the limiting factor to how large d is. Otherwise, $\sin(\beta/2)$ is bounded above and below by constant multiples of β .

Since g intersected the extended box of d , and f intersected B , we also have $\text{dist}(f \cap \text{ex}(d), g \cap \text{ex}(d)) \leq 3\sqrt{3}h(d)$. Hence we have the result of the lemma for d . But since the balance condition caused b^1 to split after d was split, b^1 must be smaller than d , and we have $h(b^1) \geq 2h(d)$.

The second case is if b^1 was split because it itself was crowded. If b^1 intersects $\partial I(\text{ex}(B))$ as well as $\partial I(B)$, we see that $h(b) \geq h(B)$ as before. Otherwise, b is a special case of a box d in the interior of $I(\text{ex}(B))$ in the previous case, so that the proof of the last case proves the desired result. ■

Lemma 4 (Corollary) *If d is an edge protected box such that $I(d)$ is contained in $I(\text{ex}(B))$ and further $I(d) \cap \partial I(\text{ex}(B)) = \emptyset$, then $h(d) \geq k \cdot \beta \cdot h(B)$, where β is the sharpest interior angle an edge in B makes with any superface of v .*

Proof. This was almost proved in the proof of the previous lemma. For the completion of the proof of the corollary, we need to consider a box strictly interior to $I(\text{ex}(B))$ that contains an edge, but whose parent was split because of the balance condition and not because it was edge crowded. We note that the last time its ancestor was split because it was crowded it had the desired size bound from the proof of the theorem. If the balance condition splits this ancestor, it must first split an adjacent non-crowded box. Once this is done, the children of this adjacent box form a buffer of two same sized non-crowded boxes between the ancestor and crowded boxes. The balance condition can never penetrate the inner box of this buffer by splitting crowded boxes adjacent to the outer box. Also, splitting other boxes adjacent to the ancestor only once will not cause the ancestor to split again, so we may assume that there are two layers of same sized boxes surrounding the first generation children of the ancestor, and so the ancestor is never split again.

Hence we have the result with a constant k half as large as for the case in the lemma where the box was split because it was edge crowded. The proof of the corollary for a protected box that does not actually contain an edge but is protected because it is adjacent to a box containing an edge follows with a constant one quarter as large, from the balance condition and the fact that it is adjacent to a box where the result holds. ■

Lemma 5 *If b^1 was split during the facet box stage, then $h(b) \geq k \cdot \alpha \cdot h(B)$.*

Proof. This is the most difficult of the three lemmas comprising the proof of Theorem 1. Many cases may be handled in the same way as in the edge lemma. Difficulties arise when considering a box that contains two facets that meet at an edge, and this is where the corollary to the previous lemma is crucial. The dependence on α in the lemma arises from the possibility of small angles in the following discussion: If all of the angles considered below are large, the factor of α may be eliminated from the result. As such, the proof concentrates on the possibility that the angles considered below are small. Certain anomalies occur in the following discussion when the angles under consideration are large. Except where noted, these anomalies are easily handled, and the resulting tedious discussion omitted. In this proof there will be a series of constants $k_1, k_2, \dots$. These labels are independent of the labels of other such constants appearing elsewhere.

If b^1 was split because of the balance condition, consider the crowded facet box d that caused us to split b^1 . In fact, we consider d to be the last descendent of this box that was split because it was facet crowded.

If $I(d)$ does not intersect $I(\text{ex}(B))$ or intersects $\partial I(\text{ex}(B))$ the same argument as in the last lemma gives the desired result.

So it remains to consider the case where d is contained in the strict interior of $I(\text{ex}(B))$. Let f be the facet contained in d , and g the closest interfering facet in $\text{ex}(d)$. If f and g meet only at v and do not share an edge, then the same argument as in the last lemma, substituting faces with the same symbols, proves this lemma.

We are left with considering a facet crowded box d containing a facet f , whose closest interfering facet g shares a common edge e with f . See Figure 6. We seek to obtain a lower bound on the distance between the two facets f and g , where for distance paths are restricted to outside of protected boxes but inside $I(\text{ex}(B))$. We bound this distance by a constant times $\alpha h(B)$. Then since the box d can only be facet crowded if both f and g intersect $I(\text{ex}(d))$, the box sizes are at least a constant fraction of this distance.

We first make an argument that shows we need only consider a box adjacent to a box where the corollary to the last lemma holds. As one travels parallel to e to boxes d farther from the vertex v , but still inside $I(\text{ex}(B))$, the distance within $I(\text{ex}(d))$ between f and g increases. So d is at least half the size of some box adjacent to B containing f . We have “half” and not

equality because of possible odities in the alignment of the edge with the coordinate (octree) directions leading to nonuniformity in which boxes are protected. We now consider the first box d^1 in the facet phase, such that d^1 contains f and is adjacent to both B and a box B_e , where B_e is an edge cone box protected because it contains e or is adjacent to a box containing e .

We now describe how badly d^1 can be crowded, that is we seek to describe the distance between f and g . Because edges are well surrounded by protected boxes, the distance from e to the boundary of B_e shared with d^1 is at least a constant fraction of $h(B_e)$. We recall that since B_e is an edge protected box, it is by definition not part of the extended box of d^1 . Hence, the closest point of $f \cap \text{ex}(d^1)$ to e is on the shared boundary of B_e and d^1 . So we have that the distance from $f \cap I(\text{ex}(d^1))$ to $e \cap I(B)$ is at least $kh(B_e)$.

Let e_2 be the face making sharpest interior angle β with e at v . The ratio of the size of B_e to B is bounded above by a constant times this angle as in Lemma 4. Combining this with the distance results of the last paragraph, the angle at v between any point of $f \cap I(\text{ex}(d^1))$ and face e is at least a constant factor k_1 times the angle β from e to e_2 , as in Figure 6 left.

Let F be an imaginary facet intersecting e and e_2 . It may be that F is coincident with either f or g . In Figure 6 left F is coincident with f , in Figure 6 right F is distinct from f and g . If e_2 is a facet, we take F to intersect it forming an edge that makes angle β with e . Note that F passes between g and f in $I(\text{ex}(d^1))$.

The two right triangles at the top of Figure 6 left are similar. First, if we assume that β is less than $\pi/4$, then the ratio of their sizes is no more than $k_1/2$. If $\beta \geq \pi/4$, then it follows from the definition of crowded that $h(B_e)$ is at least a constant fraction of $h(B)$, and the result of the lemma follows immediately from the fact that e is well centered in B_e and f and g make angle at least α at e . Second, we let β'_g be the angle between $F \cap e_2$ and the half plane of g with boundary e . Note that the larger triangle's leg between F and g is at least $k_2 \cdot \beta'_g \cdot h(B)$. Combining these two results shows that the smaller triangle's leg between F and g is at least $\frac{k_1 k_2}{2} \cdot \beta'_g \cdot h(B)$.

In Figure 6 the smaller triangle's leg between f and g is completely inside B_e and makes a right angle with $b \cap \partial I(B)$, so that the the distance from $g \cap I(\text{ex}(d^1))$ to $F \cap I(\text{ex}(d^1))$ is at least the length of the smaller triangle's leg. If the orientation of B_e with respect to f and g is such that this leg is not entirely inside B_e , then this distance may be at most a constant k_3

Figure 6: The distance between two superfacets of an edge inside an extended vertex box $I(\text{ex}(B))$ but outside an edge box is at least $k\alpha h(B)$.

smaller than the leg's length, because the facet that e is well centered in B_e implies that the angle g can make with $\partial I(B_e)$ is bounded below by a constant. Hence $\text{dist}(g \cap I(\text{ex}(d^1)), F \cap I(\text{ex}(d^1))) \geq k_4 \cdot \beta'_g \cdot h(B)$ for constant $k_4 = \frac{k_1 k_2 k_3}{2}$. Similarly for f (recall that Figure 6 left is the case where F is coincident with f , so β'_f happens to be zero). We will prove below that at least one of β'_f or β'_g is at least a constant factor times an interior angle of P . Assume that β'_g is at least $k\alpha$ for some constant k . Substituting this in the above yields $\text{dist}(g \cap I(\text{ex}(d^1)), F \cap I(\text{ex}(d^1))) \geq k_5 \cdot \alpha \cdot h(B)$.

Since F passes between f and g , we have a reverse triangle inequality of distances between sets, that is $\text{dist}(f \cap I(\text{ex}(d^1)), g \cap I(\text{ex}(d^1))) \geq \text{dist}(f \cap I(\text{ex}(d^1)), F \cap I(\text{ex}(d^1))) + \text{dist}(g \cap I(\text{ex}(d^1)), F \cap I(\text{ex}(d^1)))$. By the last paragraph, one of these two distance summands must be at least $k_5 \cdot \alpha \cdot h(B)$. Hence

$$\text{dist}(f \cap I(\text{ex}(d^1)), g \cap I(\text{ex}(d^1))) \geq k_5 \cdot \alpha \cdot h(B).$$

Recall our assumption that d^1 was split because g interfered with f . In particular, f and g are contained in $I(\text{ex}(d^1))$, so $h(d^1)$ is at least a constant fraction of the distance between f and g in $I(\text{ex}(d^1))$, that is $h(d^1) \geq k k_5 \cdot \alpha \cdot h(B)$. From the definition of facet crowded, as soon as box d^1 and its

descendants are split down to a size that is a constant fraction of this distance, the descendants are not facet crowded. Hence d^1 's descendants are at least a constant fraction of this distance, and by construction so is d, b^1 , and b . That is, for some constants k and k' ,

$$h(b) \geq k' \cdot h(d^1) \geq k \cdot \alpha \cdot h(B).$$

The proof is complete except for showing that one of β'_f or β'_g is within a constant factor of an interior angle of P . We first show that the angle β'_f is an angle interior to P , but not necessarily an “interior angle” between P faces as defined in Section 2. Similarly for β'_g . If $\beta'_g + \beta'_f \geq \pi/4$, then the proof is trivial. Otherwise, we may draw a cone C with apex at v , axis along e , and $F \cap e_2$ on its boundary. See Figure 6 right for a spherical cross section of this cone, which is nearly a “top” view of Figure 6 left. Since by assumption the P face making the smallest interior angle with e is e_2 , the only P faces visible to e within this cone are f , g , and e_2 . The angle at v between $F \cap e_2$ and the intersection of this cone with f is an interior angle. But this angle is just β'_f . Similarly for g .

We finally show that one of β'_f or β'_g is within a constant factor of an interior angle between P faces. If e_2 is a facet that meets f at an edge, then the angle β'_f between $e_2 \cap F$ and $f \cap C$ may be arbitrarily small, and not bounded by alpha. We show that even if e_2 is a facet meeting f at an edge and also g at an edge, both angles β'_f and β'_g may not be small. See Figure 6 right. The proof is as follows. If one of the angles is greater than half of β , there is nothing to prove. Otherwise they are both less than half β , but then they are a constant fraction of the angle γ_f between $e_2 \cap F$ and $e_2 \cap f$ and the angle γ_g between $e_2 \cap F$ and $e_2 \cap g$, respectively. The sum of these angles is either a true interior angle between the two P edges of $f \cap e_2$ and $g \cap e_2$ meeting at v , or if it is not interior, there is a P edge inside the triangular cone formed by f , g , and e_2 but outside the circular cone C , such that this P edge makes an interior angle with f or g smaller than this sum. In either case, this sum is at least α .

For a final note in this proof, it is surprising that the result depends on β' and that there is no explicit dependence on β , where $\beta' = \beta'_f + \beta'_g$. What is hidden is the close relationship between β' and β , which may be shown by the following mental exercise. For simplicity assume that e_2 is a subedge of f , i.e. $F = f$. Fix the angle made by the facets f and g at e at a value less than

$\pi/4$. Then β' depends linearly on β . However, if β is fixed, then β' depends linearly on the angle between f and g at e . That is, provided all angles are small compared to $\pi/4$, β' is bounded above and below by the product of β and the angle between f and g at e . In this sense, β' acts as a lower bound on β , and the explicit dependence on β is subsumed by dependence on β' . It is important to realize that β' is (within a constant factor) of a true interior angle, so that the box size has dependence on α , and not α^2 . ■

This sequence of lemmas concludes the proof of Theorem 1. ■

4.2 Edge cone boxes

Theorem 2 *Let B be a box with edge cone apex v . For any box b adjacent to B ,*

$$h(b) \geq k \cdot \alpha \cdot h(B),$$

where k is a constant.

Proof. We again consider when b was last split. We know that B was not protected in the vertex centering phase, since it is an edge protected box.

If b was last split in the vertex crowded phase, regardless of whether B is protected or not, the balance condition gives us $h(b) \geq h(B)/2$. Similarly if b was split in the edge crowded phase.

So we are left to consider if b was last split in the facet crowded phase. Let f and g be the two superfacets of e . Let I be the union of the embeddings of all of the boxes protected because of e or its vertices. Then since e is well centered in its protected boxes, the distance between a point of $f \cap \partial I$ and $g \cap \partial I$ is at least $k \cdot \beta \cdot h(B)$, where β is the angle between f and g at e . We may now repeat the argument of Lemma 3, with facets f and g meeting at e taking the place of edge f and face g meeting at v . That is, a rewording of the proof that edge boxes adjacent to a vertex box have the correct size proves that facet boxes adjacent to an edge box have the right size. ■

4.3 Facet Cone Boxes

Theorem 3 *Let B be a box with facet cone apex F . For any box b adjacent to B , $h(b) \geq h(B)/2$.*

Proof. We know that B was not protected in the vertex and edge protecting phases. If b was last split in any phase, the balance condition gives us $h(b) \geq h(B)/2$. ■

If a box is not adjacent to a protected box, then it is empty, and the balance condition provides us with $h(b) \geq h(B)/2$. We may now summarize.

Theorem 4 *Let B be any box of the octree whose embedding contains a point of P , and b an adjacent box. Then*

$$h(b) \geq k \cdot \alpha \cdot h(B),$$

where k is a constant.

5 Warping

All of the boxes we will now consider are leaf boxes of the octree. If a box is adjacent to boxes that are smaller, then we replace large faces of the box with the smaller faces of the adjacent boxes where they intersect. For example, if a box is completely surrounded by boxes one-half its size, then each of the six sides of the box is now four square facets with a common vertex and shared edges, and no longer strictly speaking a cube. We *warp* the faces of a box away from faces P when necessary to achieve good aspect ratio when we triangulate. We warp box faces that are close to P -faces, where the definition of “close” is below.

The warping consists of moving box vertices. We consider adjacent boxes to share faces of their common boundary, so that when we warp a box face, there is a change in all of the adjacent boxes that contain that face. The warping rules are selected so that no inconsistencies arise.

Note that a protected vertex box (after the centering step described in Section 3.3.1) is always adjacent to boxes the same size or smaller. This means that, without loss of generality, we do not have to triangulate and warp facets of a protected vertex box; instead we can triangulate and warp the facets of the adjacent boxes.

A special note is needed in the case of protected edge boxes when the superfacets of the edge make an angle close to 360° . If an edge box B is adjacent to two boxes b and c , such that b and c are not adjacent to each

Figure 7: The special case of sharing a vertex with two boxes.

other, but $I(b) \cap I(c)$ is non-empty, then there may be two duplicate copies of a face that should be shared by one face of $I(B)$.

See Figure 7. In this case, we retain the two distinct copies of the same face. Each face is warped separately for b and c . The portion of the face for b is retained during the triangulation of $P \wedge b$ but ignored when considering $P \wedge c$. Similarly for c .

The purpose of warping is to make sure no P face comes too close to a box face. Note that P vertices are already guaranteed to be well separated from box faces because of the vertex centering procedure described in Section 3.3.1. Thus, we only need to warp for P edges and P faces.

The warping is done in two passes. The first warping pass is for box edges close to P edges. We say that a box edge E is *close* to a P edge F if the distance from E to F is less than $h/8$. Here, h is the size of the smallest box sharing the box edge E . Note that the boxes containing E are all edge-protected, since they are all adjacent to a box containing F . This means that the sizes of these boxes are all within a factor of two from each other.

For every close edge E , we move it away from F as follows. We move every point on the edge, including the endpoints, by distance $h/8$ in the direction that is orthogonal to E and F , oriented away from F . (If E and F happen to be parallel, we move the points of E in the direction orthogonal to E and coplanar with F .)

Even though the boxes after warping have complicated shapes, we still let the “size” of a box b be its prewarped edge length, and we continue to use notation $h(b)$ for this size.

The second warping pass is for box vertices close to P facets. We say that a box vertex v is *close* to a P facet F if the distance between them is at most $h/16$, where h is the size of the smallest box containing v . For such a vertex, we move it distance $h/16$ away from the facet in the direction orthogonal to the facet.

In general, because the octree algorithm separates cones, there is never an inconsistency in these rules that could cause a vertex or edge to be close to two P facets. There is one exception to this, namely, the case of a protected edge box containing two facets that make a small interior or exterior angle. If the two facets make a reflex angle, and are both close to the same vertex, then we are in the special sharing case described before: there are actually two copies of the vertex, and each is warped separately for the appropriate facet. See Figure 7. If the facets make an acute angle, and are both close to a vertex of the protected edge box, then we project the vertex orthogonally onto the plane halfway between the facets (i.e., the plane that bisects the edge cone).

Whenever we warp a vertex in the second warping pass, we carry its adjacent edges with it. Thus, box edges continue to be straight segments between box vertices during the warping. The box facets, however, are no longer well defined because the bounding edges are no longer coplanar.

We make a new definition as follows. Let an I vertex be the intersection of a P edge with a box facet, or the intersection of a P facet with a box edge, or a box vertex.

In the last paragraph we referred to the intersection of a P edge with a box facet; the box facet is no longer well defined after warping. Accordingly, we let the I vertex be the intersection of the edge with the prewarped box facet.

Lemma 6 *After the warping, no two distinct I vertices x, y are closer than $\min(t, h/16)$. Here h is the size of the smallest box containing x, y . Variable t is defined only for the case when x, y are on two facets making a sharp interior angle in a protected edge box, or when one is a box vertex in between two such facets. In this case, $t = k_1 h a$ where a is the angle between the facets, and k_1 is a constant.*

Proof. This proof has numerous cases, all similar. For the sake of brevity, we prove the lemma in two cases only.

Figure 8: One case in the proof of Lemma 6.

We consider the case that x and y both arise as the intersection of a P facet F with box edges E_1, E_2 . (This is one of the more complicated cases of the lemma.) If E_1, E_2 are disjoint edges then the lemma is clear. Points x, y before the first warping pass are separated by at least distance h if they lie on disjoint edges. After the first warping pass x, y must be separated by at least $3h/4$. After the second warping pass, x, y must be separated by at least $5h/8$.

Next, consider the other case that E_1, E_2 have a common box vertex v . (In the third case that E_1, E_2 are the same edge, then x, y are not distinct.) This case has two subcases, namely, v is not considered close to F during the second warping pass, or that v is considered close. This case is illustrated in Figure 8.

First, observe that in either subcase, the angle $\angle xvy$ is at least 60° ; this angle is originally 90° , and the warping steps cannot reduce it more than 30° .

Now, if v is not close to F during the second warping pass, this means that the altitude of triangle xvy (from v to the xy leg) is at least $h/16$, which means by the angle bound in the last paragraph that the base of this triangle, that is, the xy leg, is at least $h/16$ long.

The same argument holds if v is close to F .

Another case of the lemma is the case of a protected edge box, where x is the intersection of a P facet F with a box edge E , and y is a box vertex. First, we know that, as a result of the first warping pass, y has distance at least $h/8$ from the P edge E in the protected edge box, where h is the size of the protected edge box, or half that size (if the protected edge box is adjacent

Figure 9: Another case in the proof of Lemma 6.

to a box half its size). Now, suppose a point on F is distance less than $h/16$ from y after the first warping pass. Then we either move y distance $h/16$ from F , or distance t , where t is half the distance from z_1 to z_2 . Here, z_1, z_2 are the endpoints of a path though y orthogonal to the bisector of the two P facets in the box. See Figure 9. But since y after the first warping pass is distance at least $h/8$ from E , the same holds for z_1, z_2 . This means that the distance between z_1 and z_2 is at least $2 \sin(a/2) \cdot h/8$, and the distance from y after the second warping pass is at least $\tan(a/2) \cdot h/8$. This formula has a lower bound of $k_1 ha$. ■

6 Two dimensional triangulation

After warping, we triangulate facets of boxes. This is a preliminary step to the three dimensional triangulation algorithm of the next section. Suppose two boxes b, b' are adjacent, and suppose their intersection is a facet S . If we assume that $h(b) \geq h(b')$, then S is a facet of b' . Under these circumstances, we always triangulate S by considering it a facet of b' , not b . (If b, b' are the same size, we break ties arbitrarily). This means that we can always assume that when triangulating a square S , there is no box vertex in the interior of S . There are four box vertices at the corners of S , and there may be additional box vertices at points along edges of S .

We now describe the triangulation of a box facet S . Note that after warping, the points of S are not necessarily coplanar, although the points

are nearly coplanar because the warping distances are small.

The triangulation of a box facet S breaks down into four cases, and cases C and D have two subcases. The cases depend on which P faces pass through the facet. Note that we only have to triangulate the portion of S interior to P , which is bounded by a closed path of line segments. Call this path of line segments π . We can also think of π as a “perturbed polygon” (since the vertices are nearly coplanar). The six subcases are illustrated in Figure 10; the regions π are shaded. The points on the boundary of π are all I vertices.

Case A, no P faces pass through the facet S . In this case, we put in a new point v at the center of the prewarped box facet, and we connect v to every segment along the boundary of S .

Case B, two P facets F, G and their common edge E pass through facet S . Let v be the I vertex where E passes through the prewarped version of facet S . Then we connect v to all segments on the boundary of π , triangulating π .

Case C, one P facet F passes through (two edges of) the facet S . Let x, y be the two points where F passes through two edges of S , and let v be the midpoint of x, y . Let c be the center of S (before warping). Let h be the edge length of S (before warping).

Subcase C1, v is within $h/4$ of c . Then we connect v to every segment along the boundary of π . This divides the region into triangles.

Subcase C2, v is further than $h/4$ from c . Then we connect c to every segment along the boundary of π . This divides polygonal region π into triangles and quadrilaterals. Then we insert a diagonal (arbitrarily) into each quadrilateral, triangulating π .

Case D, two P facets F, G pass through facet S , but no P edge. Let x, y and x', y' be points where F, G respectively cross through the boundary of S . Let v, v' be the midpoints of these segments. Let c be the center of S (before warping). Let h be the edge length of S (before warping).

Subcase D1, v or v' is within $h/4$ of c . Say v is within $h/4$ of c . Then we connect v to every segment along the boundary of π . This divides π into triangles.

Subcase D2, v and v' are both further than $h/4$ from c . Then we connect c to every segment along the boundary of π . This divides π into triangles and quadrilaterals. Then we insert a diagonal (arbitrarily) into each quadrilateral, triangulating π .

We claim that this triangulation has the following property.

Figure 10: The six cases for triangulating a box facet.

Lemma 7 *Let S be a facet of a box b , triangulated with the above algorithm. Let h be the size of S , and let h' be the size of the smallest box adjacent to S . Let T be a triangle in this triangulation. If two facets pass through S , let t be as in the last lemma. Then the radius of the largest inscribed circle in T is at least $k_2 \min(h', t)$, where k_2 is a constant.*

Proof. Standard geometry tells us the radius of the largest inscribed circle in a triangle $T = \triangle xyz$ is bounded below by

$$k_3 \cdot |xy| \cdot \min(|\angle zxy|, |\angle zyx|),$$

where k_3 is a constant. This fact is used below.

This proof breaks down into cases. First, we take the case that T has three vertices pqv , such that p, q are vertices on the boundary of π , and v is as in the above cases. This is the type of triangle we get in all cases except C2 and D2. First, observe that the distance from p to q is at least $\min(h', t/2)$; this follows from the preceding lemma. Thus, we only need a constant lower bound on $\angle vpq$ and $\angle vqp$. But these angle bounds are immediate, because

Figure 11: Triangulating a quadrilateral in Case C2 (closeup) .

v is near the center of the facet, and p and q are along the boundary of the facet.

A slightly different analysis is needed for the triangle $vx'y'$ arising in Case D1, where x' and y' are vertices arising from the intersection of the P facet G and S . We know from Lemma 6 that the distance between two I vertices, one from $F \cap S$ and one from $G \cap S$, is at least t . In fact, any point of $F \cap S$ is at distance at least kt from $G \cap S$, from the fact that the common edge between F and G is far from S after warping. To be technically correct, since the vertices of S are not actually coplanar, $F \cap S$ refers to any triangulation of S intersected with F , and similarly for G . Hence the distance from v to x' is at least kt , which gives us an edge length, and the distance from v to $G \cap S$ is also at least kt , which gives us angle bounds.

A different kind of triangle arises from splitting quadrilaterals in Cases C2 and D2. For each of the quadrilaterals in these cases, we can see that the two opposite edges e_1 and e_2 have length at least $k_4 \cdot \min(h', t/2)$. See Figure 11. For e_1 , this follows from the same argument as in the last paragraph. For e_2 , this follows because facet F in S is well separated from the center c . Therefore, the length of e_2 is at least a constant fraction of the length of e_1 . Similarly, it is clear that edges e_3 and e_4 have length at least $k_4 \cdot \min(h', t/2)$.

Similarly, we can show that the angles made by edges e_3, e_4 with e_1, e_2 are bounded below by constants. Again, this follows because c is centered, whereas e_1, e_2 are away from the center.

This is enough to show that the inscribed circles of the triangles obtained as a result of inserting a diagonal are at least a constant multiple of

$\min(h', t/2)$. ■

Theorem 5 *Let B be a box. For any triangle T on a facet of B , let r be the radius of the largest inscribed circle. Then $r \geq k \cdot \alpha \cdot h(B)$, where k is a constant.*

Proof. This is an immediate consequence of the preceding lemma and Theorem 4. In particular, h' in the previous lemma is at least $k\alpha h(B)$ by Theorem 4. Moreover, t in the previous lemma was shown to be at least $k \cdot \alpha \cdot h(B)$ in the proof of Lemma 6. ■

General results have been obtained for two-dimensional triangulations with guaranteed inscribed-circle radius bounds. Bern, Edelsbrunner, Eppstein, Mitchell, and Tan [1991] have a result concerning optimal two-dimensional triangulation of polygons, assuming new points cannot be introduced. Bern, Dobkin, and Eppstein [1991] have similar results in the case that new points can be introduced. We have not been able to incorporate these algorithms in the present version of our triangulation algorithm because we need to introduce new points, but the new points have to be introduced in a carefully controlled fashion.

7 Three dimensional triangulation

We now describe how to form tetrahedra from the triangles in the last section. We triangulate on a box by box basis. The details of how we triangulate depends on a case analysis of what is contained in a box, and how it intersects the box. However, the general principle is to find a central vertex, and then form one tetrahedron or prism for each triangle in the box by taking the convex hull of this vertex and the triangle. The organization of the argument is very similar to the two-dimensional triangulation in the last section.

7.1 Empty box tetrahedra

We first describe how to triangulate in three dimensions a box b containing no P faces. We take as a central vertex v the centroid of b . (The prewarping centroid may be used). For every triangle on the surface of the box we form a tetrahedron by taking the convex hull of v and the triangle.

7.2 Vertex box tetrahedra

For a box b containing a vertex of P , the three-dimensional triangulation is particularly easy. We take as the central vertex v the P vertex itself. We form tetrahedra by taking the convex hull of v each triangle on the surface of b . The vertex centering phase ensures that the distance from the central vertex to the triangles is at least a constant times the box size.

7.3 Edge box tetrahedra

We now describe the next case, that of triangulating in three dimensions the protected edge boxes that contain edges. Let edge E be the edge contained in box b . We let v be the midpoint of $E \cap I(b)$. We now connect v to every triangle on the surface of b , creating tetrahedra.

7.4 Facet box tetrahedra

Consider a box b containing one P facet F . We introduce a central vertex v at the centroid of the prewarped box. If v is within distance $h(b)/4$ to F , then we move v to the closest point of F to v .

We now have three cases. The first case is if v lies in the interior of P , but not on F . Here we form tetrahedra by taking the convex hull of every triangle of F and v . This requires the creation of a triangulation on the intersection of F with b , which is a polygonal region. This polygonal region is nearly convex (it would be convex if the triangles on each the surface of the box were coplanar). To form a triangulation of $F \cap I(b)$, we pick a centrally located point on the $F \cap I(b)$ and connect it to all the edges. Using similar arguments as in the previous section, it is not hard to show that the largest inscribed radius of any triangle in this two-dimensional triangulation is at least a constant fraction of h' , where h' is the size of the smallest box adjacent to b .

The second case is that v lies on F . We take the central vertex v , and proceed as if it were a vertex of P . That is, for each triangle of $I(b)$, we generate a tetrahedron by taking the convex hull of it and v .

The third case is if v lies outside of P . We form tetrahedra by taking the convex hull of v and every triangle of $I(b)$. These tetrahedra are clipped by F into “prisms” with nonparallel sides, and a triangular top and bottom.

Zero, one or two of the vertices of the top may also be vertices of the bottom. If two vertices are shared, then the prism is a tetrahedron. If one vertex is shared, the prism is split into two tetrahedra by introducing a facet between the shared vertex, and one distinct vertex of the top and one distinct vertex of the bottom. If no vertices are shared between the top and bottom, the prism is split into three tetrahedra by introducing two facets. One facet is between two of the top vertices and the opposite bottom vertex. The other facet is between the opposite bottom vertex, one of the other bottom vertices, and the opposite top vertex.

7.5 Two-facet box tetrahedra

Perhaps the most complicated case is the case of a box b containing two facets F_1, F_2 adjacent to an edge E , but the edge itself is not in the box. Let v be the centroid of box b (prewarped). If v is within distance $h(b)/4$ from F_1 or F_2 , then we move it to the closest point on either of these facets (break ties arbitrarily). Say F_1 is closer.

There are several possible arrangements for v, F_1, F_2 . For example, v could be on F_1 , or it could be between F_1, F_2 , or it could be outside P , with F_1 facing away from v and F_2 facing towards P .

For each facet F_i , ($i = 1, 2$) we triangulate $F_i \cap I(b)$ provided that F_i does not contain v and faces v . This means that either one or both of F_i is triangulated. We use the same algorithm described in the facet box case for triangulating the facet.

Now we connect v to all the triangles on the boundary of b that are interior to P . We also connect v to all the triangles of F_1, F_2 .

If v is contained in F_1 , or if v is between F_1, F_2 , then we have now divided the box into tetrahedra. If v is outside P with F_1 facing away, then by connecting v to the triangles of F_2 we have divided the interior of P into tetrahedra and prisms. We now split the prisms into tetrahedra as in the facet box case.

8 Aspect ratio of tetrahedra in $\triangle^{OCT}$

Our triangulation is optimal in two respects. First, the maximum aspect ratio among tetrahedra of our triangulation is optimal up to a constant mul-

tiple. Second, compared to all other triangulations of fixed aspect ratio, our triangulation has the minimum number of tetrahedra up to a constant multiple. In this section and the next section we establish the first optimality property, focusing on the optimality of the aspect ratio.

Recall that the tetrahedra we form are either the convex hull of a central vertex and a triangle that does not contain the central vertex on the surface of a cube or a facet of P , or else a portion of a prism that is divided into three tetrahedra. We now want to argue about the aspect ratio of all tetrahedra arising in the algorithm of the previous section.

There are three types of tetrahedra arising in the triangulation. A *Type A* tetrahedron has one vertex v centrally located in the box, and the other vertices on a box face. This type of tetrahedron arises in the cases of triangulating a box with a vertex, a box with an edge, or a box with a facet in which the vertex v in the last section lies near the center of the box, is in P , and is not close to the second facet (if the box has two facets). Also some of the tetrahedra arising from the case of two facets in a box and v on one of the facets are of this type.

A *Type B* tetrahedron arises only in the two-facet case in the last section, when the vertex v is on one facet, and the triangle it is joined to lies on the other P facet. These tetrahedra arise in cases analogous to one of the triangles in case $D1$ of Figure 10.

A *Type C* tetrahedron arises from a vertex v in the last section outside P . This happens only with boxes with one or two facets and no edges. This causes $P \wedge b$ to be divided into prisms, and then each prism is divided into tetrahedra. These tetrahedra arise in cases analogous to triangles in cases $C2$ and $D2$ in Figure 10.

Intuitively, all three types of tetrahedra have aspect ratio at most $1/\alpha$, because they involve a base (the base being a triangle) whose inscribed radius is between $k_1\alpha h(b)$ and $k_2h(b)$, and the apex is well-centered over the base, and is distance between $k_3\alpha h(b)$ and $k_4h(b)$ from the base. Here, $k_1, \dots, k_4$ are constants.

The aspect ratio computations are very technical and involved. We first prove the aspect ratio bounds for Type A tetrahedra using a construction of an inscribed sphere. The proofs for Type B and C tetrahedra follow from a different analysis of the same construction. A Type A tetrahedron has a base size between $k_1\alpha h(b)$ and $k_2h(b)$, and has a height at least $k_5h(b)$.

Consider a particular box b and its central vertex v , and a Type A tetra-

hedron t_3 in this box. In this section we let subscripts denote dimension, so that t_2 is one of the triangles on the surface of $P \wedge b$, and t_3 is the tetrahedron formed by the convex hull of t_2 and v .

Lemma 8 *The smallest containing sphere, S_3 , of t_3 has radius R_3 at most $kh(b)$, where k is a constant.*

Proof. The smallest containing sphere of t_3 has radius less than the radius of a sphere that encloses the entire box. If the box was unchanged by the warping step, we see that a sphere with radius equal to one half a diagonal of the box will do. If the box was warped, then it may be enclosed in a concentric box with side lengths $5/4$ times that of the original box by the definition of “close.” Thus we can take $k = 5\sqrt{3}/4$. ■

This concludes the analysis of the smallest sphere containing t_3 . Next, we analyze the largest inscribed sphere of t_3 by actually constructing an inscribed sphere. Consider the triangle t_2 on $\partial(P \wedge b)$. We form a new curved triangle u_2 on the surface of a sphere as follows. Let a be the closest point of t_2 to v . It does not matter if a is a vertex of t_2 or not, although in the figures of this section a is a vertex. The intersection of t_3 and the sphere centered at v with radius $\text{dist}(v, a)$ is the curved triangle u_2 . We let u_3 be the convex hull of u_2 and v . Note that u_3 is contained in t_3 , and has three flat facets, each one contained in a facet of t_3 that intersects v . This is illustrated in Figure 12. We will construct a sphere inscribed in u_3 instead of t_3 .

Lemma 9 *The distance from v to its base t_2 is at least $h(b)/8$.*

Proof. This is because, by assumption for the case, t_3 is a Type A tetrahedron. Therefore, v is centered in b . ■

Lemma 10 *Let u_2 be defined as in the last paragraph. Then the radius of the largest inscribed curved circle in u_2 is a constant fraction of the radius of the largest inscribed circle in t_2 .*

Proof. Let C be the largest inscribed circle in t_2 . We project this circle onto sphere S_3 . This yields an ellipse, whose minor axis is the product of the original inscribed radius and a factor depending on the angle between t_2 and a tangent plane to u_2 . We claim that the angles between v and any edge

Figure 12: The flat (t_3) and curved (u_3) tetrahedra.

of t_2 are bounded below by constants. This is because, before warping, the vertices of t_2 were coplanar with a box facet, and v is at least distance $h/4$ from that facet, where h is the edge length of the box. After t_2 is warped, since the warping distances are small, this can only change its inclination with respect to the original facet by a small constant angle that is arrived at via a case analysis (which we omit).

This same angle bound applies to the angle between t_2 and the tangent plane. ■

Lemma 11 (Corollary) *If v is within $k\alpha h(b)$ of the plane of t_2 , then the radius of the largest inscribed curved circle in u_2 is a constant times α times the radius of the largest inscribed circle in t_2 .*

Proof. The key observation is that v is still within a constant of the box size to the farthest point of the triangle. The angle between v and an edge of t_2 is bounded below by a term linear in the ratio of the distances from v to the plane of t_2 and from v to the farthest point of t_2 . In the lemma this was a constant, here in the corollary it is bounded by α . This will be used only for the analysis of Type B tetrahedra. ■

Figure 13: A sphere tangent to three facets of a tetrahedron.

In a series of lemmas we derive a lower bound for the radius of the largest sphere inscribed in t_3 by the radius of the largest curved circle inscribed in u_2 . Let s_3 denote the largest sphere inscribed in u_3 , r_3 its radius, and c_3 its center. Similarly, we let s_2 denote the largest circle inscribed in u_2 , r_2 its radius, and c_2 its center.

Lemma 12 *The sphere s_3 is tangent to u_3 at points equidistant from v .*

Proof. Pick a facet f of u_3 adjacent to v . Consider the triangle formed by v , the point of tangency between s_3 and f , and the center of s_3 . Note that this is a right triangle. Observe that the length of the hypotenuse and the length of one leg depends only on the position of s_3 and not on on which facet was chosen. Therefore, the lengths of the other leg (the leg between v and the points of tangency) is the same regardless of which facets is selected. See Figure 13. This theorem applies to arbitrary tetrahedra with any vertex chosen for v . ■

Lemma 13 *There is a sphere, s'_3 , that is tangent to u_3 at exactly the same three points as s_2 is tangent to u_2 .*

Proof. We call these points of tangency a, b, e . We construct s'_3 . Note that s'_3 is not contained in u_3 or t_3 , but will serve as a guide for finding an inscribed sphere. We first determine uniquely the center, c' , of s'_3 . Consider the line $\overline{vc_2}$. Now c' is the unique point on $\overline{vc_2}$ on the opposite side of c_2 as v , such

that $\angle vac'$ is right. Additionally, since $\overline{ac_2}$ is perpendicular to the boundary of u_2 , $\overline{ac'}$ is also perpendicular to the facet, f , of u_3 containing v and a in the plane tangent to the vertex sphere at a . Hence, $\overline{c'a}$ is perpendicular to the facet f . Similarly for b and e . The hypotenuse-leg theorem gives us congruent triangles, and we see that the distances from c' to each of a, b, e are equal. So, c' is indeed the center of a sphere with the desired properties, s'_3 . ■

Lemma 14 (Corollary) *Any sphere tangent to the three faces of u_3 containing v has center on $\overline{vc'}$, and points of tangency on $\overline{va}$, $\overline{vb}$ and $\overline{ve}$.*

We now seek to find a sphere inscribed in u_3 , which is almost as large as the largest inscribed sphere. Let d be the distance from v to the flat triangle t_2 . We see that we can choose an inscribed sphere tangent to the three facets of u_3 meeting at v , with radius $\tilde{r}$, and distance from center to v equal to $\tilde{l}$, such that

$$d = \tilde{l} + \tilde{r}.$$

See Figure 14 for a picture of the two dimensional case of this construction.

Let θ be the angle between $\overline{vc'}$ and $\overline{va}$. It is obvious from the figure that $\tilde{l}/\tilde{r} = \csc(\theta)$, so that

$$\tilde{r} = \frac{d}{1 + \csc(\theta)}.$$

We note that the largest inscribed sphere has radius at least that of this inscribed sphere, $r_3 \geq \tilde{r}$. Note that $0 < \theta < \pi/2$. Hence $\frac{1}{1+\csc(\theta)} = \frac{\sin(\theta)}{\sin(\theta)+1} > \sin(\theta)/2 \geq k\theta$. Hence

$$r_3 \geq kd\theta.$$

Also from Lemma 10 and Figure 14,

$$\theta = r_2/d.$$

Recall the result of Theorem 5,

$$r_2 \geq k\alpha h(b).$$

Hence $r_3 \geq k\alpha h(b)$, and we have proved the following theorem for Type A tetrahedra.

Figure 14: Constructing a large inscribed sphere of u_3 .

Theorem 6 (Aspect ratio of Δ^{OCT}) *Let B be a box. For any tetrahedron t inside B , the aspect ratio of t is at most k/α , where k is a true constant, and α is the sharpest interior angle of P .*

This theorem is the main result for this section, which the preceding lemmas prove for Type A tetrahedra. We now wish to establish this theorem for the remaining types of tetrahedra, Types B and C.

We consider Type B tetrahedra briefly because the analysis is so similar to that for Type A tetrahedra. We perform the same construction as for Type A tetrahedra, defining d and θ as before. Consider a Type B tetrahedron, the convex hull of vertex v on P facet F and triangle T on the other P facet G in box b . If the distance from v to G is at least $kh(b)$, then the analysis of Type A tetrahedra suffices. Otherwise, we have $d \geq k\alpha h(b)$ as in Lemma 6. As before, the construction gives $r_3 \geq kr_2$. Because d may be closer to the plane of T , the result of Lemma 10 does not hold, and instead we use the result of Lemma 11. This is sufficient because we have a stronger bound on the shortest altitude of T than when considering Type A tetrahedra, namely $alt(T) \geq kh(b)$.

We now consider Type C tetrahedra. These tetrahedra arise from prisms. Since the central vertex v is well centered in the box, and no P facet is close to it (and no P edges are in the box), the angle γ between the shortest segment from a triangle T to v and the normal to the plane of T is no more than a constant slightly larger than $\pi/4$, in contrast with the Type B tetrahedra. We will make several constructions as for Type A tetrahedra, only we will choose different vertices for “ v ”. However, our choice of “ v ” will be related to the box’s central vertex, so that Lemma 10 will hold for Type C tetrahedra. Hence we have $r_3 \geq kr_2 \geq k\alpha h(b)$, and we now need only

specify the constructions.

Consider the various types of prisms as in Section 7.4, having a triangular top and bottom, and one, two, or three quadrilateral sides depending on how many vertices of the top are shared by the bottom. Prisms formed by clipping the convex hull of v and a P triangle have zero shared top and bottom vertices, and will be considered last. Otherwise, if two vertices of the top of a prism are shared with the bottom, then the prism is a tetrahedron. We take the top to be a box triangle. We now perform the same construction as for Type A tetrahedra, with the non-shared bottom vertex as “ v ”.

If one vertex of the top is shared with the bottom, then the prism is split into two tetrahedra. Define the top as in the last type of prism. The bottom is a projection of the first onto a P facet, with center of projection v far from the top and bottom. Recalling that γ is bounded by a constant, the bottom altitude is at least a constant fraction of the top altitude, $k\alpha h(b)$. For the tetrahedron that is the convex hull of the top and a non shared bottom vertex, we chose “ v ” to be the bottom vertex. For the tetrahedron that is the convex hull of the bottom and a non shared top vertex, we chose “ v ” to be the top vertex.

If no top vertex is shared with the bottom, then the prism is split into three tetrahedra. We let the top be the box or P triangle. As in the previous case, the bottom triangle has altitude at least a constant fraction of the top altitude, $k\alpha h(b)$. The analysis of the one shared vertex case is sufficient for the tetrahedra with the top as a facet and the tetrahedra with the bottom as a facet.

For the third tetrahedron, the one that is the convex hull of an edge of the top and an edge of the bottom, we chose the bottom vertex opposite the two top vertices as “ v ”. The triangle of the tetrahedron opposite t in this subcase arises from introducing a diagonal on one of the quadrilateral sides of the prism. If both the top and bottom are subsets of P facets, or in general if the box is a protected edge box not containing an edge, the P edge is at least $kh(b)$ away from t , and hence t has altitude at least $k\alpha h(b)$. If the top is a subset of a box facet, then t has altitude at least $kh(b)$. In both the case where the top is a subset of a P face and in the case where the bottom is a subset of a box face, the angle between the planes of the top and bottom is between 0 and a constant slightly larger than $3\pi/4$, by the fact that neither plane is close to the central vertex, the central vertex is outside of P , and both planes pass through the box. For the definition of angle between the

planes, we chose the normal to the bottom pointing towards the interior of the prism, and the normal to the top pointing towards the exterior of the prism. In particular, the angle between the shortest segment from t to “ v ” and the normal to t is bounded by a constant, so that again Lemma 10 holds.

9 Aspect ratio tetrahedra in $\triangle^*$

In the last section we established an upper bound on the aspect ratio of our triangulation. In this section we show that we are a constant factor from optimal, by analyzing the aspect ratio of the optimal triangulation.

We wish to show that $1/\alpha$ acts as a lower bound on the aspect ratio of any triangulation of P , where α is the smallest interior angle of P . Note that if $\alpha \geq \pi/2$, then in any triangulation of P there is trivially a tetrahedron with aspect ratio at least k/α for some constant k , since aspect ratios are always greater than 1. Hence, we need only prove the following theorem

Theorem 7 *Let P have smallest interior angle α . If $\alpha < \pi/2$, any three dimensional triangulation of P has a tetrahedron with aspect ratio at least $1/\sin(\alpha)$.*

Before starting the proof, we state the following simple geometric lemma, whose proof is omitted.

Lemma 15 *Let Q be a plane in $\mathbb{R}^3$, and let x be a point on Q . Let r_1, r_2 be two rays starting from point x such that plane Q separates them (or one or both r_1, r_2 may be contained in Q). Then the angle between r_1 and Q is bounded above by the angle between r_1 and r_2 .*

Proof of Theorem 7. We first show that for any three dimensional triangulation of P , there is some tetrahedron of the triangulation with angle at most α between two edges or an edge and a facet not containing the edge. Consider the faces F and G of P that intersect at p , with the angle between F and G equal to α . Note that we assume that F and G can see each other in the sense of Section 2. Recall that we may always chose F to be an edge. If F makes angle α with any edge, we let that edge be G , otherwise G is facet. In either case there is some segment in the relative interior of G along the ray in G making angle α with F . See Figure 15. We let G_l be this segment,

most half of the altitude between p' and H , where as before p' is the vertex of F' that is not p . This altitude is $\sin(\alpha')/|F'|$, where $|F'|$ denotes the length of edge F' . Hence we have an upper bound on the size of the largest sphere inscribed in Δ_p^* , and we now find a lower bound on the size of the smallest sphere containing Δ_p^* . Since a containing sphere must have diameter at least the longest edge, we have that the smallest containing sphere has radius at least $|F'|/2$. Combining these two shows that the aspect ratio of Δ_p^* is at least $1/\sin(\alpha')$, which is at least $1/\sin(\alpha)$ as desired. ■

Let T be a three dimensional triangulation. Let $A(T)$ be its aspect ratio, defined to be the maximum over all tetrahedron $t \in T$, of the aspect ratio of t . We may now state the theorem concerning the optimality of the aspect ratio of our triangulation of P .

Theorem 8 (Aspect ratio optimality)

$$A(\Delta^{OCT}) \leq kA(\Delta^*)$$

where k is a true constant, independent of P and α , and Δ^{OCT} is the triangulation of P arising from our algorithm, and Δ^* is any other triangulation of P .

10 Optimality of the cardinality of Δ^{OCT}

In this section we prove that the number of tetrahedra in our triangulation is no more than a constant factor larger than any other triangulation with a bounded aspect ratio. The reasoning in this section is as follows. We first prove that any triangulation with bounded aspect ratio, which we denote Δ^* , has certain geometric properties concerning how fast the sizes of the tetrahedra can change. We then derive lower bounds on how small the tetrahedra of our triangulation Δ^{OCT} can be. We then compare the two sets of bounds to show that our triangulation is within a constant factor of optimal.

Conformal. Any triangulation Δ of P that we consider in this paper must be *conformal* to P . This means that the following condition must hold. Let x be a point on a P face F . Let F' be the lowest dimensional face of Δ containing x . Then $F' \subset F$. It is easy to see that the triangulation we construct has this property; vertices of P are covered by vertices of Δ^{OCT} ,

and edges of P are covered by edges of Δ^{OCT} . A more general, but equivalent definition of conformality is the following lemma, whose proof is obvious from the definition.

Lemma 16 *If a triangulation Δ is conformal to P , then the following condition is satisfied. Let F, G be two faces of P , and x, y two points on F, G . Let F', G' be the lowest dimensional faces of Δ containing x and y . Then $F' \cap G' \subset F \cap G$. In particular, if F, G are disjoint, so are F', G' .*

In this section there will be a series of positive constants $c_1, c_2, \dots$ depending only on the maximum aspect ratio A of the triangulation under consideration. For example, the maximum number of faces incident to a vertex of Δ^* is bounded by a constant c_1 .

Characteristic length function. We define $f_T : P \rightarrow \mathbb{R}$ to be the *characteristic length function* of a triangulation Δ^T . This function is defined as follows. If x is a point of P , then we define $f_T(x)$ to be the longest edge among all tetrahedra that contain x .

The next group of lemmas focus on a triangulation Δ^* of bounded aspect ratio and its characteristic length function f_* . We will show below that the longest edge of any tetrahedron containing x is bounded by a constant times the longest edge of any other tetrahedron containing x , so that any tetrahedron containing x can be used to bound $f_*(x)$ above and below.

Lemma 17 *Suppose $x, y \in P$, and suppose the two tetrahedra defining $f_*(x)$ and $f_*(y)$ are T_1 and T_2 (in particular, $x \in T_1$ and $y \in T_2$). If T_1 and T_2 share a common edge, then*

$$c_2 \cdot f_*(y) \leq f_*(x) \leq c_3 \cdot f_*(y).$$

Proof. Let e denote the common edge. Note that because of the aspect ratio bound of Δ^* , the ratio of the longest to shortest side of any simplex in Δ^* is bounded above and below. Hence

$$|e| \leq f_*(x) \leq c_4 \cdot |e|.$$

This also holds for y . ■

Wedge. A *wedge* is a set of tetrahedra, $T_1, T_2, \dots, T_k$, that share a common vertex.

Note that the bounded aspect ratio of Δ^* bounds the smallest interior angle of any simplex in Δ^* , so that the number of simplices in a wedge is bounded by a constant c_1 .

Lemma 18 *Given a bounded aspect ratio triangulation Δ^* , two points x, y , and any two faces F, G containing x, y respectively, define $H = F \cap G$. If $H \neq \emptyset$, then $f_*(x) \geq c_5 \cdot f_*(y)$.*

Proof. Note that H is a face of Δ^* . Let v be a vertex of H . We consider the tetrahedra of a maximal wedge with common vertex v . Since no vertex of P is degenerate, between any two tetrahedra there is a sequence of tetrahedra of the wedge such that each shares an edge with its successor. Since wedges consist of a bounded number of tetrahedra, the function f_* on any tetrahedron is within a constant multiple of f_* on any another by Lemma 17. This wedge includes F and G . Finally, there are tetrahedra T_1, T_2 , such that T_1 contains x and determines $f_*(x)$ (i.e., the longest edge of T_1 is equal to $f_*(x)$), and similarly T_2 determines $f_*(y)$. Then T_1 and F have a common vertex (since they both contain x) and therefore the lengths of their edges are within a constant of each other by the same argument. The same holds for T_2 and G . ■

Geodesic distance. We let $\text{dist}_P(F, G)$ denote the *geodesic distance* in P between two closed subsets F, G of P . In other words, this is the length of the shortest path in P between a point of F and a point of G . This distance is always at least as great as the Euclidean distance.

As an example of this notation, we have the following lemma. We omit the proof of the lemma, which follows from the meaning of aspect ratio for a tetrahedron.

Lemma 19 *Let T be a tetrahedron of a bounded aspect ratio triangulation, let x be a vertex of T , and let F be a subface of T not containing x . Then $\text{dist}_P(x, F) \geq c_6 \cdot f_*(x)$.*

Similarly, the following lemma follows from elementary trigonometry applied to bounded aspect ratio.

Figure 16: An illustration of the conditions stated in Lemma 20.

Lemma 20 *Let T be a tetrahedron with bounded aspect ratio. Let F be an edge or facet of T , and let x be a point on F whose distance from any subface of F is at least t . Let y be a point lying on any facet G of T , such that G does not contain F . Then the distance from y to x is at least $c_7 t$.*

Proof. An illustration of the claim made by this lemma is in Figure 16 in the case that F is a facet. If G is a subface of F , then the lemma is trivial (we could take $c_7 = 1$). Otherwise, let v be the point on the boundary of F closest to x . (Note that if F is an edge, v will be one of its endpoints). Consider the triangle $\triangle xvy$. Let $\angle FvG$ denote the interior angle between F and G . Then $\angle xvy \geq \min(\angle FvG, \pi/2)$. But $\angle FvG$ is bounded below by a constant c_8 depending on the aspect ratio of the tetrahedron. Moreover, the vx leg of this triangle has length at least t by assumption. This means that the distance from x to y is at least $t \sin c_8$. ■

We now wish to establish the following theorem about the distance between faces in a triangulation. The theorem has two cases, depending on whether the faces have a common point or not.

Theorem 9 *Let F, G be two faces of Δ^* . Let $H = F \cap G$. Let x, y be two points such that $x \in F$, $y \in G$. Then*

1. *If $H \neq \emptyset$ and $\text{dist}_P(x, H) \geq t$, then $\text{dist}_P(x, y) \geq c_9 t$.*
2. *If $H = \emptyset$, then $\text{dist}_P(x, y) \geq c_{10} f_*(x)$.*

Proof. This proof is broken down into eight cases, depending on whether we are in Case 1 or 2 or the theorem, and depending on whether F is a vertex, edge, triangle or tetrahedron. The cases are labeled A1 to D2. Note that if we are in Case 1 of the theorem, we can assume that H is a proper subset of both F and G . If H is not a proper subset, say $F = H$, then $\text{dist}_P(x, H) = 0$, and the claim is trivial. Similarly, if $G = H$ then $y \in H$ so $\text{dist}_P(x, y) \geq \text{dist}_P(x, H)$.

In each of the upcoming cases, let p be the shortest path in P from y to x .

Case A1, F is a vertex and $H \neq \emptyset$. This is a trivial case described in the previous paragraphs, since H must be the same vertex as F .

Case A2, F is a vertex and $H = \emptyset$. Let T be the first tetrahedron on p such that T contains F . Then T must be entered on a face F' of T not containing vertex F (since G does not contain F). Thus, from Lemma 19, we know that $\text{dist}_P(F, F') \geq c_6 f_*(x)$.

Case B1, F is an edge and $H \neq \emptyset$. Then G is an edge or a triangle, and H is a vertex v common to F and G . Note that t is bounded above by $f_*(x)$ since t cannot be longer than the length of F . Let w be the other endpoint of F . We take two subcases: either x is within distance $c_7 f_*(w)/2$ of w or it is further than this distance. If x is within $c_7 f_*(w)/2$ of w , then we apply Case A2 to get a lower bound of $c_6 f_*(w)$ on the distance from y to w . This means that the distance from y to x must be at least $c_6 f_*(w)/2$, which is at least $c_{11}t$. (Note that $f_*(w) \geq f_*(x)$ since any tetrahedron containing x will also contain w .)

The other subcase is that x is distance at least $c_6 f_*(w)/2$ from w and at least t from v . Let T be the first tetrahedron on the path p that contains edge F . Say that tetrahedron is entered at point y' on face F' . Then F' does not contain F . By Lemma 20 the distance from y' to x is at least

$$c_7 \min(c_6 f_*(w)/2, t).$$

Thus, in all cases, the distance from F to G is at least $c_{12}t$.

Case B2, F is an edge and $H = \emptyset$. This is handled with exactly the same argument as Case B1, namely, we observe that either x is within $c_6 \cdot f_*(w)/2$ of an endpoint w of F , in which case the distance bound can be reduced to A2, or else x is further than this distance from either endpoint, in which case we use Lemma 20. Again, in this case, we get a distance bound of $c_{12}f_*(x)$ on the distance from x to y .

Case C1, F is a triangle and $H \neq \emptyset$. Let us take two subcases:

Subcase C1(a), H is an edge of triangle F . This means that G is a second facet containing edge H . Then we first check whether x is within distance $\min(c_{12}, 1)t/3$ from a point y' on one of the other two edges E_1, E_2 of triangle F . If so (say y' is on E_1), then y' is at least distance $2t/3$ from H , so we use Case B1 (applied to E_1 and G) to get a lower bound of $2c_{12}t/3$ on the distance from y' to H , and this yields a lower bound of $c_{12}t/3$ on the distance from x to H . Otherwise, x is further than distance $\min(c_{12}, 1)t/3$ from all three edges of F . Then we look at the first tetrahedron T on path p that has F as a facet. This tetrahedron must be entered by a point y' not on F . Then Lemma 20 tells us that the distance from y' to x is at least $c_7 \min(c_{12}, 1)t/3$.

Subcase C1(b), H is a vertex of triangle F . Let E_1, E_2 be the two edges of F containing vertex H , and let E_3 be the third edge. If x is within distance $\min(c_{12}, 1)t/3$ from a point y' on either E_1 or E_2 , we can prove the theorem as in Subcase C1(a) by applying the argument of Case B1 to E_1 and G , or to E_2 and G . Another possibility is that x is within distance $c_{12}f_*(x)/2$ from a point $y' \in E_3$, in which case we can apply Case B2 to E_3 and G , since they are disjoint. This gives a lower bound of $c_{12}f_*(x)$ on the distance from y to G . Otherwise, the distance from x to any of the edges of F is bounded below by either $\min(c_{12}, 1)t/3$ or $c_{12}f_*(x)/2$. Let T be the first tetrahedron on p that has F as a facet. Then T is entered from a point y' not on F . Finally, Lemma 20 gives us a lower bound of the form

$$c_7 \min(c_{12}f_*(x)/2, \min(c_{12}, 1)t/3)$$

on the distance from x to y .

Cases D1–D2, F is a tetrahedron. These cases are not needed by the upcoming theorems, so we omit them and leave them to the reader. ■

Theorem 10 *There exist constants c_{13} and c_{14} dependent on A such that for all $x, y \in P$,*

$$f_*(x) \leq \max(c_{13} \cdot f_*(y), c_{14} \cdot \text{dist}_P(x, y)).$$

Proof. If x and y are in a common tetrahedron of Δ^* , then the result follows from Lemma 18.

Otherwise, let $\overline{xy}$ denote the shortest path in P from x to y . Label the triangles of Δ^* crossed by $\overline{xy}$ by $F_i, i = 1, 2, \dots, m$. We assume in

this proof that $\overline{xy}$ intersects each F_i in a single point; the degenerate case of exact alignment by F_i and $\overline{xy}$ can be analyzed by considering a slight perturbation to $\overline{xy}$. Note any two consecutive triangles of the sequence are in a common tetrahedron. Let x_i denote $F_i \cap \overline{xy}$. By Lemma 18 we know that $f_*(x) \leq c_5 \cdot f_*(x_1)$ and $f_*(x_m) \leq c_5 \cdot f_*(y)$.

Define

$$i^* = \max\{i : F_i \cap F_1 \neq \emptyset\}.$$

If $i^* = m$, then x_1 and x_m are in a common wedge, and $f_*(x_1) \leq c_5 \cdot f_*(x_m)$. Combining these inequalities shows that $f_*(x) \leq c_{13} \cdot f_*(y)$. The other case is that $i^* < m$, and we note that $\text{dist}_P(x_1, x_{i^*+1}) \geq c_{10} \cdot f_*(x_1)$ by Theorem 9. Moreover, $\text{dist}_P(x, y) \geq \text{dist}_P(x_1, x_{i^*+1})$. Combining these inequalities gives the result that $f_*(x) \leq c_{14} \cdot \text{dist}_P(x, y)$. ■

We now observe the following characterization of our triangulation Δ^{OCT} .

Lemma 21 *Let $x \in P$ be a point lying in octree box b . Then $f_{OCT}(x) \geq k_1 \alpha h(b)$, where k_1 is a constant.*

Proof. Let T be the tetrahedron of Δ^{OCT} lying in b containing x . Then each edge of T on the boundary of $P \wedge b$ has length at least $k_1 \alpha h(b)$ as determined in previous sections. A case analysis verifies that every tetrahedron we construct in b has an edge on the boundary of $P \wedge b$. This includes those tetrahedra arising from prisms. ■

We now have a lemma about geodesic distances in boxes. Note that in a crowded box b , it is possible to have two points x, y in b such that $\text{dist}_P(x, y)$ is arbitrarily larger than $h(b)$. For example, P might be shaped like a coil inside b . There are two cases, however, when we can get bounds on geodesic distances.

Lemma 22 *Let b be an uncrowded box in the octree, and let x, y be two points in $P \wedge b$. Then*

$$\text{dist}_P(x, y) \leq 2\sqrt{3}h(b).$$

Proof. Recall that in a box, b , $P \wedge b$ is a single component C , so x, y are both in this component. Either there is a vertex, edge or facet cone in the box, or $b \subset P$. In the latter case, the distance from x to y in b is the same as the distance from x to y in P , which is bounded by $\sqrt{3}h(b)$.

The other case is that b contains a cone. Let z be a point on the apex of the cone. Then x and y are both connected to z by a straight-line path in P , so $\text{dist}_P(x, z) \leq \sqrt{3}h(b)$ and $\text{dist}_P(y, z) \leq \sqrt{3}h(b)$. ■

We wish to establish a series of lemmas that prove that the characteristic length function at a point of Δ^* is bounded by a constant times the size of the octree box containing that point in our algorithm. The first lemma involves a construction that is useful for each of the remaining lemmas. The remaining lemmas prove the result for a point in a protected vertex box, protected edge box, protected facet box, and finally for an unprotected (empty) box.

Lemma 23 *Let b_1 be a box split in the vertex, edge or facet phase, and x a point of P face F in b_1 . Then there exists a point y of P face G such that*

1. *G does not contain F , and*
2. *$\text{dist}_P(x, y) \leq 6\sqrt{3}h(b_1)$.*

Proof. If b_1 is split because of the balance condition, let b_l be the crowded box that caused it to split. There may be a long chain of boxes that split to lead to a splitting of b_1 ; say the chain is $b_2, \dots, b_l$. In this chain, b_l is crowded, and it causes b_{l-1} to split, which causes b_{l-2} to split, $\dots$ which causes b_1 to split. Without loss of generality, we can assume that all of the boxes $b_1, \dots, b_{l-1}$ are not crowded. For example, if $b_{l'}$, $l' < l$ is crowded, then we can replace l with l' .

The sizes of these boxes must be decreasing geometrically, so that $h(b_i) = 2^{i-1}h(b_1)$, because otherwise the balance condition would not be enforced from b_{i+1} to b_i . This means that the Euclidean distance from a point in b_l to a point in b_1 is at most $2\sqrt{3}h(b)$. In other words, $I(\text{ex}(b_j)) \subset I(\text{ex}(b_1))$ for all $j = 1, 2, \dots, l$. Recall that the balance condition is only enforced between adjacent boxes, and neighboring boxes are adjacent only if the smaller box contains a point of the extended box of the larger box. Hence, the single connected component $P \wedge \text{ex}(b_1)$ contains $P \wedge \text{ex}(b_j)$ for all $j = 1, 2, \dots, l$. This line of reasoning is also correct in the case that b_1 is itself crowded, in which case we take $l = 1$.

Figure 17: An example path in the proof of Lemma 23. Here all P vertices represent edges, and edges represent facets, except for the P vertex $J = z$. Note b_3 is crowded and b_1 is split by the balance condition.

Recall that x is a point on a P face F in $P \wedge b_1$. We wish to find a face G close to F and x that does not contain F . We now have two cases, depending on why b_1 was split.

Case 1. The first case is that b_1 is crowded because a face interferes with F , and this face is visible to x . In this case, we let G be the closest interfering face to F , and y the closest point of G to x . Then we conclude G does not contain F , and $\text{dist}_P(x, y) \leq 4\sqrt{3}h(b_1)$.

Case 2. The second case consists of one of three subcases. This first subcase is that b_1 is split because of the balance condition. The second subcase is that b_1 is crowded because some face in b_1 is interfered with, but this face is not F . The third subcase is that b_1 is crowded because some face interferes with F , but this face is not visible to x . In the second and third subcases $l = 1$. In any subcase, since b_l is crowded, we may let J denote a P face in b_l such that some P face in $\text{ex}(b_l)$ interferes with it. It may be that J is F . Let H be the face interfering with J such that the geodesic distance between H and J in $P \wedge \text{ex}(b_l)$ is minimum. Since $P \wedge \text{ex}(b_1)$ is a single component containing $P \wedge b_l$, there is a shortest path entirely in $P \wedge \text{ex}(b_1)$ from x to a point z of J in b_l visible to H . This path may be considered as a chain of maximal length edges, from x to x_1 , x_1 to x_2 , $\dots$ x_m to z . A cross-sectional view of an example path is given in Figure 17.

If the smallest dimensional face containing x_1 is not F , we choose G to be this face, and y to be x_1 . Our choice of G as the smallest dimensional face containing y ensures that G does not contain F . Since there is always a

straightline path in $\text{ex}(b_1)$ from x to $x_1 = y$, we conclude that $\text{dist}_P(x, y) \leq 3\sqrt{3}h(b_1)$.

Otherwise the smallest dimensional face containing x_1 is F , that is x_1 is in the relative interior of F . Because the path construction uses maximal length edges, the only way this can happen is if x_1 is actually z . Hence there is a straightline path on F from x to z , and J is contained in F . See Figure 18. We choose G to be H , the face interfering with J , and y to be the closest point of H to z . Since G does not contain J , and J is contained in F , it follows that G does not contain F . Since there exist straightline paths from $x \in b_1$ to $z \in b_1$ to $y \in b_l$, with $h(b_l) \leq h(b_1)$, we conclude that $\text{dist}_P(x, y) \leq 6\sqrt{3}h(b_1)$. ■

Lemma 24 *For any vertex v of P in a box b , $f_*(v) \leq c_{15}h(b)$.*

Proof. Since the protecting and merging steps for a protected vertex box change the size of any box by at most a factor of two, we can assume b is a box at the end of a vertex phase. Let b_1 be the parent of b .

We apply Lemma 23 with $x = v$ in box b_1 . For any vertex, if another face does not contain it, the two are disjoint. Hence the point y on P face G is disjoint from $F = v$ and from the construction $\text{dist}_P(v, y) \leq 4\sqrt{3}h(b)$. We apply Theorem 9 to conclude $f_*(v) \leq \frac{4\sqrt{3}}{c_{10}}h(b)$. ■

Lemma 25 (Corollary) *For any point p in a box b_v protected for vertex v , $f_*(p) \leq c_{16}h(b_v)$.*

Proof. In a protected box Euclidean distance to the apex is the same as geodesic distance, so $\text{dist}_P(p, v) \leq \sqrt{3}h(b_v)$. From Theorem 10 and Lemma 24, $f_*(p) \leq \max(c_{13} \cdot f_*(v), c_{14} \cdot \text{dist}_P(p, v)) \leq \max(c_{13}c_{15}, \sqrt{3}c_{14})h(b)$. ■

Lemma 26 *For any point p of P edge F in a box b , $f_*(p) \leq c_{17}h(b)$.*

Proof. As in the vertex case, Lemma 24, since the splitting steps within a protecting step subdivide any box at most once, we may assume that the parent b_1 of b was split in the vertex or edge phase. Assume that the box containing p at the end of the octree algorithm is not vertex protected, since otherwise Lemma 25 applies.

We apply the Lemma 23 with $x = p$. Unlike the vertex case, Lemma 24, we may have that G and F intersect nontrivially. To complete the proof

Figure 18: An illustration of an instance of Case 2 of the proof of Lemma 23 leading to the case in the proof of Lemma 26 where F' and G' have a common vertex v .

we need to consider the particular triangulation Δ^* . Let F' and G' be the smallest dimensional faces in Δ^* such that $x \in F' \subset F$ and $y \in G' \subset G$. If F' and G' are disjoint, we can apply Theorem 10 as in Lemma 24 and the proof is complete. Otherwise, by conformality, $F' \cap G'$ is a P vertex v . Let b_v be the box containing v at the end of the vertex phase. See Figure 18. Since v is centered in b_v , and by assumption p is outside b_v , $\text{dist}_P(p, v) \geq h(b_v)/8$. But $h(b_v) \geq c_{15}f_*(v)$ by Lemma 24. Hence by Theorem 9 Case 1, $\text{dist}_P(p, y) \geq 8c_9c_{15}f_*(v)$. Recalling from Lemma 23 that $\text{dist}_P(p, y) \leq 6\sqrt{3}h(b_1)$ completes the proof. ■

Lemma 27 (Corollary) *For any point p in a box b_F protected for edge F , $f_*(p) \leq c_{18}h(b_F)$.*

Proof. The proof is identical to the proof of Lemma 25. ■

Lemma 28 *For any point p of P facet F in a box b , $f_*(p) \leq c_{19}h(b)$.*

Proof. As in the previous lemmas, we may assume that the parent b_1 of b was split in the vertex or edge phase, and not a protecting step. Assume that the box containing p at the end of the edge phase, is not vertex or edge protected, since otherwise Lemma 25 or Lemma 27 applies.

We apply Lemma 23 with $x = p$. As in Lemma 26, F and G may not be disjoint, so we must consider the particular triangulation Δ^* . Let F'

and G' be the smallest dimensional faces in Δ^* , such that $x \in F' \subset F$ and $y \in G' \subset G$.

If F' and G' are disjoint, we can apply Theorem 10. Otherwise, by conformality, $F' \cap G'$ is a P vertex or a subset of a P edge. Let v be the closest point of $F' \cap G'$ to p . Let b_v be the box containing v at the end of the edge phase. Since v is well centered in protected boxes, and p is not in a protected box by assumption, $\text{dist}_P(p, v) \geq h(b_v)/8$. Depending on whether b_v is vertex or edge protected, we rely on Lemma 25 or Lemma 27 to show $h(b_v) \geq \min(c_{16}, c_{18})f_*(v)$. Hence, by Theorem 9 Case 1, $\text{dist}_P(p, y) \geq 8c_9 \min(c_{16}, c_{18})f_*(v)$. Recalling from Lemma 23 that $\text{dist}_P(p, y) \leq 6\sqrt{3}h(b_1)$ completes the proof. ■

Lemma 29 (Corollary) *For any point p of a box b_F protected for facet F $f_*(p) \leq c_{20}h(b_F)$.*

Proof. The proof is identical to the proof of Lemma 25. ■

Lemma 30 *For any point p in an unprotected box b (containing no P faces), $f_*(p) \leq c_{21}h(b)$.*

Proof. As in the previous lemmas, we may assume that the parent b_1 of b was split in the vertex, edge, or facet phase. Either b_1 is split by the balance condition, or b_1 is actually crowded, but no P face is in the octant of its child b . By imitating the path construction of Lemma 23, with the only difference being that x is not on any P face, we may identify as G the first face we meet on a path from p to a face in the crowded box b_l . That is, we may show that $\text{dist}_P(p, y) \leq 3\sqrt{3}h(b_1)$, and $h(b_y) \leq h(b_1)$, for some point y on face G in box b_y at the end of the octree generation algorithm. Depending on whether b_y is a vertex, edge or facet protected box, we apply one of the previous three corollaries, Lemma 25, Lemma 27 or Lemma 29 to show $\max(c_{16}, c_{18}, c_{20})f_*(v) \leq h(b_y)$. We apply Theorem 10 to p , so that $f_*(p) \leq \max(c_{13}f_*(y), c_{14}\text{dist}_P(p, y)) \leq \max(c_{13}/\min(c_{16}, c_{18}, c_{20}), 3\sqrt{3}c_{14})h(b)$. ■

We may now summarize the results of the previous lemmas.

Theorem 11 *Let $x \in P$ lie in a box b . Then $f_*(x) \leq c_{22}h(b)$, where c_{22} is a constant depending on the aspect ratio bound of Δ^* .*

The following result combines the bounds derived for f_{OCT} and f_* .

Theorem 12 *For all polytopes P and constants A , there exists a constant c' , dependent only on A and α , such that for any $z \in P$ and any triangulation Δ^* of P with maximum aspect ratio A ,*

$$f_{OCT}(z) \geq c' f_*(z).$$

Proof. This follows directly from Lemma 21 and Theorem 11 ■

Theorem 13 *For all polytopes P and constants A , there exists a constant c'' , dependent only on A and α , such that*

$$|\Delta^{OCT}| \leq c'' |\Delta^*|,$$

where $|\Delta^*|$ denotes the number of tetrahedra of Δ^* and similarly for $|\Delta^{OCT}|$.

Proof. Let Δ^T be a triangulation with bounded aspect ratio. We derive some inequalities that depend on aspect ratio; these inequalities apply to either Δ^* or Δ^{OCT} . By the definition of f_T and bounded aspect ratio A , there are constants c_{23} and c_{24} depending only on A such that

$$c_{23} f_T(x)^3 \leq \text{vol}(T_i) \leq c_{24} f_T(x)^3$$

for all x in the interior of T_i and for any tetrahedron T_i .

Thus

$$c_{23} \leq \int_{T_i} \frac{1}{f_T(x)^3} dx \leq c_{24}$$

for every tetrahedron. Summing this chain of inequalities over all tetrahedra in the triangulation, we get

$$c_{23} \cdot |\Delta^T| \leq \int_P \frac{1}{f_T(x)^3} dx = \sum_i \int_{T_i} \frac{1}{f_T(x)^3} dx \leq c_{24} \cdot |\Delta^T|.$$

This holds for all triangulations, so in particular for Δ^{OCT} and Δ^* . By Theorem 12

$$c_{23} \cdot |\Delta^{OCT}| \leq \int_P \frac{1}{f_{OCT}(x)^3} dx \leq \int_P \frac{1}{(c' f_*(x))^3} dx \leq c_{24}/(c')^3 \cdot |\Delta^*|.$$

■

11 Conclusions

An important open question is the running time. We demonstrated a running time bound of $O(\gamma n \log n)$ for our algorithm, which is inferior to Bern, Epstein and Gilbert's running time of $O(n \log n + \gamma)$ for two dimensional regions. However, we have recently found out by private communication with those authors that this bound does not actually hold for the algorithm that they propose for two-dimensional nonconvex polygons. Therefore, it is not clear what is the best possible running time for triangulation of nonconvex polyhedral regions in either two or three dimensions.

It would be interesting to generalize our algorithm to work for higher dimensional regions. Also, the non-degeneracy conditions on the input region could be relaxed.

Another open question concerns optimizing the tetrahedra for properties other than aspect ratio. For example, is there an algorithm for triangulating a three-dimensional nonconvex polyhedral region to maximize inscribed sphere radius of the tetrahedra?

Finally, we have plans to implement this algorithm. The two-dimensional analog of this algorithm has been implemented by the first author in C++.

Acknowledgement

Thanks to Paul Chew of Cornell for discussing various triangulations with us. Thanks to Marshall Bern, David Eppstein and John Gilbert of Xerox for help with their earlier work. Thanks to Joe Mitchell of Cornell for helping with the plane sweep algorithm.

References

- I. Babuška and A. K. Aziz [1976], On the angle condition in the finite element method, *SIAM J. Num. Anal.* 13:214–226.
- M. Bern, D. Dobkin, and D. Eppstein [1991], Triangulating polygons without large angles, preprint.
- M. Bern, H. Edelsbrunner, D. Eppstein, S. Mitchell, and T. S. Tan [1991], Edge-insertion for optimal triangulations, preprint.
- M. Bern, D. Eppstein, J. Gilbert [1990] Provably Good Mesh Generation,

- Xerox Palo Alto Research Center. Also, *Proc. 1990 Symposium on Foundations of Computer Science*.
- M. Bern and D. Eppstein [1991], Mesh Generation and Optimal Triangulation, preprint.
- G. Carey, M. Sharna, and K. Wang [1988], A Class of Data Structures for 2-D and 3-D adaptive mesh refinement, *Int. J. Numer. Meth. in Engr.* 26:2607–2622.
- B. Chazelle, L. Palios [1989], Triangulating a Nonconvex Polytope, Proc. 5th ACM Symposium on Computational Geometry. 393–399.
- L. P. Chew [1989a], Constrained Delaunay Triangulation, *Algorithmica* 4:97–108.
- L. P. Chew [1989b], Guaranteed-Quality Triangular Meshes, C. S. Cornell, TR 89–983.
- H. Edelsbrunner [1987], *Algorithms in combinatorial geometry*, Springer Verlag, Berlin.
- J. Hauser and C. Taylor [1986], *Numerical Grid Generation in Computational Fluid Dynamics (Proceedings)*, Pineridge Press, Swansea, U.K.
- G. L. Miller, S.-H. Teng, and S. A. Vavasis [1991], A Unified Geometric Approach to Graph Separators, *Proc. 32nd Symp. Foundations of Comp. Sci.*
- G. L. Miller and W. Thurston [1990], Separators in Two and Three Dimensions, *Proc. 22nd Symp. on Theory of Computing*, 300–309.
- G. L. Miller and S. A. Vavasis [1990], Density Graphs and Separators, Department of Computer Science, Cornell, Tech Report 90-1169, and to appear, Symposium on Discrete Algorithms.
- W. F. Mitchell [1987], A Comparison of Adaptive Refinement Techniques for Elliptic Problems, Department of Computer Science, University of Illinois at Urbana-Champaign, Report No. UIUCDCS-R-87-1375.
- W. F. Mitchell [1988], Unified Multilevel Adaptive Finite Element Methods for Elliptic Problems, Department of Computer Science, University of Illinois at Urbana-Champaign, Report No. UIUCDCS-R-88-1446.
- D. Moore, J. Warren [1990], Multidimensional Adaptive Mesh Generation, Department of Computer Science, Rice University, Rice COMP TR 90-106.
- F. Preparata and M. Shamos [1985], *Computational Geometry: An Introduction*, Springer Verlag, New York.