Mesh Generation With Provable Quality Bounds

Scott A. Mitchell
Ph.D Thesis

93-1327
February 1993

Department of Computer Science
Cornell University
Ithaca, NY 14853-7501

MESH GENERATION WITH PROVABLE QUALITY BOUNDS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Scott A. Mitchell

January 1993

Mesh Generation with Provable Quality Bounds

Scott A. Mitchell, Ph.D.

Cornell University 1993

We consider the problem of generating a triangulation of provable quality for two and three dimensional polyhedral regions. That is, we seek a triangulation, allowing additional vertices called Steiner points, such that the triangular or tetrahedral elements have a bound on their shape. In three dimensions we also seek an upper bound on the number of tetrahedra in the triangulation. These triangulation algorithms find application in mesh generation for finite element methods. The polyhedral region must be bounded and well defined, but may have holes of arbitrary complexity.

In three dimensions, we assume there are no restrictions on where we may place Steiner points. Our triangulation is optimal in the following two senses. First, our triangulation achieves the best possible aspect ratio up to a constant factor, which is our bound on element shape. 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.

In two dimensions, we assume the restriction that we may introduce Steiner points on the polygon's interior, but not on its boundary. Of all triangulations satisfying this restriction, our triangulation has the maximum minimum angle, up to a constant factor. This is the first known algorithm for this problem with provably optimal element shape. The algorithm solves several subproblems of mesh generation.

Biographical Sketch

Scott Alan Mitchell was born on April 1, 1966 in Madison, Wisconsin. He enjoyed a happy childhood full of bicycling the backroads, climbing the Sierra mountains, canoeing the Wisconsin River, traveling to Australia, and walking with his family. He married his high school sweetheart, Carissa, an aspiring graphic designer. She dealt with the rigors of a spouse in graduate school gracefully.

After receiving his bachelor's degree in Applied Math, Engineering, and Physics from the University of Wisconsin-Madison in May, 1988, he packed up his wife and his bicycles and moved to Ithaca New York. There he had hopes of bicycling up the hills of the finger lakes region and of completing his doctorate in Applied Math at Cornell University. While the former goal was met almost immediately, the latter was realized in January 1993.

Along Scott's journey through graduate school, he and his wife were joined by a very special little boy, Evan, with whom Scott is carrying on the Mitchell tradition of long walks and blank stares.

For my wife, Carissa.

Acknowledgements

I thank Steve Vavasis for his wonderful guidance during my stay at Cornell. I truly appreciate his willingness to make time for me and help with even the most droll of subproblems. His honesty has made him a pleasure to work with.

I thank Joe Mitchell of Cornell and SUNY Stony Brook for sparking my interest in computational geometry through his well planned courses. I also must thank him for introducing me to Steve Vavasis. Joe's weekly work sessions were always stimulating. I thank those who participated in them over the years, Paul Chew, Christine Piatko, Robert Freimer, Klara Kedem, Erik Wynters, Estie Arkin, Bob Connelly and Bob's many visitors. I also thank Robert Freimer for continuing the sessions in Joe's absence.

I thank my fellow Applied Math students at Cornell. Their support through the rigorous boot camp of analysis and topology courses was invaluable. I thank Dolores Pendell for sparing me from many administrative hassles. I appreciate the picnics in the parks and parties at John Guckenheimer's. I will miss our dilapidated but homey offices in Sage Hall. I would especially like to thank classmates Rick Wicklin, Keith Saints, Patrick Worfolk, Jianguo Liu and Greg Henry, as well as Russ Qualls (civil engineering) for making Cornell bearable through the difficult times.

I thank everyone at the Xerox PARC CSL for a productive yet restful summer 1991 in California. Thanks to Marshall Bern for his continued interest in my work.

Concerning the technical development of this thesis, I thank Steve Vavasis for his indispensable involvement with the three dimensional algorithm. I thank Paul Chew for discussing various triangulations with me, and for his insights concerning the development of the two dimensional algorithm. Thanks to Marshall Bern and John Gilbert of Xerox PARC and David Eppstein of UC Irvine for help with their earlier work. Thanks to Joe Mitchell of SUNY Stony Brook for helping with the plane sweep algorithm.

Table of Contents

List of Figures

Chapter 1

Introduction

Triangulation of polyhedral regions is a fundamental geometric problem for numerical analysis. In particular, if one wishes to solve an elliptic boundary 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, and in two dimensions the triangle. Thus, a *triangulation* of the domain is required. Triangulations are also desired for other applications as well, such as solid modeling, functional interpolation, and computer graphics.

For numerical stability in the finite element method, it is necessary that the elements have some bound on their shape. For most applications it is agreed upon that the closer an element is to equilateral the better. However, which measure of an element is best is unclear, even restricted to finite element mesh generation. For information about element shape in numerical analysis, see Babuška and Aziz

[1976]. We have chosen to consider the minimum interior angle for two dimensions. That is, we seek a triangulation whose minimum angle is as large as possible. For three dimensions, we consider what we call *aspect ratio,* the ratio of the radius of the smallest sphere containing the tetrahedron, to the radius of the larges sphere contained in the tetrahedron. That is, we seek a triangulation whose aspect ratio is as small as possible. These two measures are very related, as will be described in Chapter 3.

One question that often arises when designing geometric software is how the polyhedral regions are represented. One may define such a region as the regularized intersection and union of half spaces. This is known as the "computational solid geometry," or CSG, representation. Also, one may define the region in terms of the lattice of boundary faces, and which side of the $d-1$ dimensional faces is the interior of the region. This is known as the "B-rep". The composition of half spaces is mathematically easier to describe. However, it has several algorithmic drawbacks, most notably having to traverse a tree of possibly large depth to identify specific faces of the boundary. We have chosen to use the boundary face lattice representation for our algorithms.

This thesis presents two triangulation algorithms, one for three dimensions and one for two dimensions. Both algorithms generate a triangulation for a nonconvex bounded polyhedral domain with holes.

1.1 The three dimensional algorithm

For the three dimensional algorithm, we allow ourselves the freedom to introduce additional vertices, called Steiner points, in the region's interior, and also on its boundary. The primary goal of the algorithm is to achieve good aspect ratio. However, care must be taken not to introduce too many faces, as the number of tetrahedra in the triangulation is also important. For example, the running time of a finite element method grows with the number of elements. Some very complicated and large objects, such as an entire airplane, require hundreds of thousands of elements and mere disk space can become a constraint.

We are able to achieve both of our goals. That is, the triangulation is optimal in two respects. First, the best possible aspect ratio is achieved for the tetrahedra, up to a constant factor. Second, the number of tetrahedra is within a constant factor of the best possible for any triangulation with bounded aspect ratios. This second constant factor depends on the smallest interior angle of the input region.

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 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. The key idea is that we isolate the vertex in a cube. Then, we triangulate within the cube by first triangulating in two dimensions the boundary of the cube, and then taking the convex hull of these triangles and the vertex to form tetrahedra.

Our running time bounds are probably not the best possible; this is discussed further at the end of Chapter 4. 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. This is achieved by using cube sizes that depend on the geodesic distance between disjoint faces of the input. We are able to use geodesic distance by *duplicating* boxes when the intersection of a box with the input forms several connected components. 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 before, however, to analyze the construction of a triangulation. We believe the proof method is particularly powerful and has great potential. For example, it has been adopted by Rupert[1992], whose two dimensional algorithm does not even use quadtrees.

Our technique is based on an octree partitioning of the domain. That is, we use a three dimensional grid of cubes (boxes) to divide up to domain and isolate disjoint faces of the boundary of the region. The boxes containing the boundary must be deformed away from cubes, and all boxes are triangulated. This technique has been used before, for example, by Carey, Sharna and Wang [1988].

Our algorithm generates an unstructured mesh. That is, the sizes of the cubes is allowed to vary. Structured meshes, where the size of the elements are fixed, have several advantages. See Bern and Eppstein [1992]. However, they are ill suited to input regions that have both very small and very large faces, or where the equation being solved by a finite element method has a greatly varying gradient. In both cases, to get the required resolution one must have an inordinate number of elements. We have chosen to consider problems where it is advantageous to vary the element size.

The sizes of the cube, and hence the tetrahedra, is completely determined by the geometry of the input. For some applications, one may wish there to be smaller elements in a portion of the region. For example, for using finite element methods to analyze the temperature of an object, it may be known that a certain point contains a heat source. Since the gradient of temperature will vary greatly near that point, to achieve an accurate solution small elements are needed near that source. All one needs to do to insure smaller elements near that source is to specify that the boxes may be no larger than a particular size near the source. That is, the user may define a function that takes as an argument a point of the input region, and returns the size of the largest desired box containing that point. Our algorithm could refer to this function to easily comply with the user's wishes.

Section 4.1.7 summarizes the algorithm.

1.2 The two dimensional algorithm

The problem of triangulating a two dimensional polygon with holes with minimum angle within a constant factor of optimal and simultaneous cardinality bounds has been solved by by Bern, Eppstein and Gilbert [1990]. When no Steiner points are allowed, the triangulation that maximizes the minimum angle is the Delaunay triangulation. The numerous algorithms for the Delaunay triangulation are summarized in Bern and Eppstein [1991].

The two dimensional problem we consider strikes a middle ground between arbitrary Steiner points and no Steiner points. That is, we seek to triangulate a polygon with minimum angle within a constant factor of optimal, where we are allowed Steiner points on the polygon's interior, but not its boundary. This algorithm may find application in a number of mesh generation subproblems.

This algorithm allows us to triangulate intersecting regions independently. One application is that it may be used to allow standard mesh generation algorithms to accept degenerate input. Degenerate faces such as edges that have the interior of the input region on both sides are often desired to "probe" a numeric solution at specific locations. Such an edge is also useful on the boundary between two differently doped regions in the analysis of semiconductors. We could take a degenerate edge, and fatten it in one direction into a rectangular hole of finite width. Then any mesh generation algorithm could be used to triangulate the region outside the hole. Our no boundary Steiner point algorithm can then be used to triangulate the hole, respecting the mesh already generated.

Consider another related example. In our three dimensional algorithm, we triangulate each two dimensional facet of the intersection of a cube and the input region before introducing tetrahedra. It does not matter if in this step we introduce Steiner vertices in the interior of the facet. However, if we introduced Steiner points on an edge of the boundary of the facet, that edge is shared by facets of other cubes. We may have already triangulated that facet, in which case we have ruined its triangulation by introducing a vertex in the middle of an edge of another triangle. Thus a triangulation that introduces Steiner points only on the regions interior is desired.

An important feature of the algorithm that makes it useful is that we give explicit bounds on the minimum angle of the triangulation in terms of geometric features of the input, namely edge lengths and certain angles. Consider again the example in the previous paragraph of triangulating the two dimensional intersection of the cubes of our three dimensional algorithm. If we could give explicit bounds on the geometric features of the intersection of our cubes with the input region, then we would not have to resort to the tedious case by case triangulation presented in Section 4.4. We do not take this approach because we need a triangulation that maximizes the minimum height instead.

The algorithm may also find application in mesh refinement. That is, suppose a mesh exists for a given polygon, but after running a finite element method we discover that the mesh is too coarse near a heat source. We may then "erase" the mesh in a neighborhood of the heat source. Then a finer mesh could be generated strictly inside this neighborhood, and our algorithm could be used to triangulate the gap between the outer and inner mesh. Care must be taken in selecting the inner mesh, as in general it is not possible to generate an arbitrarily fine inner mesh while retaining the outer mesh and still have minimum angle within a constant

Figure 1.1: The optimal min angle triangulation of a polygon without Steiner points (left) and with interior Steiner points (right).

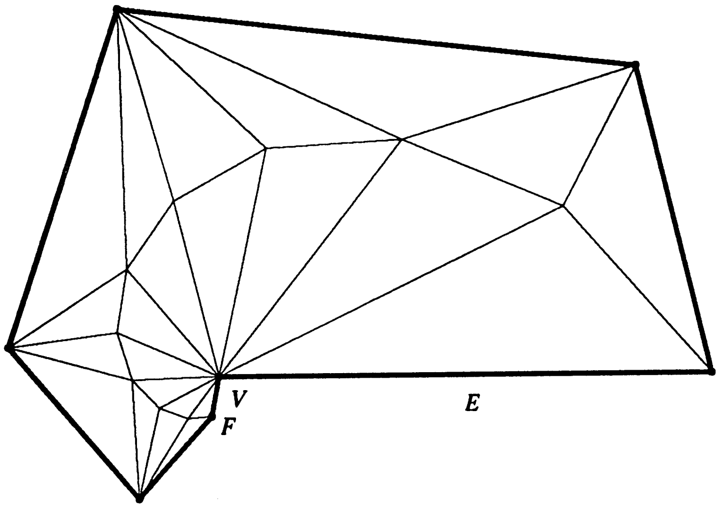

Figure 1.2: The optimal min angle may be much better than linear in the ratio of the length of an edge to its closest interfering point.

factor of optimal. The advantage for this method is that a mesh for the entire region does not need to be recalculated.

The algorithms that do not add any Steiner points are inadequate for these tasks. The four point example of Figure 1.1 shows that for a polygon the optimal min angle when Steiner points are not allowed may be much worse than the optimal min angle when Steiner points are allowed on the interior.

A key question is what features of a general polygon P determine its optimal min angle A when Steiner points are allowed on the interior of P, but not its boundary. Bern Eppstein Gilbert [1990] considers the problem where Steiner points are allowed everywhere in P. There the distance between an edge and a point on a closed face disjoint from the edge (an *interfering point*) determines how small triangles must be in order to have aspect ratio within a constant factor of optimal. For determining A, because we are not allowed to add Steiner points on ∂P, it may be impossible to make triangles this small. For example, in Figure 1.1, there will always be a triangle with an edge that is the bottom long edge of P. Indeed, as we vary the distance between the top and bottom edges of Figure 1.1, the optimal aspect ratio changes proportionally (up to a constant factor).

Hence the optimal min angle for this problem should depend on the maximum ratio of the length of an edge E to its distance to an interfering point. However, the optimal min angle is actually more closely dependent on the angle that the interfering point makes with the closer vertex V of E. That is, if this angle is large, it may be possible to triangulate with min angle much more than the interfering point distance to edge length ratio.

For an intuitive reason for why this is so, consider Figure 1.2. The closest interfering point to edge E is the vertex of F opposite V. If we vary the length of edge F, the ratio of the length of E to the closest interfering point changes proportionally. Suppose we also keep the same combinatorial triangulation as we vary F, but allow ourselves to change the position of the Steiner points. As there are ten similar triangles with vertex V, if we halve the length of F, we may also change the lengths of the edges containing V, so that the resulting triangulation has min angle much more than half of that of the original triangulation. In fact for this example we could construct a triangulation with min angle about $2^{-1/10}$ times that of the original triangulation.

1.3 Thesis overview

In Chapter 2 we briefly review some relevant literature. In Chapter 3 we give more details on aspect ratio and minimum angle.

In Chapter 4 we present the three dimensional triangulation algorithm. In Section 4.1 we discuss the formation of the octree. In Section 4.2 we demonstrate a relation between the boxes of the octree. In Section 4.3 to Section 4.5 we discuss the steps necessary to produce a triangulation given an octree. In Section 4.6 to Section 4.8 we provide the optimality proof. We conclude the three dimensional algorithm in Section 4.9.

In Chapter 5 we present the two dimensional triangulation algorithm. It has three main steps. First in Section 5.1 we introduce triangles around the boundary of P, leaving an untriangulated portion Q. Second in Section 5.2 we describe the formation of another polygon R inside Q, and in Section 5.3 we triangulate R using a standard mesh generator. Third, in Section 5.4 we describe how to triangulate between Q and R. The optimality proof relies on a sequence of lemmas and theorems found in each of these sections. We conclude the two dimensional algorithm in Section 5.5.

1.4 Some definitions.

We now describe some assumptions about the input and triangulations that apply to both algorithms.

Face. A *face* of the input, 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 a three dimensional input as a three-dimensional face. Note that the containment relation induces a partial order on faces.

Conformal. All of the triangulations we consider will be *conformal.* If we consider a triangulation as a lattice of faces, then being conformal means that two faces intersect in a face of the lattice or not at all. For example, there may be no vertices in the middle (relative interior) of an edge or triangle. In addition, we only consider triangulations that are conformal to the input, as defined in Section 4.8.

Nondegenerate. We assume that the input region P is nondegenerate in the following senses. Every facet has the interior of the region on exactly one side. 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. For the three dimensional algorithm we also require that every one-dimensional edge is incident on exactly two two-dimensional faces. These assumptions may be dropped in future work.

Chapter 2

Other work

This chapter will review selected previous results concerning triangulations. We focus our attention on maximizing the minimum angle. For a comprehensive survey of the computational geometry relevant to finite element mesh generation, see Bern and Eppstein [1991].

2.1 Computational geometry literature

The first question one might wish to settle is the existence of a triangulation. Preparata and Shamos [1985] have shown that indeed one can triangulate a planar straight line graph in optimal $\theta(n \log n)$ time, without adding any Steiner points. However, there is no bound on the shape of a triangle produced by their algorithm.

One of the first algorithms considered that provides a quality bound to triangle shape is one that produces the well known Delaunay triangulation (DT). For a point set, the DT is defined as the subgraph of the complete graph, such that each edge has a circle through its vertices containing no other input vertex. Delaunay [1934] shows that it may also be described as the planar dual of the Voronoi diagram or Dirichlet tessellation, Aurenhammer [1991] has surveyed the literature of Delaunay triangulations and Voronoi diagrams. Certain steps must be taken if the input vertices are degenerate (co-circular) to insure that the DT is indeed a triangulation.

The notion of the DT may be extended to the constrained Delaunay triangulation (CDT) for planar straight line graphs (PSLG), a category which includes polygons with holes. For a PSLG, the CDT may be defined as the subgraph of the complete graph on the vertices of the input. An edge $\overline{ab}$ is in the CDT if the vertices are visible to each other, and there is a circle through a and b such that no other input vertex c visible to a or b is contained in the circle. We say that two points are visible to each other if the edge between them does not cross an input edge. Lawson [1977] has shown that the CDT maximizes the minimum angle in the triangulation. Actually, the CDT lexicographically maximizes the angle vector, the list of angles of the triangulation from smallest to largest. See Lee [1978] or Lee and Lin [1986] Numerous optimal $O(n \log n)$ algorithms have been proposed for generating the CDT, which are nicely summarized in Bern and Eppstein [1991].

The notion of a CDT can also be used for mesh generation where Steiner points are allowed. In the $O(n^2)$ algorithm of Chew[1989b], the edges of a PSLG are divided into edges of about equal length. The CDT of these smaller edges is computed. While there exists a poorly shaped triangle, we add the center of the circle through the vertices of the triangle and recompute the CDT. This results in triangles with angles between 30 and 120 degrees. The algorithm works with a slightly restricted class of input, and there are no optimality bounds on the cardinality of the triangulation.

Recently Rupert[1992] has developed a mesh generation algorithm based on the constrained Delaunay triangulation that overcomes these shortcomings. This result has simultaneous size and quality bounds, similar to those of the earlier work of Bern, Eppstein and Gilbert [1990] and the three dimensional algorithm presented in this thesis and previously in Mitchell and Vavasis [1992]. Unlike those works, Rupert [1992] allows the simultaneous triangulation of the inside and outside of a polygon. Rupert[1992] is much easier to implement and in practice produces a triangulation with fewer triangles than Bern, Eppstein and Gilbert [1990].

Despite the power of the CDT in two dimensions, it is unlikely that the CDT will be very useful in developing three or higher dimensional algorithms. Rupert [1992] comments on the infeasibility of extending the results to higher dimension. There are two main problems. The first is that no (constrained) Delaunay triangulation may exist for general three dimensional input (Schönhardt [1928]), but with careful placement of Steiner points it is conceivable that this may be overcome. The second more fundamental problem is that Delaunay tetrahedra may have terrible shape: Even when the angles between edges are large, the angle between a face and an edge as defined in Chapter 4 may be arbitrarily small. This can happen when four vertices are nearly coplanar. That is, the results of Lee [1978] and Lee and Lin [1986] are peculiar to planar graphs. Rajan[1991] shows that CDT minimizes the maximum radius of a smallest containing sphere in $\mathbb{R}^d$. We note aspect ratio is dependent on the ratio of the radius of the smallest containing sphere to the radius of the inscribed sphere (Bern [1992], personal correspondence).

One way generate a triangulation with guaranteed element shape is to lay down a regular subdivision of the plane, and then perturb this subdivision slightly to conform to the input. The subdivision must be fine enough to isolate geometric features of the input. See for example Baker, Gross and Rafferty [1988], and the linear time algorithm described in Chew [1989b]. These triangulations may have output size much larger than optimal, however.

In higher dimensions, Chazelle and Palios [1989] consider the problem of generating a triangulation with optimal size, but are not interested in angle bounds.

The first two dimensional algorithm to simultaneously guarantee optimal bounds on triangle quality and mesh cardinality was Bern, Eppstein and Gilbert [1990]. They overcome the mesh cardinality limitations of using a regular subdivision of the plane by using a quadtree. See Carey, Sharna and Wang [1988] for a discussion of quadtrees and octrees. Bern, Eppstein and Gilbert [1990] solves many triangulation problems, including that of optimally triangulating a point set of arbitrary dimension. The general outline for each method is similar, we describe that of the algorithm that triangulates a two dimensional polygon with no small angles and optimal cardinality: Start with a large square containing the entire input. Then, recursively divide any box containing disjoint faces of the input into four smaller boxes. Element quality is retained by balancing the quadtree, forcing adjacent boxes to differ in size by no more than a factor of two. Before forming the final triangulation, faces of the quadtree that are poorly aligned with the input are perturbed. Individual boxes are triangulated independently, The mesh size optimality proof makes use of the minimum angle of the CDT of the input.

Hierarchical decomposition methods are not limited to quadtrees and octrees. Using a triangular or tetrahedral decomposition often makes the final triangulation step easier, although the decomposition process can be more complicated. Moore

[1992] discusses the properties simplices must have in order to be useful for a recursive subdivision of space. See also Moore and Warren [1990a]. For adaptive mesh generation for box-shaped regions see also Moore and Warren [1990b]. Mitchell [1987] and [1988] consider the problem of adaptive mesh generation of complex regions in regard to finite element error bounds.

2.2 Geometric modeling

In many mesh generation algorithms geometric modeling plays an important role. For example, in the three dimensional algorithm presented in this thesis, we must repeatedly find the intersection between an octree box and the input polygon. Often mesh generation algorithms use many of the functions found in geometric modelers, such as intersection. Also, complicated input regions are usually generated using a CAD system with a geometric modeler. A geometric modeler is also useful to view to resulting mesh, especially in three dimensions. Sapidis and Perucchio [1989] review existing mesh generation implementations, and look for robust extensions to three dimensions. They search in vain for a fully automatic mesh generation algorithm that may be incorporated into an existing CAD system, but emphasize the desirability of such an algorithm. Here automatic mesh generation refers to algorithms that require no user input other than the initial polyhedral region.

2.3 Engineering literature

The engineering literature is rich with finite element mesh generation algorithms. See for example, the conference proceedings edited by Hauser and Taylor [1986].

There are many techniques that do not directly involve computational geometry, such as conformal mapping. Often even the engineering literature that does involve computational geometry differs significantly from the results developed by the computational geometry community. Many of these differences arise merely from style of presentation. For example, most engineering results rely on empirical evidence rather than mathematical proofs to demonstrate the quality of the mesh. With sufficient effort, algorithms from the computational geometry community could provide empirical evidence, and in some cases proofs could be derived for algorithms from the engineering community.

Sometimes the differences go beyond this, however, and stem from what is considered a respectable problem in the two communities. For example, in the engineering literature there are many cases where reasonable element shape is only achieved for very well shaped input, no mathematical proof of element shape is possible. What is considered well shaped input is usually defined by the practical problems that arise, rather than by lower bounds on the degree of vertices and other mathematical means. Considerable effort has been made to develop heuristics that improve element shape in practice, but have no provable guarantee of success.

Another example is that in the engineering literature, many algorithms have been presented that are useful for particular equations that one wishes to solve, on particularly well shaped domains. As one might expect, these issues have been largely ignored by computational geometers, who prefer to focus on the geometry of the input regions. Most computational geometry literature focuses on arbitrarily complex domains. Also, most computational geometry results have ignored constant factors in mesh size, and often element shape as well. These constant factors can be important in implementations, and so the computational geometry literature is limited from an engineering standpoint.

One avenue of cooperation between these fields is the following. A computational geometry result, such as one presented in this thesis, can be used to generate a mesh. But, either during or after this process, well developed heuristics from the engineering literature can be applied to improve the mesh quality. As long as the heuristics do not reduce the quality of the mesh, we would still have a provably good algorithm. Hopefully, the restrictions of the computational geometry result would not hinder the heuristics from developing a good mesh in practice. Such a result would be especially valuable for very large inputs, where verifying and improving the quality of heuristic results can be tedious for practitioners.

2.3.1 Buratynski [1990]

Buratynski [1990] presents several ideas that were useful in developing the three dimensional algorithm of this thesis. It is also useful as an example of the computational geometry limitations of a heuristic result. This is not a criticism of that work, but an illustration of differences in intent.

In Buratynski [1990], the three dimensional input region is assumed to be represented by its boundary. The method is octree based. Similar to our three dimensional algorithm, in both the octree generation and final triangulation steps input faces are considered on a hierarchal basis. That is first vertices, then edges and finally facets. In typical mesh generation algorithms, the octree boxes are deformed to the input. In this method, however, the input geometry is deformed to the boxes (called bricks by Buratynski [1990]), and then the inverse transform is performed to both the input and the box to recover the original input. This approach is easier in that there is less to consider when forming the final triangulation of a box intersected with the input. However, mathematically proving a lower bound on element shape using these transformations would be difficult.

This method works very well in practice on the class of input it is designed for, namely injection plastic molds. However, it is unclear if it can handle high (> 12) degree vertices, so there are correctness concerns from a computational geometry viewpoint. The analysis of the quality of the mesh does not consider some factors that we consider important in this thesis, but probably are not important for the class of input the algorithm was designed for. That is, there are no mathematically proved lower bounds on element quality in terms of the input geometry. There are limitations on the size of the boxes that can contain a particular input edge, which limits how close to optimal the cardinality of the mesh may be. Indeed, a lower bound on the mesh size is not considered.

The algorithm is automatic in the sense that the user does not specify individual boxes, although there are user specified starting parameters. These parameters place some limits on how coarse the mesh may be. Also, there are points in the algorithm that the program must stop and ask the user for additional parameters if it encounters unusual features of the input. In light of these user inputs, it is doubtful that nontrivial mathematical bounds on mesh cardinality and element shape exist. However, these inputs appear to be valuable practitioners.

Chapter 3

Element shape

In two dimensions we seek a triangulation whose minimum angle is as large as possible, up to a constant factor. In three dimensions, we seek a triangulation whose aspect ratio is as small as possible, up to a constant factor. Here the minimum angle of a triangulation is the minimum of the angles of its triangles, and the aspect ratio of a triangulation is the maximum of the aspect ratios of its tetrahedra.

In three dimensions, in order to argue about the aspect ratio of triangulations of P, we wish to define the angle between any two faces of P or 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°. This is also how we define interior angle in two dimensions.

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" in Chapter 4 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 we define to be 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 (circumscribing sphere). 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 in Chapter 4. This is also easy to show in two dimensions. Thus, up to constant factors, aspect ratio and the reciprocal of min angle are equivalent measures. In Chapter 4 we will make heavy use of this fact because the min interior angle is so much easier to analyze.

Sphere ratio. Often another measure goes by the name of aspect ratio. It is what we will call *sphere ratio.* The sphere ratio of a d dimensional simplex is the ratio of the radius of the d dimensional sphere through the vertices of the simplex, to the radius of the largest sphere contained in the simplex. While related to aspect ratio, the sphere ratio is not equivalent, even up to constant factors.

Bern [1992] (personal correspondence) has proven that in d dimensions the sphere ratio of a simplex is between A and A^2, where A is the aspect ratio defined above. Bern and Teng [1992] (personal correspondence) have extended this result to show the equivalence of aspect ratio with several measures. For example, A is bounded above and below by constant factors times the ratio of the largest distance between a pair of supporting hyperplanes to the smallest distance between a pair of supporting hyperplanes, and also by constant factors times the inverse of smallest solid angle.

A previously known consequence is that if a triangle has no large angles, its

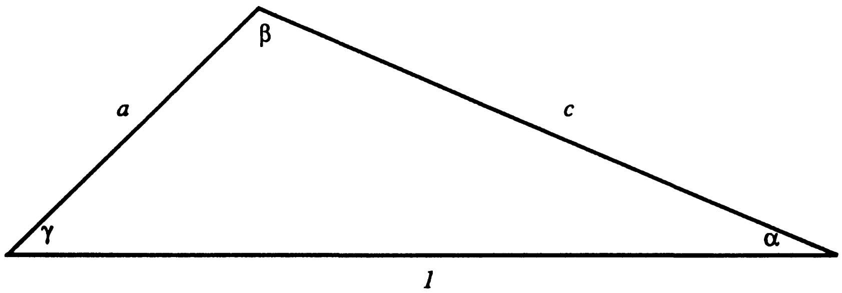

Figure 3.1: The sides and angles of a general triangle.

sphere ratio and aspect ratio are within constant factors of each other.

Sphere ratio is unwieldy, even in two dimensions. We present the following approximation of sphere ratio for two dimensions.

Theorem 1 *Let S be the sphere ratio of a triangle, α its smallest angle, and β its largest angle. Then*

$$\frac{3}{8}\frac{1}{\sin\alpha\sin\beta} \leq S \leq 4\frac{1}{\sin\alpha\sin\beta}.$$

Proof. Without loss of generality, assume that the longest edge of the triangle, which is opposite β, has length 1. Let the shortest edge, which is opposite α, have length a. Denote the remaining angle by γ, and the length of the opposite edge by c. See Figure 3.1.

We first bound r, the radius of the inscribed circle. From the hypotenuse leg theorem of elementary trigonometry, the center of the inscribed circle is at the intersection of the angle bisectors. Hence, from Figure 3.2, we have $r = l_\gamma \tan\frac{\gamma}{2}$, $r = l_\alpha \tan\frac{\alpha}{2}$, and $l_\gamma + l_\alpha = 1$. Eliminating l_γ and l_α from these equations, we have

$$\frac{1}{r} = \cot\frac{\alpha}{2} + \cot\frac{\gamma}{2}.$$

Since by assumption $\alpha \leq \gamma$, we have $\frac{1}{2}\tan\frac{\alpha}{2} \leq r \leq \tan\frac{\alpha}{2}$. Since $0 \leq \alpha \leq \pi/3$, we

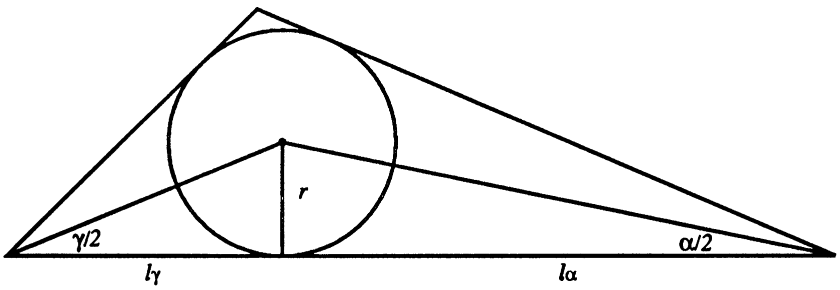

Figure 3.2: The radius of the inscribed circle.

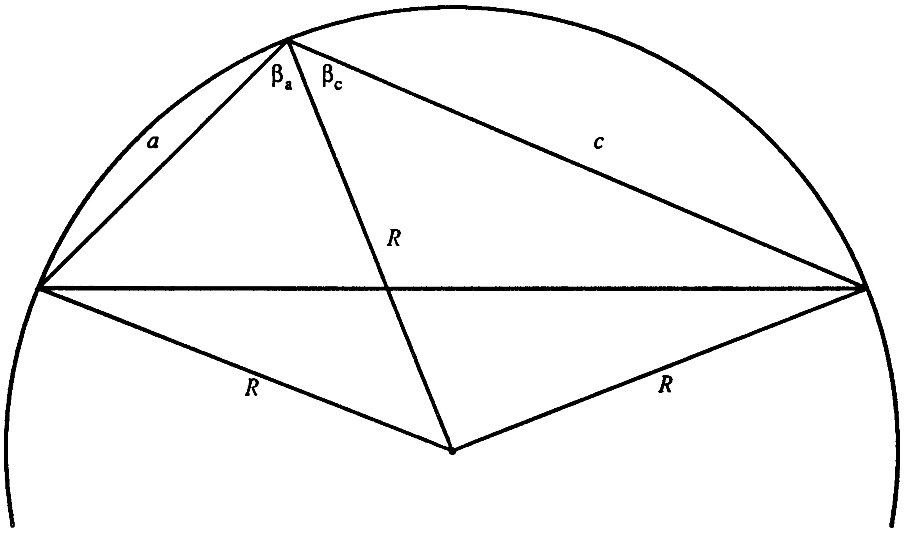

Figure 3.3: The radius of the circumcircle.

have $\frac{1}{2}\sin\alpha \leq \tan\frac{\alpha}{2} \leq \frac{2}{3}\sin\alpha$. Hence

$$\frac{1}{4}\sin\alpha \leq r \leq \frac{2}{3}\sin\alpha.$$

We next examine R, the radius of the circumscribing circle. see Figure 3.3. Drawing the segments from the center of the circumscribing circle to each of the vertices of the triangle, we obtain two triangles. The first is a similar triangle with two sides of length R and the remaining side of length a. We define β_a to be the angle at a base vertex of this triangle. The second triangle is also a similar triangle, with two sides of length R, but the remaining side is of length c. We define β_c to be the angle at a base vertex of this second triangle. Then

$$R = \frac{c}{2\cos\beta_c}.$$

We wish to convert the above equality involving β_c and c into inequalities involving only β. Write $\beta = \pi - \epsilon$, with $0 < \epsilon < 2\pi/3$. Note $\beta_a + \beta_c = \beta$.

We obtain an upper bound on R involving β as follows. Since $c > a$, we have $\beta_c > \beta_a$, and $\beta_c \leq \frac{\beta}{2} = \frac{\pi}{2} - \frac{\epsilon}{2}$. Hence $\cos\beta_c \geq \cos(\frac{\pi}{2} - \frac{\epsilon}{2}) = \sin\frac{\epsilon}{2} \geq \frac{\sin\epsilon}{2} = \frac{\sin\beta}{2}$. Since c is the length of the second longest side of the triangle T, we have $c \leq 1$. Hence

$$R \leq \frac{1}{\sin\beta}.$$

We obtain a lower bound on R involving β as follows. Since β_a is the base angle of a similar triangle, $\beta_a \leq \frac{\pi}{2}$. Hence $\beta_c = \beta - \beta_a \geq \frac{\pi}{2} - \epsilon$. Hence $\cos\beta_c \leq \cos\frac{\pi}{2} - \epsilon = \sin\epsilon = \sin\beta$. Since c is the length of the second longest side of the triangle T, from the triangle inequality we have $c \geq \frac{1}{2}$. Hence

$$R \geq \frac{1}{4\sin\beta}.$$

Combining the bounds on R with those on r, we obtain the following bounds on the sphere ratio $S = \frac{R}{r}$ of

$$\frac{3}{8}\frac{1}{\sin\alpha\sin\beta} \leq S \leq 4\frac{1}{\sin\alpha\sin\beta}.$$

∎

Chapter 4

Quality mesh generation in three dimensions

We show how to triangulate a three dimensional polyhedral region with holes. We assume there are no restrictions on where we may place Steiner points. Our triangulation is optimal in the following two senses. First, our triangulation achieves the best possible aspect ratio up to a constant factor. 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.

4.1 Subdivision of P into cubes, the octree

4.1.1 Data structures and definitions

Our triangulation is denoted $\triangle^{OCT}$. 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.

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, $\mathrm{ex}(b)$, such that $I(\mathrm{ex}(b))$ is the cube concentric with $I(b)$ and with each dimension expanded by a factor of 5. We use the notation $P \wedge \mathrm{ex}(b)$ to denote the component of $P \cap I(\mathrm{ex}(b))$ that contains $P \wedge b$. The extended box contains only the P faces of $P \wedge \mathrm{ex}(b)$. Note that $\mathrm{ex}(b)$ 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. We say that box b_1 is *adjacent* to box b_2 if they are neighbors and there is a point of P common to both, i.e. if $P \wedge b_1 \cap P \wedge b_2 \neq \emptyset$. 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. We say b_1, b_2 are *balance-adjacent* if they are neighbors and there is a point of P common to one of the boxes and the extended box of the other, i.e. if $(P \wedge \mathrm{ex}(b_1)) \cap (P \wedge b_2) \neq \emptyset$ or if $(P \wedge \mathrm{ex}(b_2)) \cap (P \wedge b_1) \neq \emptyset$. Boxes that are adjacent are balance-adjacent, but not all balance-adjacent boxes are adjacent, see for example Figure 4.3.

Balance Condition. If we split a box, we immediately split other boxes to maintain the following invariant called the *balance condition*: No box is balance-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 an adjacent box is bounded above by a constant that depends linearly on $1/\alpha$, as we prove in the next section.

4.1.2 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.

4.1.3 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, $F_1, F_2, \ldots, F_k$, 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 $F_1, F_2, \ldots, F_k$ are exactly the P faces incident upon $P \wedge \mathrm{ex}(b)$.

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 $P \wedge \mathrm{ex}(b)$.

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 $P \wedge \mathrm{ex}(b)$ *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 $\mathrm{ex}(b)$ contains a face that interferes with v.

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.

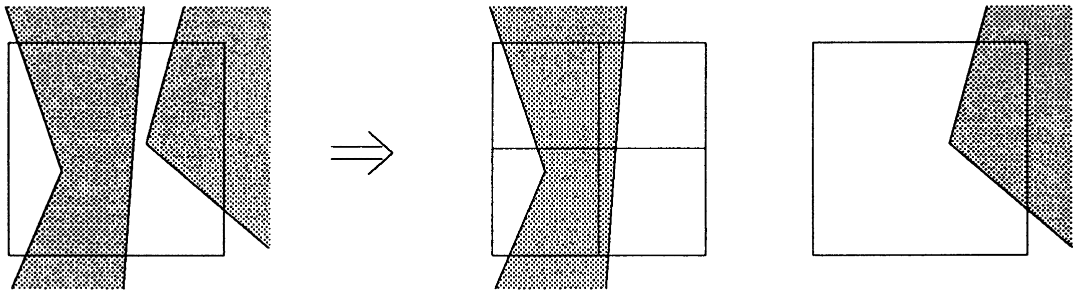

Figure 4.1: Here a box is duplicated for two components, and one duplicate box is split because it is crowded.

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

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 4.3 and later sections.

We first make the following observation concerning the boxes at the conclusion of the vertex phase. Let b be a box containing a P vertex v. We claim that any box balance-adjacent to b has size at least that of b. We denote the size of b with the notation $h(b)$. Let b_1 be balance-adjacent to b. The balance condition ensures that $h(b_1)$ is either equal to $h(b)$, $2h(b)$, or $h(b)/2$. Hence we need only prove that $h(b_1) \neq h(b)/2$.

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 $b_1', \ldots, b_k'$ such that b_k' is vertex crowded, b_i' is balance-adjacent to b_{i+1}' (for $i = 1, \ldots, k-1$), and such that $h(b_i') = 2h(b_{i+1}')$. 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 $I(\mathrm{ex}(b))$, no matter how large k is, because $\mathrm{ex}(b)$ 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 $P \wedge \mathrm{ex}(b)$, 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 balance-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 balance-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 balance-adjacent to a box b with a vertex to be split.

Consider a setup of boxes that might cause a box b_1 balance-adjacent to a box b to be split by propagation of the balance condition. Suppose b_1 is a box balance-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 balance-adjacent b_k such that b_k is the same size as $\hat{b}$, and b_k is balance-adjacent to a box b_{k-1} four times as large, which

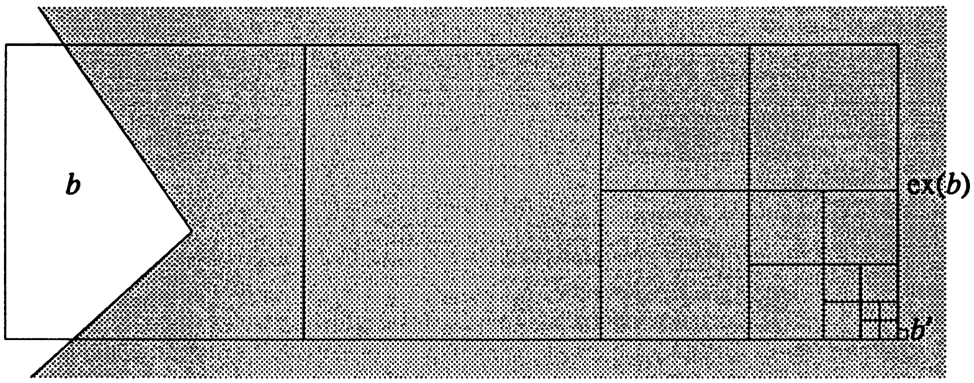

Figure 4.2: No matter how small we split a box b', if it is outside of ex(b) then the balance condition will not split b.

is balance-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 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 balance-adjacent to boxes of the same size. Consider such a box b. Without loss of generality, b has 26 balance-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 balance-adjacent to 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

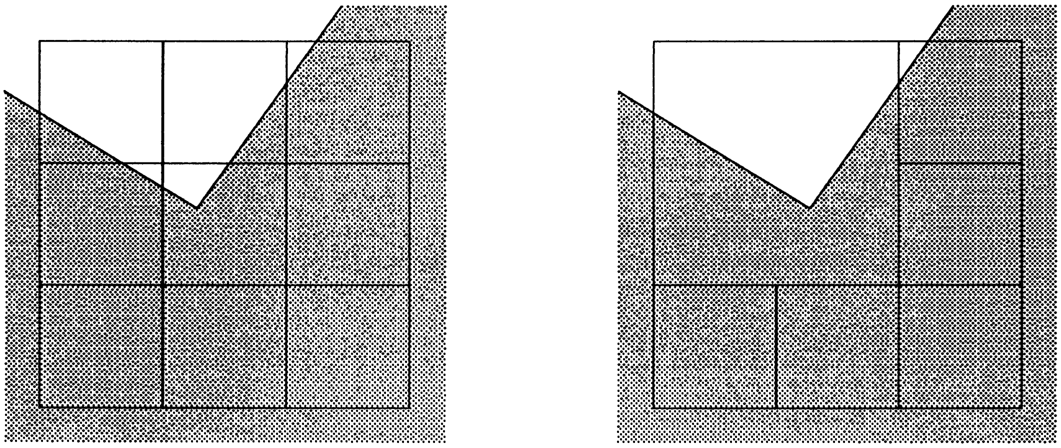

Figure 4.3: Uniting boxes around a vertex box in two dimensions.

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 4.3 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 balance-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.

4.1.4 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 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 extended 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.

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 balance-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 4.3) 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 balance-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 balance-adjacent to a box b' containing an edge E, then b cannot be balance-adjacent to a box containing any other edge E'.

We always protect every box containing a P edge, and every box balance-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.

4.1.5 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 balance-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 extended 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 $P \wedge \mathrm{ex}(b)$.

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 $P \wedge \mathrm{ex}(b)$. 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.

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 balance-adjacent boxes and enforcing the balance condition. This ensures that a box is balance-adjacent to at most one facet cone apex. Thus, in the warping and triangulation step to follow in Section 4.3, we may consider the boxes containing a facet without reference to any other facets, except for boxes balance-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 balance-adjacent to them.

4.1.6 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 $P \cap I(\mathrm{ex}(b))$, which is necessary for determining whether boxes are crowded, and when to propagate the balance condition.

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

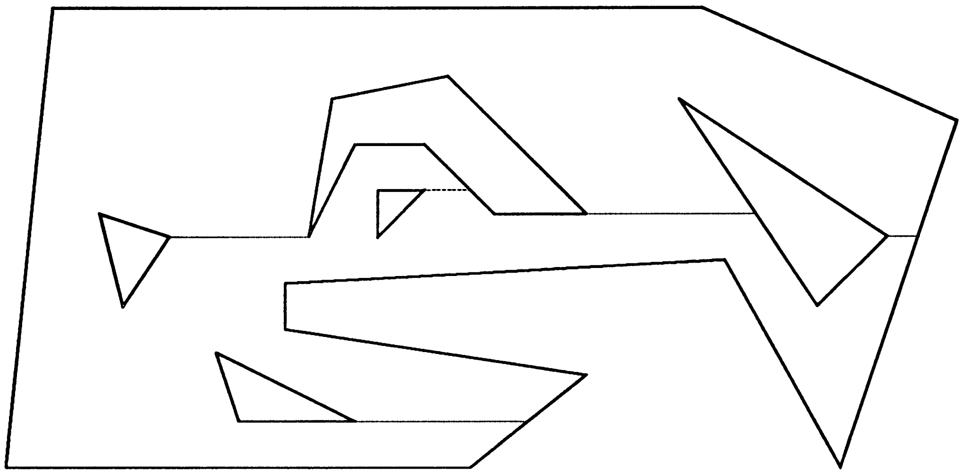

Figure 4.4: Floaters and tethers for a two dimensional polygon.

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 4.4 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 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 literature. The reader unfamiliar with these tools should consult Edelsbrunner [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 4.5.

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

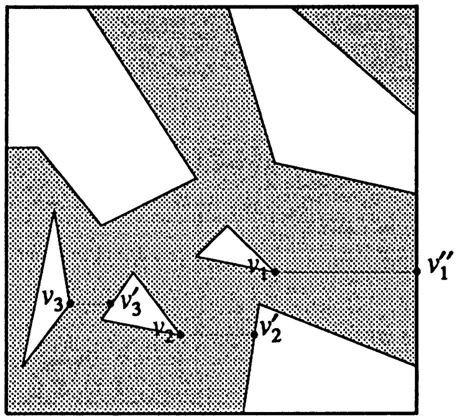

Figure 4.5: Determining the components of a box. Sweeping the boundary determines two components, and following tethers determines that all internal floaters belong to the same component, and there are no internal components.

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 $P \wedge \mathrm{ex}(b)$ for testing crowdedness. This is done by running the component-determination algorithm on $\mathrm{ex}(b)$, and then saving the component of $P \cap I(\mathrm{ex}(b))$ that contains a point in $P \wedge b$. This also allows us to determine which neighboring boxes are balance-adjacent to b.

4.1.7 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' balance-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 balance-adjacent boxes are split to further separate edges from each other. The balance condition is enforced.

7. Every box containing an edge and its balance-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 balance-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 4.9 for more remarks on this bound.

4.2 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 3. 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.

4.2.1 Vertex cone boxes

Theorem 2 *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 3.*

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 $I(\mathrm{ex}(B))$ or intersects $\partial I(\mathrm{ex}(B))$ the balance condition gives the result: Now matter how small a box is that intersects $\partial I(\mathrm{ex}(B))$, a box that intersects $\partial I(B)$ will be split to a size no smaller than half the distance from $\partial I(B)$ to $\partial I(\mathrm{ex}(B))$. This is very similar to Figure 4.2. If $I(d)$ actually lies outside of $I(\mathrm{ex}(B))$, then the same bound holds.

Subcase 2. Otherwise $I(d)$ lies strictly interior to $I(\mathrm{ex}(B))$. 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 $\partial I(\mathrm{ex}(B))$. If d is split down to a size one half that of this buffer box, for each of the children d_1 of d, $I(\mathrm{ex}(d_1))$ is a subset of $I(\mathrm{ex}(B))$. 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 $f \cap \mathrm{ex}(d)$ 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 a balance-adjacent non-crowded box. Once this is done, the children of this balance-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 balance-adjacent to the outer box. Also, splitting other boxes balance-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 balance-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 balance-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 2. 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 $k_1, k_2, \ldots$. 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 4.6. 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 $f \cap \mathrm{ex}(d^1)$ to e is on the shared boundary of B_e and d^1. So we have that the distance from $f \cap I(\mathrm{ex}(d^1))$ 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 $f \cap I(\mathrm{ex}(d^1))$ and face e is at least a constant factor k_1 times the angle β from e to e_2, as in Figure 4.6 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 4.6 left F is coincident with f, in Figure 4.6 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 $I(\mathrm{ex}(d^1))$.

The two right triangles at the top of Figure 4.6 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

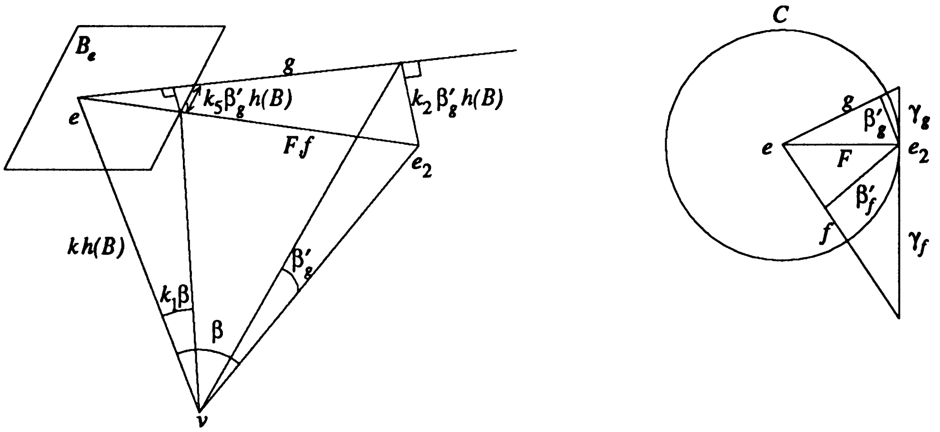

Figure 4.6: The distance between two superfacets of an edge inside an extended vertex box $I(\text{ex}(B))$ but outside an edge box is at least $k\alpha h(B)$.

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 4.6 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 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 4.6 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 $\mathrm{dist}(g \cap I(\mathrm{ex}(d^1)), F \cap I(\mathrm{ex}(d^1))) \geq k_5 \cdot \alpha \cdot h(B)$.

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

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

Recall our assumption that d^1 was split because g interfered with f. In particular, f and g are contained in $I(\mathrm{ex}(d^1))$, so $h(d^1)$ is at least a constant fraction of the distance between f and g in $I(\mathrm{ex}(d^1))$, that is $h(d^1) \geq kk_5 \cdot \alpha \cdot h(B)$. 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 3. 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 4.6 right for a spherical cross section of this cone, which is nearly a "top" view of Figure 4.6 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 4.6 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 2. ∎

4.2.2 Edge cone boxes

Theorem 3 *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. ∎

4.2.3 Facet Cone Boxes

Theorem 4 *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 5 *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.

4.3 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

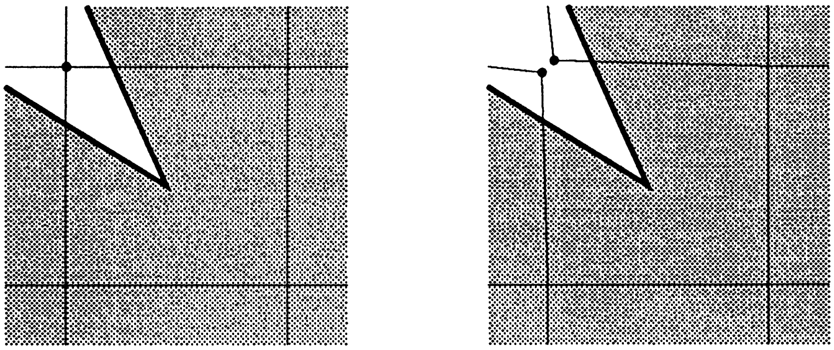

Figure 4.7: The special case of sharing a vertex with two boxes.

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 4.1.3) 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 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 4.7. 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 4.1.3. 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 4.7. 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 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, $t = k_1 ha$ 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.

We consider the case that x and y both arise as the intersection of a P facet F

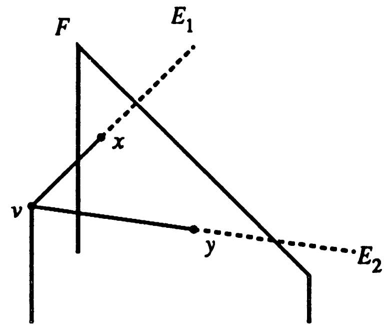

Figure 4.8: One case in the proof of Lemma 6.

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 4.8.

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

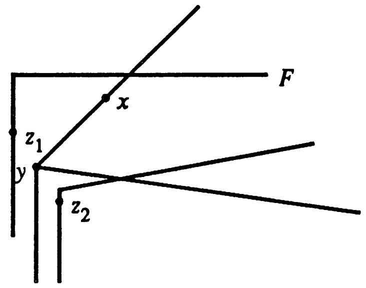

Figure 4.9: Another case in the proof of Lemma 6.

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 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 4.9. 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 $k_1 ha$. ∎

4.4 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 4.10; 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.

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

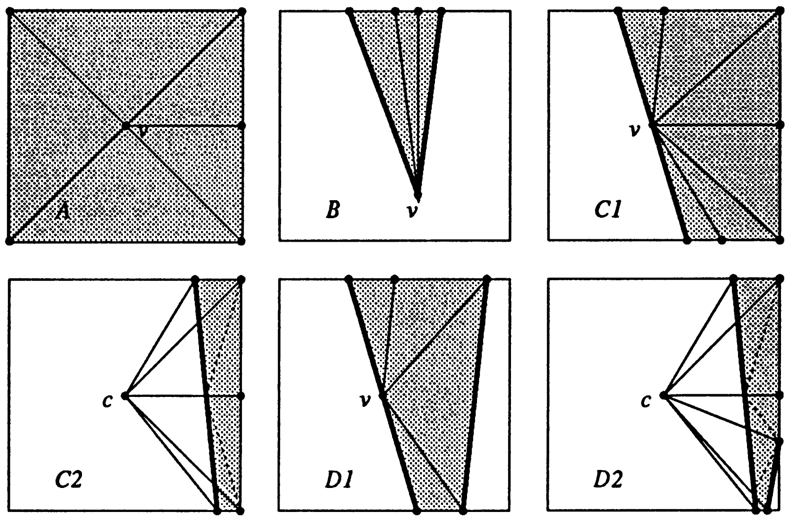

Figure 4.10: The six cases for triangulating a box facet.

triangle $T = \Delta xyz$ 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 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 $F \cap S$

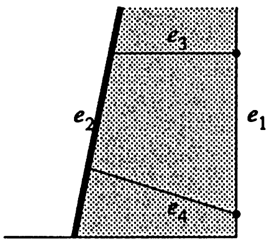

Figure 4.11: Triangulating a quadrilateral in Case C2 (closeup) .

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 4.11. 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 6 *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 5. In particular, h' in the previous lemma is at least $k\alpha h(B)$ by Theorem 5. 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.

4.5 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.

4.5.1 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.

4.5.2 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.

4.5.3 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.

4.5.4 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.

4.5.5 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.

4.6 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 multiple. 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 4.10.

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 4.10.

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, $k_1, \ldots, k_4$ 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 tetrahedron 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

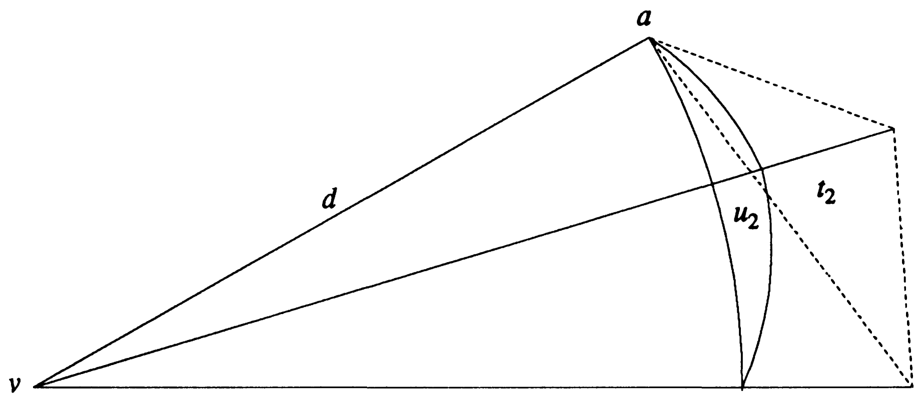

Figure 4.12: The flat (t_3) and curved (u_3) tetrahedra.

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 4.12. 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 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. ∎

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.

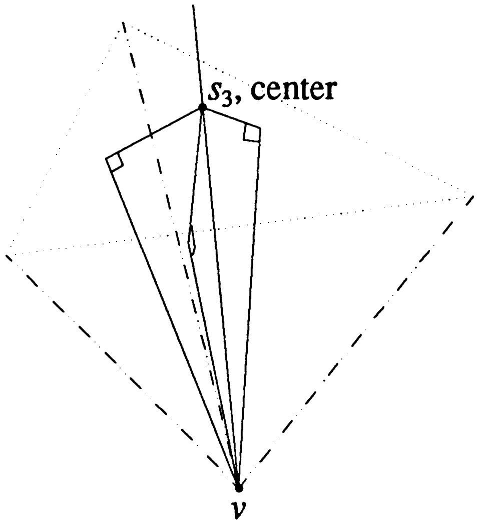

Figure 4.13: A sphere tangent to three facets of a tetrahedron.

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 4.13. 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 chose 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 4.14 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

$$r_3 \geq kd\theta.$$

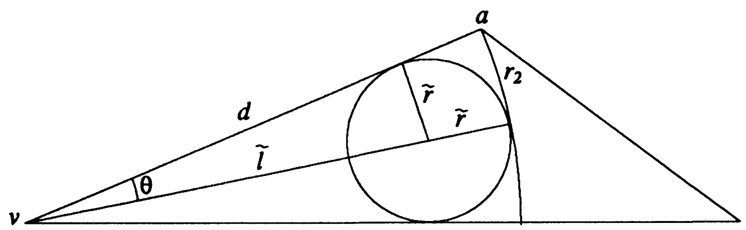

Figure 4.14: Constructing a large inscribed sphere of u_3.

Also from Lemma 10 and Figure 4.14,

$$\theta = r_2/d.$$

Recall the result of Theorem 6,

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

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

Theorem 7 (Aspect ratio of $\triangle^{OCT}$) *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 chose 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 4.5.4, 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.

4.7 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 8 *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 8. 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 3. 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 4.15. We let G_l be this segment, and note that p is contained in G. Now, let G_m be a segment joining a point on F and a point on G_l so that G_l, G_m, and F make a triangle and so that F and G_m make a right angle. Let p' be the intersection of F and G_m. Let p'' be another point on G_m close to p'. Given a three dimensional triangulation $\triangle^*$, we may chose G_m close enough to p and p'' close enough to p', such that every tetrahedron of $\triangle^*$ containing p'' also contains p and p'.

Let $\triangle^*_p$ be one of these tetrahedron containing p''. Then $\triangle^*_p$ contains p, p', and also has an edge, say F', that is coincident with a subsegment of F. Without loss of generality, G_m is moved so that p' is a vertex of $\triangle^*_p$ terminating edge F'. Let G' be the face of $\triangle^*_p$ opposite p'. Then the angle made by G' and F' is no larger than α, because G' is on a plane separating F' and G_l as in the last lemma.

We now show that this small angle leads to a large aspect ratio for $\triangle^*_p$. We first bound the altitude of $\triangle^*_p$ at p', which bounds the radius of the largest sphere inscribed in $\triangle^*_p$. If G' is a facet then we chose $H = G'$, otherwise we chose H as the unique facet of $\triangle^*_p$ containing edge G' but not edge F'. That is, we let H be the facet of $\triangle^*_p$ that intersects F' at exactly p. Then the angle α' between F' and the plane of H is at most α. It is clear that since tetrahedra are convex, the radius of the largest sphere inscribed in $\triangle^*_p$ is at most half of the altitude between p' and

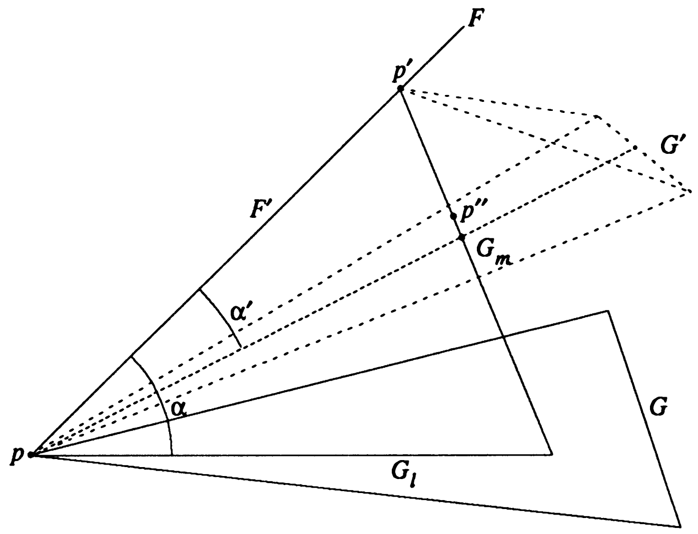

Figure 4.15: Finding the tetrahedron with an angle at most α.

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 $\triangle_p^*$, and we now find a lower bound on the size of the smallest sphere containing $\triangle_p^*$. 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 $\triangle_p^*$ 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 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 9 (Aspect ratio optimality)

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

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

4.8 Optimality of the cardinality of $\triangle^{OCT}$

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

Lemma 16 *If a triangulation $\triangle$ 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 $\triangle$ 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 $c_1, c_2, \ldots$ 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 $\triangle^*$ 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 $\triangle^T$. 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 $\triangle^*$ 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 $\triangle^*$, the ratio of the longest to shortest side of any simplex in $\triangle^*$ 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, $T_1, T_2, \ldots, T_k$, that share a common vertex.

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

Lemma 18 *Given a bounded aspect ratio triangulation $\triangle^*$, 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 $\triangle^*$. 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 $\mathrm{dist}_P(F, G)$ 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* $\mathrm{dist}_P(x, F) \geq c_6 \cdot f_*(x)$.

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

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 4.16 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 Δxvy. Let $\angle FG$ denote the interior angle between F and G. Then $\angle xvy \geq min(\angle FG, \pi/2)$. But $\angle FG$ 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.

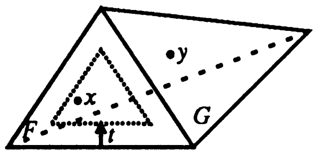

Figure 4.16: An illustration of the conditions stated in Lemma 20.

Theorem 10 *Let F, G be two faces of $\triangle^*$. 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 $\mathrm{dist}_P(x, H) \geq t$, then $\mathrm{dist}_P(x, y) \geq c_9 t$.*
2. *If $H = \emptyset$, then $\mathrm{dist}_P(x, y) \geq c_{10} f_*(x)$.*

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 $\mathrm{dist}_P(x, H) = 0$, and the claim is trivial. Similarly, if $G = H$ then $y \in H$ so $\mathrm{dist}_P(x, y) \geq \mathrm{dist}_P(x, H)$.

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 $\mathrm{dist}_P(F, F') \geq c_6 f_*(x)$.

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

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

Proof. If x and y are in a common tetrahedron of $\triangle^*$, 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 $\triangle^*$ crossed by $\overline{xy}$ by $F_i, i = 1, 2, \ldots, m$. 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 $\mathrm{dist}_P(x_1, x_{i^*+1}) \geq c_{10} \cdot f_*(x_1)$ by Theorem 10. Moreover, $\mathrm{dist}_P(x, y) \geq \mathrm{dist}_P(x_1, x_{i^*+1})$. Combining these inequalities gives the result that $f_*(x) \leq c_{14} \cdot \mathrm{dist}_P(x, y)$. ∎

We now observe the following characterization of our triangulation $\triangle^{OCT}$.

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 $\triangle^{OCT}$ 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 $\mathrm{dist}_P(x, y)$ 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*

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

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 $\mathrm{dist}_P(x, z) \leq \sqrt{3}h(b)$ and $\mathrm{dist}_P(y, z) \leq \sqrt{3}h(b)$. ∎

We wish to establish a series of lemmas that prove that the characteristic length function at a point of $\triangle^*$ 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. $\mathrm{dist}_P(x, y) \leq 6\sqrt{3}h(b_1)$.

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 $b_2, \ldots, b_l$. In this chain, b_l is crowded, and it causes b_{l-1} to split, which causes b_{l-2} to split, ... which causes b_1 to split. Without loss of generality, we can assume that all of the boxes $b_1, \ldots, b_{l-1}$ 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, $I(\mathrm{ex}(b_j)) \subset I(\mathrm{ex}(b_1))$ for all $j = 1, 2, \ldots l$. Recall that the balance condition is only enforced between balance-adjacent boxes, and neighboring boxes are balance-adjacent only if the smaller box contains a point of the extended box of the larger box. Hence, the single connected component $P \wedge \mathrm{ex}(b_1)$ contains $P \wedge \mathrm{ex}(b_j)$ for all $j = 1, 2, \ldots l$. This line of reasoning is also correct in the case that b_1 is itself crowded, in which case we take $l = 1$.

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 $\mathrm{dist}_P(x, y) \leq 4\sqrt{3}h(b_1)$.

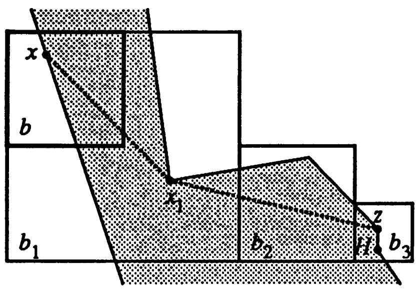

Figure 4.17: An example path in the proof of Lemma 23. Here all P vertices represent edges, and edges represent facets, except for the P vertex $J = z$. Note b_3 is crowded and b_1 is split by the balance condition.

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, $\ldots x_m$ to z. A cross-sectional view of an example path is given in Figure 4.17.

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 4.18. 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 in 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 10 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 11 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, $f_*(p) \leq c_{17}h(b)$.*

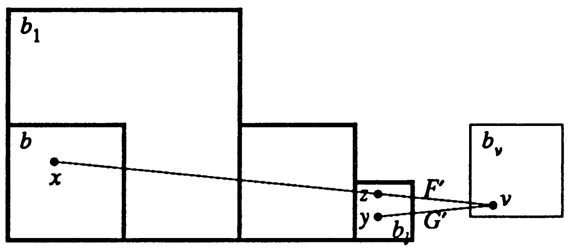

Figure 4.18: An illustration of an instance of Case 2 of the proof of Lemma 23 leading to the case in the proof of Lemma 26 where F' and G' have a common vertex v.

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 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 11 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 4.18. 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 10 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 $\triangle^*$. Let F' and G' be the smallest dimensional faces in $\triangle^*$, such that $x \in F' \subset F$ and $y \in G' \subset G$.

If F' and G' are disjoint, we can apply Theorem 11. 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, $\mathrm{dist}_P(p, v) \geq h(b_v)/8$. 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 10 Case 1, $\mathrm{dist}_P(p, y) \geq 8c_9 \min(c_{16}, c_{18}) f_*(v)$. Recalling from Lemma 23 that $\mathrm{dist}_P(p, y) \leq 6\sqrt{3}h(b_1)$ 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), $f_*(p) \le c_{21}h(b)$.*

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 $\mathrm{dist}_P(p, y) \le 3\sqrt{3}h(b_1)$, and $h(b_y) \le h(b_1)$, 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 $\max(c_{16}, c_{18}, c_{20})f_*(v) \le h(b_y)$. We apply Theorem 11 to p, so that $f_*(p) \le \max(c_{13}f_*(y), c_{14}\mathrm{dist}_P(p, y)) \le \max(c_{13}/\min(c_{16}, c_{18}, c_{20}), 3\sqrt{3}c_{14})h(b)$. ∎

We may now summarize the results of the previous lemmas.

Theorem 12 *Let $x \in P$ lie in a box b. Then $f_*(x) \le c_{22}h(b)$, where c_{22} is a constant depending on the aspect ratio bound of $\triangle^*$.*

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

Theorem 13 *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 $\triangle^*$ of P with maximum aspect ratio A,*

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

Proof. This follows directly from Lemma 21 and Theorem 12 ∎

Theorem 14 *For all polytopes P and constants A, there exists a constant c'', dependent only on A and α, such that*

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

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

Proof. Let $\triangle^T$ be a triangulation with bounded aspect ratio. We derive some inequalities that depend on aspect ratio; these inequalities apply to either $\triangle^*$ or $\triangle^{OCT}$. 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

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

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

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

This holds for all triangulations, so in particular for $\triangle^{OCT}$ and $\triangle^*$. By Theorem 13

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

∎

4.9 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. I implemented the two-dimensional analog of this algorithm in C++ while at Xerox PARC.

Chapter 5

Maxmin angle triangulation with no boundary Steiner points

We show how to triangulate a two dimensional polygon with holes. We assume the restriction that we may introduce Steiner points on the polygon's interior, but not on its boundary. Of all triangulations satisfying this restriction, our triangulation has the maximum minimum angle, up to a constant factor. This is the first known algorithm for this problem with provably optimal element shape. The algorithm solves several subproblems of mesh generation.

The largest possible minimum angle of a triangulation of a polygon P is dependent on the geometry of the polygon. Consider an edge $E = \overline{VU}$ of P, and W a point of a face of P disjoint from E (interfering with E). Let r be the ratio of the distance from W to the closer vertex V of E to the length of E, or $L(E)$. Let $\omega = \angle UVW$. Then we show that any triangulation of P has min angle at most

$$h(E, W) = O(\omega / \max(1, -\log r)).$$

Hence an upper bound on the max min angle of a triangulation of P is $g(P)$, the minimum of $f(E, W)$ over all E and W of P.

The algorithm we describe in this chapter constructs a triangulation with minimum angle at least a constant factor times $g(P)$. Thus our algorithm is optimal, up to a constant factor.

5.0.1 Two dimensional algorithm overview

First, we triangulate part of P by introducing sequences of triangles called *wedges* at vertices of P. The triangles have min angle within a constant factor of the optimal min angle of a triangulation of P. We call the remaining untriangulated polygon Q. We may triangulate Q using more traditional techniques, and still achieve min angle within a constant factor of optimal.

So that we do not need to introduce vertices on the boundary of Q, we *shrink* Q to the polygon R, where R closely resembles a small copy of Q lying inside Q. We triangulate R with the two dimensional analog of the algorithm in Chapter 4. This triangulates R with min angle within a constant factor of optimal, but also introduces Steiner points on the boundary of R.

Lastly, we *match* R to Q, that is we triangulate the region of P between Q and R using the edges of Q as a guide. No new vertices on the boundary of either Q or R are introduced in this last step.

The proof of the optimality of our algorithm is a step-by-step argument that we never introduce a triangle with an angle smaller than a constant times the optimal.

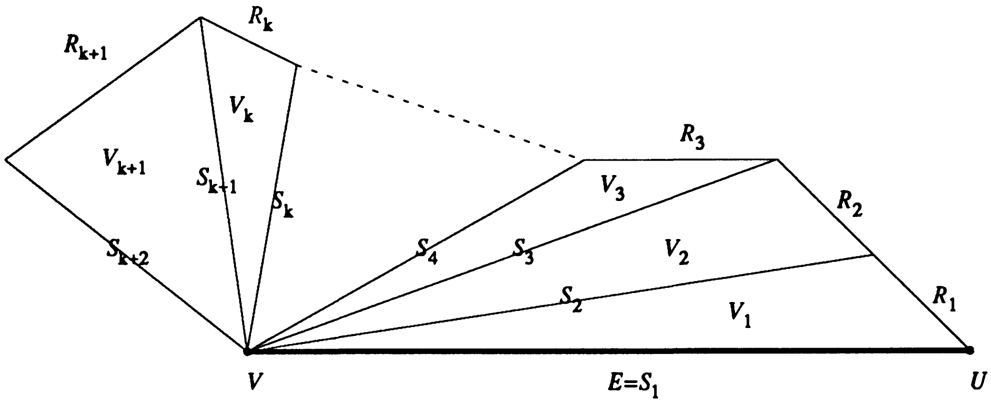

Figure 5.1: The triangles V_i, spokes S_i, and rim edges R_i of a wedge at V.

5.1 Constructing wedges around vertices

In this section we begin the triangulation of P. We form triangles that share an edge or vertex with ∂P. We call the portion of P which remains to be triangulated Q. The boundary of Q is completely composed of faces from the triangles we introduce in this section. That is, no triangle of a later section will contain an edge or vertex of ∂P. We first provide motivation for how we should triangulate.

The following definitions differ slightly from those given in Chapter 4.

Wedge. A *wedge* is a sequence of triangles such that all triangles contain a common vertex, and consecutive triangles contain a common edge. The common vertex V is called the vertex of the wedge, and the wedge is often called the *wedge at* V. The edges containing V are called the *spokes* of the wedge, and are ordered as in Figure 5.1. The edges opposite V are called the *rim edges* and are similarly ordered.

Interfere. For a closed face F of a polygon, we say that the closed polygon face G interferes with F if F and G are disjoint. We also say that any point of G interferes with F.

Length. We use the notation $L(E)$ to donote the length of the edge E.

Our goal is to produce a triangulation of P with max min angle at least a constant factor times the optimal min angle. In order to get a lower bound on the optimal min angle, we consider a class of triangulations. Any triangulation will have a wedge at every vertex V of P. Let E be a P edge containing V, and W a P vertex interfering with E. We consider an arbitrary triangulation from the class of triangulations that have an edge from V to W. We produce an upper bound on the min angle of the wedge at V with first spoke E and last spoke $\overline{VW}$. We do this by constructing an "optimal" wedge at V. The optimal wedge is closely related to the "A-optimal" wedge that we define below.

A-optimal wedge. We define the *A-optimal wedge* to be the wedge that has optimal min angle for the problem of constructing a wedge at V subject to only the constraints that its first spoke is E and its last spoke is $\overline{VW}$. Note that such a wedge may intersect the other faces of ∂P arbitrarily. Since, by assumption, in the triangulation under consideration there is a wedge at V that also satisfies these simple spoke constraints, the min angle of the A-optimal wedge is a lower bound on the min angle of this triangulation. Our optimal wedge is a wedge consisting of similar triangles that has min angle within a constant factor of the min angle of the A-optimal wedge. Note that Theorem 18 below shows that there really is an A-optimal wedge, and not an infinite sequence of wedges with increasing minimum angle. To summarize, we have proved the following theorem.

Theorem 15 *Let P vertex W interfere with P edge $E = \overline{VU}$. Then among all those wedges at V with first spoke E and last spoke $\overline{VW}$, there is a wedge with optimal min angle called the* A-optimal wedge. *The min angle of any triangulation* *of P that has edge $\overline{VW}$ is at least that of the A-optimal wedge.*

For simplicity we restrict our attention to the case that $L(E) \geq L(\overline{VW})$. To get a stronger grasp on what the triangles of the A-optimal wedge are like, we have the following three theorems.

Theorem 16 *Let θ be the angle at V for a triangle T of an A-optimal wedge in Theorem 15. Then*

$$\theta \leq \pi/2.$$

Proof. Suppose that the two spokes of T are of equal length (this is the limiting case). Let γ be the angle at the other vertices of T. If $\theta > \pi/3$, then $\gamma < \theta$, and the min angle of T is given by $(\pi - \theta)/2$. Otherwise $\theta \leq \pi/3$ and θ is the min angle of T. For $\theta \geq \pi/2$ the min angle of T is smaller than the min angle for the triangle with angle at V half that of T. Hence, if T had angle at V greater than $\pi/2$, we could divide it into two and improve the min angle, a contradiction.

If the spokes of T are not the same length, then creating two triangles by bisecting the angle at V with an appropriate length spoke improves the min angle even more. ∎

This bound allows us to write an approximation of the min angle of the optimal wedge in terms of a and b.

Theorem 17 *Let A be the min angle of a triangle T of an A-optimal wedge at V in Theorem 15. Let a be the ratio of the length of the shorter spoke to the longer spoke of T, and θ the angle of T at V. Then*

$$k_1 a\theta \leq A \leq k_2 a\theta,$$

where k_1 and k_2 are constants.

Proof. Let the longer spoke have length 1 and the shorter spoke length a. Let the rim edge have length c. Let γ be the angle at the vertex of the longer spoke that is not V. So, either $\gamma = A$ or $\theta = A$.

Suppose $A = \theta \leq \gamma$. Then, from the law of sines, $c \leq a$. But from the triangle inequality $a + c \geq 1$, so that $a \geq 1/2$. By definition $a \leq 1$, so $1a\theta \leq \theta \leq 2a\theta$.

Otherwise $A = \gamma \leq \theta$. Then, from the law of sines, $\sin\gamma = (a/c)\sin\theta$, and $c \geq a$. Since $c + a \geq 1$, we have $c \geq 1/2$. But also, from Theorem 16, we have that $c \leq \sqrt{a^2+1} \leq \sqrt{2}$. Hence $(a/\sqrt{2})\sin\theta \leq \sin(A) = \sin\gamma \leq 2a\sin\theta$. Noting that γ is bounded above by $\pi/2$ shows that the sines are within a constant factor of their arguments, and completes the proof. ∎

It is surprising that we may also derive a constant upper bound on the ratio between successive spokes.

Theorem 18 *Let a be the ratio of the lengths of the shorter spoke to the longer spoke of a triangle T in an A-optimal wedge in Theorem 15. Then*

$$1/4 \leq a \leq 1.$$

Proof. That $a \leq 1$ follows from its definition. To show $a \geq 1/4$, consider the min angle approximation $a\theta$. Assume the longer spoke has unit length. Divide T into two similar triangles by introducing a new spoke along the angle bisector of θ. Let the length of the new spoke be $a' = \sqrt{a}$. Then the new triangles have min angle approximation $a'\theta' = \sqrt{a}\theta/2$. This is an improvement over $a\theta$ if $a \geq 1/4$. Using the actual min angle instead of the approximation shows a greater improvement. Hence if a was less than $1/4$, then adding more spokes would improve the min angle of the A-optimal wedge. But since by definition the A-optimal wedge cannot be improved, this is a contradiction. ∎

Optimal wedge. We now define an *optimal wedge* to be the wedge that has the max min angle *approximation* $a_T\theta_T$ from Theorem 17, subject to $\theta_T \leq \pi/2$. Contrast this with the definition of an A-optimal wedge, which is the wedge that has the max min angle. Formally, the optimal wedge for a P edge $E = \overline{VU}$ and P vertex V, and interfering point W, is the wedge at V that has maximum minimum $a_T\theta_T$ over all triangles T of the wedge, subject to the constraint that the wedge has first spoke E, last spoke $\overline{VW}$, and $\theta_T \leq \pi/2$ for all triangles T of the wedge. We consider the optimal wedge because it is easier to analyze than the A-optimal wedge.

Theorem 19 *All triangles in an optimal wedge are similar, with spoke lengths following a geometric progression. The parameters a, the ratio of successive spoke lengths, and θ, the angle of a triangle at V, determine the shape of an optimal wedge.*

Proof. It is sufficient to show this for any two consecutive triangles T and U in an optimal wedge at V. Suppose that the number of spokes has been determined, and also the lengths of all spokes except for the spoke shared by both T and U. We now show that in order to maximize the minimum angle approximation of T and U, we should chose this shared spoke so that T and U are similar, and these three spokes form a geometric progression. Let a_T be the ratio of the length of the longer spoke of T to the length of the shorter spoke, and θ_T the angle of T at V. Similarly define a_U and θ_U.

We first prove that in an optimal wedge $a_T\theta_T = a_U\theta_U$. Suppose $a_T\theta_T < a_U\theta_U$, then we could decrease θ_U and increase θ_T by rotating the shared spoke about V, and the resulting max min angle approximation would increase, a contradiction.

Next we have three cases corresponding to possible ranges of the length of the shared spoke in the optimal wedge. Let the shared spoke have length s_s. Let the spoke unique to T have length s_T, and the spoke unique to U have length s_U. Without loss of generality assume that $s_T/s_U = r < 1$. Let the sum of the angles of T and U at V be ω.

Case 1, $s_T \le s_s \le s_U$. Then the problem can be stated as

$$\begin{aligned} &\text{Maximize } a_T\theta_T \\ &\text{s. t. } a_T\theta_T = a_U\theta_U \\ &a_Ta_U = r \\ &\theta_T + \theta_U = \omega \\ &\theta_T \le \frac{\pi}{2},\ \theta_U \le \frac{\pi}{2}. \end{aligned}$$

If we ignore the last constraint, we can use the other constraints to solve for θ_T in the objective function. The problem reduces to

$$\text{Maximize } \sqrt{r\theta_T(\omega - \theta_T)},$$

which has solution $\theta_T = \omega/2$. By assumption, the number of spokes have been determined, so that $\omega \le \pi$. Hence this solution satisfies the constraint that we ignored. Solving for the other variables leads to $\theta_U = \omega/2$ and $a_T = a_U = \sqrt{r}$, which shows that triangles T and U are similar. Moreover $s_s = \sqrt{r}\, s_U$, and $s_T = \sqrt{r}\, s_s$, so the spokes lengths are in a geometric series. The solution value is $\sqrt{r}\,\omega/2$.

Case 2, $s_s \le s_T$. Then it follows that $s_U \ge s_s$, and the problem can be stated as

$$\text{Maximize } a_T\theta_T$$

$$\text{s. t. } a_T\theta_T = a_U\theta_U$$

$$a_U/a_T = r$$

$$\theta_T + \theta_U = \omega$$

$$\theta_T \le \frac{\pi}{2},\ \theta_U \le \frac{\pi}{2}.$$

Using similar techniques, this has optimal solution at $\theta_T = \omega/(1+r), \theta_U = \omega r/(1+r), a_T = 1$ and $a_U = r$. The solution value is $\omega/(1+r)$. The solution is the same as that of Case 1 for $r = 1$. For the case that $r < 1$, the solution value is inferior to that of Case 1, and so this will not be the optimal solution.

Case 3 $s_s \ge s_U$. Then if follows that $s_s \ge s_T$, and the problem can be stated as

$$\text{Maximize } a_U\theta_U$$

$$\text{s. t. } a_T\theta_T = a_U\theta_U$$

$$a_T/a_U = r$$

$$\theta_T + \theta_U = \omega$$

$$\theta_T \le \frac{\pi}{2},\ \theta_U \le \frac{\pi}{2}.$$

By symmetry with Case 2, the solution here will always be inferior to the Case 1 solution for $r < 1$, and identical to it for $r = 1$. ∎

This allows us to write the optimal wedge for $E = \overline{VU}$ and W as the solution to a simple optimization problem. Let $r = L(\overline{VW})/L(E) \le 1$, the ratio of the length of the last spoke to the length of the first spoke. Let $a \le 1$ be the ratio of the size of a triangle to the one preceding it. Let ω be the angle between $\overline{VW}$ and E, that is $\omega = \angle WE = \angle WVU$. Let the angle at V of any triangle be θ. Then the number of triangles in the wedge is ω/θ. Our constraint that the first spoke is E and the last is $\overline{VW}$ may be written as $a^{\omega/\theta} = r$. Thus the optimal wedge is the solution to:

$$\text{Maximize } a\theta$$
$$\text{s. t. } a^{\omega/\theta} = r$$
$$\omega/\theta \geq 1, \text{ integer}$$

In light of Theorem 17 (and its dependence on Theorem 16), the optimal wedge has min angle within a constant factor of the solution value of the A-optimal wedge if $\theta \leq \pi/2$. Let us consider the relaxed problem in the case that $r < 1$:

$$\text{Maximize } a\theta$$
$$\text{s. t. } a^{\omega/\theta} = r$$

Since $r < 1$, we have that $a < 1$, and the constraint may be rewritten $\theta = \omega \log a / \log r$. Hence we may rewrite the problem as simply

$$\text{Maximize } a \log a \frac{\omega}{\log r}$$

From calculus this has solution $a = \mathrm{e}^{-1}$, and solution value

$$\frac{\omega}{-\mathrm{e} \log r}.$$

The corresponding value of $\theta = -\omega/\log r$. Noting that the minimum angle is always positive, and considering separately the cases when θ is greater than and less than $\pi/2$, this is sufficient to show that a constant factor times $\frac{\omega}{-\mathrm{e}\log r}$ is an upper bound on the min angle of the A-optimal wedge for E and W.

However, for $r < \mathrm{e}^{-1}$, we wish to find a wedge with min angle at least a constant factor times this upper bound. In general, in the solution to the relaxed problem

$\omega/\theta = -\log r > 0$ is non-integer. However, in the case that $r \leq e^{-1}$, we take the relaxed problem solution, and decrease θ to the nearest value such that ω/θ is an integer. We then increase a so that $a^{\omega/\theta} = r$ is once again satisfied. Since $r \leq e^{-1}$ implies $\omega/\theta \geq 1$, at worst we halve θ, and a is increased, so that there is a solution to the unrelaxed problem with value at least one half the relaxed solution. That is, the unrelaxed problem has solution value at least

$$\frac{\omega}{-2e\log r}.$$

Moreover, any feasible solution to the unrelaxed problem with ω/θ greater than some value has min angle worse than this feasible solution, hence there is an optimal wedge and not an infinite sequence of wedges approaching an optimal value.

In the case that $r > e^{-1}$, we merely note that no triangle can have angle at V larger than ω. Hence, no wedge at V can have min angle more than ω. In this case it is easy to construct a wedge at V, consisting of no more than four similar triangles, with θ at least a constant times ω and no more that $\pi/2$, and with a bounded by a constant. Thus we have proved the following theorem.

Theorem 20 *For every edge E and interfering vertex W, there exists an optimal wedge consisting of a finite number of triangles. The min angle of the optimal wedge is $O(\omega/\max(1, -\log r))$. The minimum angle of the optimal wedge is bounded above and below by constant factors times the minimum angle of the corresponding A-optimal wedge.*

We narrow the definition of wedge slightly to differentiate between wedges that we construct and wedges of arbitrary triangulations.

Spiral. A wedge that we construct that consists of similar triangles is called a spiral.

Optimal wedges are spirals. In the following discussion, we will be concerned with spirals only when E is the longest edge of the spiral containing V, and the interfering point W defining the spiral is closer to the vertex of the spiral than to the other vertex of E. The first condition will always be satisfied when $\text{dist}(V, W) \leq L(E)$. The P edge E will always be the first edge of the spiral.

Theorem 21 *Given P edge $E = \overline{VU}$ and interfering point W, if* $\text{dist}(V, W) \leq L(E)$ *and* $\text{dist}(V, W) \leq \text{dist}(U, W)$*, then any triangulation must have min angle at most a constant times that of the optimal wedge with first spoke E and last spoke $\overline{VW}$.*

Proof. Note that W is any point interfering with E, and not necessarily a P vertex. If $\overline{VW}$ is an edge of the given triangulation, then Theorem 20 applies. So there is only something to prove in the case when there is *no* triangulation edge from V to W. Note that there will still be a wedge at V, but not one with last spoke $\overline{VW}$. Let the sequence of triangles $V_1, V_2, \ldots V_k$ be the wedge at V such that E is the first spoke and W is in the infinite sector defined by the spokes of V_k. See Figure 5.2. The outline of the proof follows.

We transform the wedge at V in a given triangulation to one that includes the edge $\overline{VW}$. We use different transformations depending on whether $k > 1$ or $k = 1$. For $k > 1$, there is one step that at most halves the minimum angle, while all other steps do not decrease the minimum angle. For $k = 1$, the transformation does not increase the minimum angle. In either case, the transformed wedge must have minimum angle at most that of the A-optimal wedge, by definition. Hence any given triangulation must have min angle at most twice that of the A-optimal wedge. The proof is completed by noting that the A-optimal wedge has min angle

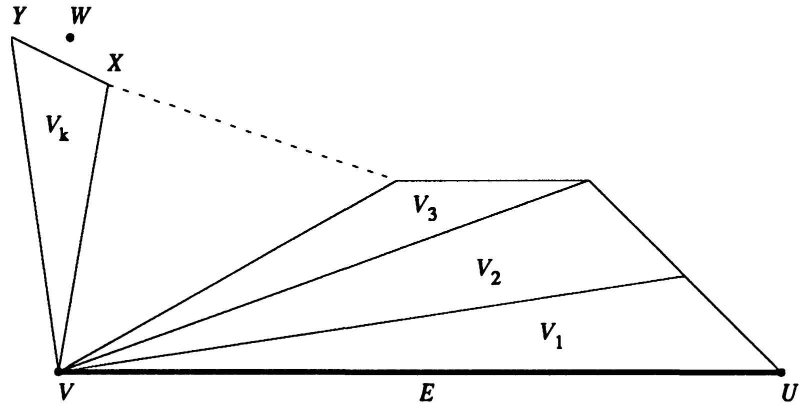

Figure 5.2: The wedge at V in a given triangulation where the sector of the last triangle $V_k = \triangle YVX$ contains the interfering vertex W.

at most a constant factor larger than the min angle of the optimal wedge. We now describe the transformations.

If $k > 1$, then let X be the vertex opposite V of the kth (second to last) spoke. We have two cases which require different transformations.

Case 1, $L(\overline{VX}) \leq 2\text{dist}(V, W)$. In this case we will ignore V_k and consider transformations to the wedge $V_1, V_2, \ldots V_{k-1}$. By definition, the wedge $V_1, V_2, \ldots V_{k-1}$ has min angle at most that of the A-optimal wedge at V for last spoke $\overline{VX}$, and hence at most a constant factor times that of the corresponding optimal wedge.

If we place a vertex X' at the midpoint of $\overline{VX}$, the optimal wedge at V with last spoke $\overline{VX'}$ has min angle approximation at least half that of the optimal wedge at V with last spoke $\overline{VX}$. This can be seen by noting that we have just halved r in the optimization problem leading to Theorem 20. Now, moving X' to W increases r and increases ω in the constraint $a^{\omega/\theta} = r$, so that the optimal wedge for last spoke $\overline{VX'}$ has min angle approximation smaller than that of the optimal wedge for last spoke $\overline{VW}$. Hence, the min angle of the given triangulation is at most a constant times the min angle of the optimal wedge at V with last spoke

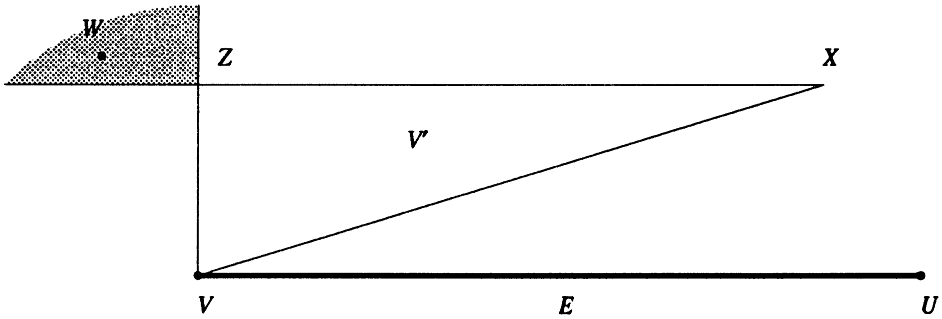

Figure 5.3: The first subcase of where an interfering vertex W may lie relative to the spiral triangle V'.

$\overline{VW}$.

Case 2, $L(\overline{VX}) > 2\text{dist}(V, W)$. We consider a series of transformations of the wedge $V_1, V_2, \ldots V_k$ that changes it into a wedge that contains W. We replace V_k by the triangle $V' = \triangle XVZ$, where the angle of V_k and V' at X are the same, and the angle of V' at Z is $\pi/2$. We wish to show that the min angle of V' is not smaller than the min angle of V_k. By assumption the altitude of V_k and V' at V is at most $\text{dist}(V, W) < L(\overline{VX})/2$. Hence the angle γ at X is at most $\pi/6$, and the smallest angle of V' occurs at X. Since the angle at X of the two triangles is the same, the min angle of V' is at least the min angle of V_k.

We have two subcases. The first subcase is when W is not in the infinite sector $\angle XVZ$. See Figure 5.3. We note that the min angle of the wedge $V_1, V_2, \ldots V_{k-1}, V'$ is at most a constant times that of the optimal wedge at V with last spoke $\overline{VZ}$. Since the angle at X was unchanged when we replaced V_k by V', and W is in the sector at V defined by V_k, but not inside V_k itself, it is always the case that line $\overline{ZX}$ separates W from V. Since Z is the closest point of line $\overline{ZX}$ to V, we have that $\text{dist}(V, W) \geq \text{dist}(V, Z)$, and $\angle WVU \geq \angle ZVU$. Referring to the constraints leading to Theorem 20 shows that the min angle approximation of the optimal

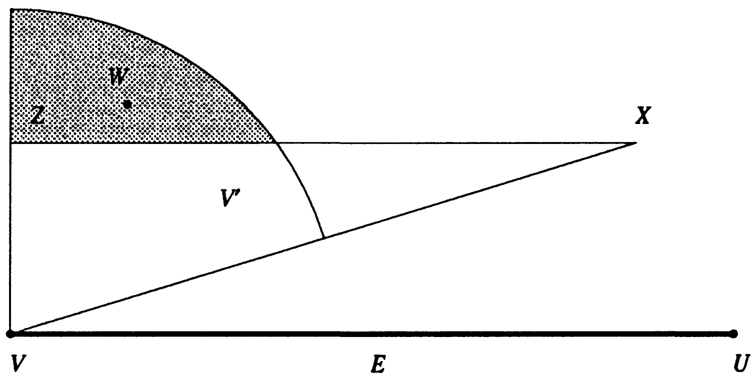

Figure 5.4: The second subcase of where an interfering vertex W may lie relative to the spiral triangle V'.

wedge at V with last spoke $\overline{VW}$ is more than the min angle approximation of the optimal wedge at V with last spoke $\overline{VZ}$. Hence we have proved the theorem for this subcase.

The second subcase is when W is in the infinite sector $\angle XVZ$. Since also $\text{dist}(W, V) \leq \text{dist}(X, V)/2$, we note that W is in the shaded region of Figure 5.4. Recall that the smallest angle of V' occurs at X. We transform V' by moving moving Z to W. From the law of sines, the new angle at V is at least the old angle at X. Also the new angle at X is at least the old angle at X. Hence the min angle of V' is not decreased by this transformation. We have constructed a wedge $V_1, V_2, \ldots V_{k-1}, V'$ with min angle at least that of the original wedge at V. This wedge $V_1, V_2, \ldots V_{k-1}, V'$ contains W as the last vertex so its min angle is at most that of the A-optimal wedge at V to W, and at most a constant factor times that of the optimal wedge.

If $k = 1$, then let Y be the third vertex of V_1. We transform V_1 into V' by intersecting it with the half plane containing U defined by the line through $\overline{VW}$. The angle at U is the same for both V_1 and V'. By assumption W is closer to V

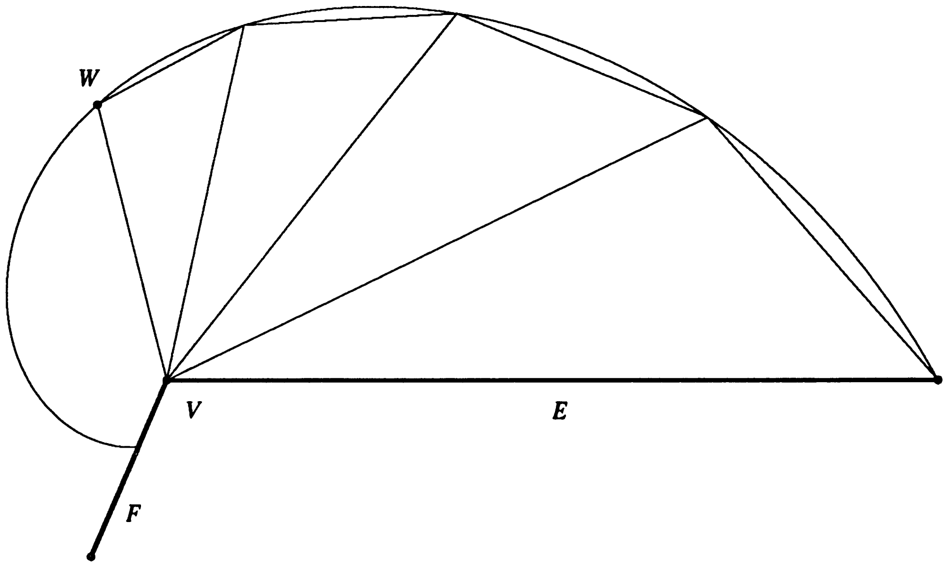

Figure 5.5: The continuous spiral for an optimal wedge.

than to U, so the angle at V of V' is at least the angle at U. The angle at the new vertex Y' of V' is greater than the angle at Y of the old triangle. Hence V' has min angle at least that of V_1. Next, the optimal wedge for E and Y' has min angle at least a constant factor times that of V', which is at least that of V_1. By examining the constraints leading to Theorem 20 we see that this optimal wedge has r smaller than r for the optimal wedge for E and W. Hence the optimal wedge for E and W has min angle at least a constant factor times that of the optimal wedge for E and Y', which is at least that of V_1, and the proof is complete. ∎

Continuous spiral. An optimal wedge for P edge E and interfering point W defines a and θ as in Theorem 19. We define the *continuous spiral* for E and W to be the $\{\rho, \phi\}$ polar curve $\rho = L(E)a^{\phi/\theta}$, where the origin is at V, the $\phi = 0$ axis is aligned with E. We allow ϕ to vary between 0 and the angle that the other P edge containing V makes with E. See Figure 5.5. Note that the spoke vertices of the optimal wedge lie on the continuous spiral. All other points of the optimal wedge lie "inside" the continuous spiral. We say that set Y lies *inside* a continuous spiral Z if for all $y = \{\rho_y, 0 \leq \phi_y \leq 2\pi\} \in Y$, there is a $z = \{\rho_z, \phi_z\} \in Z$ such that $\phi_z = \phi_y$ and $\rho_y \leq \rho_z$.

Shrinking factor and optimal spiral. We may rewrite a continuous spiral as $\rho = L(E)(a^{1/\theta})^{\phi}$. We define $a^{1/\theta}$ to be the *shrinking factor* of the continuous spiral. The shrinking factor will always be less than one for the continuous spirals we consider. The shrinking factor and $L(E)$ defines a continuous spiral. The optimal wedge that defines a continuous spiral has max min angle approximation for all spirals with spoke vertices on a continuous spiral, angle at V no more than $\pi/2$, and the fixed last spoke. Hence, we may think of a continuous spiral as defining an optimal wedge, except that we don't know the angle of the last spoke of the optimal wedge. So we distinguish as the *optimal spiral*, the corresponding optimal wedge whose last spoke vertex has ϕ that of the other P edge containing V. The optimal spiral for a continuous spiral has min angle approximation at least a constant times that of any optimal wedge that defines the continuous spiral. This fact follows from examining the rounding we did when establishing the optimal spiral from the relaxed optimization problem solution leading the Theorem 20.

In the following discussion, we will take a given continuous spiral, transform it by increasing its shrinking factor, and consider the optimal spiral it defines. Hence we wish to explicitly write how the min angle approximation of the optimal spiral depends on the shrinking factor of the continuous spiral.

If we ignore where the last spoke of an optimal spiral must lie, then we may write the corresponding optimal spiral in terms of a shrinking factor s. That is, we may write it as the solution to

$$\begin{aligned} &\text{Maximize } a\theta \\ &\text{s. t. } a^{1/\theta} = s, \\ &\qquad \theta \le \frac{\pi}{2}. \end{aligned}$$

If we ignore the last constraint, this has solution $a = \mathrm{e}^{-1}$ and $\theta = -1/\log s$, and solution value $-1/\mathrm{e}\log s$. The constraint that $\theta \le \pi/2$ is only important when s is large (close to 1), in which case we can find a spiral with min angle within a constant factor of the angle between the two P edges meeting at V. Imposing the constraint of where the last spoke must lie, we may find a feasible solution with larger a, and θ no less than a constant fraction of the above solution value. Hence we have proved the following theorem.

Theorem 22 *Let ω be the angle between two P edges E and F meeting at a vertex. Given a continuous spiral for edge E with shrinking factor s, there is a corresponding optimal spiral with min angle approximation bounded above and below by constants times* $\min(\omega, -1/\mathrm{e}\log s)$.

We now begin to define curves that will contain our first set of triangles in the triangulation of P. We start with the continuous spiral defined in the following theorem, and transform it in the next definition.

Theorem 23 *Given P edge $E = \overline{VU}$, then there is a continuous spiral F called the* empty spiral, *such that*

1. *The only points of ∂P inside F are closer to U than to V, or are in the P edges containing V.*

2. *The min angle of the optimal spiral defined by F is at least a constant factor times the optimal min angle of P.*

Proof. Each optimal wedge at V for edge E defined by an interfering point in turn defines a continuous spiral. Of these continuous spirals, let F be the one with smallest shrinking rate. Note F will be inside all the other continuous spirals. Each point interfering with E is a vertex of the optimal spiral it defines, and hence lies on the corresponding continuous spiral. Hence we may shrink F by an infinitesimal amount so that 1 is satisfied. Item 2 follows from Theorem 21 and Theorem 22.

If there is no interfering point close enough to V to define an optimal wedge, then we take S to be the spiral with shrinking factor 1, that is $\rho = L(E)$. The corresponding optimal spiral has min angle at least a constant factor times the angle between the P edges meeting at V. ∎

Smallest spiral. For each empty spiral F we define a continuous spiral called the *smallest spiral.* Suppose the empty spiral is defined by $\rho = L(E)s^{\phi}$. Then the smallest spiral is defined by $\rho = \frac{L(E)}{4}(s^4)^{\phi}$. The smallest spiral shrinks four times as fast as the empty spiral. Also, the distance from V to the point at $\rho = 0$ in the smallest spiral is only one quarter the length of E.

Theorem 24 *For a smallest spiral S, the optimal spiral defined by it is called the* smallest optimal spiral *and denoted T. Let A be the min angle approximation of T.*

1. *A is at least a constant factor times the min angle of the optimal spiral for the empty spiral at V for E.*

2. For any rim edge J of T, let d be the distance from J to a smallest optimal spiral for another vertex. Then $d/L(J) \geq kA$ for some constant k.

Proof. The first item follows directly from the definition of the smallest spiral and Theorem 22.

For the second item, we develop a series of lemmas. We first establish the length of a rim edge.

Lemma 31 *For a point $Y = \{\rho_Y, \phi_Y\}$ on a rim edge J,*

$$L(J) \leq k_1 \rho_Y,$$

for some constant k_1.

Proof. Suppose Y is a spoke vertex. The length of the larger rim edge J containing Y is $L(J) = \rho_Y \sin\theta/\sin\gamma$, where θ and γ are defined as for a triangle of an optimal wedge. Since $a \geq \mathrm{e}^{-1}$, γ is never more than a constant factor smaller than θ, and $L(J) \leq k_1 \rho_Y$, for some constant k_1.

If Y is not actually the vertex of J, then let Z be the vertex of J making larger angle with E. Since $\theta \leq \pi/2$, ρ_Y is at least $\rho_Z \cos\frac{\pi}{4}$. The bound on $L(J)$ in terms of ρ_Z from the last paragraph completes the proof. ∎

We next establish the distance from a point of S to any other smallest spiral. The proof is not very direct. Let F be the empty spiral defining S. For a point Z of S with ϕ_Z in a certain range we are able to show that it is closer to E than to F by a constant factor, and its distance to F is reasonably large. If this were true for all Z of all smallest spirals regardless of ϕ_Z, then since all other faces of P are outside of F, this would show that the two smallest spirals for disjoint P edges are well separated. However, it is *not* true for points with ϕ_Z in the remaining range, and we must argue directly that any other smallest spiral must be far away from these points. Let s be the shrinking factor of F.

Lemma 32 *For a point Z in S with $\phi_Z \geq \pi/2$,*

$$\text{dist}(Z,F) \geq 3\rho_Z = 3\text{dist}(Z,E).$$

Proof. For any point W of F, the triangle inequality gives us $\text{dist}(Z,W) \geq \rho_W - \rho_Z$. For $\phi_Z \geq \pi/2$, we have $\rho_Z \leq \frac{L(E)}{4}s^{4(\pi/2)} = \frac{L(E)}{4}s^{2\pi}$. But $\rho_W \geq L(E)s^{2\pi}$. ∎

This generalizes to

Lemma 33 *For a point Z of S,* $\text{dist}(Z,W) \geq 3\rho_Z$ *for all points W of F with $\phi_W \leq 4\phi_Z$.*

We are also able to show that points Z with small ϕ_Z are closer to E than to F.

Lemma 34 *For a point Z with $\phi_Z \leq \pi/3$,*

$$\text{dist}(Z,F) \geq k\frac{\rho_Z}{\max(1,-log(s))},$$

for some constant k, and

$$\text{dist}(Z,F) \geq \frac{2}{\sqrt{3}}\text{dist}(Z,E).$$

Proof. We first find the point W at angle ϕ_W such that $\rho_W = 2\rho_Z$. This is equivalent to $s^{\phi_W} = s^{4\phi_Z}/2$, which has solution $\phi_W = 4\phi_Z + \frac{1}{-\log s}$.

It might be the case that no such W exists. That is, perhaps the solution for ρ_W above is larger than 2π. Then for all W, $\rho_W \geq 2\rho_Z$. Then the triangle inequality yields $\mathrm{dist}(W, Z) \geq \rho_Z$. But also $\mathrm{dist}(Z, E) = \rho_Z \sin \phi_Z \leq \frac{\sqrt{3}}{2}\rho_z$, and we have proved the lemma for this Z.

So assume there is such a W. The reasoning of the last paragraph proves the lemma if we restrict F to the points W' of F such that $\phi'_W \leq \phi_W$.

It remains to consider the distance between Z and any point W' with $\phi_{W'} \geq \phi_W$. This distance is at least the distance between Z and the ray at V at angle ϕ_W. Let $\delta = \phi_W - \phi_Z$. If $\delta \geq \pi/2$, then this distance is ρ_Z and the lemma follows for all such points W' of F.

Otherwise, $\delta \leq \pi/2$. Note that $\mathrm{dist}(Z, W') \geq \rho_Z \sin \delta$, or

$$\mathrm{dist}(W', Z) \geq \rho_Z \sin(3\phi_Z + \frac{1}{-\log s}).$$

Here $\sin \delta \geq \frac{2}{\pi}\delta$, and we may reduce this to

$$\mathrm{dist}(W', Z) \geq \frac{2}{\pi}\frac{\rho_Z}{-\log s},$$

which proves the first assertion of the lemma.

For the second assertion note that $\mathrm{dist}(Z, E) = \rho_Z \sin \phi_Z$. If $\phi_Z = 0$, the second assertion is trivial. Otherwise $\phi_Z > 0$ and

$$\frac{\mathrm{dist}(W', Z)}{\mathrm{dist}(Z, E)} \geq \frac{\sin 3\phi_Z}{\sin \phi_Z}.$$

Recall that we are in the case $\delta \leq \pi/2$, so in particular $\phi \leq \pi/6$ and the ratio on the right hand side of the above is at least 2. ∎

In the previous two lemmas, we considered the distance from a point Z to F. However, there may be points interfering with E inside F, but closer to the other vertex of E than to Z. The following lemmas proves bounds similar to those of the previous lemmas for this case.

Lemma 35 *If point W interferes with E and is inside S, then for all Z of S*

$$\mathrm{dist}(Z, W) \geq \mathrm{dist}(Z, V),$$

and

$$\mathrm{dist}(Z, W) \geq \sqrt{3}\mathrm{dist}(Z, E).$$

Proof. By the definition of smallest spiral, $\rho_Z = \frac{L(E)}{4} s^{4\rho_Z}$. Hence $\rho_Z \leq L(E)/4$. The only possible W interfering with E but inside S are closer to the other vertex of E than to Z. So W lies on the side of the perpendicular bisector of E opposite V. Hence $\mathrm{dist}(Z, W) \geq L(E)/2$ for $\phi_Z \geq \pi/2$, and $\mathrm{dist}(Z, W) \geq L(E)(2 - \cos\phi_W)/4$ otherwise. Thus $\mathrm{dist}(Z, W) \geq L(E)/4 \geq \rho_Z$ and the first inequality is proved.

To prove the second inequality we consider two ranges of ϕ_Z. For $\phi_Z \geq \pi/2$ we have $\mathrm{dist}(Z, W) \geq L(E)/2 \geq 2\mathrm{dist}(Z, E)$. For $\phi_Z \leq \pi/2$ we have $\mathrm{dist}(Z, E) = \rho_Z \sin\phi_Z \leq \frac{L(E)}{4} \sin\phi_Z$. For $\phi_Z = 0$ the second inequality is trivial, otherwise

$$\frac{\mathrm{dist}(Z, W)}{\mathrm{dist}(Z, E)} \geq \frac{2 - \cos\phi_W}{\sin\phi_Z}.$$

This has minimum value $\sqrt{3}$ at $\phi_Z = \pi/3$. ∎

Lemma 36 *For point Z of S, and point Z_1 of S_1, if one of the following holds*

1. $\phi_Z \leq \pi/3$, *or* $\phi_Z \geq \pi/2$, *or* Z_1 *is inside* S.

2. $\phi_{Z_1} \leq \pi/3$, *or* $\phi_{Z_1} \geq \pi/2$, *or* Z *is inside* S_1,

then

$$\mathrm{dist}(Z, Z_1) > 0.07 \max(\mathrm{dist}(Z, E), \mathrm{dist}(Z_1, E_1)).$$

Proof. Assume 2 holds. Let G be the closest point of E to Z, and similarly define G_1.

The easy case is if 1 also holds. In this case we may relabel so that $\mathrm{dist}(Z,G) \geq \mathrm{dist}(Z_1,G_1)$. Note that G_1 interferes with E, so from Lemma 32, Lemma 34, or Lemma 35, we have that $\mathrm{dist}(Z,G_1) \geq \frac{2}{\sqrt{3}}\mathrm{dist}(Z,G)$. But then the triangle inequality gives $\mathrm{dist}(Z,Z_1) \geq \mathrm{dist}(Z,G_1) - \mathrm{dist}(Z_1,G_1) \geq (\frac{2}{\sqrt{3}} - 1)\mathrm{dist}(Z,G)$.

It remains to consider when 1 does not hold, that is when $\pi/3 \leq \phi_Z \leq \pi/2$ and S_1 is not inside S. Suppose the lemma were false and $\mathrm{dist}(Z,Z_1) \leq 0.07\mathrm{dist}(Z,G)$. From Lemma 32, Lemma 34 or Lemma 35, we have $\mathrm{dist}(Z_1,G_1) \leq \frac{\sqrt{3}}{2}\mathrm{dist}(Z_1,G)$. Hence $\mathrm{dist}(Z,G_1) \leq \mathrm{dist}(Z_1,G_1) + \mathrm{dist}(Z,Z_1) \leq \frac{\sqrt{3}}{2}\mathrm{dist}(Z_1,G) + 0.07\mathrm{dist}(Z,G)$. But $\mathrm{dist}(Z_1,G) \leq \mathrm{dist}(Z_1,Z) + \mathrm{dist}(Z,G) \leq (1+0.07)\mathrm{dist}(Z,G)$. So $\mathrm{dist}(Z,G_1) \leq (\frac{\sqrt{3}}{2}(1+0.07) + 0.07)\mathrm{dist}(Z,G) < \mathrm{dist}(Z,G)$.

Hence we have that G_1 lies inside the circle at Z with radius $\mathrm{dist}(Z,E)$. Since additionally $\phi_Z \leq \pi/2$, we have that the angle that G_1 makes with E at V is no more than $\phi_Z + \pi/2 \leq 2.5\phi_Z$. Note G_1 interferes with E. Hence either G_1 is not inside S, and from Lemma 33 we have that $\mathrm{dist}(V,G_1) \geq 4\mathrm{dist}(Z,V)$, or G_1 is inside S and from the proof of Lemma 35 we have $\mathrm{dist}(V,G_1) \geq 2\mathrm{dist}(Z,V)$. Hence $\mathrm{dist}(Z,G_1) \geq \mathrm{dist}(V,G_1) - \mathrm{dist}(Z,V) \geq \mathrm{dist}(Z,V) \geq \mathrm{dist}(Z,E)$. But in the last paragraph we just showed that $\mathrm{dist}(Z,G_1) < \mathrm{dist}(Z,E)$, a contradiction, and hence $\mathrm{dist}(Z,Z_1) > 0.07\mathrm{dist}(Z,G)$.

Since we want a bound in terms of $\max(\mathrm{dist}(Z,E),\mathrm{dist}(Z_1,E_1))$, it remains to consider the case that $\mathrm{dist}(Z_1,G_1) \geq \mathrm{dist}(Z,G)$. Using the same arguments as in the easy case described above when both 1 and 2 were true, we can show that $\mathrm{dist}(Z,Z_1) \geq (\frac{2}{\sqrt{3}} - 1)\mathrm{dist}(Z_1,G_1)$. ∎

Only one lemma remains to fill in the case when both suppositions 1 and 2 of the previous lemma are false.

Lemma 37 *For point Z of S, and point Z_1 of S_1, if $\pi/3 \leq \phi_Z \leq \pi/2$, Z_1 is not inside S, $\pi/3 \leq \phi_{Z_1} \leq \pi/2$, and Z is not inside S_1, then*

$$\mathrm{dist}(Z, Z_1) \geq 2\max(\rho_Z, \rho_{Z_1}).$$

Proof. Suppose $\rho_Z \geq \rho_{Z_1}$. We consider the angle ϕ_{V_1} that V_1 makes with E at V. If $\phi_{V_1} \leq 4\phi_Z$, then Lemma 33 applies and $\mathrm{dist}(V, V_1) \geq 3\rho_Z$. Hence from the triangle inequality $\mathrm{dist}(Z, Z_1) \geq \mathrm{dist}(V, V_1) - \mathrm{dist}(V, Z) - \mathrm{dist}(V_1, Z_1) \geq 2\rho_Z$ and the lemma is proved restricted to such V_1.

It remains to consider V_1 such that $\phi_{V_1} \geq 4\phi_Z$. From the previous paragraph, it is obvious that we may also restrict our attention to Z_1 such that $\mathrm{dist}(V, Z_1) \leq L(E)/2$. This avoids certain complications when considering geodesic distance below. Define $\delta = \phi_{V_1} - \phi_Z$. Note $\delta \geq \pi$. The main consequence of this is that the shortest geodesic path from Z to V_1 passes through V, and so $\mathrm{dist}(Z, V_1) = \rho_Z + \mathrm{dist}(V, V_1)$. This and the triangle inequality form the basis of the proofs in the cases below.

Define ϕ_V to be the angle V makes with E_1 at V_1. We have two main cases depending on the orientation of the interior side of E_1 relative to that for E.

The first case is when the orientation or handedness of the interior side of E_1 with respect to V_1 is *the same* as the handedness of E with respect to V. See Figure 5.6. We have two subcases.

The first subcase is if $\phi_V \geq \phi_{Z_1}$. Then the shortest geodesic path from Z_1 to Z passes through V, and so has distance at least ρ_V.

The second subcase is if $\phi_V \leq \phi_{Z_1}$. But then since V interferes with V_1, we have $\rho_{Z_1} \leq \rho_V$. Hence $\mathrm{dist}(Z, Z_1) \geq \mathrm{dist}(Z, V_1) - \mathrm{dist}(Z_1, V_1) \geq \rho_Z$.

The second case is when the handedness of the interior side of E_1 with respect

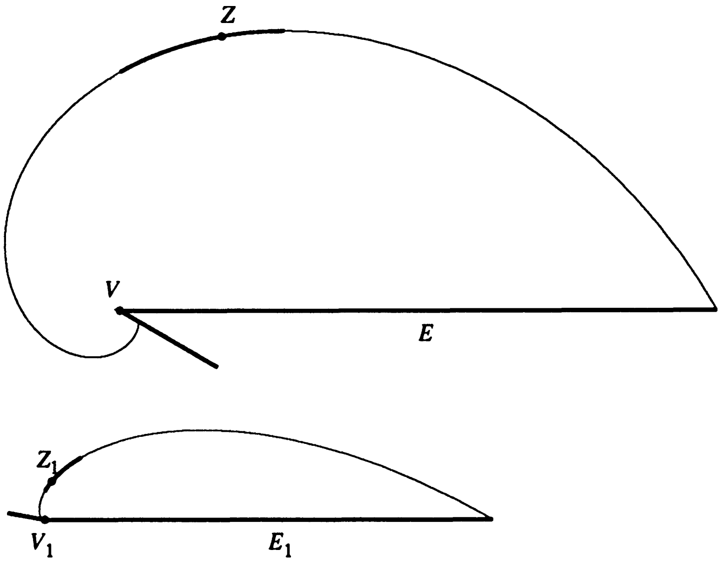

Figure 5.6: The interior side of E relative to V is the same as the interior side of E_1 with respect to V_1.

to V_1 is *opposite* the handedness of E with respect to V. See Figure 5.7. If $\phi_V \geq \pi$, then the fact that $\phi_{Z_1} \leq \pi/2$ insures that $\phi_V - \phi_{Z_1} \geq \pi/2$. Hence the closest point to Z on the ray at V_1 at angle ϕ_{Z_1} is V_1, so $\text{dist}(Z, Z_1) \geq \rho_Z$.

It remains to consider the subcase that $\phi_V \leq \pi$. But then $\phi_V \leq 4\phi_{Z_1} \leq 4\pi/3$, so that $\rho_{Z_1} \leq 4\rho_V$. Hence by the triangle inequality, $\text{dist}(Z, Z_1) \geq \text{dist}(Z, V_1) - \text{dist}(V_1, Z_1) \geq \rho_Z + \rho_V - \rho_V/4 \geq \rho_Z$. ∎

Thus we have succeeded in showing that for a point Z of a smallest spiral S, its distance to a point of any other smallest spiral is at least a constant factor times $\rho_Z A$, where A is the min angle of the corresponding smallest optimal spiral. In particular, this implies that the curves of two smallest spirals S and S' for different P vertices do not intersect. Since the point at angle 0 of S is not inside S', no point of S' is inside S.

Now consider a point Y of a rim edge J of the smallest optimal spiral T. Let

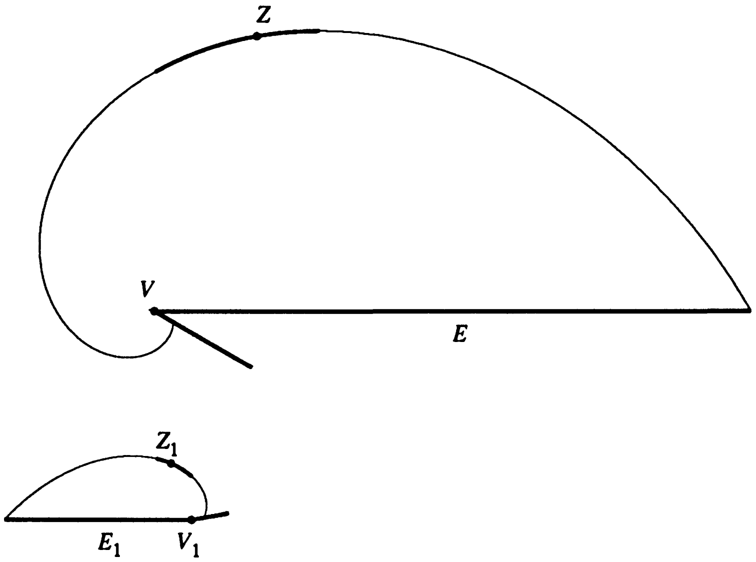

Figure 5.7: The interior side of E relative to V is opposite the interior side of E_1 with respect to V_1.

J_1 and J_2 be the vertices of J with $\phi_{J_1} < \phi_{J_2}$. Note $\rho_Y \leq \rho_{J_1}$. Clearly, the closest point Z of S to Y has $\rho_{J_2} \leq \rho_Z \leq \rho_{J_1}$. Recalling that ρ_{J_1} and ρ_{J_2} are within a constant factor of one another, we see that the distance from Y to a smallest spiral for another vertex is at least a constant times $\rho_Y A$. Its distance to the corresponding smallest optimal spiral is at least this distance. Recalling Lemma 31 completes the proof of the second item of the theorem.

This completes the proof of the Theorem 24. ∎

5.1.1 Constructing a complete wedge around a vertex

We now describe how we construct a wedge around a P vertex V that includes both of the P edges containing V.

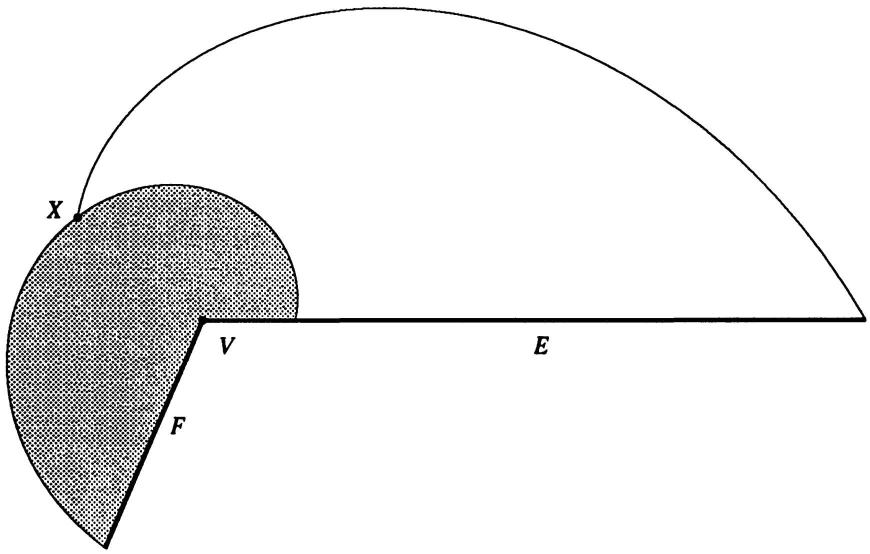

Figure 5.8: Two smallest spirals for the same vertex V meet at X.

Welding two spirals at a vertex

There a smallest spiral for vertex V and each of the two edges containing it. Our first task is to "weld" these two smallest spirals.

Let S and T be the two smallest spirals for E and F at V, where $L(E) \geq L(F)$. Let X be the point common to the boundaries of both S and T as in Figure 5.8. It may be that both S and T have shrinking factor 1, in which case they are coincident. In this case we simply ignore T and consider S to be the welded spiral.

Hence we may assume that X is unique. If X lies on E we discard S, and if X lies on F we discard T, and consider the remaining spiral to be the welded spiral.

Otherwise, we now form the smallest optimal spiral S' for S subject to the last spoke being $\overline{VX}$. That is, in Theorem 22 we take $\omega = \phi_X$ instead of $\omega = \phi_F$. Here ϕ_F denotes the angle between F and E at V. Similarly define T'. See Figure 5.9. From Theorem 22 we see that S' has min angle at least a constant factor times what is optimal for P, unless ϕ_X is less than half the angle θ of the smallest optimal spiral when we take $\omega = \phi_F$. We now show how to fix S' in this case. Should this case arise for T', it is fixed analogously.

First note from Theorem 22 that this case can only happen when $\phi_X < \phi_F/2$. This implies that S' consists of just one triangle. This also implies that this case cannot arise simultaneously for S' and T'. Hence the min angle of T' is within a constant factor of what is optimal for P. If the min angle of S' is within a constant factor of that for T', there is nothing that needs to be done. Otherwise, let Y be the second to last spoke vertex of T', and Z the second to last spoke vertex of S'. We move X to the point X' along the line $\overline{ZX}$ such that $\overline{VX'}$ bisects $\angle YVZ$. This transformation keeps T' and S' inside T and S. The angle of S' at Z is unchanged.

The transformation at most halves the angle of T' at V. Because of the constant lower bound on a for both S' and T', the transformation decreases a for T' by at most a constant factor. Also because of the constant upper and lower bounds on a for both S' and T', the min angle of S' is at least a constant factor times the min angle of T'. Hence the min angles of S' and T' are both within a constant factor of the optimal achievable for P, and we have fixed the spirals for this case.

We always make an additional transformation to guarantee a large angle between the rim edges meeting at X, in order to make Theorem 26 true. We replace the triangles of S' and T' that contain X by the three triangles of Figure 5.10. That is, we replace $\overline{VX}$ by two spokes, $\overline{VX'}$ and $\overline{VX''}$, with a rim edge $\overline{X'X''}$ between them. See Figure 5.10. The length of the new spokes is the same as that of the old spoke $\overline{VX}$. For each of the new triangles, the angle at V is one third of the angle at V between the second to last spokes of S' and T'. This makes at

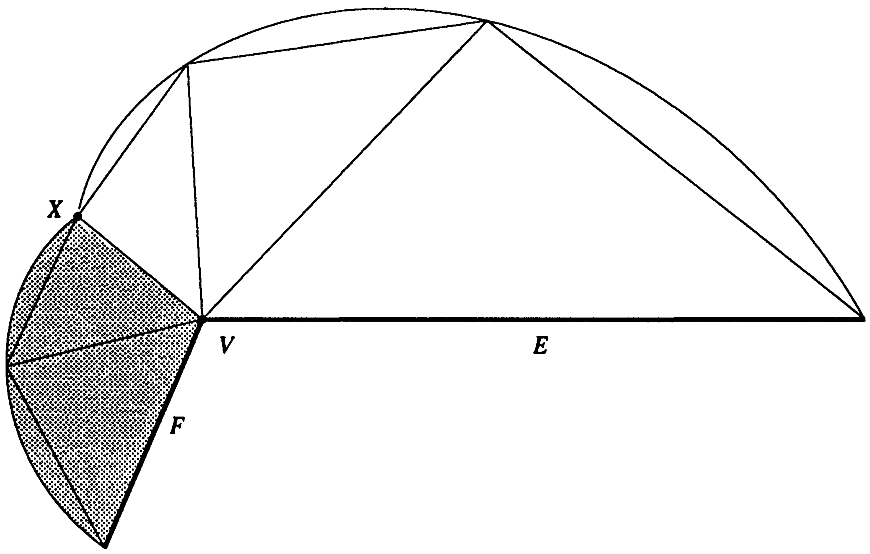

Figure 5.9: Welding the two smallest optimal spirals for V.

most a constant factor decrease in the min angle of S' and T'. This also insures that Theorem 28 below is true.

Welding two spirals for an edge

The wedge at V is complete except for the triangles containing P edges E and F. We consider E, the procedure for F is analogous. We rotate the first spoke N of S' (the one contained in E), uniformly distributing the angle change over all S', until the angle N makes with E is the angle of S'. This at most halves the min angle of S'. The edge E also contains another P vertex U. We complete the same steps in the construction of the complete wedge for U as we have for V before continuing. Call the first spoke in that construction M.

Recall that $L(N) = L(M) = L(E)/4$. Without loss of generality, we may assume that $\angle NE \leq \angle ME$. Let X be the other vertex of M, and Y the other vertex of N as in Figure 5.11. We introduce the edges $\overline{YU}$ and $\overline{YX}$ into our construction. See Figure 5.11.

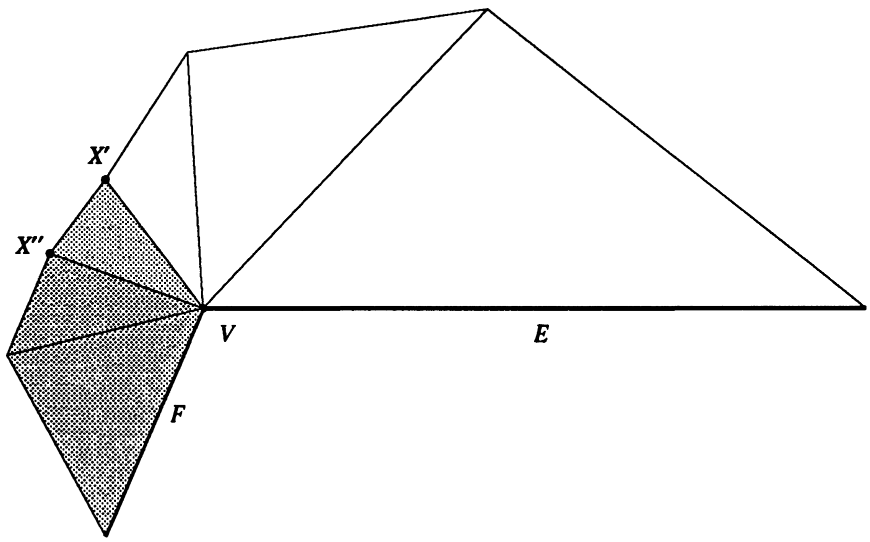

Figure 5.10: Replacing $\overline{VX}$ by the two spokes $\overline{VX'}$ and $\overline{VX''}$. This is the final transformation for welding the two smallest optimal spirals at V.

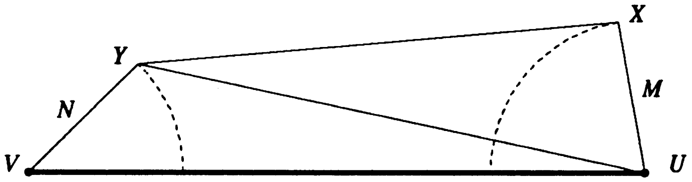

Figure 5.11: Welding the two wedges for one P edge and its two vertices V and U.

It remains to analyze the min angle of the two new triangles we introduced. For $\triangle YUV$, we note that the ratio of the length of E to the length of N is 4, and its angle at V is the same as that of (the transformed) S'. Hence, its min angle is within at least a constant factor times optimal.

For $\triangle YXU$, from the law of sines,

$$\frac{L(YU)}{\sin \angle YVU} = \frac{L(N)}{\sin \angle YUV}.$$

Since $L(YU) \geq \frac{3}{4}L(E)$, we have that $\sin \angle YUV \leq \frac{1}{3} \sin \angle YVU$. Since by assumption $\angle XUV \geq \angle YVU$, we have $\angle XUY = \angle XUV - \angle YUV \geq \angle YVU - \frac{1}{3}\angle YVU =$

$\frac{2}{3}\angle YVU$. So the min angle of $\triangle YXU$ is at least a constant times that of $\triangle YVU$, and hence also at least a constant times that of the optimal triangulation.

This completes the construction of a complete wedge about any vertex of P. We form the complete wedge about all vertices of P. This triangulates a portion of P; the region left untriangulated we call Q, and the triangulated region we call Q^c. The boundary of Q is composed of edges of wedges opposite the vertex of the wedge. Combining the min angle theorems of this section yields the following.

Theorem 25 *The min angle of Q^c is at least a constant factor times the min angle of any other triangulation of P.*

For the following sections, we also have a bound on the smallest interior angle of Q.

Theorem 26 *The interior angle at any vertex of Q is at least* 0.39π.

Proof. We must analyze three types of vertices, depending on what types of edges contain it. The interior angle is at least π for a vertex contained in two rim edges of a spiral. This follows from the fact that spoke lengths follow a geometric series and from the upper bound on the angle between successive spokes given by Theorem 16.

The interior angle is at least $\pi/2$ for a vertex contained in one rim edge of a spiral and one edge $\overline{X'X''}$ as in Section 5.1.1. This is the limiting value as $L(\overline{X'X''})$ approaches zero and the angle at V between successive spokes approaches zero. See Figure 5.12 for an example of this construction at a small interior angle of P.

The interior angle is at least 0.39π for a vertex contained in an edge between two spirals for the same P edge but different P vertices, as in Section 5.1.1. The

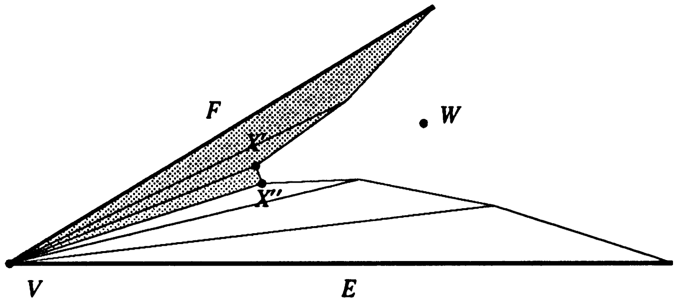

Figure 5.12: The interior angle of Q where we welded two smallest spirals for one P vertex and two different P edges is at least $\pi/2$, regardless of the interior angle at V.

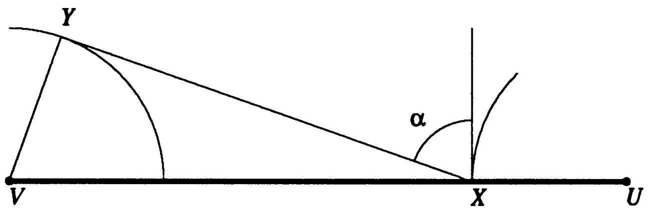

Figure 5.13: The smallest interior angle α of Q may occur where we welded the two smallest spirals for one P edge.

limiting case is when the angle for the spiral at one of the P vertices approaches zero, as illustrated in Figure 5.13. The smallest interior angle α is thus at least $\pi/2 - \arcsin(1/3) \geq 0.39\pi$. ∎

Theorem 27 *The interior angle at a vertex of Q is at most $3\pi/2$.*

Proof. Consider the interior angle β_Q between successive rim edges for the same smallest spiral. These edges lie in similar triangles of a (transformed) smallest spiral. So the interior angle of Q at that vertex is equal to 2π minus the sum of the two angles of the similar triangles that do not occur at V. Hence β_Q is equal to π plus the angle at V of a (transformed) smallest optimal spiral. But the angle at V of any smallest optimal spiral is at most $\pi/2$. Hence $\beta_Q \leq 3\pi/2$.

It is clear that the interior angle at rim vertices where we welded two smallest spirals for the same P vertex is smaller than the case just analyzed. Consider the interior angle ω_Q at the first rim vertex, where we welded the two smallest spirals for a P edge. The limiting case is for the smallest optimal spiral S' left with the larger angle at V. Consider Figure 5.13, which shows the limiting case as $\angle XUV$ approaches 0. For S' we have $a \leq \mathrm{e}^{-1}$, but $L(\overline{VY}) = L(E)/4$ and $L(\overline{VX}) \geq L(E)/2$. Hence $\angle VYX$ is at least the angle of a smallest optimal spiral vertex at a triangle of S'. From the analysis of the last paragraph, we see that ω_Q is at most 2π minus the angle of S' at V, and hence $\omega_Q \leq 3\pi/2$. ∎

We now turn to establishing the last theorem of this section which will find application in the next section.

Delta. For a closed face E of a polygon, we define $\delta(E)$ to be the minimum over all closed faces $F \subseteq E$ of the geodesic distance from F to the closest face W interfering with F.

Theorem 28 *For any edge E of ∂Q, $\delta(E)/L(E)$ is at least a constant times $A(Q^c)$, where $A(Q^c)$ denotes the minimum angle among all triangles we have introduced between P and Q.*

Proof. From Theorem 24, the ratio of the length of an edge of ∂Q to its distance to a spiral for another P edge is at most $A(Q^c)$. The ratio of $L(E)$ to its distance to a vertex or edge of the same spiral is at most $1/a$, where a is the ratio between consecutive spokes of the spiral. This is bounded by the constant e. A constant bound also holds for the edge introduced between the first spokes of two spirals for the same P edge in Section 5.1.1.

We now consider the theorem for an edge E that is the last rim edge of a spiral at V whose closest interfering point is a vertex of a different smallest spiral at V. That is, consider when we matched the two smallest spirals for the two edges containing a P vertex in Section 5.1.1. Let E and F be the rim edges in the last triangles of the spirals. Let $E = \overline{XY}$ and $F = \overline{XZ}$. Let ω_E be the angle at V opposite E, and ω_F the angle at V opposite F.

Without loss of generality, we seek to bound $\mathrm{dist}(Z,E)/L(E)$. Note $\mathrm{dist}(Z,E) \geq L(\overline{VX}) \sin y\omega_F$. Since by construction $\omega_E = \omega_F$, $\mathrm{dist}(Z,E) \geq L(\overline{VX}) \sin(\omega_E)$. Since the first spoke of a spiral is the longest, and $\omega_E \leq \pi/2$, we have that $L(\overline{VY}) \geq L(E)$. Hence $\mathrm{dist}(Z,E)/L(E)$ is at least the min angle of the triangle containing E. ∎

5.2 Shrinking the polygon Q to R

We have already triangulated Q^c, it remains to triangulate Q. In order to triangulate Q, we first shrink it. That is, we form another polygon R, contained in Q, that is approximately what one would get by moving all the edges of Q towards its interior.

For every edge E of Q we introduce a *shrunken* edge F parallel to E in the interior of Q, at distance $\delta(E)/4$ from E. The placement of the vertices of F depends on the interior angles at the vertices of E. For every vertex V of Q, we introduce shrunken vertices according to the following two cases.

Case 1. If $\angle V \leq 3\pi/4$, we introduce a single shrunken vertex. This vertex is

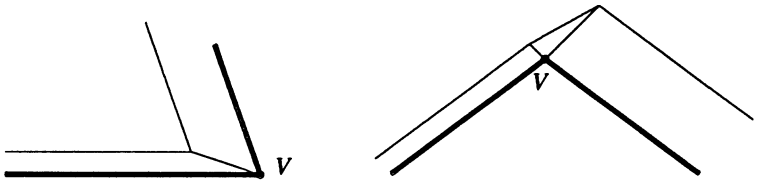

Figure 5.14: Shrinking Q to R near a vertex V with interior angle $\leq 3\pi/4$ (left) and interior angle $\geq 3\pi/4$ (right).

placed at the intersection of the shrunken edges for the Q edges containing V.

Case 2. If $\angle V \geq 3\pi/4$, then we introduce two shrunken vertices and a shrunken edge between them. Consider a ray from V at $\pi/4$ radians clockwise from the bisector of the interior angle at V. We introduce a shrunken vertex at the intersection of this ray and the shrunken edge for the clockwise Q edge at V. Similarly, we introduce a shrunken vertex at the intersection of the ray $\pi/4$ radians counter-clockwise from the interior angle bisector at V and the shrunken edge for the counter-clockwise Q edge at V. We introduce a shrunken edge between the two shrunken vertices we just introduced. See Figure 5.14.

Delta for R and Q. We extend the definition of the delta function to R and Q. That is, for a face F of R we define $\delta(F)$ in terms of the geodesic distance in R between F and its subfaces and interfering faces of R. The following theorems relating δ and the lengths of edges are important for later sections.

Theorem 29 *Let F be an edge shrunken for Q edge E. Let V be a vertex of E, and W the corresponding vertex of F. Then the angle ϕ at V between $\overline{VW}$ and E is at least a constant times $\delta(E)/L(E)$.*

Proof. If we are in Case 2 of the above construction, that is the angle between E and the other Q edge containing V is at least $3\pi/4$, then the angle ϕ is at least

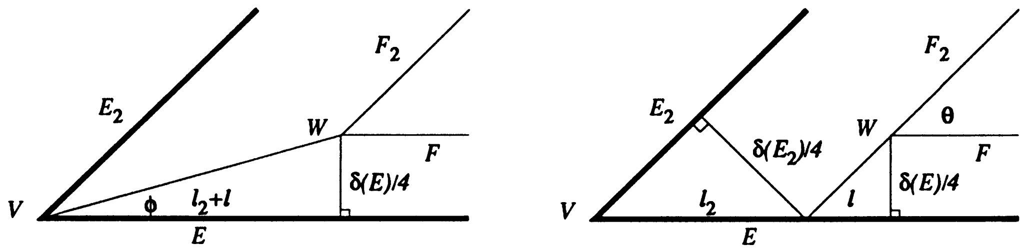

Figure 5.15: Bounding ϕ, the angle between $\overline{VW}$ and E.

$\pi/8$.

Otherwise, we are in Case 1 of the above construction. Let E_2 be the other Q edge containing V. Let $\theta = \angle EE_2$. Define the distances l and l_2 as in Figure 5.15. We have that $l = \frac{\delta(E)}{4\tan\theta}$ and $l_2 = \frac{\delta(E_2)}{4\sin\theta}$. Hence $l + l_2 \geq k_1\frac{\delta(E)+\delta(E_2)}{\theta}$ for some constant k_1. Since $\tan\phi = \frac{\delta(E)}{4(l+l_2)}$, we have the bound that

$$\phi \geq k_2\frac{\theta\delta(E)}{\delta(E)+\delta(E_2)}.$$

If $\delta(E_2) \leq \delta(E)$, then $\phi \geq k_2\theta/2$. Let us ignore the fact that θ is bounded by a constant for polygon Q, and prove the theorem for any polygon. In any polygon, we may get a bound for $\delta(E)$ in terms of θ as follows. Suppose $L(E) \geq L(E_2)$. The vertex of E_2 opposite V is disjoint from E, and its distance to E is $L(E_2)\sin\theta$. Similarly, if $L(E) \leq L(E_2)$, then note that the vertex of E opposite V is disjoint from E, and its distance to E_2 is $L(E)\sin\theta$. Since this vertex is a subface of E, this distance bounds $\delta(E)$. Hence $\delta(E) \leq \min(L(E), L(E_2))\sin\theta$. That is, $\theta \geq \delta(E)/L(E)$, and $\phi \geq k_2\frac{\delta(E)}{2L(E)}$.

Otherwise $\delta(E) \leq \delta(E_2)$, so that $\phi \geq k_2\frac{\theta\delta(E)}{2\delta(E_2)}$. But from the last paragraph, we may also claim that $\delta(E_2) \leq \min(L(E), L(E_2))\sin\theta$. Substituting this expression for $\delta(E_2)$ in the equation bounding ϕ allows us to conclude the theorem ∎

Theorem 30 *For any edge F of R arising from shrinking edge E of Q, $\delta(F) \geq \delta(E)/2$. Also, for any vertex or edge F of R arising from shrinking vertex V of Q, $\delta(F) \geq \delta(V)/2$.*

Proof. We first show that the shrunken edges are close to the corresponding Q edges in terms of δ. This leads to the fact that shrunken edges for two disjoint Q edges are far from one another in terms of δ. We next note that the shrunken edges for Q vertices have a large length in terms of δ, and are close to the corresponding Q vertex. Finally we note that shrunken vertices for Q vertices with small interior angle are far from other faces. This allows us to conclude that that shrunken edges have bounded δ even when considering all faces of R.

Recall that a shrunken edge F is introduced parallel to a Q edge E at distance $\delta(E)/4$. The angle between a shrunken edge F and the ray from one of its endpoints through the corresponding Q vertex V is at least $\pi/4$. Hence any point on a shrunken edge F for Q edge E is at distance at most $\frac{\sqrt{2}}{4}\delta(E)$ from the closest point of E. For an edge F shrunken for vertex V, we note that since $\delta(V)$ is at least δ for any edge containing V, the distance from a point of F to V is at most $\frac{\sqrt{2}}{4}\delta(V)$. The triangle inequality allows us to conclude that two shrunken edges E and F for disjoint Q edges are at least distance $(1-\frac{1}{\sqrt{2}})\delta(E)$ from each other.

For any shrunken vertex W arising from a Q vertex V with interior angle greater than $3\pi/4$, the distance from W to V is at most $\frac{\sqrt{2}}{4}\delta(V)$. Hence the distance from any shrunken edge for a Q edge E to a shrunken edge for a disjoint Q vertex is at least $(1-\frac{1}{\sqrt{2}})\delta(E)$. Also, since the shrunken edge for a vertex V with large interior angle subtends an angle of $\pi/2$ at V, its length is at least the distance from V to the further shrunken edge for a Q edge containing V. Consider a Q vertex V with edges E_1 and E_2 meeting at an interior angle larger than $3\pi/4$. Let F_1 and F_2 denote the shrunken edges arising from E_1 and E_2. Then the distance from F_1 to F_2 is at least $\max(\delta(E_1), \delta(E_2))$.

Using the above arguments, we can also show that for a shrunken vertex W arising from a Q vertex V with edges E_1 and E_2 meeting at an interior angle greater than $3\pi/4$, $\delta(W) \geq \min(\delta(E_1), \delta(E_2))$.

In order to bound δ for a face, it remains to consider the distance from that face to a Q vertex V with edges E_1 and E_2 meeting at interior angles less than $3\pi/4$. Let F_1 and F_2 denote the shrunken edges for E_1 and E_2 and W the shrunken vertex for V. Note that the shrunken edges meet at the same angle as the Q edges. Since this angle is less than π, to determine δ for a shrunken edge that is disjoint from F_1 and F_2 it is not necessary to consider its distance to W. However, it may be the case that the closest disjoint face to F_1 is the other vertex of F_2. Hence we need to consider the length of F_1 and F_2. See Figure 5.16. Assume that the closest face disjoint to E_1 is X, the vertex of E_2 that is not V, and similarly assume that the closest face disjoint from E_2 is Y, the vertex of E_1 that is not V. Assume that E_1 and E_2 are of equal length, and the angle at V is acute. Introduce F_2 first. We see that the line through F_2 cuts off from E_1 an edge of length $3/4$ times that of E_1. If we now introduce F_1, we see that F_1 is of length $3/4$ that of E_1. If a similar situation occurs at the vertex of E_1 that is not V, we see that F_1 can be as short as $1/2$ the length of E_1. Similarly for F_2. If there is some other face of Q that is the closest face disjoint from E_1 rather than X, then this only makes F_1 longer. If the angle at V is not acute, then $1 - \frac{1}{\sqrt{2}}$ times the length of E_1 is a weak lower bound on the length of F_1. We omit many tedious cases and symmetry arguments,

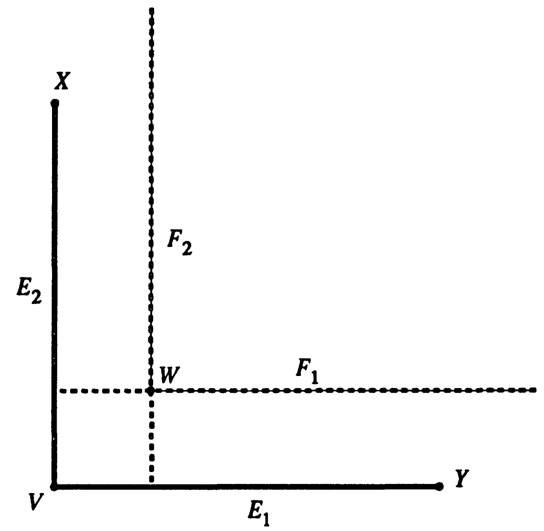

Figure 5.16: Bounding the lengths of shrunken edges F_1 and F_2 meeting at an angle $\leq 3\pi/4$.

and conclude that a lower bound on δ for a shrunken vertex W for Q vertex V is $\delta(V)(1 - \frac{1}{\sqrt{2}})$. ∎

5.3 Triangulation of R

We now triangulate the shrunken polygon R, adding Steiner points to its boundary. We use the two dimensional analog of Chapter 4. Alternatives are discussed in Section 5.5.1.

A feature of this triangulation algorithm is that the size of a quadtree box containing an input edge is at least a constant c_1 times δ for that edge. Also, the edge of a triangle contained in an input edge is at least a constant c_2 times the size of the quadtree box containing the triangle. Hence a shrunken edge F is subdivided by the triangulation of R into edges of length at least $c_1 c_2 \delta(E)$.

Balanced. We also note that the quadtree is balanced: Any quadtree box containing an R edge is at least $1/2$ as large as any adjacent box. Any quadtree box containing an R vertex is at least a constant c_3 times as large as any adjacent box. Hence, any two adjacent triangulation edges on the boundary of R differ in size by at most a factor of $1/(c_3c_2)$.

The constants c_1, c_2 and c_3 all depend on the smallest interior angle of R. However, the smallest interior angle of R is equal to the smallest interior angle of Q, and Theorem 26 bounds this by the constant 0.39π. Hence c_1, c_2 and c_3 are all bounded below by true constants, namely k_1, k_2 and k_3.

In the upcoming section we will refer to the size of the quadtree box containing a given triangulation edge G as $q(G)$.

5.4 Matching R to Q

We now show how to triangulate the portion of Q outside of R. We do this by matching an edge or vertex of Q with its shrunken edges to form a trapezoid or triangle. Each such figure is triangulated with min angle at least a constant times $\delta(F)/L(F)$ for some R edge F. We add Steiner points to a figure's interior and to its sides, but not to any edge of R or Q. Adjacent figures share sides, so that care must be taken when adding Steiner points to the sides.

Match faces. We introduce an edge from each Q vertex to the corresponding vertices of R. We call these edges *match edges.* This partitions $Q \setminus R$ into *match trapezoids* and *match triangles.* There is one trapezoid for each edge of Q, and one triangle for each Q vertex with interior angle greater than $3\pi/4$. For each trapezoid, one of the parallel edges is a Q edge and the opposite R edge is divided into several smaller edges by the quadtree triangulation of R. Similarly for the

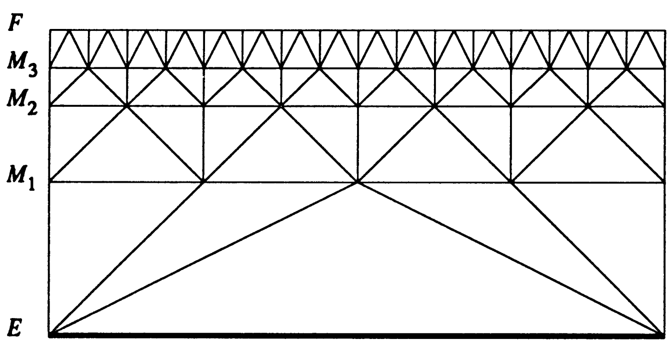

Figure 5.17: Triangulating a special case trapezoid with layers of boxes.

triangles; one vertex is a Q vertex, and the opposite R edge is subdivided into several smaller edges.

5.4.1 Triangulating a match trapezoid

Consider a particular trapezoid T, with Q edge E and R edge F. Let d be the distance from F to E, and recall that d is bounded above and below by constant factors times $\delta(E)$. Consider a set of lines parallel to E and F, the first one halfway between E and F, the second halfway between the previous line and F, etc. That is, consider the set of lines $M_1, M_2, \ldots M_n$ parallel to F, at distance $d/2, d/4, \ldots, d/2^n$ from F, where $n = \lfloor \log(L(E)/d) \rfloor$.

We triangulate T by first introducing squares with two sides parallel to F, where those sides lie on these parallel lines. We then triangulate each square independently, without adding vertices on the boundary of the square. The min angle of this triangulation will be at least a constant factor time A, the optimal min angle of any triangulation of Q.

A special case

Before describing the general algorithm and proving this min angle bound, we consider an easier special case. First recall from Section 5.3 that in general $q(X)$ is at least a constant k_1 times d for any point $X \in F$. So suppose that $q(X) = k_1 d$ for all $X \in F$. Then we divide M_n into edges of length $2k_1 d$, and triangulate as in Figure 5.17. The distance between M_n and F is $d/2^n \geq d\delta(E)/L(E)$. The min angle of a triangle is within a constant factor of the ratio of the radii of the smallest containing sphere to the inscribed sphere. The containing sphere for any triangle has radius within a constant factor of $k_1 d$. The inscribed sphere has radius within a constant factor of $d\delta(E)/L(E)$. Hence the min angle of the triangles we introduce is within a constant factor of $\delta(E)/(k_1 L(E))$. Note that this holds for the triangles intersecting a match edge, since the angle between a match edge and F is at least A, the optimal min angle of any triangulation of R. We have drawn Figure 5.17 as if this angle was $\pi/2$. We may repeat this process for the lines $M_2, M_3, \ldots M_n$. Each time we triangulate with min angle at least a constant times $\delta(E)/(k_1 L(E))$. We will subdivide M_1 into a constant number of edges (linearly dependent on k_1), each of length at least a constant times the length of E. Hence we may triangulate between M_1 and E by introducing an edge from each vertex of M_1 to the closer vertex of F (and one more edge for the remaining trapezoid). The resulting min angle is within a constant factor of $\delta(E)/(k_1 L(E))$. We have drawn Figure 5.17 in the case when M_1 is divided into four segments.

Drawing squares

We now show how to modify the special case to handle the general case when F is not subdivided uniformly by the quadtree. Recall that the quadtree is balanced as described in Section 5.3.

We define the *level* of an edge to be defined in terms of the size of the quadtree box containing it, but first we need to make a special definition of size for boxes containing a vertex of F. The quadtree is not necessarily balanced by a factor of two at boxes containing a vertex of F. We treat a quadtree box b_V containing a vertex V of F as if it had the same size as the quadtree box b_F, where b_F is the unique box that shares a point of F with b_V. Since the ratio of the true sizes of these quadtree boxes is bounded by a constant k_3 from Section 5.3, this will affect the min angle of our triangulation by at most a constant factor.

Level. We now define the level of a triangulation edge so that the smallest level edges have level 1, the second smallest level 2, the third smallest level 3, etc. That is, the edges arising from the ith smallest quadtree boxes have level i. Formally, the level of an edge G in F is $\log_2(sq(G))$, where $s = 2/\min_{X \in F} q(X)$. We order the edges from the leftmost vertex of F to the rightmost vertex of F, and place them in a list S.

The algorithm for drawing the squares is iterative, having steps $1, 2, \ldots n$, where as before $n = \lfloor \log(L(E)/d) \rfloor$, and is the number of parallel guide lines M_i. At the i step we draw squares around edges of level at most i. For each square, its edge parallel to and furthest from F will lie on line M_{n-i+1}. This edge G is then considered as an edge of level $i+1$ for consideration by the next iteration step. We say that the square is also of level $i+1$. The edges we drew the square around are consecutive edges in S. We remove that subsequence from S and insert G in its place.

We now describe exactly which edges we draw squares around in iteration step

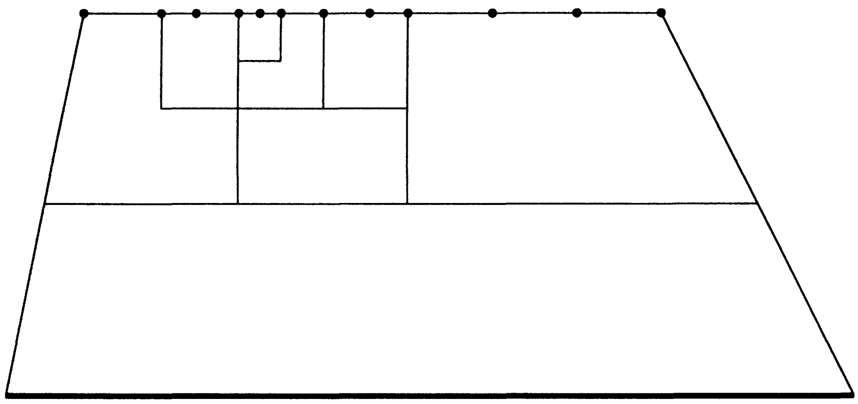

Figure 5.18: Drawing squares as a preliminary step to triangulating a general match trapezoid.

i. Identify each maximal length subsequence C of consecutive edges in S, where each edge in C has level at most i. Denote the cardinality of C by m. If $m > 1$, then for $j = 0, 1, \ldots \lfloor m/2 \rfloor$, we draw a square around the $2j$th and $2j+1$th edge of C. That is, we draw a square around each pair of edges in the subsequence starting with the first two, and skipping the last if the subsequence is of odd cardinality. The exception to this rule is if $m > 1$ and the last edge of C is the last edge of S. In that case, the last square we draw is drawn around three edges, namely edges $m-2, m-1$ and m of C. See Figure 5.18.

This concludes the definition of the square drawing algorithm. We now consider the shape of the squares we have drawn.

Elongation. We define the *elongation* of a triangulation edge to be 1. We define the *elongation* of a square b (and its corresponding edge placed in S) to be

$$e(b) = \frac{\sum_{G \in b} 2^{l(G)}}{2^{l(b)}},$$

where G ranges over all triangulation edges contained in b, and $l(x)$ denotes the level of x. If all the subsequences encountered by the algorithm were of even length, then each square would have elongation exactly 1.

It is easy to show that the elongation of b may also be expressed as

$$e(b) = \sum_{H \in b} e(H) 2^{l(H)-l(b)},$$

where H ranges over the edges that b is immediatly drawn around, both triangulation and square edges. It is this formula that will be used in the lemmas below.

We seek constant upper and lower bounds on the elongation of a square. We first derive the upper bounds.

Lemma 38 *Any edge not containing the rightmost vertex of F has elongation at most* 1.

Proof. We prove this by induction. Before the first iteration step begins, the lemma is true by definition of elongation for a triangulation edge. Consider a square b drawn in the ith iteration step. Now b is drawn around two edges. Moreover, neither of these edges contains the rightmost vertex of F, and hence by induction have elongation at most 1. Thus the elongation of b is at most $1/2 + 1/2 = 1$. ∎

Lemma 39 *Any edge containing the rightmost edge of F has elongation at most* 2.

Proof. Again we prove this by induction. Before the first iteration step begins, the edge containing the rightmost vertex of F has elongation 1 by definition. Before the ith iteration step, there is exactly one edge in S that contains the rightmost edge of F. If we do not draw a square around this edge in the ith iteration step, then after this step this edge is a triangulation edge of level at least i and has elongation 1. If we do draw a square around this edge, then it is drawn around at most three edges, each of which can have level at most i. Thus from the previous lemma the square has elongation at most $1/2 + 1/2 + 2/2 = 2$. ∎

We now derive the lower bounds. First we need the following two lemmas about the algorithm.

Lemma 40 *If an edge of level i is not drawn around in iteration step i, then it is a triangulation edge.*

Proof. If an edge G of level i is not drawn around in iteration step i, then it is adjacent in S to an edge H of level $i+1$. Since $i+1$ is greater than the current iteration step, i, H is a triangulation edge. By the balance condition, all triangulation edges adjacent to an edge of level $i+1$ have level at least i. A triangulation edge adjacent to H is contained in G, hence G contains a triangulation edge of level i. But by assumption G is a level i edge, and so it cannot contain any other triangulation edges, and is a single triangulation edge. ∎

Lemma 41 *An edge of level i is drawn around on iteration step i or $i+1$.*

Proof. Suppose an edge G is not drawn around on iteration step i. We have two cases.

The first case is if G does not contain the rightmost vertex of F. Then G's successor H in S is of level greater than i. By the previous lemma G is a single triangulation edge, and H is a triangulation edge of level exactly $i+1$. Hence G will be in a maximal subsequence in iteration step $i+1$, but will not be the leftmost edge in the subsequence. So in iteration step $i+1$ the algorithm will draw a square around G.

The second case is if G does contain the rightmost vertex of F. The only way we would not draw a square around such an edge is if the maximal subsequence containing it has length 1. But then by the same arguments as the last paragraph, it is preceded by a triangulation edge H of level $i+1$. Hence G will be in a maximal subsequence of length at least 2 in iteration step $i+1$. So in iteration step $i+1$ the algorithm will draw a square around G. ∎

We may now prove a lower bound on the elongation of an edge.

Lemma 42 *Any edge has elongation at least* $1/2$.

Proof. We prove this by induction. The lemma is trivially true before the first iteration step. Consider an edge H drawn around by a square at iteration step i. From the previous lemma, H may be of level i or $i-1$.

Suppose H is of level i. By induction H has elongation at least $1/2$, so $e(H)2^{l(H)-l(b)} \geq 1/4$.

Suppose H is of level $i-1$. But by Lemma 40, such an edge H is a triangulation edge and must have elongation 1. Hence $e(H)2^{l(H)-l(b)} = 1/4$. Every square is drawn around at least two edges, so the square will have elongation at least $1/4 + 1/4 = 1/2$. ∎

We now describe how the algorithm is modified when the match edges are not perpendicular to E and F: If one vertex of the square lies on a match edge, we use the match edge for a side of a "square."

After we describe how to triangulate the squares, we will wish to prove that the min angle of the triangles is at least a constant factor times A, the optimal min angle of any triangulation of Q. To that end, we have the following two lemmas concerning the angle that a match edge may make with an R edge F that it intersects.

Lemma 43 *The angle between a match edge and F is at least $\pi/4 + kA$, where k is a constant and A is the optimal min angle of any triangulation of Q.*

Proof. Suppose we are in Case 2 of the construction of R. Theorem 27 gives an upper bound of $3\pi/2$ on the interior angle at a vertex of E. From Case 2 of the construction of R in Section 5.2, the angle between a match edge and E is at most $\frac{3\pi}{4} - \frac{\pi}{4} = \frac{\pi}{2}$. Hence the angle that F makes with a match edge is at least the supplement of this, or $\frac{\pi}{2}$.

If we are in Case 1 of the construction of R the situation is not as nice. Here the angle at E may be as great as $3\pi/4 - k_1\delta(E')/\delta(E)$ for some constant k_1, where E' is the other R edge intersecting E at the vertex under consideration. The ratio of δ's is bounded by a constant times the optimal min angle, or k_2A. The angle that F makes with a match edge is at least the supplement of the above, or $\pi/4 + k_1k_2A$. ∎

Lemma 44 *The angle between a match edge and F is at most $\pi - kA$.*

Proof. If we are in Case 2 of the construction of R, we have that the angle between E and the match edge is at least $3\pi/8 - \pi/4 = \pi/8$. Hence the angle between F and that match edge is at most the supplement of this, or $7\pi/8$.

If we are in Case 1 of the construction of R the situation again is not as nice. Here the angle at E may be as small as $k_1\delta(E)/\delta(E')$ for some constant k, where E' is the other R edge intersecting E at the vertex under consideration. The ratio of δ's is bounded by the optimal min angle times a constant, that is k_2A. The

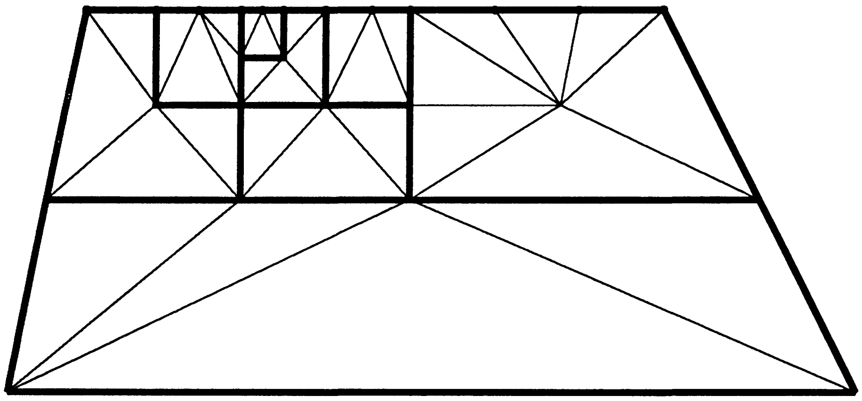

Figure 5.19: Triangulating the squares with the help of a center vertex.

angle that F makes with a match edge is at least the supplement of the above, or $\pi - k_1 k_2 A$. ∎

Triangulation of squares

We now describe how to triangulate the squares. We triangulate between E and M_1 as in the special case described above.

Center vertex. We then triangulate each square by taking the convex hull of the *center vertex* of the square and each edge of the boundary of the square as in Figure 5.19. For squares that were drawn around two edges in S we may chose one of the vertices of these edges for the center vertex. In particular, if the rightmost edge is of level i, then we use its leftmost vertex as the center vertex. Otherwise, the leftmost edge is of level i and we use its rightmost vertex as the center vertex.

For the squares drawn around three edges in S, it is easy to prove that there are six cases depending on whether these three edges are triangulation edges or edges of a drawn square. These squares are triangulated as illustrated in Figure 5.20. The center vertices in squares without any squares inside them is the midpoint

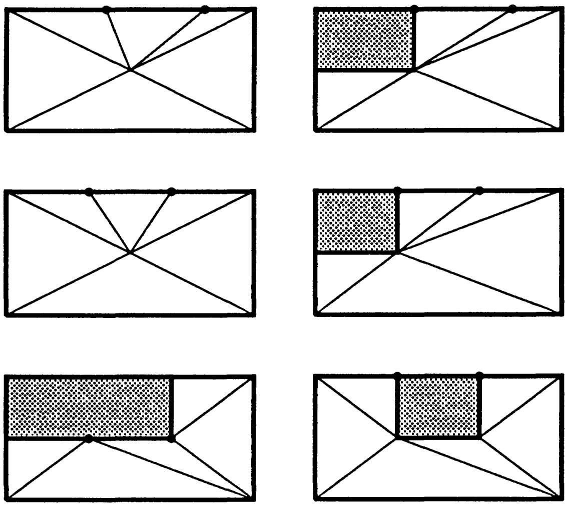

Figure 5.20: How to triangulate a square drawn around three edges.

of a horizontal segment drawn halfway between the bottom and top edges. Note that the leftmost and rightmost vertical edges may have vertices in their interior, arising from smaller adjacent squares. In that case we introduce an edge from that square vertex to the opposite vertex of the triangle that contains it in Figure 5.20.

A special situation can arise for the first square drawn around a triangulation edge containing a vertex of F. For two trapezoids T_1 and T_2 sharing a match edge F, the number of levels n may be different. Moreover, the level of the edges G_1 and G_2 containing a vertex of F may be different for T_1 and T_2. Thus the first square b_1 drawn around G_1 may be smaller than the first square drawn around G_2. That is, the side of b_1 that is contained in F may have vertices in its interior, arising from the vertices of squares drawn in T_2. We triangulate as in Figure 5.21 when b_1 is drawn around only triangulation edges. In general, we triangulate by introducing edges between the match edge vertices and the central vertex. The extension to squares drawn around three edges is analogous. We show in Lemma

Figure 5.21: There may be many $(1/A)$ square vertices in a match edge for constant shaped squares, or at most a constant number of square vertices for wide and short $(1/A)$ squares.

45 that there are strict bounds on the distance between the match vertices and the central vertex, so that the min angle of triangles in b_1 is at least a constant factor times A.

Min angle of triangles in squares

We now set out to prove that our triangulation of the squares is within a constant factor of what is optimal.

Lemma 45 *Any triangle of a square without a match edge for a side and not intersecting E has min angle at least a constant times $\delta(E)/L(E)$.*

Proof. Consider the triangles that contain an edge G parallel to F. Since elongation is between $1/2$ and 1 for G, we have $k_1 L(E)2^{l(G)-n-1} \leq L(G) \leq k_1 L(E)2^{l(G)-n}$. The vertex V opposite G makes altitude at least $d2^{l(G)-n-3}$, and no more than $d2^{l(G)-n}$. These bounds follow from the bounds on elongation. The exact altitude value depends on whether the square has level $l(G), l(G)+1$, or $l(G)+2$. In all cases this altitude lies inside the triangle. Hence the angle of the triangle at the vertices of G is at least

$$\arctan \frac{d2^{l(G)-n-3}}{k_1 L(E)2^{l(G)-n}} = \arctan \frac{d}{8k_1 L(E)},$$

which is within a constant factor of $\delta(E)/L(E)$. Also the angle of the triangle at V is at least

$$\arctan\frac{k_1 L(E) 2^{l(G)-n-1}}{d 2^{l(G)-n}} = \arctan\frac{k_1 L(E)}{2\delta(E)},$$

which is bounded below by constant.

Consider a triangle that contains an edge G perpendicular to F. Let b be the square containing triangle. The length of G is within constant factors of $d2^{l(b)-n-2}$. The altitude at the vertex opposite G is within a constant factor of $k_1 L(E) 2^{l(b)-n-1}$, and hence the triangle has min angle at least $\arctan(\delta(E)/L(E))$, which completes the proof. ∎

Lemma 46 *Any triangle intersecting E has min angle at least a constant times $\delta(E)/L(E)$.*

Proof. As shown before, M_1 is subdivided into at most a constant number of triangulation edges, where the constant factor depends linearly on k_1. The distance between E and M_1 is $d/2$. But the lower bound on the angle that a match edge makes with E may be written in terms of $\delta(E)/L(E)$. Hence the length of M_1 is at least a constant times the length of E. The constant upper bound on the angle that a match edge may make with E, and the fact that $d < L(E)/2$ insures that $L(M_1) < 2L(E)$. Thus from the construction each triangulation edge in M_1 has length at least $c_1 L(E)$ and at most $c_2 L(E)$. A lower bound on the worst triangle that may be formed is shown in Figure 5.22. Thus any triangle has min angle at least

$$\arctan(\frac{d}{L(E)}) - \arctan(\frac{1}{1+2c_1}\frac{d}{L(E)}).$$

Since $d/L(E) < 1/2$, we have that the above is at least $0.8d/L(E)$. ∎

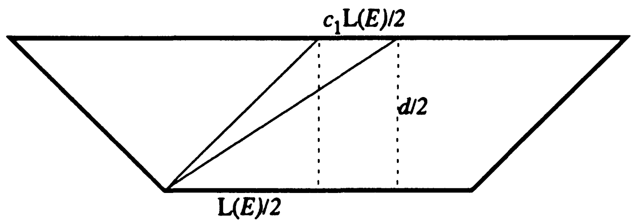

Figure 5.22: The worst case of a triangle containing a vertex of E.

We now consider triangles in squares that are on the boundary between two trapezoids.

Lemma 47 *Any triangle in a square with a match edge for a side has min angle at least a constant times A.*

Proof. On the boundary between two trapezoids, the subdivision of the match edge for one trapezoid may not correspond to the subdivision for the other. However, from Lemma 40 we see that this can only happen inside the first squares drawn around the R triangulation edge containing the R vertex of the match edge.

For triangles of a square that is not drawn around the triangulation edge containing the match edge R vertex, there is no vertex in the side of the square. Hence the arguments of the previous two lemmas are sufficient to show that the triangulation of the square is at least a constant factor times A.

We now consider the squares drawn around the triangulation edge containing the match edge R vertex. The algorithm to triangulate R centers vertices of R in a quadtree box. That is, the two triangulation edges G and G' containing a R vertex V and contained in R edges F and F' are contained in the same quadtree box. Thus the square drawing algorithm considers $q(G)$ and $q(G')$ to be within a constant factor of one another. We have two cases.

The first case is if the lengths of E and E' are within constant factors of one another, then the levels of G and G' are within constant factors of one another. Hence from Lemma 41 the square drawn around G has at most a constant number of vertices in its match edge side arising from squares around F'. Similarly for G'.

The second case is if there are more than a constant number of subdivision vertices in its match edges, we wish to show two things. First we wish to show that the box has constant height to width ratio in terms of Euclidean distance. Second we wish to show that $L(F)/L(F')$ is bounded below by A. Thus we may triangulate as in Figure 5.21 and achieve min angle at least A. Note that in the case that the match edge makes small angle with F and all the edges drawn around are triangulation edges we would instead introduce edges from the match vertices to the lower right vertex of E in Figure 5.21.

We now provide the details for the second case. Edges F and F' arising from shrinking edges of the same smallest spiral have length ratio bounded by a constant. This follows from the constant bound on the ratio between similar triangles of the same smallest spiral. So it remains to consider edges arising from shrinking two edges E and E' that were used to weld two spirals. Let us consider the case that $E = \overline{X'X''}$ was used to weld two spirals for the same vertex. We have already shown in Theorem 28 that $L(E)/L(E') \geq kA(Q^c)$ for some constant k. However, in this case E is at most a constant times the lengths of the two spokes containing its vertices, which are equal in length. Also, these spokes are at most one quarter of $\delta(V)$, where V is the P vertex they contain. Hence $\delta(E)$ is at least a constant factor times $L(E)$. We now consider the case that E' was used to weld two spirals for the same P edge, and E is the first rim edge of a spiral at V. Again from

Theorem 28 we have $L(E)/L(E') \geq kA(Q^c)$ for some constant k. However, the only case that $L(E)/L(E')$ is less than a constant is when the smallest optimal spiral consists of one triangle. That is, when the interior angle at V is small and the shrinking rate of the smallest spiral is not too small in relation to it. But as in the previous welding case, then E is between two spokes of (nearly) equal length, and hence $\delta(E)$ can be no less than a constant times $L(E)$. ∎

5.4.2 Triangulating a match triangle

Triangulating a match triangle T is no more difficult that triangulating a match trapezoid. The key observation is a bound on the lengths of the two match edges of a triangle T. Let V be the Q vertex of T. Let F be the R edge of T opposite V. Let H_1 and H_2 be the R edges intersecting F, and N_1 and N_2 the match edges from Q to F. Assume $L(N_1) \geq L(N_2)$. Note $L(N_1)$ is within a constant factor of $\delta(H_1)$. Similarly for N_2. Since the vertex of H_2 opposite V interferes with H_1, we have that $\delta(H_1) \leq L(H_2)$. Hence $L(N_1) \leq L(H_2)$, and we conclude that $L(N_1)/L(N_2) \leq L(H_2)/\delta(H_2)$.

Thus, we may proceed as in the trapezoid case. We introduce lines parallel to F, spaced as before. Recall that the angle of T at V is $\pi/2$. Hence, by the last paragraph, the angle between these lines and a match edge is bounded by a constant times $L(H_2)/\delta(H_2)$. Hence we may build squares as in the trapezoid case, with min angle at a match edge at least a constant factor times $\min(\delta(H_1)/L(H_1), \delta(H_2)/L(H_2))$, which is at least a constant factor times A.

The algorithm used to triangulate the squares is the same as in the trapezoid case. We triangulate between M_1 and V by introducing an edge from W to each of the square vertices of M_1. The proof of Lemma 46 shows that these triangles have min angle at least a constant times A.

5.5 Conclusions

5.5.1 Why the two dimensional analog of Chapter 4?

The details of the method used in Section 5.3 to triangulate R is not important for the correctness of the other parts of the algorithm. Any two dimensional mesh generation algorithm that generates triangles with guaranteed min angle or aspect ratio may be used, as long as it satisfies the two mentioned features:

1. The length of a triangulation edge G contained in an R edge E must be at least a constant times δ for that edge.

2. The ratio of the lengths of two intersecting triangulation edges must be bounded by a constant.

Any triangulation algorithm with bounded aspect ratio would have to satisfy the second condition, although the constant could be quite large.

The method of Bern, Eppstein and Gilbert[1991] satisfies the second condition, and provides us with a quadtree to guide Section 5.4, if the input has no small interior angles that need to be cut off. This is the case for R because of the lower bound on the smallest interior angle given by Theorem 26. The first condition is not satisfied by their algorithm for general input. Normally, the duplication of boxes in Chapter 4 is needed to guarantee this. The method is akin to Reimann sheets in Seidel[1988]. However, the fact that we shrunk Q to R by moving edges a constant fraction of δ towards the interior of R shows that generating small triangles in one part of R will not necessarily generate small triangles in another part of R close in Euclidean distance but far in geodesic distance. Hence Bern, Eppstein and Gilbert[1991] would satisfy the first condition for R.

The algorithm of Rupert[1992] will also satisfy both conditions. This is startling because it is based on Delaunay triangulations and is not quadtree based. That algorithm recursively divides edges in half, so that the ratio of the lengths of two intersecting triangulation edges contained in an edge of R is at most 2. If Rupert[1992] was used, the algorithm presented in Section 5.4 would have to be modified to handle the fact that there is no quadtree to guide us.

5.5.2 Running time

We now provide an upper bound on the running time of our algorithm, in the process describing how we identify the smallest optimal spirals and compute the delta function.

Let n denote the number of edges of P. We start our algorithm by computing the *medial axis* of P in optimal time $O(n \log n)$. See Duda and Hart [1973] and Kirkpatrick [1979]. This allows us to identify on average a constant number of faces that have points that could define the smallest optimal spiral for a given vertex and edge. Thus, after the medial axis is computed, we may identify the smallest spirals for all edges and vertices of P in total time $O(n)$.

Producing the triangles of the wedges about a vertex (Q^c) takes time linear in the number of triangles produced. The medial axis of Q may be used to efficiently compute the delta function for the faces of Q. This may take time at most $O(m \log m)$, where m is the number of faces of Q. After the medial axis of Q is computed, R takes time $O(m)$ to produce. Note that $m > n$, and m may be as large as on the order of the cardinality of the final triangulation.

Let t be the time to triangulate R, and γ_1 the number of triangulation edges on the boundary of R. Using geometric series, it is easy to prove that triangulating between Q and R produces $O(\gamma_1)$ triangles. Triangulating between Q and R takes time linear in the number of triangles produced. To achieve this bound in the triangulation algorithm of Section 5.4, at each iteration level we use the previous iteration's maximal subsequences to guide our determination of the current iteration's maximal subsequences. Thus we have the following.

Theorem 31 *The running time of the algorithm is $O(\gamma \log \gamma + t)$, where γ is the number of edges of the final triangulation, and t is the time taken to triangulate R.*

As the algorithm is described, using the two dimensional analog of Chapter 4 to triangulate R, t is no worse than $O(\gamma^2 \log \gamma)$. So, t is the limiting factor in the running time of this algorithm. See Section 4.9 for more comments on t.

5.5.3 Cardinality

The cardinality of the triangulation is not optimal. The construction of the wedges around the vertices of P triangulating Q^c produces an optimal number of triangles, up to a constant factor. However, triangulating R with any algorithm designed to maximize the minimum angle, allowing Steiner points arbitrarily in R, is bound to produce too many triangles for certain input.

5.5.4 Open problems

Concerning triangulations with only interior Steiner points, there are several other measures for which algorithms may be useful. As mentioned in the introduction, maximizing the minimum height would be useful for the three dimensional mesh generation algorithm. The algorithm as described bounds sphere ratio in light of Theorem 1 in Chapter 3, but is not within a constant factor of optimal. Other possible measures include minimizing the maximum angle.

Is there an algorithm for triangulating a three dimensional polytope with maximum minimum angle? The difficulty in extending this work to three dimensions is that we must define an optimal spiral at a vertex of possibly high degree. The boundary of a three dimensional spiral may be a complicated surface, and not a simple curve as presented here.

References

F. Aurenhammer [1991], Voronoi diagrams- a survey of a fundamental geometric data structure. *ACM Computing Surveys* 23, 345-405.

I. Babuška and A. K. Aziz [1976], On the angle condition in the finite element method, *SIAM J. Num. Anal.* 13:214–226.

B.S. Baker, E. Gross and C.S. Rafferty [1988], Non-obtuse triangulation of polygons. *Disc. and Comp. Geom. 3 (1988)* 147-168.

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.

E. K. Buratynski [1990], A fully automatic three-dimensional mesh generator for complex geometries, *Int. J. for Num. Meth. in Engr.* vol 30 931-952.

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.

B. Delaunay [1934], Sur la sphère vide, *Bull. Acad. Sci. USSR(VIII),* Classe Sci. Mat. Nat., 793-800.

R. Duda and P. Hart [1973], *Pattern Classification and Scene Analysis,* Interscience, New York.

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.

D. Kirkpatrick [1979], Efficient computation of continuous skeletons, *Proc. 20th IEEE Annual Symp. on Foundations of Comput. Sci.,* pp. 18-27, Oct. 1979.

C.L. Lawson [1977], Software for C^1 surface interpolation, in *J. Rice*, ed., Mathematical Software III, Academic Press 161-194.

D.T. Lee [1978], Proximity and reachability in the plane. Coordinated Science Laboratory, Univ. Illinois, Tech.Rep. R-831.

D.T. Lee and A. Lin [1986], Generalized Delaunay triangulation for planar graphs. *Disc. and Comp. Geometry 1 (1986)*, 201-217.

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.

S. A. Mitchell and S. A. Vavasis [1992], Quality Mesh Generation in Three Dimensions, *Proc. 8th Symp. on Computational Geometry*, 212-221.

D. Moore [1992], Simplicial Mesh Generation with Applications, C.S. Cornell, PhD thesis, May 1992.

D. Moore and J. Warren [1990a]. Adaptive mesh generation I: packing space. Dep. of Computer Science, Rice University, Tech. Report TR 90-106.

D. Moore and J. Warren [1990b], Multidimensional Adaptive Mesh Generation, Department of Computer Science, Rice University, Rice COMP TR 90-106.

F.P. Preparata and M.I. Shamos [1985], *Computational Geometry: An Introduction*, Springer Verlag, New York.

V.T. Rajan [1991], Optimality of the Delaunay Triangulaion in R^d. In *Proc. 7th ACM Symp. Comp. Geometry* 357-363.

J. Rupert [1992], A New and Simple Algorithm for Quality 2-Dimensional Mesh Generation, Computer Science Division (EECS), UC Berkeley, Report No. UCB/CSD 92/694.

N. Sapidis and R. Perucchio [1989], Advanced techniques for automatic finite element meshing from solid models. *Computer Aided Design,* 21 4 248-253.

E. Schönhardt [1928], Uber die Zerlegung von Dreieckspolyedern in Tetraeder. *Math. Ananalen* 98 309-312.

R. Seidel [1988], Constrained Delaunay triangulations and Voronoi diagrams with obstacles, Technical Report 260, Inst. for Information Processing, Graz, Austria.