Delaunay Refinement Algorithms

A Thesis
Presented to
The Academic Faculty

by

Steven Elliot Pav

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Department of Mathematical Sciences
Carnegie Mellon University
May 2003

For Richard and Nancy.

ACKNOWLEDGEMENTS

It would be difficult to acknowledge those people who have contributed to this work without either making an omission or exceeding the proper page limits:

First, partial funding for this work was provided by National Science Foundation Grants DMS–0208586 and CCR–9902091.

I could not have completed my education without the constant support and encouragement of my family: my father and mother, my two brothers, sisters-in-law, and three young ones to whom "Uncle Steve" is synonymous with "the bogeyman."

I thank my advisor, Noel Walkington, for three years of guidance and support. Much of this work was developed in (nearly) innumerable discussions with Noel and Gary Miller. Many of the insights for this work are theirs, while the deficiencies are mine only.

I have much profitted from discussing the topic with Guy Blelloch and Clemens Kadow. The meshing algorithms described herein were coded with components developed by Umut Acar, Guy Blelloch, Franklin Chen, Herb Derby, David Garmire, Robert Harper and Aleksandar Nanevski, as part of the PSciCo project. That my code runs at all is only through their efforts. Many of the figures (and some of the results) herein were generated with KSEG, written by Ilya Baran. Computational support was provided by Florin Manolache. Much deserved thanks go to the department secretaries.

I enjoyed a summer working with Kokichi Sugihara at Tōkyō University. I thank the NSF and MEXT for offering such a program. I am grateful to Momoko Ono and the Ono family for their friendship and hospitality. I suspect my admission to Carnegie Mellon was due to the special efforts of William Zeimer, for which I am grateful. My knowledge of triangle centers dates to Douglas Hofstadter's CAT-DOG class. The constellations in my math sky all have 17 yellow pigs, er, stars, thanks to David Kelly.

I thank the colleagues with whom I suffered through my first year at Carnegie Mellon: Caner Kazançı, Adam Speight and Kate "Skater" Trapp. John Krueger, Kerry Ojakian and Ksenija Simic round out the fresh-math crew. That there has been no gunplay in our cramped office over the years is a lasting tribute to my officemates, Chad Brown and Matthew Szudzik. My compadres worldwide bear mention for couches slept-upon, phone calls, and art swaps: Daniel Seifert, Stephen Markacs, Paul Ford, Joy Patterson, Heather Rowe, Ian Rothwell and Mark Bishop. If this thesis is defended, I will surely celebrate with Jessi Berkelhammer, Abie Flaxman, Devon Yates, and Critical Mass Pittsburgh.

I cannot thank enough the ladyfriend, Linda Shipley, who *is* the law as far as I know it.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS OR ABBREVIATIONS

Ω	The domain of the input, $(\mathcal{P}, \mathcal{S})$, *i.e.*, the convex hull of $\mathcal{P}$.
θ^*	A lower bound on the angle subtended by input segments, with $0 < \theta^* \leq \pi/3$.
α	The minimum angle in a mesh.
κ	The output angle parameter for the Delaunay Refinement Algorithm.
$\lvert x - y \rvert$	The Euclidian distance between the points x and y.
$\mathrm{lfs}\left(x\right)$	The local feature size at the point x with respect to an understood input, $(\mathcal{P}, \mathcal{S})$.
β, ρ	Constants relevant to the Encroachment Sequence Bound Hypothesis and the generalized good grading theorem.
$\mathrm{d}_1\left(x\right)$	The distance from x to the nearest distinct feature of an understood input, $(\mathcal{P}, \mathcal{S})$.
$\mathrm{lfs}'\left(x\right)$	The local feature size at the point x with respect to an augmented input, $(\mathcal{P}', \mathcal{S}')$.
γ	Constant describing the possible scaling of a Bounded Reduction Augmenter or Feature Size Augmenter.
$\hat{\kappa}$	The output angle parameter for the Adaptive Delaunay Refinement Algorithm.
η	Roughly the ratio of the distance from which a point can cause a segment to split to the radius of that segment, $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$.
μ	The grading constant for split segments under the Adaptive Delaunay Refinement Algorithm. Normally $\mu = \mathcal{O}\left(\frac{\eta}{\theta^*}\right)$.
$\zeta\left(\theta^*\right)$	The optimality constant for input with lower angle bound θ^*. Roughly, this constant reflects how much worse, with regard to number of Steiner points, the Delaunay Refinement Algorithm is than any other mesher with the same output angle guarantees.

GLOSSARY

Local Feature Size — A positive Lipschitz function determined by a set of points and a set of segments, $(\mathcal{P}, \mathcal{S})$. The local feature size of a point x is the radius of the smallest circle centered at x which touches two disjoint elements from the set of points and segments.

Planar Straight Line Graph — Abbreviated PSLG, a set of points and a set of segments, $(\mathcal{P}, \mathcal{S})$, such that two segments intersect at most at an endpoint, the endpoints of all segments are in the collection of points, and points only intersect segments at the ends.

SUMMARY

In this thesis, the problem of planar mesh generation with quality bounds is considered. A good mesh generation algorithm should accept an arbitrary planar straight line graph and output a Delaunay Triangulation of a set of points (all the input points plus some acceptably small number of Steiner Points) which conforms to the input and has no small or large angles. Ruppert's algorithm for solving this problem is reanalyzed: it is shown that this algorithm terminates for a larger class of input than previously proven. The analysis removes a deficiency of Ruppert's original exposition and analysis: the requirement that input segments meet at nonacute angles.

An "adaptive" variant of the algorithm is introduced. This simple variant deals with small input angles adaptively, producing meshes where most triangles have no small angles, except where conformality to the input prevents this. The analysis shows that small output angles of the mesh are not much smaller than nearby small input angles, and that large angles are entirely prevented. As in Ruppert's algorithm, the output meshes are "well-graded" in the sense that the size of edges in the output is bounded by some function of the input.

Lastly, a certain measure on triangulations with bounded minimum angle is considered. This measure is similar to that analyzed by Mitchell. The new measure allows an asymptotic improvement in the optimality claims of the algorithms considered.

CHAPTER I

INTRODUCTION

"There are only two kinds of math books. Those you cannot read beyond the first sentence, and those you cannot read beyond the first page." –*C.N. Yang*

1.1 Motivation

Physical systems are often modelled by a collection of partial differential equations. The solution to such a collection is a function which satisfies the equations on the domain of the problem, and exhibits some prerequisite continuity and boundary conditions. Typically the class of potential solutions is an infinite dimensional vector space, call it $\mathcal{U}$. The method of finite elements selects some finite dimensional subspace, $\mathcal{U}_h \subset \mathcal{U}$, and finds an member of this space which approximates the actual solution to the equations. Finding this approximate solution is performed automatically, by computer, and the task is subdivided into determining the basis vectors of $\mathcal{U}_h$, constructing some (typically sparse) system of linear equations, then solving the system, which yields the approximate solution [24, 36, 4].

The space $\mathcal{U}_h$ usually consists of truncated polynomials of a certain degree with some local support; it is convenient to construct this space by partitioning the original problem domain, call it Ω, into combinatorially similar *elements*, then defining the basis vectors in terms of this partitioning. Two dimensional domains are usually decomposed into triangles or quadrilaterals; for three dimensional domains, tetrahedra or hexahedra. The partition is called a *mesh*, and the objective of the study of meshing is to reconcile how the choice of meshes affects the success of the finite element method. A decomposition of a physical domain is illustrated in Figure 1.

The classical analysis of the finite element method shows that the approximation error is bounded by some function of the element size. This matches the intuition that somehow finer meshes should produce more accurate results, though at the computational cost of solving larger linear systems. However, this analysis always assumes a bound on the *aspect ratio* of mesh elements [24]. The aspect ratio is some measure of the shape of the elements, and a number of roughly equivalent definitions exist [21, 44].

The celebrated work of Babuška and Aziz shows that in two-dimensional triangular meshes, large angles (near π) can yield poor interpolation accuracy [2, 23]. Small (and large) angles in meshes are known to generate poorly conditioned linear systems for the finite element method, leading to poor performance of iterative solvers [1, 44].

Many meshing algorithms pursue the strategy of producing meshes with no small angles, since this also eliminates large angles. In some situations no such mesh may exist; for example when the boundary of the domain Ω contains a small angle the mesher cannot hope to partition Ω without including a small angle. In the case that small angles are unavoidable, a good mesher should at least avoid large angles. The remainder of this thesis, then, is devoted to the fairly simpleminded task of dividing two dimensional domains into triangular meshes, of modest cardinality, while avoiding small and large angles. This is actually a rather small subset of meshing; the reader interested in a broader view is encouraged to consult the survey paper by Bern and Eppstein [6].

1.2 Meshing

Many meshing algorithms generate Delaunay or "nearly" Delaunay meshes. Delaunay methods are preferred because they have a sound theoretical footing, there are many algorithms for generating Delaunay Triangulations, not a few of which have been implemented and are commonly available, and Delaunay Triangulations have a number of attractive and well-studied properties [37, 15, 18, 40, 9, 47, 22]. In particular, the Delaunay Triangulation is known to maximize the minimum angle over all meshes on a given set of coplanar points. Moreover, of all possible meshes, it lexicographically maximimizes the vector of all angles of the triangulation (modulo degeneracies) [15, 33].

The Constrained Delaunay Triangulation is a generalization of the Delaunay Triangulation to the case of a set of input segments and a set of input points. It has the property of maximizing the minimum angle over all triangulations on the set of points which include the set of segments as edges [26, 11].

Most meshing algorithms do not simply output the (Constrained) Delaunay Triangulation of the input, as this can lead to rather small angles, even when they are absent from the input. Rather a typical mesher will add *Steiner Points* to the input points, sometimes dividing input segments into subsegments, and return a mesh which is the (Constrained) Delaunay Triangulation of this augmented set of points. A common means of implementing this is by generating some mesh on the input, then incrementally examining the properties of the current mesh, and introducing Steiner Points to locally improve mesh quality. As illustrated in Figure 1, there is usually some tradeoff between number of Steiner Points added and overall mesh quality.

Chew introduced the first Delaunay Refinement algorithm with provable quality bounds [12]. The algorithm accepts a *planar straight-line graph,* and, under some modest conditions on the input, returns a Constrained Delaunay Triangulation in which no angle is smaller

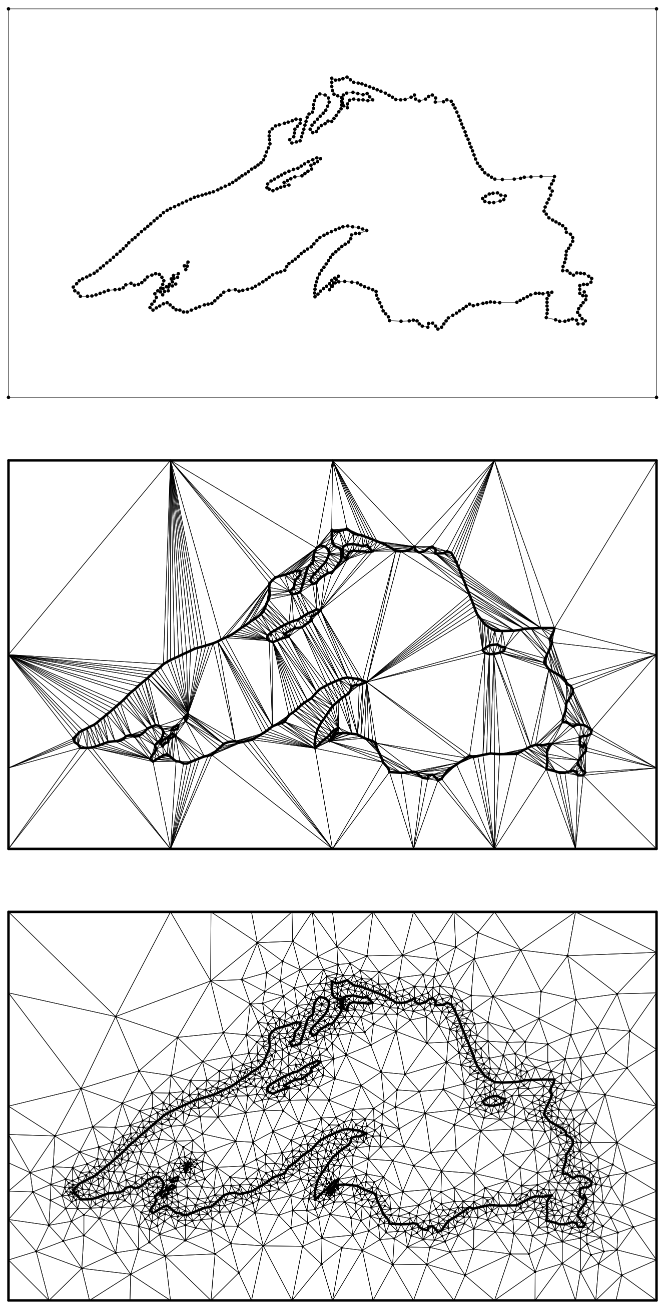

Figure 1: A "famous" domain from the world of meshing, Lake Superior, is shown at top. In the middle, a Delaunay mesh of the domain is shown, with many small angle triangles. The minimum angle in the mesh is around $0.9°$, and the maximum angle is around $177°$. At bottom, a high quality Delaunay mesh of the domain is shown; all angles are greater than $14°$, and less than $125°$. In general one can expect some tradeoff between mesh quality and cardinality of the point set. Input edges are shown in bold in the bottom two figures.

than $\pi/6$. The method is notable for its simplicity: the algorithm maintains a Constrained Delaunay Triangulation of the current set of points, then adds to this set the circumcenter of any triangle in the mesh with circumradius larger than some parameter, h, which is user-provided; when no such triangle exists, the algorithm outputs the mesh. The main drawback of the method is that the meshes are uniform–all edges in the output mesh have length between h and $2h$. By itself, this isn't damning, but it implies that the number of triangles in the mesh is $\Omega\left(A/h^2\right)$, where A is the area of the input domain, Ω. This may be an unacceptably large number of elements.

Ruppert developed a Delaunay Refinement algorithm which is more input-sensitive [39]. The algorithm constructs a Conforming Delaunay Triangulation of the input, *i.e.*, a mesh which is the Delaunay Triangulation of a set of points including the input points, and in which each input edge is the union of output edges. We will describe the algorithm in greater detail in the following chapters. Roughly it guarantees that input edges are represented by adding midpoints of edges as necessary, and removes triangles with small angles by adding their circumcenters to the set of maintained points. The algorithm can guarantee that no angle in the output mesh is smaller than some parametrizable $\kappa < \arcsin \frac{1}{2\sqrt{2}} \approx 20.7°$. Unlike in Chew's method, the edges of the mesh are well-graded, that is, short edges can be "blamed" on nearby parts of the input which are close to one another. Moreover, it comes with some optimality guarantee: the number of Steiner Points added by the algorithm is within a constant of the number added by *any* mesher which generates a mesh on the input with minimum angle κ. This optimality constant depends only on the angle κ.

A few holes remained in Ruppert's original analysis, however. The most serious drawback of the original exposition was that it required input angles meet at nonacute angles. This restriction was loosened to a $\pi/3$ lower bound on input angles [41, 30]. It was known that the algorithm could fail (*i.e.*, run indefinitely) for input with small angles, but that this wasn't always the case. Ruppert's original paper suggests a heuristic for splitting input segments to prevent the algorithm from diverging, but no proof is given of correctness [39]. This heuristic cannot be employed without some other modification of the algorithm, as small input angles could lead to vertices in the mesh with out-edges of highly varying length, which clashes with the fact that in a mesh with no angle larger than κ, the ratio of the lengths of edges emanating from a vertex is bounded by at most $(2\cos\kappa)^{\frac{2\pi}{\kappa}}$.

Shewchuk's "Terminator" algorithm incorporates Ruppert's idea for dealing with small angles, and adds some modification for subverting the rules of Ruppert's method when they could lead to a sequence of decreasing mesh edges [43]. The proof of termination is laudable for its elegance; the main idea of the proof–and of the algorithm–is to ensure that any newly created mesh edge can either be directly blamed on some input features, or is not shorter than some (possibly distant) previously created mesh edge. Thus the algorithm terminates with no mesh edge smaller than some h, which is determined solely by the input. This property, however, points to an obvious failing of the analysis, *e.g.*, that there is no theoretical guarantee of good-grading. Though this may not be observed in practice, the algorithm might generate meshes where all edges have length $\Theta(h)$, and thus have $\Omega\left(A/h^2\right)$ elements.

The output angle guarantees of the Terminator are also unsatisfactory; the analysis shows only that no angle of the output is smaller than $\arcsin\left[\sin\left(\frac{\theta^*}{2}\right)/\sqrt{2}\right]$, where $\theta^* \leq \pi/3$ is a lower bound on input angle. The proof of this lower bound is tight, in that it provides an example where an angle of size $\arcsin\left[\sin\left(\frac{\theta^*}{2}\right)/\sqrt{2}\right] + \epsilon$ could persist in the final mesh. Moreover, no maximum angle bound is given, other than the naïve one of $\pi - 2\arcsin\left[\sin\left(\frac{\theta^*}{2}\right)/\sqrt{2}\right]$, which deteriorates when θ^* is small.

Shewchuk shows that his algorithm can be modified to remove angles which are smaller than $\pi/6$ and are "far" from input features [43]. This partly addresses a long-standing gap between the theoretical and actual performance of Ruppert's algorithm; starting with Ruppert [39], users of the algorithm have noted that it can be run in practice with κ chosen as large as $\pi/6$, even though the theoretical upper bound on the number of Steiner Points given by Ruppert diverges as κ approaches $\arcsin \frac{1}{2\sqrt{2}}$.

Ruppert's original optimality constants are far too large to be meaningful; for the case where $\kappa = \pi/9$, Ruppert proved a constant of 1.81×10^{25}. Mitchell studied a measure on triangulations, showing that it is inversely related to the minimum angle of the mesh [31]. Use of this measure improves the optimality constants by several orders of magnitude; for the case where $\kappa = \pi/9$, an optimality constant of 1.1×10^6 can be shown.[1] Such a guarantee, however, is useless. Fortunately the algorithm in practice performs much better than the theory can guarantee.

Chew refined his original algorithm to accept a user-defined grading function, and to split triangles that do not satisfy the grading function, or which have angles smaller than $\pi/6$. The algorithm returns a mesh which is constrained Delaunay, although some input segments have been split. The algorithm relies on the operation of removing Steiner Points from the mesh, which marks a departure from his first algorithm and Ruppert's algorithm which only add Steiner Points. Chew claims some vague optimality result; this claim could be interpreted, roughly, to mean that the algorithm produces a mesh with a number of Steiner Points within a constant of that in any mesh which also satisfies the user-defined grading function. No formal statement or proof of optimality is made. Moreover, the

[1] Mitchell claims a value of 6.3×10^5 which is clearly erroneous. I have attributed this to his use of his Theorem 10 when computing the constants, and not his Theorem 11, as is appropriate [31].

algorithm, in the form presented by Chew, requires that no angle of the input is less than $\pi/3$ [13].

Shewchuk analyzes Chew's second algorithm, removing the user-defined grading function. He shows that it can produce meshes which are graded to the input, as those produced by Ruppert's algorithm, and in which no angle is smaller than a parametrizable $\kappa < \arcsin 5^{-1/2} \approx 26.57°$. Given this proof, the optimality results of Ruppert apply, giving an affirmation of Chew's claim [43]. Again, this result relies on a $\pi/3$ lower bound on input angle.

Shewchuk entertains a bewildering array of modifications which can be applied to standard Delaunay Refinement algorithms [43]. The first is the replacement of "diametral circles" with "diametral lenses;" roughly the idea is to make the algorithm less likely to split an edge because a Steiner Point is too close to an input segment. In so doing, he is able to increase the output guarantee of his Terminator algorithm to $\arcsin\left[\left(\sqrt{3}/2\right)\sin\left(\frac{\theta^*}{2}\right)\right]$; applying this modification, however, can sacrifice the Delaunay property, and the output meshes are only guaranteed to be Constrained Delaunay. As mentioned above, Shewchuk shows that a bound of $\kappa < \pi/6$ can be applied to triangles "far" from input features, without sacrificing good grading. He also discusses "range-restricted segment splitting," which is a means of modifying how input edges are split in Ruppert's algorithm to get a $\pi/6$ output angle bound; this modification requires an input bound of $\pi/3$, and has no grading guarantee.

Ollivier-Gooch and Boivin formulate a grading function which is controlled by two user-defined parameters. They use this function in a modified version of Ruppert's algorithm in the same way a grading function is used in Chew's second algorithm [38]. They also extend this work to the more general case where the input consists of vertices and curves [8].

Miller has recently formulated a Delaunay Refinement algorithm which admits a timing analysis. This analysis shows that the algorithm can run in time $\mathcal{O}(m \log m)$ plus some geometric factor which depends on the input, where m is the number of output points. [28]

Delaunay Refinement has also been generalized to three dimensions. Ruppert considers the matter briefly and without optimism; he recognized that his method of proving good grading could only guarantee a bound on a weak quality measure in three dimensions [39]. Indeed, the Delaunay Refinement method bounds the "radius-edge" ratio of simplices; in two dimensions this is equivalent to a bounded classical aspect ratio, but in three and higher dimensions, the two ratios need not be related. Known generalizations of Ruppert's Algorithm to three dimensions are notorious for producing meshes with "slivers," which are tetrahedra with nearly coplanar vertices, and which have a fine radius-edge ratio, but a poor classical aspect ratio. Meshes with bounded radius-edge ratio may be sufficient for some applications [27], but in general a bound on classical aspect ratio is desired. Practitioners often rely on a postprocessor to remove slivers from meshes produced by the Delaunay Refinement method [10, 17].

Shewchuk proves termination and good grading of a three-dimensional generalization of Ruppert's algorithm; it requires that input segments make angles no smaller than $\pi/3$, that no input segment meets a plane at an angle smaller than $\arccos \frac{1}{2\sqrt{2}} \approx 69.3°$, and, essentially, that adjoining input planes have nonacute dihedral angle. The algorithm is shown to produce well-graded meshes. However, in three dimensions, this does not imply size optimality. As mentioned above, output meshes may contain slivers [41]. Miller *et al.*, and Ollivier-Gooch and Boivin also discuss three-dimensional Delaunay Refinement algorithms, with results similar to Shewchuk's [30, 29, 38].

While this thesis focuses mostly on the Delaunay Refinement method, it would be negligent not to give a brief survey of other advances in guaranteed-quality mesh generation. Many of the meshing algorithms outside the Delaunay Refinement framework rely on quadtree and other anisotropic methods.

Baker *et al.* describe an algorithm that produces meshes with no angle smaller than $\arctan \frac{1}{3} \approx 18.43°$, and no angle larger than $\pi/2$. The algorithm comes without guarantees on the number of Steiner Points added [3]. The technique involves a grid overlaying the input which is used to generate Steiner Points.

Bern *et al.* produced the first meshing algorithm with quality bounds and a size-optimality guarantee. The algorithm matches the $\arctan \frac{1}{3}$ lower bound of Baker *et al.*, but comes with a guarantee on the number of Steiner Points added. The algorithm runs in (optimal) $\mathcal{O}(n \log n + m)$ time, where m is the output size [7]. The algorithm is based on quadtrees: essentially, space is recursively divided into squares until each square contains at most one input feature and neighboring squares are not of widely varying size, then the mesh is constructed from these squares. As presented the algorithm is more of a theoretical victory, since output meshes have an unacceptably large number of Steiner Points, even for relatively simple inputs. Practical mesh generators based on this algorithm must rely heavily on heuristics to reduce the output cardinality. Moreover, meshes produced by this method display their Cartesian ancestry. The quadtree-method has been generalized to higher dimensions by Mitchell and Vavasis [32].

Bern *et al.* developed an algorithm for generating nonobtuse triangulations of polygons with (optimal) $\mathcal{O}(n)$ triangles. The method uses disc packing to subdivide the polygon into regions which are then meshed. No lower angle bound is provided, and the algorithm only works on polygons, possibly with polygonal holes–no generalization to planar straight-line graphs has been discovered. The output of this algorithm is often visually aberrant–triangle size in the output mesh is not closely tied to the local feature size of the input [5].

1.3 This Thesis

This thesis reanalyzes the Delaunay Refinement Algorithm of Ruppert; it is shown that the algorithm terminates for a wider class of input than previously suspected. Without any alteration, the method is shown to be applicable to input with minimum angle as small as $\pi/4$, subject to a condition on the angles about every input point. This result comes with a non-trivial loss in output quality. It is shown that Ruppert's algorithm with no modification will terminate with good grading for input with arbitrarily small angles if the input satisfies a certain condition regarding the lengths of segments which share a common endpoint. In this case, an output bound of $\arcsin\left[\sin\left(\frac{\theta^*}{2}\right)/\sqrt{2}\right]$ is possible, where $\theta^* \leq \pi/3$ is a lower bound on the input angle. It is shown that arbitrary input can be put in the latter form by the addition of a few augmenting points, while keeping the optimality result (albeit with an extra factor in the constant). Ruppert's original strategy of splitting on concentric circular shells performs this augmenting procedure on an "as-needed" basis, thus it is shown that Ruppert's original algorithm with concentric shell splitting can be applied to input with arbitrarily small angles and produces well-graded output with the same output guarantees as Shewchuk's Terminator.

An alteration of the algorithm, the Adaptive Delaunay Refinement Algorithm, is also analyzed. This algorithm simply redefines the bad-triangle test of Ruppert's algorithm to avoid edges which are too short compared to the local feature size. Assuming a slightly more restrictive condition on edge-lengths than above, this algorithm produces well-graded meshes on input with arbitrary lower angle bound, θ^*. All angles in the output mesh are at least $\arcsin 2^{-7/6} \approx 26.45°$, except those at or opposite an input angle of size θ: such output angles are shown to be at least $\arctan\left(\frac{\sin\theta}{2-\cos\theta}\right)$. Thus the output contains no angle smaller than $\min\left\{\arcsin 2^{-7/6}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$. Moreover, in spite of the potential of arbitrarily small output angles if θ^* is small, the analysis guarantees that no output angle is *larger* than around 137.1°. Again, Ruppert's method of concentric shell splitting is shown to put input into the required form. It is shown that this variant is entirely unnecessary when θ^* is modest. That is, Ruppert's algorithm with concentric shell splitting employing output angle parameter $\kappa < \min\left\{\arcsin 2^{-7/6}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$ has the same performance guarantees as the Adaptive Delaunay Refinement Algorithm. Thus when θ^* is greater than about 36.53°, the adaptive variant is not necessary.

The work of Mitchell on optimality measures is also expanded. An integral over the edges of a triangulation with bounded minimum angle is related to the minimum angle. This allows an improvement of Ruppert's optimality proof where segment midpoints and triangle circumcenters are considered separately. This also allows an asymptotic improvement in the optimality constants for the abovementioned algorithms.

CHAPTER II

DELAUNAY TRIANGULATIONS AND DELAUNAY REFINEMENT

"The truth is rarely pure, and never simple." –*Oscar Wilde*

2.1 The Delaunay Triangulation

An introduction to Delaunay Triangulations is provided here, though more skillful and complete expositions exist elsewhere [37, 15, 18]. Although the theory easily extends to $\mathbb{R}^n$, we will work entirely in two dimensions. We suppose the existence of a finite set of "sites," $\mathcal{S}$, which are points in $\mathbb{R}^2$.

Definition 2.1.1 (Delaunay Triangulation). An edge or polygon, $\mathcal{P}$, which has corners $C = (s_0, s_1, \ldots, s_l)$, is said to have the *(Strong) Delaunay Property* with respect to $\mathcal{S}$ if there is some 2-dimensional open ball, $\mathcal{B}$, such that s_i is on the boundary of $\mathcal{B}$ for $i = 0, 1, \ldots, l$, and such for every $s_j \in \mathcal{S} \setminus C$, s_j is not in (the closure of) $\mathcal{B}$.

We let $\mathcal{D}(\mathcal{S})$ be the collection of edges (s_i, s_j) with the Strong Delaunay Property. Then the *Delaunay Triangulation* is the graph $(\mathcal{S}, \mathcal{D}(\mathcal{S}))$. We implicitly identify the graph with its embedding in $\mathbb{R}^2$, and identify the edges with line segments.

A set of points in $\mathbb{R}^2$ and the associated Delaunay Triangulation are shown in Figure 2.

Balls of the type mentioned above play a special part in the following analysis, making the following definition convenient.

Definition 2.1.2 (Circumcircle). Given three noncollinear sites, s_0, s_1, s_2, in $\mathbb{R}^2$, the *circumcircle* of the triangle on these sites is the unique circle passing through all three sites. The *circumcenter* of the triangle is the center of its circumcircle.

Given two distinct sites, s_0, s_1, in $\mathbb{R}^2$, a circumcircle of the edge (s_0, s_1) is any circle with the edge as a chord. The *diametral circle* of the edge is the circle with the edge as a diameter.

This description of the Delaunay Triangulation might better be termed a *pre-triangulation* since in the presence of degeneracies, the embedded graph will not decompose space into triangles. This is illustrated in Figure 3, which shows the Delaunay Triangulation of a set of five cocircular sites in $\mathbb{R}^2$. Each of the edges $(A, B), (B, C), (C, D), (D, E), (E, A)$ has

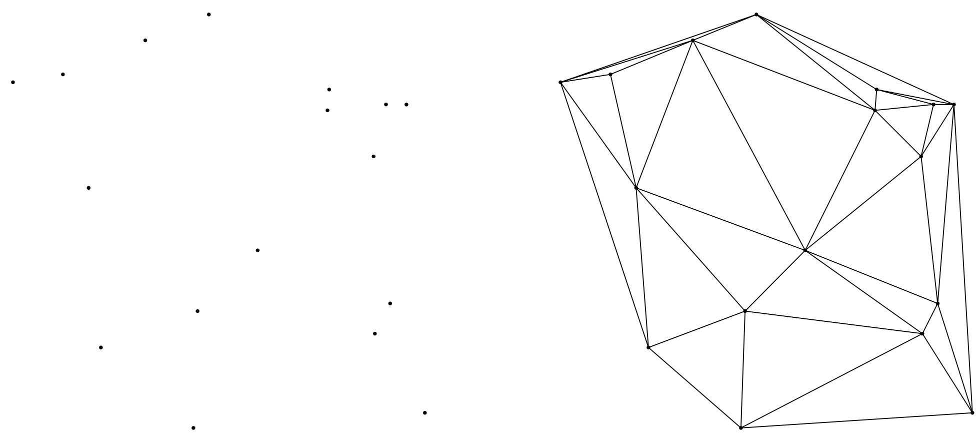

(a) A Set of Points in the Plane

(b) The Delaunay Triangulation of these Points

Figure 2: In (b), the Delaunay Triangulation of the points of (a) is shown.

the Strong Delaunay Property, but there is no edge cutting through the pentagon $ABCDE$. In fact, this “triangulation” contains no triangles at all, rather one pentagon.

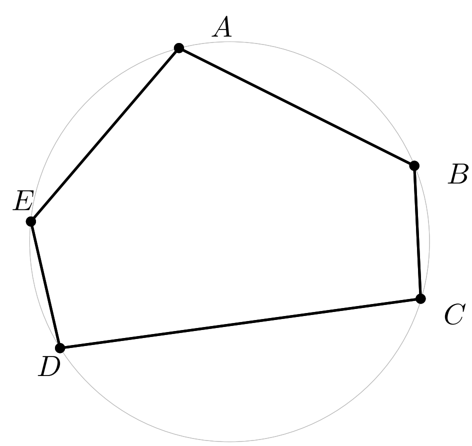

Figure 3: The Delaunay Triangulation of a set of cocircular sites, $\mathcal{S} = \{A, B, C, D, E\}$, is shown. The outer edges of the pentagon $ABCDE$ all have the Strong Delaunay Property and thus are in $\mathcal{D}(\mathcal{S})$, but no edge through the pentagon has the Strong Delaunay Property. The “triangulation” then contains no triangles.

Theoreticians and practitioners of meshing generally agree that degeneracies are an unavoidable nuisance. A number of strategies exist for dealing with degeneracy. At one end of the spectrum is the approach of rewriting the common geometric predicates to either symbolically perturb input in a consistent manner so that no input is degenerate, or to better detect degeneracy via higher precision calculation [19, 48, 41]. In the latter case, the algorithm must still be able to produce correct results when degeneracy has indeed been detected. In either case, the rewritten predicates may be unacceptably slow or too difficult to implement. At the other end of the spectrum are "topological" methods, which assume imperfect predicates, but attempt to at least output topologically consistent results [45, 16].

In this thesis, the common strategy of ignoring degeneracy is employed to simplify the exposition. The algorithms described herein can be used to generate a Delaunay Triangulation as defined above, after which degeneracies can be resolved in some way consistent with the input and optimizing output angles. Alternatively, they may be combined with some of the abovementioned methods for dealing with degeneracy to produce an actual triangulation.

2.2 The Meshing Problem

The meshing problem is described in terms of the input to the algorithm and the expected conditions on the output. The input to the mesher is defined as follows:

Assumption 2.2.1 (Input). The input to the meshing problem consists of a finite set of points, $\mathcal{P} \subseteq \mathbb{R}^2$, and a set of segments $\mathcal{S}$ such that

(a) the two endpoints of any segment in $\mathcal{S}$ are in $\mathcal{P}$,

(b) any point of $\mathcal{P}$ intersects a segment of $\mathcal{S}$ only at an endpoint,

(c) two segments of $\mathcal{S}$ meet only at their endpoints, and

(d) the boundary of the convex hull of $\mathcal{P}$ is the union of segments in $\mathcal{S}$.

Let Ω denote the convex hull of the input, and let $0 < \theta^* \leq \pi/3$ be a lower bound on the angle between any two intersecting segments of the input.

Items (a)-(c) characterize $(\mathcal{P}, \mathcal{S})$ as a Planar Straight Line Graph (PSLG); item (d) can always be satisfied by augmenting an arbitrary PSLG which does not satisfy it with a bounding polygon (typically a rectangle). The restriction that $\theta^* \leq \pi/3$ is merely for convenience; asserting a larger lower bound does not give any better results.

Assumption 2.2.2 (Output). The algorithm outputs sets of points, segments, triangles, $\mathcal{P}', \mathcal{S}', \mathcal{T}'$, respectively, satisfying:

(a) **Complex:** The output collectively forms a simplicial complex, *i.e.*, $\{\emptyset\} \cup \mathcal{P}' \cup \mathcal{S}' \cup \mathcal{T}'$ is closed under taking boundaries, and under intersection.

(b) **Delaunay:** Each triangle of $\mathcal{T}'$ has the Delaunay property with respect to $\mathcal{P}'$.

(c) **Conformality:** $\mathcal{P} \subseteq \mathcal{P}'$, and for every $s \in \mathcal{S}$, s is the union of segments in $\mathcal{S}'$.

(d) **Quality:** There are few or no "poor-quality" triangles in $\mathcal{T}'$.

(e) **Cardinality:** Few Steiner points have been added, *i.e.*, $|\mathcal{P}' \setminus \mathcal{P}|$ is small.

Given the discussion in the introductory chapter on mesh quality, one passable definition of item (d) is that there are some reasonably large constants $0 < \alpha \leq \omega \leq \frac{\pi+\alpha}{4}$ such that for every triangle $t \in \mathcal{T}'$, no angle of t is smaller than α or larger than $\pi - 2\omega$. However, such a guarantee is not consistent with conformality of the triangulation (item (c)) when the input contains angles less than α. Thus a weaker definition is that most triangles satisfy the above condition, and those that do not (a) are describably near an input angle of size θ, (b) have no angle smaller than $\theta - \mathcal{O}\left(\theta^2\right)$, and (c) have no angle larger than $\pi - 2\omega$. This definition, of course, is rigged so that it will be satisfied by some algorithm herein described.

2.3 The Delaunay Refinement Algorithm

We describe a whole class of algorithms, which we collectively refer to as "the" Delaunay Refinement Algorithm. This class contains Ruppert's original formulation [39], as well as the incremental version [30].

We suppose that the algorithm maintains a set of "committed" points, initialized to be the set of input points, $\mathcal{P}$. The algorithm also maintains a set of "current" segments, initialized as the input set, $\mathcal{S}$. The algorithm will "commit" points to the set of committed points. At times the algorithm will choose to "split" a current segment; this is achieved by removing the segment from the set of current segments, adding the two half-length subsegments which comprise the segment to the set of current segments, and committing to the midpoint of the segment. The word "midpoint" should be taken to mean one of these segment midpoints for the remainder of this work, to distinguish them from the other kind of Steiner Point, which will be called "circumcenters."

The algorithm has two high-level operations, and will continue to perform these operations until it can no longer do so, at which time it will output the committed points, the current segments and the Delaunay Triangulation of the set of committed points. For convenience, we say that a segment is "encroached" by a point p if p is inside the diametral circumball of the segment. Then the two major operations are as follows:

(Conformality) If s is a current segment, and there is a committed point that encroaches s, then split s.

(Quality) If a, b, c are committed points, the circumcircle of the triangle Δabc contains no committed point, triangle Δabc has an angle smaller than the global *minimum output angle*, κ, and the triangle's circumcenter, p is in Ω, then attempt to commit p. If, however, the point p encroaches any current segment, then do not commit to point p, rather in this case split one, some, or all of the current segments which are encroached by p.

It should be clear that if the algorithm terminates then every segment of the set $\mathcal{S}$

has been decomposed into current segments, none of which are encroached by committed points, and thus have the Delaunay property with respect to the final point set, and are thus present in the output Delaunay Triangulation. The algorithm clearly never adds any points outside Ω. The following claim insures that if the algorithm terminates, no triangle in the Delaunay Triangulation has an angle greater than the minimum output angle κ. The proof requires the notion of the radical axis [14, 25], and is somewhat tangential to our discussion of the algorithm; we include it for its completeness, if not for its clarity. The proof is essentially the same as that of Shewchuk [41, Lemma 13]

Claim 2.3.1. Suppose that no current segment is encroached by a committed point. Let Δabc be a triangle with committed corners and with the Delaunay Property with respect to committed points. Then the circumcenter of the triangle is inside Ω.

Proof. Let p be the circumcenter of the triangle. If all of a, b, c are on the diametral circle of some segment, then p is the center of the segment, and thus is in Ω. Assuming, toward a contradiction, that p is outside Ω, there is some segment s on $\partial\Omega$ which intersects the interior of one of (p, a),(p, b),(p, c). Starting from the point p, imagine a growing circle, ending with $\mathcal{C}_1$, the circumcircle of the triangle. The *radical axis* of the growing circle and $\mathcal{C}_2$, the diametral circle of s, is perpendicular to the line through the centers of $\mathcal{C}_1$,$\mathcal{C}_2$. When the growing circle first touches $\mathcal{C}_2$, the radical axis runs through the point of intersection; as the growing circle reaches $\mathcal{C}_1$, which must have a point not on $\mathcal{C}_2$, the radical axis must pass one or both endpoints of s, which are thus actually inside $\mathcal{C}_1$, a contradiction to the Delaunay Property of the triangle. $\square$

The heuristics involved with determining which operation to perform when and on which segment or poor-quality triangle are not relevant to our discussion. This is not to say that they might not affect ease of implementation, running time, cardinality of the final set of committed points, parallelizability, etc. A common heuristic (and the one chosen by Ruppert and others) is to prefer conformality operations over quality operations, which likely results in a smaller output, and which simplifies detecting that a circumcenter is outside of Ω. A description of a member of this class of algorithms would have to include some discussion of how to figure out which current segments are encroached, which triangles are suitable for removal via the quality operation, how to deal with degeneracy, etc. We do not concern ourselves with these details (though see [15, 46, 28, 30, 35, 41, 34]).

2.3.1 The Delaunay Refinement Algorithm at Work

The reader familiar with the Delaunay Refinement Algorithm may have found the above description a bit odd. Most expositions on the matter give much more detailed accounts, talking about meshes which are updated by point additions, segments which are "buried" in the mesh, etc. We ignore these details since they are irrelevant to our results. Since this may leave some readers bewildered, we clarify our description of the algorithm by illustrating a possible run of the algorithm in Figure 4.

2.4 Failures of the Delaunay Refinement Algorithm

We describe some input for which the Delaunay Refinement Algorithm will not work. These are input which either result in an infinite recursion, or in poor quality output. The consideration of these input is instructive in improving the algorithm and its analysis.

Ping-Pong Recursion: A well-known failure [39] of the Delaunay Refinement Algorithm is the "ping-pong," illustrated in Figure 5. If two input segments meet at an angle θ less than $\pi/4$, the algorithm can easily fall into an infinite recursion. One segment may be encroached by an input point on the other segment, as in Figure 5(b). The algorithm splits the encroached segment to ensure conformity, but the newly committed midpoint encroaches the other segment, as in Figure 5(c). Thus the other segment will be split. The picture has effectively been scaled in half, and the cycle of midpoints will be repeated *ad infinitum*, as in Figure 5(d). We will show that the ping-pong can only occur if the segments meet at an angle less than $\pi/4$.

Spiral Recursion: Another case that can lead to an infinite recursion is the "spiral," as illustrated in Figure 6(a). The seven segments, $\{S_i\}_{i=0}^{6}$, are each separated by $2\pi/7 > \pi/4$, (so will not cause ping-ponging), and their lengths form a geometric progression. As shown in Figure 6(b), an endpoint of S_6 encroaches on S_0. The algorithm splits S_0 to ensure conformity, but the newly committed midpoint encroaches S_1. Another midpoint is committed, which encroaches S_2. Eventually S_6 is split because it is encroached. However, the midpoint of S_6 encroaches a subsegment on S_0, as shown in Figure 6(c). The entire figure has been effectively scaled in half, so the cycle will occur again, as in Figure 6(d). Indeed, the cycle will repeat forever. We will show that, without any restriction on lengths of input segments, there is a condition concerning the angles about a given input point that prevents spirals; moreover, this condition is simple to check.

Refinement Recursion (I): Next we consider an input for which conformity alone does not cause problems, rather the choice of κ too large can cause an infinite recursion. Consider 7 input segments, $\{S_i\}_{i=0}^{6}$, each separated by $\pi/4$, (except S_6 and S_0, which are separated by $\pi/2$.), and such that the length of S_0 is 1 and the length of S_i is $2^{-i/2} + \epsilon$, for $i = 1, 2, \ldots, 6$, and some $\epsilon > 0$. If the midpoint of S_0 is added to the mesh, it will encroach S_1, which will be split, causing S_2 to be split, and eventually S_6, which has a length $\frac{1}{8} + \epsilon$, is split. This is shown in Figure 7(a). The triangle shown in Figure 7(b) (which has empty

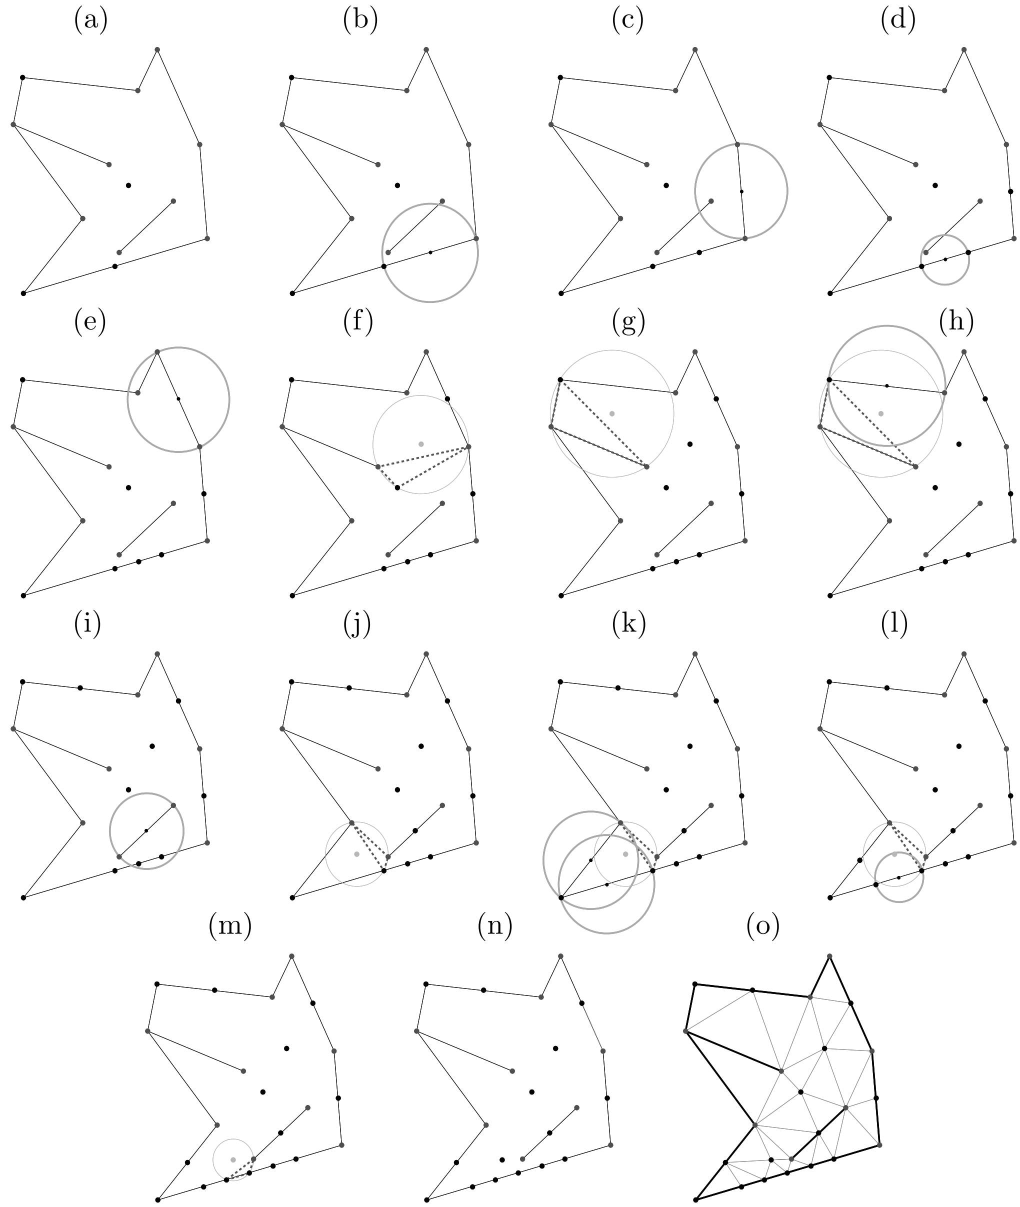

Figure 4: A possible run of the Delaunay Refinement Algorithm on the input of (a) is shown. In (b)-(e), segments are split in (Conformality) operations. In (f) a circumcenter is added in a (Quality) operation. In (g) a circumcenter is considered, but it encroaches a segment, as shown in (h); the segment is split instead. In (i) another segment is split. A circumcenter is considered in (j), but it encroaches on two segments, both of which are split. In (l) the same circumcenter is considered again, but instead another segment is split. A circumcenter is committed in (m). In (n) the final vertex and segment sets are shown, in (o) the Constrained Delaunay Triangulation is shown, with segments in bold.

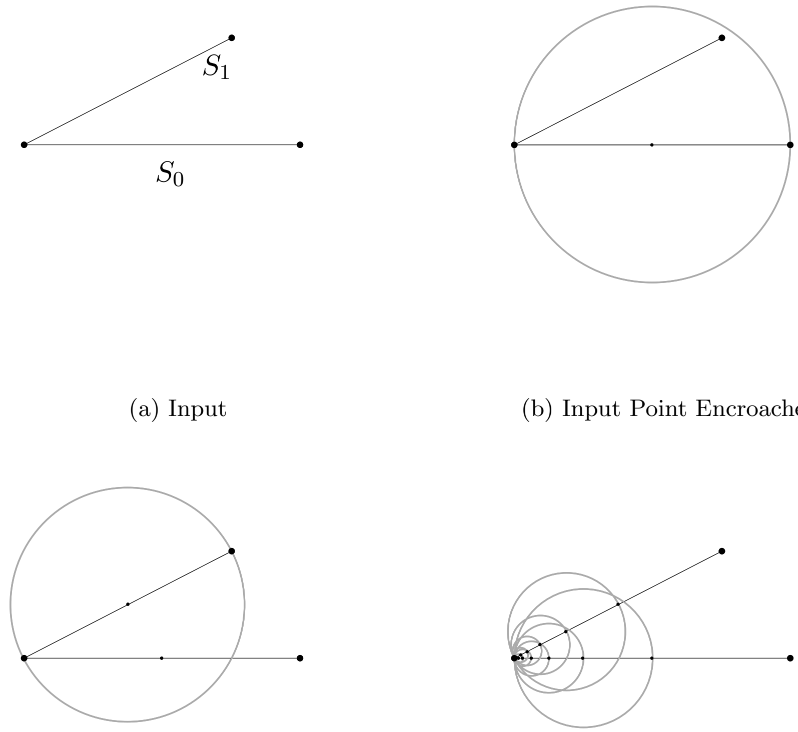

(a) Input

(b) Input Point Encroaches

(c) Midpoint Encroaches Opposite Segment

(d) Cycle Repeats

Figure 5: An input which breaks the Delaunay Refinement Algorithm is illustrated. The input is shown in (a); the segments meet at an angle of approximately 27°. The endpoint of S_1 encroaches on S_0, as in (b). The newly committed midpoint encroaches on S_1, whose midpoint is committed, as in (c). Since the figure has essentially been scaled in half, the cycle repeats, as shown in (d); the algorithm will fail to terminate for this input.

circumcircle) has an angle smaller than $\arcsin\left(\frac{1+8\epsilon}{\sqrt{73}}\right)$. If the algorithm attempts to add this triangle's circumcenter, it will encroach on a subsegment on S_0, as shown. This will cause another sequence of midpoints, as shown in Figure 7(c). Since splitting at the midpoints has merely scaled the picture by a half, another skinny triangle will be present, causing another encroachment spiral, causing another skinny triangle, etc. and the algorithm will fall into an infinite recursion.

Thus the algorithm will not terminate unless κ in this case is no greater than $\arcsin \frac{1}{\sqrt{73}} \approx 6.72°$. Clearly this minimum output angle is far smaller than the smallest angle of the input, $\pi/4$.

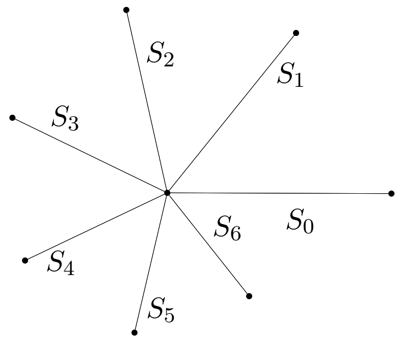

(a) Input

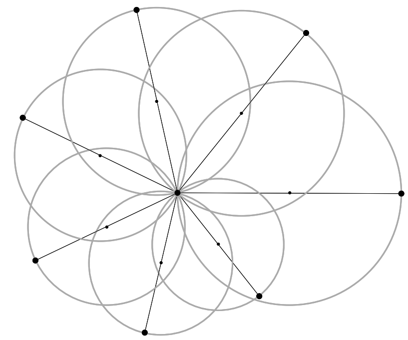

(b) Spiral of Midpoints

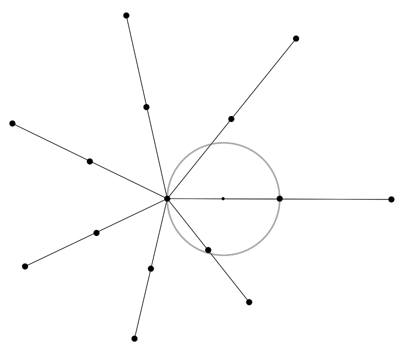

(c) Last Midpoint Encroaches First Segment

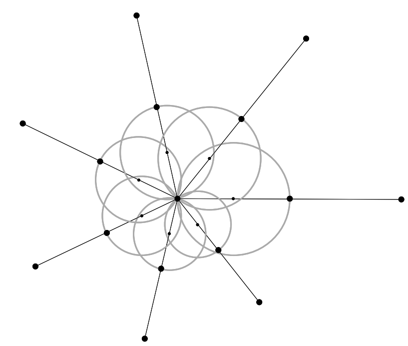

(d) Cycle Repeats

Figure 6: An input which breaks the Delaunay Refinement Algorithm is illustrated. The input is shown in (a); the segments are separated by an angle of $2\pi/7$, and their lengths form a geometric progression. The endpoint of S_6 encroaches S_0, so its midpoint is committed. This midpoint encroaches on S_1, whose midpoint encroaches on S_2, etc, as shown in (b). The newly committed midpoint of S_6 encroaches on the new subsegment on S_0, as shown in (c). Since the input has been effectively scaled in half, the cycle repeats, as shown in (d); the algorithm will fail to terminate for this input.

Refinement Recursion (II): Lastly, we consider an input for which any $\kappa > \pi/6$ will cause an infinite recursion. The input consists of two segments, separated by an angle of $\frac{7\pi}{12} - \epsilon$, with $\sqrt{2}$ the ratio of their lengths, as shown in Figure 8(a). The triangle formed by the endpoints has an angle $\pi/6 + \mathcal{O}(\epsilon)$ and its circumcenter encroaches the longer segment.

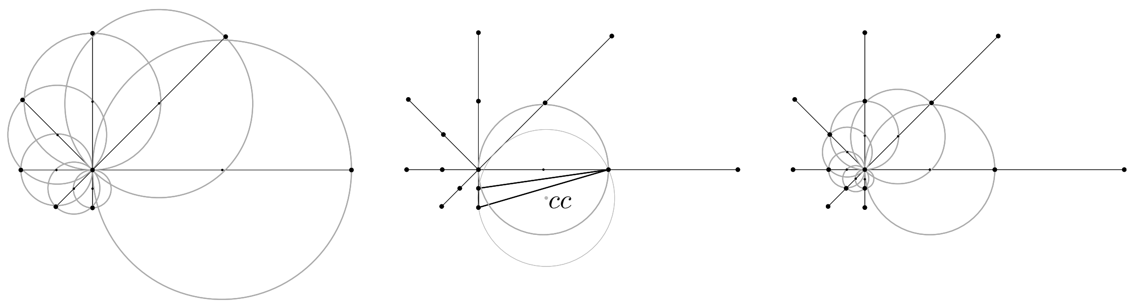

(a) Original Spiral (b) Circumcenter Encroaches (c) Cycle Repeats

Figure 7: An input which breaks the Delaunay Refinement Algorithm is illustrated. In (a) the midpoint of the longest segment has been committed, causing the midpoints of each of the segments to be committed, as each new midpoint encroaches another segment. After this sequence, the algorithm appears to have stabilized. However, as in (b), when the circumcenter, *cc*, of a skinny triangle is considered for addition to the mesh, it encroaches a diametral circle, so is rejected; the segment is split and its midpoint is committed. This causes another sequence of midpoints to be committed, as shown in (c). Since the input has been effectively scaled in half, the cycle repeats, causing an infinite recursion; the algorithm will fail to terminate for this input.

Attempted addition of the circumcenter to remove this triangle causes the long segment to be split, as in Figure 8(b). The picture will then have been effectively scaled by $\frac{1}{\sqrt{2}}$ and "flipped." That is, there will be two segments of length ratio $\sqrt{2}$, meeting at an angle of $\frac{7\pi}{12} - \epsilon$. This process will repeat as shown in Figure 8(c), and the algorithm suffers an infinite recursion.

2.5 A Generic Proof of Good Grading

It is relatively easy to prove a generic theorem concerning termination and optimality of the Delaunay Refinement Algorithm, given an assumption on the behaviour of midpoint-midpoint interactions for certain classes of input.

Some preliminary definitions and results are essential to the exposition. First there is the matter of terminology: if p is a committed point that was the midpoint of a segment, we say this segment is the "parent" segment (or parent subsegment) of p; the "radius" of a segment is half its length, while the radius associated with a midpoint is the radius of its parent segment; any segment derived from a segment $s \in \mathbb{S}$ by splitting is a "subsegment" of (or on) s; segments in $\mathbb{S}$ which share an endpoint are nondisjoint; distinct nondisjoint segments are said to be "adjoining."

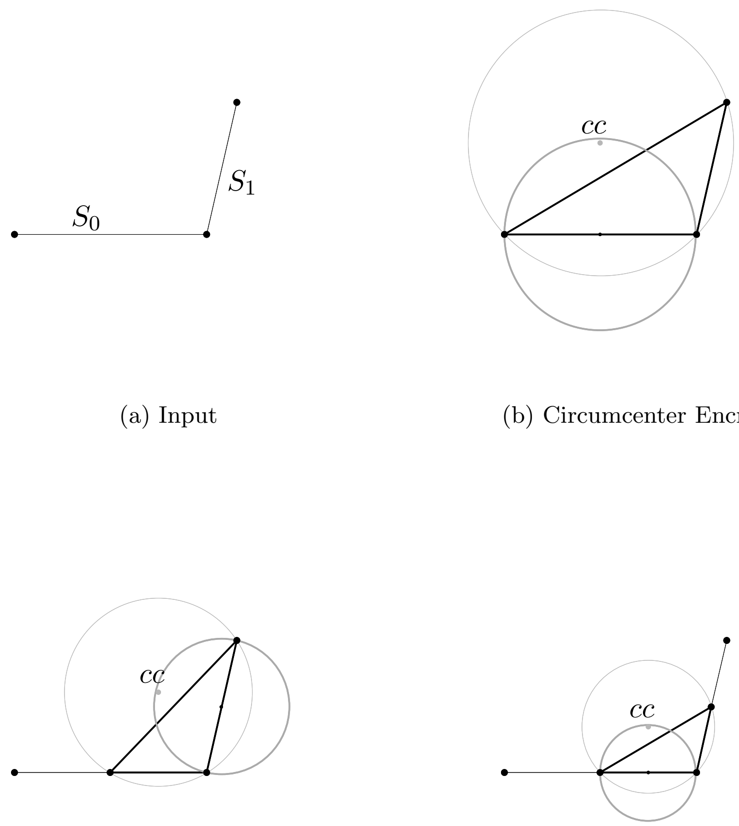

(a) Input

(b) Circumcenter Encroaches

(c) New Triangle Circumcenter Encroaches

(d) Cycle Repeats

Figure 8: An input which breaks the Delaunay Refinement Algorithm is illustrated. The input is shown in (a); the segments are separated by an angle of $\frac{7\pi}{12} - \epsilon$, and the ratio of their lengths is $\sqrt{2}$. The angle opposite S_1 is $\pi/6 + \mathcal{O}(\epsilon)$. If the triangle is to be removed, its circumcenter encroaches on the segment S_0, which is split instead, as shown in (b). The newly committed midpoint creates a triangle which is similar to the one considered for removal previously, but its circumcenter encroaches S_1, as shown in (c). Another midpoint is committed, and the picture has been scaled in half. The cycle will repeat forever, as shown in (d).

Throughout this thesis, we let $|x - y|$ denote the Euclidian distance between points x and y. For a segment S, we let $|S|$ denote the length of the segment. Local feature size is defined in terms of the input, and is the classical definition:

Definition 2.5.1 (Local Feature Size). For a point $x \in \mathbb{R}^2$, the *local feature size* at x, relative to an input PSLG, $(\mathcal{P}, \mathcal{S})$, is

$$\mathrm{lfs}\,(x) = \mathrm{lfs}_{\mathcal{P},\mathcal{S}}\,(x) = \min\left\{r > 0 \,\middle|\, \bar{B}_r(x) \text{ intersects at least two disjoint features of } \mathcal{P} \cup \mathcal{S}.\right\},$$

where $\bar{B}_r(x)$ is the closed ball of radius r centered at x. The local feature size is a Lipschitz function, *i.e.*, $\mathrm{lfs}\,(x) \leq |x - y| + \mathrm{lfs}\,(y)$.

The definition of local feature size of illustrated in Figure 9.

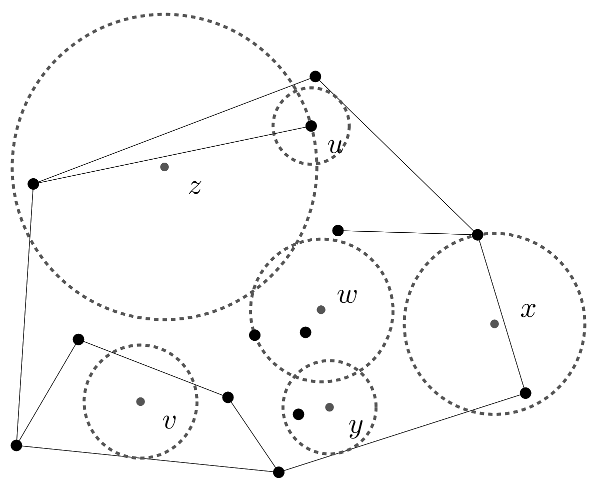

Figure 9: For a number of points in the plane, the local feature size with respect to the given input is shown. About each of the points u, v, w, x, y, z is a circle whose radius is the local feature size of the center point. The point u is an input point.

Considering sequences of segment midpoints will allow us an abstract condition for proving good grading. In particular we are concerned with encroachment sequences.

Definition 2.5.2. A sequence of segment midpoints, $\{p_i\}_{i=0}^{l-1}$ is an *encroachment sequence* if for $i = 1, 2, \ldots, l-1$, at the time p_i is committed, its parent segment is encroached by the already committed point p_{i-1}.

We cite a simple hypothesis that can be applied to analyze the Delaunay Refinement Algorithm. This hypothesis serves as an interface to the generic proof and will be applied in the following sections.

Definition 2.5.3. The **Encroachment Sequence Bound Hypothesis** is defined as follows:

There are constants $1 \leq \beta$, and $1 \leq \rho$ such that if p_{l-1} is a midpoint whose parent segment is encroached by a committed point, p_{l-2}, which is a midpoint on an adjoining input segment, then either

(a) The inequality $\text{lfs}\,(p_{l-1}) \leq \beta \left|p_{l-1} - p_{l-2}\right|$ holds, or

(b) There is an encroachment sequence, $\{p_i\}_{i=0}^{l-1}$, such that (i) the parent segment of p_0 was *not* encroached by a point on an adjoining input segment, and (ii) the inequality

$$\text{lfs}\,(p_{l-1}) \leq (\beta + \rho \frac{\text{lfs}\,(p_0)}{r_0}) \left|p_{l-1} - p_{l-2}\right|,$$

where r_0 is the radius associated with p_0, holds.

The following lemma ensures that if this hypothesis holds, and a point p is committed, then the distance from p to the nearest committed point is bounded below by the local feature size at p (divided by a constant). This establishes termination of the algorithm, and locally bounds below the length of any edge in the output triangulation. Moreover, the cardinality of the set of committed points will be at most a constant times the number of points of *any* mesh which conforms to the input and has no angle less than κ. The proof is very little different from Ruppert's good-grading proof [39].

Lemma 2.5.4 (Generic Good Grading). *Suppose that the Encroachment Sequence Bound Hypothesis holds with constants β, ρ. Then there are positive constants C_1, C_2 depending on β, ρ such that when the Delaunay Refinement Algorithm, operating with a minimum output angle, $\kappa < \arcsin \frac{1}{2\rho\sqrt{2}}$, commits or attempts to commit the point p, and the nearest already committed point to p is q, then*

- *If p is the midpoint of a subsegment, s, of radius r, then*
  - *If q is a committed midpoint on an adjoining input segment which encroached the parent segment of p, then*

$$\text{lfs}\,(p) \leq (\beta + C_1) \left|p - q\right| \leq (\beta + C_1) r.$$

  - *If q is not such a midpoint, then*

$$\text{lfs}\,(p) \leq (1 + \sqrt{2} C_2) \left|p - q\right| \leq (1 + \sqrt{2} C_2) r.$$

- *If p is the circumcenter of a triangle then*

$$\text{lfs}\,(p) \leq C_2 \left|p - q\right|.$$

Proof. We determine the sufficient conditions on the constants. To reduce verbosity, we say that q "provokes" midpoint p if q is an already committed point that encroaches the parent segment of p.

- Suppose p is the midpoint of segment s, of radius r, and q is a midpoint on a nondisjoint input segment. Using the Encroachment Sequence Bound Hypothesis, either $\mathrm{lfs}\,(p) \leq \beta\,|p-q|$, and it suffices to take C_1 nonnegative, or there is an encroachment sequence, beginning with the point p_0, and ending with the points q, p in order such that $\mathrm{lfs}\,(p) \leq (\beta + \rho\frac{\mathrm{lfs}\,(p_0)}{r_0})\,|q-p|$, and such that p_0 was not provoked by a point on an adjoining input segment.

  Using this lemma inductively, since p_0 was not provoked, the second alternative holds, and thus $\mathrm{lfs}\,(p_0) \leq (1+\sqrt{2}C_2)r_0$. So $\frac{\mathrm{lfs}\,(p_0)}{r_0} \leq (1+\sqrt{2}C_2)$. Then $\mathrm{lfs}\,(p) \leq (\beta + \rho(1+\sqrt{2}C_2))\,|p-q|$. Thus any C_1 satisfying the following inequality suffices:

  $$\boxed{\rho(1+\sqrt{2}C_2) \leq C_1.}$$

  Since q provoked p, then $|p-q| \leq r$.

- Suppose p is the midpoint of segment s, of radius r, and q is *not* a midpoint on an adjoining input segment. We consider the cases:
  - Suppose that p was not provoked by q, so q is not in the diametral circle of s. In this case p is being committed because it was encroached by a circumcenter, q', of some triangle T, which was not committed, as shown in Figure 10. Note that s must have been current when the algorithm last attempted to kill T, thus both the endpoints of s were committed at that time. Let $r_{q'}$ be the circumradius of T. Because the circumcircle of T was empty at the time it was considered for removal, the circumradius $r_{q'}$ cannot have been more than the distance from q' to the nearest endpoint of s, which, because q' encroaches the segment, is no greater than $\sqrt{2}r$. Inductively $\mathrm{lfs}\,(q') \leq C_2 r_{q'} \leq \sqrt{2}C_2 r$. Using the Lipschitz property, $\mathrm{lfs}\,(p) \leq |p-q'| + \mathrm{lfs}\,(q') \leq (1+\sqrt{2}C_2)r$. Noting that $|p-q| = r$, the result holds.
  - Otherwise suppose that p was provoked by q. Then q is an input point or a midpoint on a disjoint input feature. By definition, $\mathrm{lfs}\,(p) \leq |p-q|$. Since we assume C_2 to be positive, then $\mathrm{lfs}\,(p) \leq (1+\sqrt{2}C_2)\,|p-q|$, which suffices since $|p-q| < r$.

- If p is a circumcenter of a skinny triangle of circumradius r, then let a, b be the vertices of the shortest edge. Note that $|p-q| = r$, since by definition of the quality operation, the circumcircle of the triangle is empty when p is to be committed. If a, b are both input points, then they are disjoint and by definition $\mathrm{lfs}\,(p) \leq r$, so it suffices to take $1 \leq C_2$. Otherwise let b be the most recently committed of the two points. By definition of the quality operation and the sine rule, $|a-b| \leq 2\sin\kappa\, r$.

  Consider, inductively, the time when b was committed: since a was already committed,

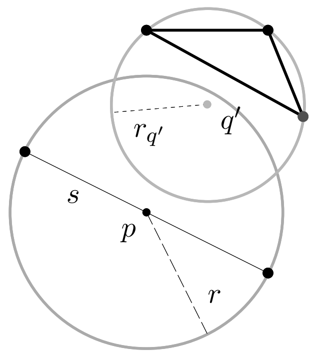

Figure 10: The "yield" case is illustrated; s is encroached by q', a circumcenter of a triangle. Rather than commit q', the segment is split, causing the midpoint p to be committed. Because the endpoints of s cannot have been inside the circumcircle of the triangle, $r_{q'}$ cannot exceed $\sqrt{2}r$.

$\operatorname{lfs}(b) \leq C|b-a|$ for some constant C. Considering all the cases in the lemma, $C = \max\left\{\beta + C_1, 1+\sqrt{2}C_2, C_2\right\} = \beta + C_1$ works. Using the Lipschitz condition

$$\operatorname{lfs}(p) \leq r + \operatorname{lfs}(b) \leq r + C|a-b| \leq (1+2C\sin\kappa)r,$$

it suffices to ensure that $1 + 2C\sin\kappa \leq C_2$.

In all it suffices to ensure that

$$\boxed{1 + 2(\beta + C_1)\sin\kappa \leq C_2.}$$

Accumulating the boxed constraints, it suffices to take

$$\begin{aligned} \kappa &< \arcsin\frac{1}{2\rho\sqrt{2}}, \\ C_1 &= \rho(1+\sqrt{2}C_2) \\ C_2 &= \frac{1+2(\beta+\rho)\sin\kappa}{1-2\rho\sqrt{2}\sin\kappa}. \end{aligned}$$

□

The lemma shows that the distance between committed points is bounded from below by the local feature size. It is also possible to show that the distance will be bounded from above, *i.e.*, if the local feature size of a committed point is small, eventually the algorithm will have committed another point nearby. The result is used in bounding the cardinality of the final point set, in Corollary 2.5.7.

Definition 2.5.5 (Distinct Input Distance). For a point $x \in \mathbb{R}^2$, the *distinct input distance* at x, relative to an input PSLG, $(\mathcal{P}, \mathcal{S})$, denoted by $\mathrm{d}_1(x)$, is the distance from x

to the nearest member of the input which is distinct from x. The function $\mathrm{d}_1(\cdot)$ is not a distance function at all; in particular it is not continuous.

Thus if x is not a point of $\mathcal{P}$ and is not on any segment of $\mathcal{S}$, $\mathrm{d}_1(x)$ is the distance from x to the nearest member of either of these sets. If x is a point of $\mathcal{P}$, $\mathrm{d}_1(x)$ is the distance to the nearest point of $\mathcal{P} \setminus \{x\}$ or segment of $\mathcal{S}$ which does not have x as an endpoint. If x is on segment $s \in \mathcal{S}$, but not an endpoint of s, then $\mathrm{d}_1(x)$ is the distance to the nearest point of $\mathcal{P}$, or segment of $\mathcal{S} \setminus \{s\}$.

Clearly $\mathrm{d}_1(x) \leq \mathrm{lfs}(x)$. Moreover, if $\mathcal{C}$ is a circle centered at x with radius $\mathrm{d}_1(x)$, then there is a point of $\mathcal{P}$ on $\mathcal{C}$, or the projection of x onto a segment of $\mathcal{S}$ is on $\mathcal{C}$.

Lemma 2.5.6. *Given an input PSLG $(\mathcal{P}, \mathcal{S})$, let $(\mathcal{P}', \mathcal{S}')$ be a PSLG such that $\mathcal{P} \subseteq \mathcal{P}'$, every $s \in \mathcal{S}$ is the union of segments in $\mathcal{S}'$, and every segment of $\mathcal{S}'$ has an empty diametral circle with respect to the set $\mathcal{P}'$.*

Then for every $p \in \mathcal{P}'$ there is some $q \in \mathcal{P}'$ such that

$$|p - q| \leq \sqrt{2}\,\mathrm{d}_1(p),$$

where $\mathrm{d}_1(\cdot)$ is with respect to $(\mathcal{P}, \mathcal{S})$.

Proof. Let $\mathcal{C}$ be the circle centered at p of radius $\mathrm{d}_1(p)$. Then either there is a point of $\mathcal{P}$ on $\mathcal{C}$, in which case let this point be q; otherwise the projection of p onto some segment distinct from any input containing p is on $\mathcal{C}$. Let this projection be p'. The point p' is on some segment $s' \in \mathcal{S}'$. Since the segment s' does not contain the point p, it must be the case that its endpoints are no farther from p' than $\mathrm{d}_1(p)$, in which case they are no farther than $\sqrt{2}\,\mathrm{d}_1(p)$ from p. □

We briefly note that by substituting the empty diametral circle condition in the previous lemma with an empty diametral lens condition [41], we can claim $|p - q| \leq 2\mathrm{d}_1(p)$; if we substitute it with a minimum angle bound of α for $(\mathcal{P}', \mathcal{S}')$ a triangulation, we can claim $|p - q| \leq \frac{1}{\sin\alpha}\mathrm{d}_1(p)$.

Corollary 2.5.7. *Suppose the hypotheses of the lemma hold for the input $\mathcal{P}, \mathcal{S}$ to the Delaunay Refinement Algorithm. Then the algorithm terminates with output $(\mathcal{P}', \mathcal{S}', \mathcal{T}')$ satisfying*

$$\left|\mathcal{P}' \setminus \mathcal{P}\right| \leq \frac{2\left(2(\beta + C_1) + 3\right)^2}{\pi} \int_\Omega \frac{1}{\mathrm{lfs}^2(x)}\mathrm{d}x = \mathcal{O}\left(\left(\frac{\beta + \rho}{1 - 2\rho\sqrt{2}\sin\kappa}\right)^2 \int_\Omega \frac{1}{\mathrm{lfs}^2(x)}\mathrm{d}x\right). \quad (1)$$

Moreover, if there is a triangulation, $\mathcal{T}''$ on a set of points $\mathcal{P}''$ that conforms to the input $(\mathcal{P}, \mathcal{S})$ and has minimum angle κ then

$$\left|\mathcal{P}'\right| = \mathcal{O}\left(\frac{1}{\kappa}\left(\frac{\beta + \rho}{1 - 2\rho\sqrt{2}\sin\kappa}\right)^2\right)\left|\mathcal{P}''\right|. \quad (2)$$

Proof. The termination result, and equation 1 will follow as in Ruppert [39]. First note that $\max\left\{\beta + C_1, 1 + \sqrt{2}C_2, C_2\right\} = \beta + C_1$, so if p is committed before q then $\operatorname{lfs}(q) \leq (\beta + C_1)|p-q|$. Thus, by the Lipschitz condition, if p, q are any two committed points, without knowing which was committed first, we can conclude that $\operatorname{lfs}(p) \leq (\beta + C_1 + 1)|p-q|$.

Suppose the algorithm runs forever. Since each major operation results in a point being committed, it must be the case that an infinite number of points are committed. Clearly no point is committed twice, and, as already noted, all points are added in the closed, bounded, and therefore compact, set, Ω. So there must be an infinite sequence converging to an accumulation point, point P. Let $\{P_i\}_{i=0}^{\infty}$ be such an infinite sequence. Then

$$\operatorname{lfs}(P) = \lim_i \operatorname{lfs}(P_i) \leq \lim_i (\beta + C_1 + 1)|P_i - P_{i-1}| = 0,$$

contradicting $\operatorname{lfs}(x) > 0$ for all $x \in \Omega$. It follows that the algorithm must terminate.

Next, as in Ruppert [39], for each $p \in \mathcal{P}' \setminus \mathcal{P}$, let B_p be an open ball centered at p, of radius $r_p = \frac{\operatorname{lfs}(p)}{2(\beta + C_1 + 1)}$. Then if q is another committed point $r_p \leq \frac{|p-q|}{2}$, so the balls are disjoint, and contain no other point of $\mathcal{P}'$. Now we claim that at least half of each B_p is in Ω. If $p \in \partial\Omega$, then, since it is not in $\mathcal{P}$, it is not a corner of $\partial\Omega$. Because the balls B_p do not intersect, the ball does not contain the endpoints of the edge that p is on. Then half of B_p is outside Ω. We call the other half of B_p the "inside" half, and show that it really is fully inside of Ω.

By Lemma 2.5.6, there is some committed point q such that $|p-q| \leq \sqrt{2}\,\mathrm{d}_1(p)$. By definition of r_p, it is the case that $r_p \leq \frac{|p-q|}{2} \leq \frac{1}{\sqrt{2}}\mathrm{d}_1(p) < \mathrm{d}_1(p)$. Thus there is no segment intersecting B_p other than the one containing p. Thus half of B_p is inside Ω.

If $p \notin \partial\Omega$, *i.e.*, p is a circumcenter or is on an internal segment, then we can repeat the above argument to show that all of B_p is contained in Ω. Then, as in Ruppert's paper, we bound

$$\begin{aligned}\int_\Omega \frac{1}{\operatorname{lfs}^2(x)}\mathrm{d}x &\geq \frac{1}{2}\sum_{p\in\mathcal{P}'\setminus\mathcal{P}}\int_B \frac{1}{\operatorname{lfs}^2(x)}\mathrm{d}x \geq \frac{1}{2}\sum_{p\in\mathcal{P}'\setminus\mathcal{P}}\frac{\pi r_p^2}{(\operatorname{lfs}(p)+r_p)^2}\\ &\geq \frac{\pi}{2\left(2(\beta + C_1 + 1)+1\right)^2}\left|\mathcal{P}'\setminus\mathcal{P}\right|.\end{aligned}$$

Algebraic cancellation reveals that $\beta + C_1 \leq \frac{\beta + \rho(1+\sqrt{2})}{1 - 2\rho\sqrt{2}\sin\kappa}$, which establishes equation 1.

For equation 2, we use Theorem 11 of Mitchell [31], which states that if a triangulation of the points $\mathcal{P}''$ respects the input and has minimum output angle κ, then

$$\left|\mathcal{P}''\right|\left(\frac{21.5}{\kappa} + 11.9\right) \geq \int_\Omega \frac{1}{\operatorname{lfs}^2(x)}\mathrm{d}x.$$

This establishes the optimality of $|\mathcal{P}'\setminus\mathcal{P}|$. We then use $\mathcal{P} \subseteq \mathcal{P}''$ to bound $|\mathcal{P}|$. $\square$

It is then sufficient to show that the Encroachment Sequence Bound Hypothesis holds when the algorithm is given a specific input. Moreover, if the ρ of the hypothesis can be identified, then the minimum output angle can be set accordingly. In the following chapter, it is shown that the hypothesis can be established for certain classes of input.

CHAPTER III

THE DELAUNAY REFINEMENT ALGORITHM

"There is no royal road to geometry." –*Euclid*

In this chapter, the Encroachment Sequence Bound Hypothesis is established for two different classes of input, namely input with θ^* as small as $\pi/4$, subject to a "cosine condition," and input with arbitrarily small θ^*, with an edge-length condition. In Chapter 4, procedures for putting arbitrary input into the latter form by the addition of a few "augmenting" points are discussed.

3.1 Encroachment Basics

The following simple claims regarding encroachment will be used ubiquitously.

Claim 3.1.1. Let (a, b) be a subsegment of an input segment which has endpoint x. Let $|x - a| < |x - b|$. Then either $x = a$ or $|a - b| \leq |x - a|$.

Proof. See Figure 11. Suppose that a is distinct from x. Then a must be a midpoint of some subsegment of radius at least $|a - b|$. However, $|x - a|$ is at least this radius, *i.e.*, $|a - b| \leq |x - a|$. □

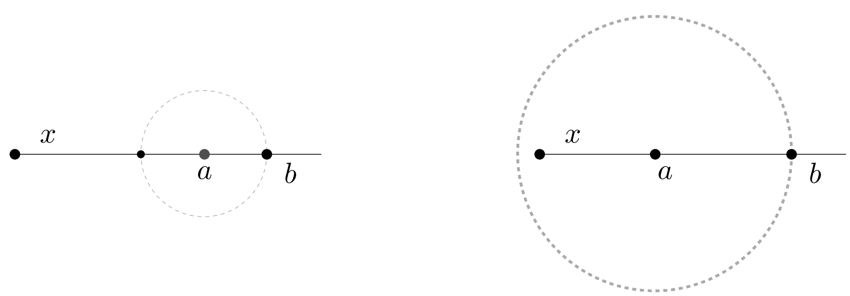

(a) Possible Case (b) Impossible Case

Figure 11: The argument of Claim 3.1.1 is shown. When (a, b) is a subsegment on an input segment with endpoint x, such that $0 < |x - a| < |x - b|$, we show that $|a - b| \leq |x - a|$, as shown in (a). The case illustrated in (b) is impossible since a would have to be the midpoint of a subsegment which actually contained the endpoint x. In both figures we show the diametral circle of the subsegment of which a is the center.

Claim 3.1.2. Let (a, b) be a subsegment of an input segment which has endpoint x. Suppose p is a point on an input segment which shares the endpoint x that encroaches on the diametral circle of (a, b). Assume that $|x - a| < |x - b|$, and let θ be the angle between the two input segments. Then $|x - a| < |x - p| \cos\theta$, and $\theta < \frac{\pi}{2}$. Also we can claim $|x - p| \leq |x - b|$. Moreover, if m is the midpoint of (a, b), and r is its radius, then $|x - m| \sin\theta < r$.

Proof. The gist of this claim is shown in Figure 12. Since p encroaches the diametral circle of (a, b), then so does it's projection onto the line containing the segment, p'. Thus $|x - a| < |x - p'|$. But $|x - p'| = |x - p| \cos\theta$. This implies that $\cos\theta$ is strictly positive, so $\theta < \frac{\pi}{2}$. Since the circle centered at x of radius $|x - b|$ contains the diametral circle of (a, b), it contains the point p, so then $|x - p| \leq |x - b|$.

Since p encroaches (a, b), the radius of the diametral circle must be at least the distance from m to the line segment containing p, which is $|x - m| \sin\theta$. □

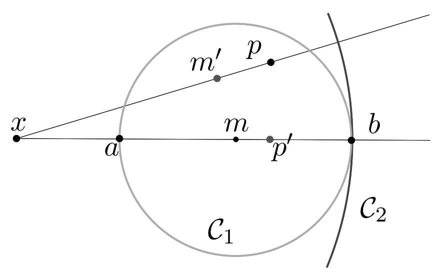

Figure 12: The argument of Claim 3.1.2 is shown. Letting $\theta = \angle axp$, by definition of sine and cosine, $|m - m'| = |x - m| \sin\theta$, and $|x - p'| = |x - p| \cos\theta$, where m', p' are projections of the points m, p, onto the opposing segment. Thus $|x - a| \leq |x - p| \cos\theta$, and the radius of the circle is at least $|x - m| \sin\theta$. Part of the circle $\mathcal{C}_2$ centered at x of radius $|x - b|$ is shown. Since $\mathcal{C}_2$ contains $\mathcal{C}_1$, the diametral of (a, b), then $|x - p| \leq |x - b|$.

Lemma 3.1.3. *Let $(x, y), (x, z)$ be two input segments sharing a common endpoint x, and subtending an angle of θ, with $\arcsin \frac{1}{3} \leq \theta < \pi/2$. Consider a subsegment, (a, b), on (x, y), and assume that $0 < |x - a| < |x - b|$. Then no point on (x, z) encroaches (a, b).*

As a consequence, if there is a subsegment (a, b) on (x, y) which is encroached by a point on (x, z), then $a = x$. Moreover, if m is the midpoint of a subsegment on (x, y) which was encroached by a point on (x, z), then $|x - m| \leq |y - m|$.

Proof. Suppose some point p on (x, z) encroaches the segment. Since $0 < |x - a|$, by Claim 3.1.1, $|a - b| \leq |a - x|$, and so $r \leq |x - a|/2$, where r is the radius of (a, b). By Claim 3.1.2, r is large compared to the distance from the midpoint of (a, b) to x, *i.e.*, $(|x - a| + r) \sin\theta < r$, where $|x - a| + r$ is the distance from x to the midpoint of the subsegment. Putting these two inequalities together we have

$$\frac{|x-a|\sin\theta}{1-\sin\theta} < r \leq \frac{|x-a|}{2},$$

which requires that $2\sin\theta < 1-\sin\theta$, or $\sin\theta < \frac{1}{3}$, contradicting the assumption that $\arcsin\frac{1}{3} \leq \theta$. $\square$

The following lemma is invaluable to our proofs.

Lemma 3.1.4. *Given two rays, R and R' from a point x with angle θ between them, suppose there is a ball of radius r with center p on ray R such that the ball does not contain x but does contain a point q of R'. Then if $\pi/4 \leq \theta < \pi/2$,*

$$\frac{|q-x|}{|p-x|} \leq \frac{|q-x|}{r} < \frac{|q-x|}{|p-q|} \leq 2\cos\theta.$$

If $0 < \theta < \pi/4$, then only the inequality

$$\frac{|q-x|}{|p-x|} < 2\cos\theta$$

can be asserted.

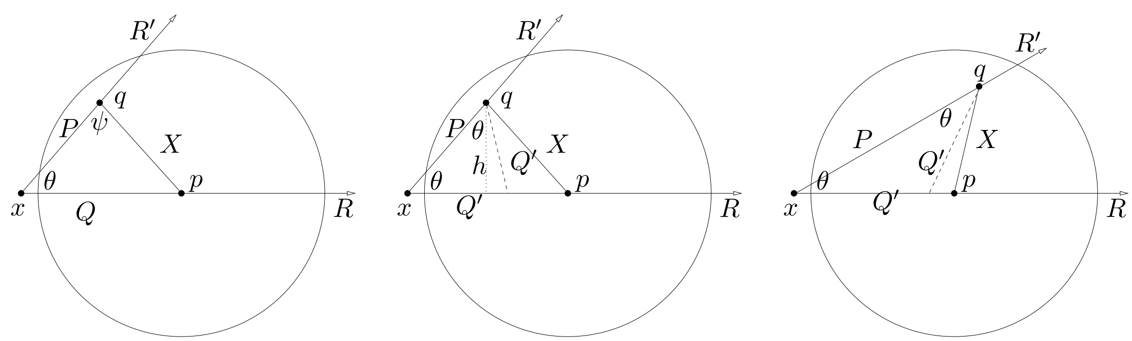

(a) as stated (b) bounded by the isosceles (c) the small angle case

Figure 13: Proof of Lemma 3.1.4; The lemma as stated is shown in (a). It can be shown that $\theta \leq \psi$, so we may draw the isosceles triangle, as in (b) with base angle θ to get the desired bound. The altitude is also drawn in (b), and both triangle legs will be to its right in the order shown. The case where $\theta < \pi/4$ is shown in (c); in this case the ordering of the legs relative to the altitude is not fixed, and only a weaker result is obtained.

Proof. Letting P, Q, X be as in Figure 13, first note that $X < r \leq Q$ because x is not inside the ball (which has radius r), but q is; thus for the "large" angle case, it suffices to show only that $P/X \leq 2\cos\theta$. Using the sine identity, we find that $\sin\theta \leq \sin\psi$, implying that

$\theta \leq \psi$. We can then draw an isosceles triangle of base angle θ and base (x, q), with side lengths Q'. In the case where $\pi/4 \leq \theta$, the apex angle of this isosceles triangle will be acute, so the altitude h, the leg Q' and the leg X are ordered left to right as shown in Figure 13(b), and thus $Q' \leq X$. Using the cosine relation it is easy to show that $P/Q' = 2\cos\theta$ and thus $P/X \leq 2\cos\theta$, as desired.

In the "small" angle case, the apex angle may be obtuse. However, since $\theta < \psi$, we have $Q' < Q$. Thus $2\cos\theta = P/Q' > P/Q$, as desired.

□

3.2 Input with Restricted Cosines

In this section we show that the Delaunay Refinement Algorithm may be applied to input with angles as low as $\pi/4$, although with a single added restriction which is simple to check. The drop from $\pi/3$ is not without a concomitant loss in output quality. In light of Lemma 2.5.4, and assuming a proper lower bound on input angles, it suffices to show that encroachment sequences are well-behaved. We start with the assumptions on our input.

3.2.1 Assumptions on the Input

Assumption 3.2.1. In addition to Assumption 2.2.1 we assume that the input satisfies the following constraints.

(a) The lower bound on input angles, θ^*, is in the range $\pi/4 \leq \theta^* \leq \pi/3$.

(b) There is some $1 \leq d \leq \min\left\{7, \lfloor \frac{2\pi}{\theta^*} \rfloor\right\}$, such that no input point has degree greater than d, where by "degree" of a vertex we mean the number of incident edges.

(c) If an input point has $k \leq d$ segments emanating from it, and $\{S_i\}_{i=0}^{k-1}$ are the k angles between the segments, and every θ_i is acute, then

$$\prod_{i=0}^{k-1} 2\cos\theta_i \leq 2. \tag{3}$$

The condition of item (c) is automatic for $k < 4$. It is elementary to prove that it holds for $k = 4$, and Corollary 3.2.3, following, shows that it holds automatically for $k = 5, 6$. The spiral recursion of Figure 6 shows that it needn't hold for $k = 7$, and it cannot be satisfied for $k = 8$ (which would require that $\theta_i = \theta^* = \pi/4$.). Note that the input shown in Figure 7 *does* conform to these assumptions.

Lemma 3.2.2. *Given the collection* $\{\theta_i\}_{i=0}^{k-1}$ *with* $\theta_i \in [0, \pi/2]$, *then*

$$\prod_{i=0}^{k-1} \cos\theta_i \leq \cos^k\left(\frac{\sum_{i=0}^{k-1}\theta_i}{k}\right).$$

Proof. First note that due to the restriction on the θ_i, that $\cos\theta_i$ is nonnegative, so the arithmetic mean of the collection $\{\cos\theta_i\}_{i=0}^{k-1}$ exceeds its geometric mean, that is

$$\sqrt[k]{\prod_{i=0}^{k-1}\cos\theta_i} \leq \frac{\sum_{i=0}^{k-1}\cos\theta_i}{k},$$

and so it suffices to prove that

$$\frac{\sum_{i=0}^{k-1}\cos\theta_i}{k} \leq \cos\left(\frac{\sum_{i=0}^{k-1}\theta_i}{k}\right).$$

This follows, however, from Jensen's Inequality, since the cosine is concave on $[0,\pi/2]$. □

Corollary 3.2.3. *Given a collection of k angles, $\{\theta_i\}_{i=0}^{k-1}$, which sum to 2π, and with $\pi/4 \leq \theta_i < \pi/2$ for $0 \leq i < k$, then if $k=5$ or $k=6$, the following relation holds:*

$$\prod_{i=0}^{k-1} 2\cos\theta_i \leq 2.$$

Proof. By the lemma it suffices to show that $\left(2\cos\frac{2\pi}{k}\right)^k \leq 2$. However, $\cos\frac{2\pi}{k} \leq \frac{1}{2}$ for $k=5,6$. □

3.2.2 Establishing the Encroachment Sequence Bound Hypothesis

We now set out to show that labouring under these assumptions, the Delaunay Refinement Algorithm will not cause infinite "ping-pong" cascades of the kind shown in Figure 5, or infinite inward "spirals," as outlined in Figure 6.

Lemma 3.2.4 (Ping Pong). *Let $(x,y),(x,z)$ be two input segments sharing a common endpoint x, and subtending an angle of θ, with $\pi/4 \leq \theta < \pi/2$. Suppose a committed midpoint, p, on (x,z) encroaches a subsegment on (x,y). Then the midpoint of this segment, m, does not encroach on any subsegment along (x,z).*

Proof. Let p be the point on (x,z) that encroached the subsegment on (x,y), as in Figure 14. Let $L=|x-p|$, let $M=|x-m|$, and let r be the radius of the segment on (x,y) that was encroached. By Lemma 3.1.4,

$$L/M \leq L/r < 2\cos\theta. \tag{4}$$

Suppose there is a subsegment on (x,z) which is encroached by m. By Lemma 3.1.3, it must be the case that x is one of the endpoints of the subsegment, since $\arcsin\frac{1}{3} < \pi/4 \leq \theta$. Let r' be the radius of this subsegment. Since this subsegment does not contain

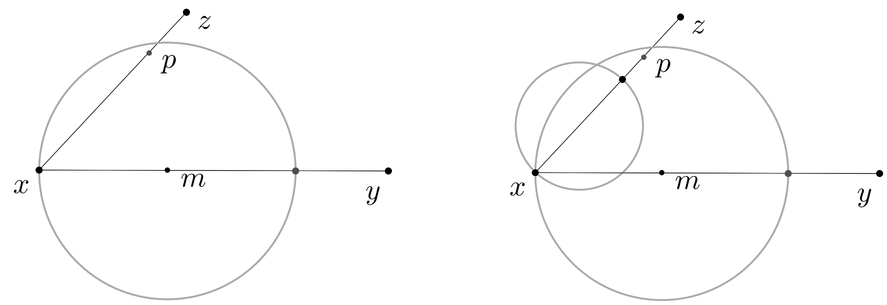

(a) "Ping" (b) No "Pong"

Figure 14: The argument of Lemma 3.2.4 is shown. The point p encroaches a subsegment along (x, y), causing its midpoint m to be committed, as shown in (a). Because the angle $\theta = \angle yxz$ is assumed to be large, by Lemma 3.1.3, it must be the case that the subsegment has x as one of its endpoints, as shown. Then the point m cannot encroach any subsegment which will ever be current along (x, z), as shown in (b). Again because of the angle condition, this subsegment would have to have endpoint x, but cannot have diameter greater than $|x - p|$. This leads to a contradiction on the cosine of θ.

the committed point p, it must be the case that $2r' \leq |x - p| = L$, or $r' \leq \frac{L}{2}$. Using Lemma 3.1.4 we have $M/r' < 2\cos\theta$. Combined with equation 4 we have

$$\frac{L}{2\cos\theta} < M < 2r'\cos\theta \leq L\cos\theta.$$

Thus $\frac{1}{2} < \cos^2\theta$, which is impossible for $\theta \in [\pi/4, \pi/2)$. $\square$

We now turn to the proof of Lemma 3.2.5 where the key geometric arguments required to establish the Encroachment Sequence Bound Hypothesis are developed.

Lemma 3.2.5. *Suppose that the input to the Delaunay Refinement Algorithm conforms to Assumption 3.2.1. Let $\{S_i\}_{i=0}^{k-1}$ be a collection of $k \leq 7$ consecutively numbered (either clockwise or anticlockwise) input segments sharing a common endpoint, x. Let θ_i be the angle between S_i and S_{i+1}, where S_k is read as S_0.*

Furthermore, let $\{p_j\}_{j=0}^{l-1}$ be an encroachment sequence such that each p_j is a midpoint on some segment from the collection $\{S_i\}_{i=0}^{k-1}$ Assume p_0 is on S_0, and p_1 is on S_1. Let r_j be the radius of the subsegment of which p_j was the midpoint. Then the following hold:

(a) $l \leq k$. Moreover, p_j is on S_j for $0 \leq j < l$.

(b) $|x - p_0| \prod_{i=0}^{j-1} \frac{1}{2\cos\theta_i} \leq |p_j - p_{j-1}| < r_j = |x - p_j|$ for $1 \leq j < l$.

(c) For $1 \leq j < l$,

$$\mathrm{lfs}\,(p_j) \leq \left(\frac{1}{\sin\theta_{j-1}} + (1 + \frac{\mathrm{lfs}\,(p_0)}{r_0}) \prod_{i=0}^{j-1} 2\cos\theta_i \right) |p_j - p_{j-1}|\,.$$

(d) $|p_{l-1} - p_0| \leq \left(1 + 2\cos\theta_{l-2} + \prod_{i=0}^{l-2} 2\cos\theta_i\right) |p_{l-1} - p_{l-2}|$.

Proof. First note that by the angle condition, $\pi/4 \leq \theta_i$, the angle between two non-consecutive input segments must be at least $\pi/2$; so a point on one of them will never encroach a subsegment along another. Thus it is the general case to assume that p_1 was on a segment consecutive to S_0, which, by convention, we have named S_1. By Lemma 3.1.3, because $\arcsin \frac{1}{3} < \theta_i$ for all i, it is the case that x is an endpoint of every encroached segment, and so $r_j = |x - p_j|$, for $j \geq 1$. Note, however, that we can only claim that $r_0 \leq |x - p_0|$.

We establish item (b) and item (a) together, inductively; By assumption, p_j is on S_j for $j = 0, 1$. Moreover, since p_0 encroached on a subsegment on S_1, causing the midpoint p_1 to be committed, by Lemma 3.1.4

$$\frac{|x - p_0|}{2\cos\theta_0} \leq |p_1 - p_0| < r_1 = |x - p_1|.$$

Now the induction step: suppose that for all j up to $m < l - 1$ that P_j is on $S_{j \mod k}$, and that

$$|x - p_0| \prod_{i=0}^{j-1} \frac{1}{2\cos\theta_{i \mod k}} \leq |p_j - p_{j-1}| < r_j = |x - p_j|.$$

We show these hold for $j = m+1$. By Lemma 3.2.4, p_{m+1} cannot be on the same segment as p_{m-1}, nor, as per above, can it be on an input segment not consecutive to the one containing p_m, that is $S_{m \mod k}$. Thus it must be on $S_{m+1 \mod k}$. By Lemma 3.1.4

$$\frac{|x - p_m|}{2\cos\theta_{m \mod k}} \leq |p_{m+1} - p_m| < r_{m+1} = |x - p_{m+1}|.$$

Using the inductive result, this gives immediately that

$$|x - p_0| \prod_{i=0}^{m} \frac{1}{2\cos\theta_{i \mod k}} \leq |p_{m+1} - p_m| < r_{m+1} = |x - p_{m+1}|.$$

So we need only show that $l \leq k$ to show that item (b) and item (a) hold. Suppose to the contrary that some point p_k, the midpoint of a subsegment along S_0 is committed due to the subsegment being encroached. (Note that since $\theta_i < \pi/2$, it must be the case that $k > 4$.) By Lemma 3.1.3, one endpoint of the subsegment must be x. It must be, then, that $r_k \leq \frac{|x - p_0|}{2}$ Putting this together with the above result we have

$$|x - p_0| \prod_{i=0}^{k-1} \frac{1}{2\cos\theta_i} < r_k \leq \frac{|x - p_0|}{2}.$$

Thus $2 < \prod_{i=0}^{k-1} 2\cos\theta_i$, which contradictions equation 3 in Assumption 3.2.1.

For item (c), we use the Lipschitz property of $\text{lfs}(\cdot)$, the initial estimate at p_0, and the above results. Let $C = \frac{\text{lfs}(p_0)}{r_0}$. Then

$$\begin{aligned}
\text{lfs}(p_j) &\leq |p_j - p_0| + \text{lfs}(p_0) \\
&\leq |x - p_j| + |x - p_0| + Cr_0 \\
&\leq |x - p_j| + (1 + C)|x - p_0| \\
&\leq |x - p_j| + (1 + C)|p_j - p_{j-1}| \prod_{i=0}^{j-1} 2\cos\theta_i.
\end{aligned}$$

The result then follows from the sine estimate of Claim 3.1.2, $|x - p_j| \sin\theta_{j-1} \leq |p_j - p_{j-1}|$.

To establish item (d), we start with the triangle inequality:

$$|p_{l-1} - p_0| \leq |p_{l-1} - p_{l-2}| + |p_{l-2} - x| + |x - p_0|.$$

By item (b), $|x - p_0| \leq \prod_{i=0}^{l-3} 2\cos\theta_i \, |x - p_{l-2}|$. By use of Lemma 3.1.4, $|p_{l-2} - x| \leq 2\cos\theta_{l-2}\, |p_{l-1} - p_{l-2}|$. Thus

$$|p_{l-1} - p_0| \leq \left(1 + 2\cos\theta_{l-2} + \prod_{i=0}^{l-2} 2\cos\theta_i\right) |p_{l-1} - p_{l-2}|.$$

□

The previous two lemmata allow us to establish the Encroachment Sequence Bound Hypothesis; Theorem 3.2.7 then follows from Lemma 2.5.4.

Corollary 3.2.6. *Suppose that the input to the Delaunay Refinement Algorithm conforms to Assumption 3.2.1. Let* $\rho = (2\cos\theta^*)^{d-1}, \beta = 3 + 2\rho$. *If* p_{l-1} *is a midpoint whose parent segment is encroached by a committed midpoint* p_{l-2} *on an adjoining input segment, then either*

(a) The inequality $\text{lfs}(p_{l-1}) \leq \beta |p_{l-1} - p_{l-2}|$ *holds, or*

(b) There is an encroachment sequence, $\{p_i\}_{i=0}^{l-1}$, *such that (i) the parent segment of* p_0 *was* not *encroached by a committed point on an adjoining input segment, and (ii) the inequality*

$$\text{lfs}(p_{l-1}) \leq (\beta + \rho\frac{\text{lfs}(p_0)}{r_0}) |p_{l-1} - p_{l-2}|,$$

where r_0 *is the radius associated with* p_0, *holds.*

Proof. Let x be the input point shared by the input segments containing p_{l-2}, p_{l-1}. Let $\{p_i\}_{i=0}^{l-1}$ be a maximal encroachment sequence ending with p_{l-2}, p_{l-1} such that each p_i is on an input segment with endpoint x. By Lemma 3.2.5, and since angles are bounded by θ^*,

$$\text{lfs}(p_{l-1}) \leq \left(\frac{1}{\sin\theta^*} + \rho + \rho\frac{\text{lfs}(p_0)}{r_0}\right) |p_j - p_{j-1}| \leq \left(\beta + \rho\frac{\text{lfs}(p_0)}{r_0}\right) |p_j - p_{j-1}|.$$

If the parent segment of p_0 was not encroached by a point on an adjoining input segment, then the second alternative holds. So suppose there was some encroaching point on an adjoining input segment. By item (d) of Lemma 3.2.5, and using the bounds on input angles,

$$|p_{l-1} - p_0| \leq (3 + \rho) |p_{l-1} - p_{l-2}| .$$

Let $y \neq x$ be the input point shared by the segment containing p_0 and the encroaching point. By Lemma 3.1.3, $|p_0 - y| \leq |p_0 - x|$. The definition of local feature size gives, $\mathrm{lfs}(p_0) \leq |p_0 - x| \vee |p_0 - y| = |p_0 - x|$. Then by item (b) of Lemma 3.2.5, $\mathrm{lfs}(p_0) \leq \rho |p_{l-1} - p_{l-2}|$. Using the Lipschitz property,

$$\mathrm{lfs}(p_{l-1}) \leq |p_{l-1} - p_0| + \mathrm{lfs}(p_0) \leq (3 + 2\rho) |p_{l-1} - p_{l-2}| = \beta |p_{l-1} - p_{l-2}| ,$$

and the first alternative holds. □

With a little bit more work, one could establish a slightly smaller value for β, namely $3 + \rho$. Throughout this work we prefer slightly shorter proofs to slightly smaller constants. Asymptotically the larger constants will make no difference.

Theorem 3.2.7. *Suppose that the input to the Delaunay Refinement Algorithm conforms to Assumption 3.2.1. Let* $\rho = (2 \cos \theta^*)^{d-1}$. *Then the Delaunay Refinement Algorithm applied with* $\kappa < \arcsin \frac{1}{2\rho\sqrt{2}}$ *terminates with good grading. Moreover, the cardinality and optimality results of Corollary 2.5.7 apply with the given* ρ*, and with* $\beta = 3 + 2\rho$.

Note that if $(2 \cos \theta^*)^{d-1} = 1$, then it must be that $\theta^* = \pi/3$, and so our analysis is redundant, *i.e.*, the classical analysis applies [30].

Note that $\rho \leq 8$, by the conditions on θ^* and d. This bound is tight, *i.e.*, some input, including the input of Figure 7 (which is covered by the theorem), will exhibit $\rho = 8$. In this case the theorem only guarantees the algorithm will terminate for $\kappa < \arcsin \frac{1}{16\sqrt{2}} \approx 2.53°$. We have already seen that if $\kappa \geq \arcsin \frac{1}{\sqrt{73}} \approx 6.72°$, the algorithm will not terminate for the input of Figure 7.

3.3 Input with Restricted Lengths

The results of the previous section are perhaps of limited applicability; the lower bound on the input angle may be too restrictive, and the minimum output angle far too small to be of practical use.

By making assumptions about the lengths of adjoining line segments, we can prove a much stronger result; In fact, the Delaunay Refinement Algorithm may handle arbitrarily small input angles, with some loss of output quality. Moreover, in contrast to the previous section, the assumptions made on the input can be forced onto a wide class of input by a "grooming" process, as shown in Chapter 4. In Chapter 6, it is shown that the grooming can be performed adaptively. These ideas are not new, rather they somehow formalize Ruppert's strategy of splitting on concentric circular shells [39].

3.3.1 Assumptions on the Input

The algorithm makes some assumptions about the lengths of adjoining input segments. Since the algorithm splits segments at midpoints, we are interested in the ratio of the lengths of adjoining segments, modulo powers of two.

Definition 3.3.1. Given an ordered pair of positive real numbers, (l_1, l_2) we define their *essential ratio* to be the unique $\xi \in [1, 2)$ such that $\frac{l_1}{l_2} = 2^k \xi$, for some integer k. We denote the essential ratio as $\operatorname{ess}(l_1, l_2)$. Note that $\operatorname{ess}(l_2, l_1) = \frac{2}{\operatorname{ess}(l_1, l_2)}$, unless $\operatorname{ess}(l_1, l_2) = 1$, in which case $\operatorname{ess}(l_2, l_1) = \operatorname{ess}(l_1, l_2) = 1$.

Definition 3.3.2. Two adjoining input segments which subtend an angle $\theta \in (0, \pi]$, with lengths l_1, l_2 are said to have *pairwise acceptable lengths* if

(a) $\pi/3 \leq \theta$, or

(b) $\pi/4 \leq \theta < \pi/3$, and either $\operatorname{ess}(l_1, l_2) = 1$, or $2 \cos\theta \leq \operatorname{ess}(l_1, l_2) \leq \frac{1}{\cos\theta}$, or

(c) $0 < \theta < \pi/4$, and $\operatorname{ess}(l_1, l_2) = 1$.

Notice that in item (b) and item (c) of Definition 3.3.2, the order of the line segments is immaterial. Also notice that for $\pi/3 \leq \theta < \pi/2$, that the condition $2\cos\theta \leq \operatorname{ess}(l_1, l_2) \leq \frac{1}{\cos\theta}$ holds automatically since $\operatorname{ess}(l_1, l_2) \in [1, 2)$. If $\pi/2 \leq \theta$, then no point on one of the segments will ever encroach any subsegment along the other.

The following assumption on the input can be satisfied by augmenting the input with at most $2|\mathbb{S}|$ points, splitting some or all of the line segments in the input which share an endpoint so they have pairwise acceptable lengths, as defined in Definition 3.3.2; see Chapter 4.

Assumption 3.3.3. In addition to those of Assumption 2.2.1 we also assume the following:

(a) Two adjoining input segments which are consecutive (*i.e.*, no segment with common endpoint is between them) have pairwise acceptable lengths, as defined in Definition 3.3.2.

3.3.2 Establishing the Encroachment Sequence Bound Hypothesis

We begin with the following claim which is applicable to the Delaunay Refinement Algorithm regardless of assumptions on input. The elementary proof by induction is omitted.

Claim 3.3.4. Let (a, b) be a current segment on an input segment (x, y). Then $\log_2 \frac{|x-y|}{|a-b|}$ is a nonnegative integer. Moreover $\frac{|x-a|}{|a-b|}$ is either zero or is an integral power of two, as is $\frac{|y-a|}{|a-b|}$, $\frac{|x-b|}{|a-b|}$, and $\frac{|y-b|}{|a-b|}$.

A geometric argument follows which helps us establish that radii don't "dwindle" in encroachment sequences. The lemma which follows is a mild improvement on Lemma 3.1.4 when the input satisfies Assumption 3.3.3

Claim 3.3.5. Given three noncollinear points, x, p, q, with $|x - q| \le |x - p|$ then $|p - q| \ge 2\,|q - x| \sin \frac{\theta}{2}$, where $\theta = \angle qxp \le \pi$.

Proof. Let $L = \frac{|x-p|}{|x-q|} \ge 1$. Using the cosine rule on Δxpq,

$$
\begin{aligned}
|p-q|^2 &= |x-p|^2 + |x-q|^2 - 2\,|x-p|\,|x-q| \cos\theta. \\
&= (1+L^2)\,|x-q|^2 - 2L\,|x-q|^2 \cos\theta \\
&\ge 2L\,|x-q|^2 - 2L\,|x-q|^2 \cos\theta \\
&= 2L\,|x-q|^2\,(1-\cos\theta),
\end{aligned}
$$

where we have used that $1 + L^2 \ge 2L$. Using $L \ge 1$, we obtain $\frac{|p-q|}{|x-p|} \ge \sqrt{2(1-\cos\theta)}$. It is a simple exercise to show that $2 \sin \frac{\theta}{2} = \sqrt{2(1-\cos\theta)}$ for $\theta \in [0, \pi]$. □

Lemma 3.3.6. *Suppose that the input conforms to Assumption 3.3.3. Let p be the midpoint of a segment which is encroached by a committed point, q, on an adjoining input segment. Let r_p be the radius associated with p, and r_q that of q. Then $r_q \le r_p$, and moreover,*

$$|p - q| \ge 2 r_q \sin \frac{\theta}{2},$$

where θ is the angle between the two input segments.

Proof. Let $(x, y), (x, z)$ be the two input segments containing, respectively, p, q. Let (a, b) be the subsegment of which p is midpoint. Let (c, d) be that for which q is midpoint. Assume that a is closer to x than b is, and assume c is closer to x than d is. It may be the case that $x = a$, or $x = c$. Consider the cases for θ, and the lengths of the two segments:

- Suppose $\pi/3 \le \theta$. Then by Lemma 3.1.4, $1 \ge 2\cos\theta > \frac{|x-q|}{r_p}$. Since $r_q \le |x - q|$, then $1 \ge \frac{r_q}{r_p}$, as desired.
- Suppose $\pi/4 \le \theta < \pi/3$. By Lemma 3.1.3, since $\arcsin \frac{1}{3} < \theta$, we know that $a = x$, and so $r_p = |x - p|$. By Lemma 3.1.4, $2\cos\theta > \frac{|x-q|}{|x-p|} = 2^l \operatorname{ess}(|x-q|, |x-p|)$, for integer l.

  By item (3.3.3) of Assumption 3.3.3, the two segments have pairwise acceptable lengths, and so either $\operatorname{ess}(|x-z|, |x-y|) = 1$, or $2\cos\theta \le \operatorname{ess}(|x-z|, |x-y|) \le \frac{1}{\cos\theta}$.

In the first case, we have $2 > 2\cos\theta > 2^l$, forcing $l \leq 0$. But then $\frac{|x-q|}{|x-p|} = 2^l \leq 1$.
In the second case, we have $2\cos\theta > 2^l \operatorname{ess}(|x-z|, |x-y|) \geq 2^l 2\cos\theta$, again forcing $l < 0$. But $\frac{|x-q|}{|x-p|} = 2^l \operatorname{ess}(|x-q|, |x-p|) \leq \frac{1}{2\cos\theta} < 1$, since $\theta < \pi/3$.
In either case, $\frac{|x-q|}{r_p} = \frac{|x-q|}{|x-p|} \leq 1$. Using $r_q \leq |x-q|$ gives the desired result $r_q \leq r_p$.

- Suppose $\theta < \pi/4$. By item (3.3.3) of Assumption 3.3.3, the two segments have pairwise acceptable lengths, and so $\operatorname{ess}(|x-z|, |x-y|) = 1$. If $a = x$, then the reasoning of the first part of the previous case applies. So assume otherwise.
By Claim 3.3.4, $\log_2 \frac{|x-y|}{|a-b|}$, and $\log_2 \frac{|x-z|}{|c-d|}$ are nonnegative integers. By the assumption on the essential length ratios, $\log_2 \frac{|x-y|}{|x-z|}$ is an integer. Thus $\log_2 \frac{|a-b|}{|c-d|} = \log_2 \frac{r_p}{r_q} = j$ is also an integer. We wish to show that it is nonnegative.
By Claim 3.1.2, $|x-a| < |x-q| < |x-b|$, so that $|x-a| < |x-c| + r_q < |x-a| + 2r_p$. Using Claim 3.3.4 shows that $k = \frac{|x-a|}{|a-b|} = \frac{|x-a|}{2r_p}$ is a nonnegative integer, as is, *mutatis mutandis*, $l = \frac{|x-c|}{2r_q}$. Thus
$$2kr_p < (2l+1)r_q < 2(k+1)r_p, \quad \text{or}$$
$$2^{j+1}k < (2l+1) < 2^{j+1}(k+1), \quad \text{and so}$$
$$\frac{2l+1}{2^{j+1}} - 1 < k < \frac{2l+1}{2^{j+1}}.$$
If j is a negative integer, then 2^{j+1} is a power of two no greater than 1; in particular it divides any integer, thus $\frac{2l+1}{2^{j+1}} = m$ is an integer. This gives the contradiction that $m - 1 < k < m$ for integer m, k. Thus j is a nonnegative integer, or $r_p \geq r_q$.

For the second part, by Claim 3.3.5, $|p-q| \geq 2(|x-q| \wedge |x-p|)\sin\frac{\theta}{2}$. Clearly $|x-p| \geq r_p \geq r_q$, and $|x-q| \geq r_q$, so the result $|p-q| \geq 2r_q \sin\frac{\theta}{2}$ holds, as desired. □

We now prove the Encroachment Sequence Bound Hypothesis for input conforming to Assumption 3.3.3.

Lemma 3.3.7. *Suppose that the input to Delaunay Refinement Algorithm conforms to Assumption 3.3.3. Then there are constants ρ, β, both at least 1, such that if p_{l-1} is a midpoint of a segment that is encroached by a committed point, which is a midpoint, p_{l-2}, on an adjoining input segment, with the two segments subtending angle θ, then either*

(a) $\operatorname{lfs}(p_{l-1}) \leq \beta |p_{l-1} - p_{l-2}|$, *or*

(b) there is an encroachment sequence, $\{p_i\}_{i=0}^{l-1}$, such that

  (i) the parent segment of p_0 was not encroached by a committed point on an adjoining input segment,

  (ii) the inequality
$$\operatorname{lfs}(p_{l-1}) \leq (\beta + \rho \frac{\operatorname{lfs}(p_0)}{r_0}) |p_{l-1} - p_{l-2}|, \tag{5}$$
  where r_0 is the radius associated with p_0, holds.

Moreover, $\rho = \frac{1}{2\sin\frac{\theta^*}{2}}$, *and* $\beta = 1 + \rho(1 + \frac{4}{\sin\theta^*})$ *suffice.*

Proof. Again it is convenient to say that q "provokes" midpoint p if q is an already committed point that encroaches the parent segment of p.

Suppose p_{l-1} is a midpoint which is provoked by a committed midpoint, p_{l-2}, on an adjoining input segment. Let the nondisjoint input segments share input point x. Construct a maximal encroachment sequence $\{p_i\}_{i=0}^{l-1}$ that is radial about x, *i.e.*, each midpoint is on an input segment with endpoint x. Let r_i be the radius associated with p_i. By Lemma 3.3.6, it is clear that $r_i \geq r_{i-1}$ for $i = 1, 2, \ldots, l-1$.

By Claim 3.1.2, since for $i > 0$, p_i is provoked by a midpoint on an adjoining input segment then $|x - p_i| \sin\theta^* < r_i$. Using this property, and the triangle inequality, gives the bound:

$$\begin{aligned}
|p_{l-1} - p_0| &\leq |p_{l-1} - p_{l-2}| + |p_{l-2} - x| + |x - p_1| + |p_1 - p_0|, \\
&\leq |p_{l-1} - p_{l-2}| + \frac{r_{l-2}}{\sin\theta^*} + \frac{r_1}{\sin\theta^*} + r_1, \\
&\leq |p_{l-1} - p_{l-2}| + \left(1 + \frac{2}{\sin\theta^*}\right) r_{l-2},
\end{aligned}$$

where we have used $|p_0 - p_1| \leq r_1$, which holds because p_0 provokes p_1, and the fact that, by Lemma 3.3.6, $r_0 \leq r_1 \leq \ldots \leq r_{l-2}$. By the same lemma we know that $r_{l-2} \leq \frac{|p_{l-1} - p_{l-2}|}{2\sin\frac{\theta}{2}}$. Thus

$$|p_{l-1} - p_0| \leq \left(1 + \frac{1}{2\sin\frac{\theta}{2}} + \frac{1}{\sin\theta^* \sin\frac{\theta}{2}}\right) |p_{l-1} - p_{l-2}|. \tag{6}$$

Moreover because $r_0 \leq r_{l-2}$, then

$$\begin{aligned}
\operatorname{lfs}(p_{l-1}) &\leq |p_{l-1} - p_0| + \operatorname{lfs}(p_0), \\
&\leq |p_{l-1} - p_{l-2}| + \left(1 + \frac{2}{\sin\theta^*}\right) r_{l-2} + \frac{\operatorname{lfs}(p_0)}{r_0} r_0, \\
&\leq |p_{l-1} - p_{l-2}| + \left(1 + \frac{2}{\sin\theta^*} + \frac{\operatorname{lfs}(p_0)}{r_0}\right) r_{l-2}, \\
&\leq \left(1 + \frac{1}{2\sin\frac{\theta}{2}} \left[1 + \frac{2}{\sin\theta^*} + \frac{\operatorname{lfs}(p_0)}{r_0}\right]\right) |p_{l-1} - p_{l-2}|, \\
&\leq \left(\beta + \rho \frac{\operatorname{lfs}(p_0)}{r_0}\right) |p_{l-1} - p_{l-2}|.
\end{aligned}$$

If p_0 was not provoked by a committed point on an adjoining input segment, alternative (b) holds for the encroachment sequence $\{p_i\}_{i=0}^{l-1}$. So suppose to the contrary that p_0 is provoked by a committed point on an adjoining input segment. By maximality of the chosen encroachment sequence, the two input segments do not share endpoint x. We've assumed they do share an endpoint, call it y. Let these two segments subtend angle ϕ. By Claim 3.1.2, $|p_0 - y| \leq \frac{r_0}{\sin\phi}$. Since $\theta^* \leq \phi < \pi/2$, then $|p_0 - y| \leq \frac{r_0}{\sin\theta^*}$.

We bound $|p_{l-1} - y|$; using equation 6, and the above estimate on $|p_0 - y|$, and because $r_0 \le r_{l-2} \le \frac{|p_{l-1}-p_{l-2}|}{2\sin\frac{\theta}{2}}$, we have

$$\begin{aligned} |p_{l-1} - y| &\le |p_{l-1} - p_0| + |p_0 - y|, \\ &\le \left[1 + \frac{1}{2\sin\frac{\theta}{2}} + \frac{1}{\sin\theta^*\sin\frac{\theta}{2}}\right]|p_{l-1} - p_{l-2}| + \frac{r_0}{\sin\theta^*}, \\ &\le \left[1 + \frac{1}{2\sin\frac{\theta}{2}} + \frac{3}{2\sin\theta^*\sin\frac{\theta}{2}}\right]|p_{l-1} - p_{l-2}|, \\ &\le \left[1 + \frac{2}{\sin\theta^*\sin\frac{\theta}{2}}\right]|p_{l-1} - p_{l-2}| \le \beta\,|p_{l-1} - p_{l-2}|. \end{aligned}$$

By Claim 3.1.2 we can bound $|p_{l-1} - x| \le \frac{|p_{l-1}-p_{l-2}|}{\sin\theta^*}$. Then by definition of local feature size,

$$\begin{aligned} \operatorname{lfs}(p_{l-1}) &\le |p_{l-1} - x| \vee |p_{l-1} - y| \\ &\le \left(\frac{1}{\sin\theta^*} \vee \beta\right)|p_{l-1} - p_{l-2}| \le \beta\,|p_{l-1} - p_{l-2}|. \end{aligned}$$

Thus alternative (a) is established. □

In Chapter 5, we will construct a proof similar to the preceding, thus it may be instructive to take note of the general idea of the proof. A sequence of midpoints $\{p_i\}_{i=0}^{l-1}$ was constructed, where point p_i has associated radius r_i. Then the following facts were key in the proof: (a) The radii are nondecreasing: $r_0 \le r_1 \le \ldots \le r_{l-1}$, (b) $|p_i - p_{i-1}|$ is within a constant of both r_i and r_{i-1}, and (c) $|p_i - x|$ is bounded by a constant times r_i.

By Lemma 2.5.4 the following is then immediate.

Theorem 3.3.8. *Suppose that the input to the Delaunay Refinement Algorithm conforms to Assumption 3.3.3. Let* $\rho = \frac{1}{2\sin\frac{\theta^*}{2}}$*. Then the Delaunay Refinement Algorithm applied with* $\kappa < \arcsin\frac{1}{2\rho\sqrt{2}}$ *terminates with good grading. Moreover, the cardinality and optimality results of Corollary 2.5.7 apply with the given* ρ*, and with* $\beta = 1 + \rho(1 + \frac{4}{\sin\theta^*})$.

Thus the Delaunay Refinement Algorithm can guarantee well graded output when using any $\kappa < \arcsin\left[\sin\left(\frac{\theta^*}{2}\right)/\sqrt{2}\right]$. This bound is the same that was achieved by Shewchuck for his "terminator" algorithm [43]. In Chapter 5, we show that the algorithm can actually use a larger value of κ.

CHAPTER IV

AUGMENTING INPUT

"What's lost upon the roundabouts we pulls up on the swings!" –*Patrick R. Chalmers*

The results of Section 3.3 and later of Section 5.4 suggest that input with small angles may be handled by Delaunay Refinement algorithms subject to some restriction on the lengths of consecutive segments. The goal of this chapter is to show that by augmenting an input with points on segments that the input can be made to conform to Assumption 3.3.3, with a bounded decrease in local feature size.

Throughout this chapter, we assume there is a set of input points and segments, $(\mathcal{P}, \mathcal{S})$, conforming to Assumption 2.2.1. We denote the local feature size relative to this input by $\operatorname{lfs}(x)$. We will augment the set of points, $\mathcal{P}$ to form a new set $\mathcal{P}'$, with all new points on segments of $\mathcal{S}$. The set $\mathcal{S}$ will be replaced by $\mathcal{S}'$, where every segment of $\mathcal{S}$ is the union of segments in $\mathcal{S}'$. We denote the local feature size relative to this pair of sets by $\operatorname{lfs}'(x)$. Loosely, we call this process an *augmentation*, and we will analyze two different augmenting strategies.

4.1 Bounded Reduction Augmenter

The first kind of augmenter that we consider breaks an input segment into pieces not too much shorter than the original. We introduce an algorithm which performs this kind of augmentation and creates input that conforms to Assumption 3.3.3. First we give it a name:

Definition 4.1.1. A γ-*Bounded Reduction Augmenter* is a procedure that takes an input $(\mathcal{P}, \mathcal{S})$, and produces an output $(\mathcal{P}', \mathcal{S}')$ such that

(a) $\mathcal{P} \subseteq \mathcal{P}'$,

(b) every segment of $\mathcal{S}$ is the union of segments in $\mathcal{S}'$,

(c) every point of $\mathcal{P}' \setminus \mathcal{P}$ and every segment of $\mathcal{S}'$ is on a segment of $\mathcal{S}$, and

(d) if $S' \in \mathcal{S}'$ is a segment on segment $S \in \mathcal{S}$, then $|S| \leq \gamma |S'|$.

Note the definition requires that $\gamma \geq 2$, otherwise the augmenting procedure is idempotent, *i.e.*, leaves $(\mathcal{P}, \mathcal{S})$ unchanged. Also note that for a γ-Bounded Reduction Augmenter that $|\mathcal{P}' \setminus \mathcal{P}| \leq \lfloor \gamma - 1 \rfloor |\mathcal{S}|$.

We now show that a γ-Bounded Reduction Augmenter does not greatly diminish local feature size.

Theorem 4.1.2. *Let* $\mathrm{lfs}(x)$ *be the local feature size on* $(\mathcal{P}, \mathcal{S})$, *and let* $\mathrm{lfs}'(x)$ *be the local feature size on* $(\mathcal{P}', \mathcal{S}')$, *where* $\mathcal{P} \subseteq \mathcal{P}'$, *every segment of* $\mathcal{S}$ *is the union of segments in* $\mathcal{S}'$, *and every point of* $\mathcal{P}' \setminus \mathcal{P}$ *and segment of* $\mathcal{S}'$ *is on a segment of* $\mathcal{S}$.

Furthermore suppose that coincident segments of $\mathcal{S}$ *meet at an angle no less than some* $\theta^* \leq \pi/2$, *and that there is a constant* $\gamma \geq 2$ *such that if* $S' \in \mathcal{S}'$ *is a segment on the segment* $S \in \mathcal{S}$, *that* $|S| \leq \gamma |S'|$.

Then there is a constant σ *such that* $\mathrm{lfs}'(x) \geq \sigma \mathrm{lfs}(x)$ *for every* $x \in \mathbb{R}^2$. *Additionally,* $\sigma = \frac{\sin\theta^*}{2\gamma - 1}$ *suffices.*

Note that although we assume in this thesis that $\theta^* \leq \pi/3$, this theorem only requires that θ^* is no greater than $\pi/2$.

Proof. We will refer to those features of $(\mathcal{P}, \mathcal{S})$ as being "input."

Let $r = \mathrm{lfs}'(x)$, and let B be the open ball centered at x of radius r. By definition there are two disjoint features touching the closure of B. Since features of $(\mathcal{P}', \mathcal{S}')$ are on features of the input, we can assume there are two (not necessarily distinct) features of the input, say X, Y that intersect the closure of B. We consider the cases:

- If X, Y are disjoint features of the input, then by definition $\mathrm{lfs}(x) \leq r = \mathrm{lfs}'(x)$, since the closure of B intersects two features of the input.
- If $X = Y$, then the closure of B contains two points of $\mathcal{P}'$ placed on some segment $(a, b) \in \mathcal{S}$. If both these points are in $\mathcal{P}$, then by definition $\mathrm{lfs}(x) \leq \mathrm{lfs}'(x)$. So assume that there is one, call it p, that is in $\mathcal{P}'$. Our assumptions on the lengths of segments in $\mathcal{S}'$ give us
$$|a - p| \vee |b - p| \leq \frac{(\gamma - 1)|a - b|}{\gamma}.$$
By definition of the local feature size, $\mathrm{lfs}(x) \leq |x - a| \vee |x - b|$. Using the triangle inequality gives $\mathrm{lfs}(x) \leq r + (|p - a| \vee |p - b|) \leq r + \frac{(\gamma-1)|a-b|}{\gamma}$. So
$$\frac{\mathrm{lfs}'(x)}{\mathrm{lfs}(x)} \geq \frac{\mathrm{lfs}'(x)}{\mathrm{lfs}'(x) + \frac{(\gamma-1)|a-b|}{\gamma}}.$$
The two terms involved in the right hand side are both positive, but only one of them depends on x. It is easy to see that the right hand side is minimized when $\mathrm{lfs}'(x)$ is minimized. But the new local feature size, $\mathrm{lfs}'(x)$ must be at least half the distance from p to the other point, *i.e.*, at least half the length of a segment on (a, b), so

$\mathrm{lfs}'(x) \geq \frac{|a-b|}{2\gamma}$. Thus

$$\frac{\mathrm{lfs}'(x)}{\mathrm{lfs}(x)} \geq \frac{\frac{|a-b|}{2\gamma}}{\frac{|a-b|}{2\gamma} + \frac{(\gamma-1)|a-b|}{\gamma}} \geq \frac{1}{1+2(\gamma-1)},$$

which suffices.

- If X, Y are non-disjoint input features, then we may assume they are segments, as the case where one is an endpoint of the other is treated in the previous case. Let the two segments be $(a,b), (a,c)$. Without loss of generality assume there is a feature of $(\mathcal{P}', \mathcal{S}')$, call it Z, that is on (a,b) and intersects the closure of B, as does the segment (a,c). Furthermore we may assume that Z is disjoint from every feature of $(\mathcal{P}', \mathcal{S}')$ which is on (a,c). Then there is a point $p \in \mathcal{P}' \setminus \mathcal{P}$ which is on (a,b) which is no farther from a than Z. See Figure 15.

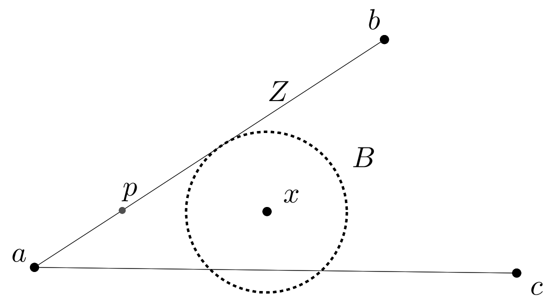

Figure 15: The "segment-segment" case for the proof of Theorem 4.1.2 is shown. The segments (a,b),(a,c) are in the input, while p is an augmenting point.

This then gives us the lower bound: $\mathrm{lfs}'(x) \geq \frac{|a-p|\sin(\theta\wedge\pi/2)}{2}$, where θ is the angle subtended by the segments $(a,b), (a,c)$. This lower bound holds because the distance from Z to (a,c) is at least the distance from a to Z times $\sin(\theta \wedge \pi/2)$. Using the bounds from the previous case, this gives $\mathrm{lfs}'(x) \geq \frac{|a-b|\sin(\theta\wedge\pi/2)}{2\gamma}$.
The local feature size (with respect to the input) of x is no greater than the distance from x to b, since b is outside the closure of B, which intersects an input feature disjoint from b, *i.e.*, the segment (a,c). Using the triangle inequality to bound the distance from x to b gives an upper bound: $\mathrm{lfs}(x) \leq r + |p-b|$. This holds because Z intersects the closure of B, but is no farther from b than p is. Thus

$$\frac{\mathrm{lfs}'(x)}{\mathrm{lfs}(x)} \geq \frac{\mathrm{lfs}'(x)}{\mathrm{lfs}'(x) + |p-b|}.$$

Again, the right hand side is minimized when $\mathrm{lfs}'(x)$ is minimized. Using the above obtained lower bound on $\mathrm{lfs}'(x)$ gives

$$\begin{aligned}
\frac{\operatorname{lfs}'(x)}{\operatorname{lfs}(x)} &\geq \frac{\frac{|a-b|\sin(\theta\wedge\pi/2)}{2\gamma}}{\frac{|a-b|\sin(\theta\wedge\pi/2)}{2\gamma} + \frac{(\gamma-1)|a-b|}{\gamma}} \\
&\geq \frac{\sin(\theta\wedge\pi/2)}{\sin(\theta\wedge\pi/2) + 2(\gamma-1)} \\
&\geq \frac{\sin(\theta\wedge\pi/2)}{2\gamma-1} \geq \frac{\sin\theta^*}{2\gamma-1}.
\end{aligned}$$

□

As Algorithm 1 we present a 5-Bounded Reduction Augmenter, which can produce augmenting input $(\mathcal{P}', \mathcal{S}')$ that conforms to Assumption 3.3.3. By Theorem 4.1.2, it follows that $\operatorname{lfs}'(x) \geq \frac{\sin\theta^*}{9}\operatorname{lfs}(x)$, where $\operatorname{lfs}'(x)$ is with respect to the augmented pair, and θ^* is a proper lower bound on input angles. Moreover, it should be clear that the procedure adds no more than $2|\mathcal{S}|$ Steiner Points.

The algorithm can be briefly described as follows: for each input point x, partition the segments with endpoint x into maximal collections such that each segment is separated from another in the collection by an angle no greater than $\pi/3$; For each such maximal collection with more than one segment, for each segment (x, y) in the collection, add an augmenting point, p, to the segment such that $0.2|x-y| \leq |x-p| \leq 0.4|x-y|$, and such that each new segment created in the maximal collection has the same length modulo a power of two.

Algorithm 1: Algorithm for making input conform to Assumption 3.3.3.
Input: The input points and segments.
Output: An augmented set of points and segments.
BoundedReductionAugment$(\mathcal{P}, \mathcal{S})$
(1) **foreach** point $x \in \mathcal{P}$ which is the endpoint of at least two segments in $\mathcal{S}$
(2) Partition the segments with endpoint x into maximal collections such that each segment is separated from another in the collection by an angle no greater than $\pi/3$.
(3) **foreach** maximal collection with more than one segment
(4) To each segment (x, y) in the collection, add an augmenting point p, such that $0.2|x-y| \leq |x-p| \leq 0.4|x-y|$, and such that each new segment (x, p) has the same length modulo a power of two.

Since this *is* a thesis, and we are not constrained by space, we include a slight variant of Algorithm 1 which puts input into the form of Assumption 5.4.1. This augmenter, presented as Algorithm 2, adds augmenting points around each vertex of degree at least two, and thus can be a bit more aggressive in adding augmenting points. This is because Assumption 5.4.1 is a more onerous assumption than Assumption 3.3.3. In the end the procedure is still a

5-Bounded Reduction Augmenter and adds no more than $2|\mathcal{S}|$ augmenting points. The output of this algorithm is shown in Figure 16.

Algorithm 2: Algorithm for making input conform to Assumption 5.4.1.
Input: The input points and segments.
Output: An augmented set of points and segments.
BoundedReductionAugment'($\mathcal{P}, \mathcal{S}$)
(1) **foreach** point $x \in \mathcal{P}$ which is the endpoint of at least two segments in $\mathcal{S}$
(2) To each segment (x, y) in $\mathcal{S}$, add an augmenting point p, such that $0.2|x-y| \leq |x-p| \leq 0.4|x-y|$, and such that each new segment (x, p) has the same length modulo a power of two.

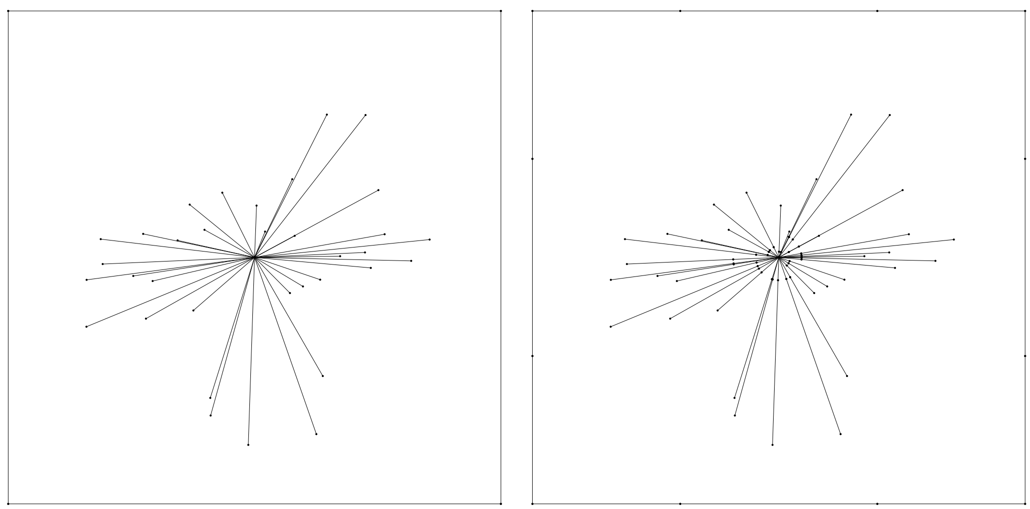

(a) Input (b) Augmented

Figure 16: The input shown in (a) was fed to Algorithm 2, resulting in the augmented input shown in (b). The augmenting points are added at different concentric circular shells around the "axis" of the segments.

We briefly argue that Ruppert's strategy of splitting on concentric circular shells is something like an as-needed 6-Bounded Reduction Augmenter. This heuristic fixes midpoint-midpoint interactions as follows: When an input segment is first to be split, it is split at its midpoint, leaving two subsegments which have one input endpoint each; when either of these two subsegments is split, they are split at the point closest to their midpoint which is at a distance 2^k from their input endpoint, for integer k; after this, all subsegments are split at their midpoints [39].

Consider the two subsegments created by the "off-center" split. Clearly neither of them is shorter than one third the length of their parent subsegment, which is one half the length of the input segment. Thus this strategy is a 6-Bounded Reduction Augmenter.

4.2 Feature Size Augmenter

There is another kind of augmenter to consider. This kind may break a segment into very small pieces, but creates no segments smaller than local feature size would suggest. Shewchuk considered such an augmenter as a grooming step in the production of Constrained Delaunay Triangulations in three dimensions [42], however he had no reason to consider the resultant loss of local feature size off of the input features.

Definition 4.2.1. For $\gamma > 0$, a γ-*Feature Size Augmenter* is a procedure that takes an input $(\mathcal{P}, \mathcal{S})$, and produces an output $(\mathcal{P}', \mathcal{S}')$ such that

(a) $\mathcal{P} \subseteq \mathcal{P}'$,

(b) every segment of $\mathcal{S}$ is the union of segments in $\mathcal{S}'$,

(c) every point of $\mathcal{P}' \setminus \mathcal{P}$ and every segment of $\mathcal{S}'$ is on a segment of $\mathcal{S}$, and

(d) if $S' \in \mathcal{S}'$, and x is any point on S', including its endpoints, then $\operatorname{lfs}(x) \leq \gamma |S'|$, where $\operatorname{lfs}(x)$ is with respect to the input $(\mathcal{P}, \mathcal{S})$.

Note that every γ-Bounded Reduction Augmenter is a γ-Feature Size Augmenter. However this is a strict containment and there is no guarantee that a Feature Size Augmenter is a Bounded Reduction Augmenter at all. To see this, consider input consisting of two parallel unit-length segments which are ϵ distance apart from each other. The local feature size for every point on one of these segments is merely ϵ, so a γ-Feature Size Augmenter may break each segment into as many as γ/ϵ subsegments. Thus there is no $\gamma > 0$ such that a γ-Feature Size Augmenter is necessarily idempotent, and there is no *a priori* upper bound on the number of subsegments created by a γ-Feature Size Augmenter.

We now show that a γ-Feature Size Augmenter does not greatly diminish local feature size.

Theorem 4.2.2. *Let* $\operatorname{lfs}(x)$ *be the local feature size on* $(\mathcal{P}, \mathcal{S})$, *and let* $\operatorname{lfs}'(x)$ *be the local feature size on* $(\mathcal{P}', \mathcal{S}')$, *where* $\mathcal{P} \subseteq \mathcal{P}'$, *every segment of* $\mathcal{S}$ *is the union of segments in* $\mathcal{S}'$, *and every point of* $\mathcal{P}' \setminus \mathcal{P}$ *and segment of* $\mathcal{S}'$ *is on a segment of* $\mathcal{S}$.

Furthermore suppose that coincident segments of $\mathcal{S}$ *meet at an angle no less than some* $\theta^* \leq \pi/2$, *and that there is a constant* $\gamma > 0$ *such that if* $S' \in \mathcal{S}'$ *and* x *is any point along* S' *then* $\operatorname{lfs}(x) \leq \gamma |S'|$.

Then there is a constant σ *such that* $\operatorname{lfs}'(x) \geq \sigma \operatorname{lfs}(x)$ *for every* $x \in \mathbb{R}^2$. *Additionally,* $\sigma = \frac{\sin\theta^*}{2\gamma+3}$ *suffices.*

Proof. We will refer to those features of $(\mathcal{P}, \mathcal{S})$ as being “input.”

Let $r = \mathrm{lfs}'(x)$, and let B be the open ball centered at x of radius r. By definition there are two disjoint features touching the closure of B. Since features of $(\mathcal{P}', \mathcal{S}')$ are on features of the input, we can assume there are two (not necessarily distinct) features of the input, say X, Y that intersect the closure of B. We consider the cases:

- If X, Y are disjoint features of the input, then by definition $\mathrm{lfs}(x) \leq r = \mathrm{lfs}'(x)$, since the closure of B intersects two features of the input.
- If $X = Y$, then the closure of B contains an entire segment of $\mathcal{S}'$, call it (a, b). By the Lipschitz condition, $\mathrm{lfs}(x) \leq |x - a| + \mathrm{lfs}(a) = r + \mathrm{lfs}(a)$. By assumption $\mathrm{lfs}(a) \leq \gamma |a - b|$. Since the segment is entirely in the closure of B, we have $|a - b| \leq 2r$. Putting these all together gives $\mathrm{lfs}(x) \leq (1 + 2\gamma) r$, *i.e.*, $\mathrm{lfs}'(x) \geq \frac{1}{2\gamma+1} \mathrm{lfs}(x)$, which suffices.
- If X, Y are non-disjoint input features, then we may assume they are segments, as the case where one is an endpoint of the other is treated in the previous cases. Let the two segments be $(a, b), (a, c)$. Without loss of generality there is a point $p \in \mathcal{P}' \setminus \mathcal{P}$ that is on (a, b) such that $(a, p) \in \mathcal{S}'$ does not intersect the open ball B. Note it may be the case that p is in the closure of B. Let z be a point of (a, b) that *is* in the closure of B. See Figure 17.

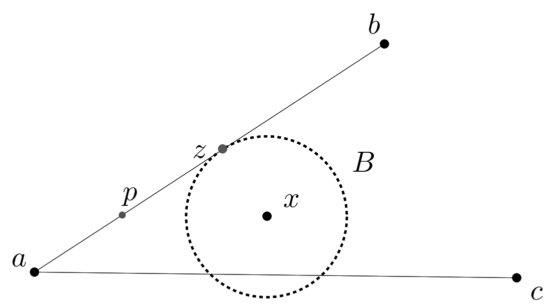

Figure 17: The "segment-segment" case for the proof of Theorem 4.2.2 is shown. The segments (a, b),(a, c) are in the input, while p is an augmenting point. The point z is on (a, b), and inside the closure of B. Although drawn this way, the point z need not be unique.

The distance from z to (a, c) is $|a - z| \sin(\theta \wedge \pi/2)$, where θ is the angle subtended by $(a, b), (a, c)$. Since the closure of B contains a point of (a, c), then the segment from z to that point is inside the closure of B and has length no greater than $2r$. That is $|a - z| \sin(\theta \wedge \pi/2) \leq 2r$. Using $\theta^* \leq \theta \wedge \pi/3$ gives $|a - z| \leq \frac{2r}{\sin \theta^*}$.

Now we use the Lipschitz condition:

$$\begin{aligned} \mathrm{lfs}(x) &\leq |x - z| + |z - a| + \mathrm{lfs}(a), \\ &\leq r + \frac{2r}{\sin \theta^*} + \gamma |a - p|. \end{aligned}$$

Noting that $|a-p| \leq |a-z| \leq \frac{2r}{\sin\theta^*}$ gives

$$\mathrm{lfs}\,(x) \leq r + \frac{2\,(1+\gamma)\,r}{\sin\theta^*} \leq \frac{1+2\,(1+\gamma)}{\sin\theta^*}r = \frac{2\gamma+3}{\sin\theta^*}\mathrm{lfs}'\,(x)\,.$$

□

As Algorithm 3 we present a $\frac{3+\sqrt{13}}{2}$-Feature Size Augmenter, which can produce augmenting input $(\mathcal{P}', \mathcal{S}')$ that conforms to Assumption 3.3.3. By Theorem 4.2.2, it follows that $\mathrm{lfs}'\,(x) \geq \frac{\sin\theta^*}{6+\sqrt{13}}\mathrm{lfs}\,(x)$, where $\mathrm{lfs}'\,(x)$ is with respect to the augmented pair, and θ^* is a proper lower bound on input angles. Moreover, it should be clear that the procedure adds no more than $2\,|\mathcal{S}|$ Steiner Points.

Algorithm 3: Algorithm for making input conform to Assumption 3.3.3.
Input: The input points and segments.
Output: An augmented set of points and segments.
FEATURESIZEAUGMENT$(\mathcal{P}, \mathcal{S})$
(1) **foreach** point $x \in \mathcal{P}$ which is the endpoint of at least two segments in $\mathcal{S}$
(2) Partition the segments with endpoint x into maximal collections such that each segment is separated from another in the collection by an angle no greater than $\pi/3$.
(3) **foreach** maximal collection with more than one segment
(4) Let l be the length of the shortest segment in the collection.
(5) To each segment (x, y) in the collection, add an augmenting point p, such that $|x-p| = \frac{2}{1+\sqrt{13}}l$.

The algorithm can be briefly described as follows: for each input point x, partition the segments with endpoint x into maximal collections such that each segment is separated from another in the collection by an angle no greater than $\pi/3$; For each such maximal collection with more than one segment, let l be the length of the shortest segment of the collection. Then for each segment (x, y) in the collection, add the augmenting point, p, to the segment such that $|x-p| = \frac{2}{1+\sqrt{13}}l$.

Claim 4.2.3. Algorithm 3 is a $\frac{3+\sqrt{13}}{2}$-Feature Size Augmenter.

Proof. There are three kinds of subsegments in $\mathcal{S}$': those with zero, one, or two endpoints in $\mathcal{P}$. We let $S' \in \mathcal{S}'$ be a subsegment of $S \in \mathcal{S}$, and consider the cases.

- Let S' have two endpoints in $\mathcal{P}$. Then $S' = S$, and by definition, for any point x on S, $\mathrm{lfs}\,(x)$ is less than $|S|$, so it is less than $\frac{3+\sqrt{13}}{2}\,|S'|$.
- Let S' have one endpoint in $\mathcal{P}$. Let $S' = (a, p)$, with $a \in \mathcal{P}$. Let $S = (a, b)$. There are two separate subcases: p could have been added "near" a or b. By "near," we mean because of a cluster of segments around a or b. First we suppose it was near a. In this

case there is some segment (a, c) of length l such that $(a, p) = \frac{2}{1+\sqrt{13}} l$. Then if x is a point on (a, p), by definition of the local feature size, and using the triangle inequality,

$$\operatorname{lfs}(x) \leq |x - a| + l \leq |p - a| + \frac{1 + \sqrt{13}}{2} |a - p|,$$

and thus $\frac{\operatorname{lfs}(x)}{|a-p|} \leq \frac{3+\sqrt{13}}{2}$. Note this is may be a gross overestimate when, for example, $c = b$.

Now suppose that p was near b. We know that $|p - b| \leq \frac{2}{1+\sqrt{13}} |a - b|$. Thus

$$|a - p| \geq \left[1 - \frac{2}{1 + \sqrt{13}}\right] |a - b| = \frac{\sqrt{13} - 1}{1 + \sqrt{13}} |a - b|.$$

If x is a point on (a, p), then $\operatorname{lfs}(x) \leq |a - b|$, so

$$\frac{\operatorname{lfs}(x)}{|a - p|} \leq \frac{1 + \sqrt{13}}{\sqrt{13} - 1} < \frac{3 + \sqrt{13}}{2}.$$

- Let S' have no endpoint in $\mathcal{P}$. Let $S' = (p, q)$, let $S = (a, b)$, and suppose that (a, p) and (q, b) are in $\mathcal{S}'$. We know that neither $|a - p|$ nor $|q - b|$ are greater than $\frac{2}{1+\sqrt{13}} |a - b|$. It follows that

$$|p - q| \geq \left[1 - \frac{4}{1 + \sqrt{13}}\right] |a - b| = \frac{\sqrt{13} - 3}{1 + \sqrt{13}} |a - b|.$$

Now given x on (p, q), using the definition of local feature size,

$$\operatorname{lfs}(x) \leq |x - a| \vee |x - b| \leq \frac{2}{1 + \sqrt{13}} |a - b|.$$

Thus

$$\frac{\operatorname{lfs}(x)}{|p - q|} \leq \frac{2}{1 + \sqrt{13}} \frac{1 + \sqrt{13}}{\sqrt{13} - 3} = \frac{2}{\sqrt{13} - 3} = 2\frac{3 + \sqrt{13}}{13 - 9} = \frac{3 + \sqrt{13}}{2}.$$

□

Again, we display an algorithm for putting input into the form of Assumption 5.4.1 merely for completeness. The astute reader will have already surmised its form. The action of the algorithm is illustrated in Figure 18.

Algorithm 4: Algorithm for making input conform to Assumption 3.3.3.
Input: The input points and segments.
Output: An augmented set of points and segments.
FEATURESIZEAUGMENT'($\mathcal{P}, \mathcal{S}$)
(1) **foreach** point $x \in \mathcal{P}$ which is the endpoint of at least two segments in $\mathcal{S}$
(2) Let l be the length of the shortest segment of the form (x, y) in $\mathcal{S}$.
(3) To each segment (x, y) in $\mathcal{S}$, add an augmenting point p, such that $|x - p| = \frac{2}{1+\sqrt{13}} l$.

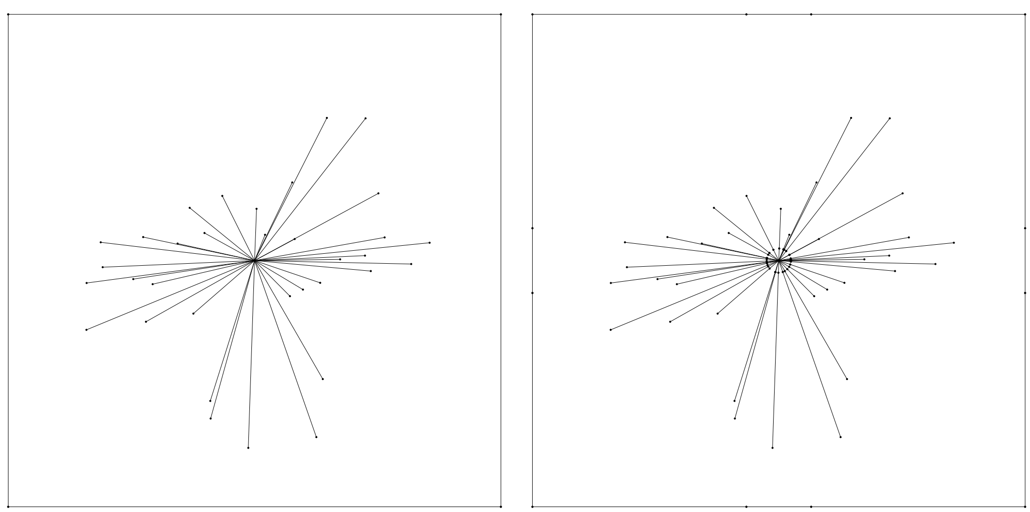

(a) Input (b) Augmented

Figure 18: The input shown in (a) was fed to Algorithm 4, resulting in the augmented input shown in (b). The augmenting points are all added at the same distance from the "axis" of the segments.

CHAPTER V

THE ADAPTIVE DELAUNAY REFINEMENT ALGORITHM

"What I tell you three times is true." –*Lewis Carroll*

5.1 Description of the Algorithm

We describe a locally adaptive algorithm which minimizes the effects of small input angles by determining quality with respect to location. The algorithm appears to ignore all poor quality triangles which are "near" small input angles; it will be shown that the ignored triangles cannot be too poor. Ultimately the algorithm will not be able to guarantee a minimum output angle greater than $\arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)$, but such small angles will only be near input segments which meet at an angle of θ^*. The algorithm will also be able to guarantee that angles "far" from input features are no less than $\arcsin 2^{-7/6}$.

We will assume that the algorithm maintains, for every midpoint, a pointer to the two input points which are endpoints of the input segment containing the midpoint. This can be done easily by slightly modifying the segment split procedure to maintain this information.

The adaptive algorithm then has two major operations. The first is the same as in the Delaunay Refinement Algorithm; the second replaces the operation (Quality) (as described in Section 2.3) by the operation (Quality′):

(Conformality) If s is a current segment, and there is a committed point that encroaches s, then split s.

(Quality′) If a, b, c are committed points, the circumcircle of the triangle Δabc contains no committed point, $\angle acb < \hat{\kappa}$, the circumcenter, p, of the triangle is inside Ω and either (i) both a, b are midpoints on distinct nondisjoint input segments, sharing input endpoint x, and $\angle axb > \pi/3$, or (ii) a, b are not midpoints on adjoining input segments, then attempt to commit p. If, however, the point p encroaches any current segment, then do not commit to point p, rather in this case split one, some, or all of the current segments which are encroached by p.

In summary, the algorithm removes angles smaller than $\hat{\kappa}$ except when the opposite edge spans a small angle in the input, in which case the small output angles are ignored. For this variant we call $\hat{\kappa}$ the *output angle parameter*; the output mesh may well contain angles smaller than $\hat{\kappa}$. We will let α be the minimum angle in the output mesh.

5.2 *Is Adaptivity Necessary?*

We here make the claim that the Delaunay Refinement Algorithm is as good as its adaptive variant when the latter runs with a small output angle parameter $\hat{\kappa}$. The claim is formalized as follows:

Claim 5.2.1. Suppose that we can guarantee that if the Adaptive Delaunay Refinement Algorithm is run with output angle parameter $\hat{\kappa}$, on any input with minimum input angle θ^* that (a) the algorithm terminates, (b) no angle of the output mesh is smaller than $\hat{\kappa}$, (c) no angle is larger than $\pi - 2\omega$, and (d) output mesh edges are graded with the local feature size by some constant.

Then if the Delaunay Refinement Algorithm is run on any input with minimum input angle θ^*, using output angle parameter $\kappa = \hat{\kappa}$, then (a) the algorithm terminates, (b) no angle of the output mesh is smaller than κ, (c) no angle is larger than $\pi - 2\omega$, and (d) output mesh edges are graded with the local feature size by the same constants as above.

Proof. The Adaptive Delaunay Refinement Algorithm only attempts to remove a Delaunay triangle if it has minimum angle smaller than $\hat{\kappa}$. Moreover, it produces meshes with no angle smaller than $\hat{\kappa}$. Then the (QUALITY′) operation could be rewritten as follows:

> (QUALITY′) If a, b, c are committed points, the circumcircle of the triangle Δabc contains no committed point, $\angle acb < \hat{\kappa}$, and the circumcenter, p, of the triangle is inside Ω then attempt to commit p. If, however, the point p encroaches any current segment, then do not commit to point p, rather in this case split one, some, or all of the current segments which are encroached by p.

This is the same as the operation (QUALITY) of the Delaunay Refinement Algorithm. □

The analysis that follows should be read with a tacit understanding that it can be applied to the Delaunay Refinement Algorithm as well, if κ is set propertly. For example, it will be shown that if an input with $\theta^* \approx 36.53°$ conforms to Assumption 5.4.1, then the Adaptive Delaunay Refinement Algorithm with $\hat{\kappa} = 26.45°$ will terminate leaving no angle in the output mesh smaller than $\hat{\kappa}$, and no angle larger than $\pi - 2\hat{\kappa}$. Then we can immediately claim that the Delaunay Refinement Algorithm (*i.e.*, Ruppert's Algorithm) with $\kappa = 26.45°$ will also terminate on the same input, and with the same grading guarantees.

Thus the adaptive variant is only necessary when θ^* is small, say smaller than about $36.53°$. When θ^* *is* small, the adaptive variant will remove small angles where this is possible, *i.e.*, away from small input angles.

5.3 Circumcenter Sequences

We here analyze sequences of triangle circumcenters, which will be used much like midpoint sequences were used in Chapter 2 to analyze the Delaunay Refinement Algorithm. The analysis of circumcenter sequences is relatively simple, but allows examination of the interaction between segment midpoints, which can be more complex.

Definition 5.3.1. A *circumcenter sequence* is a sequence of points, $\{b_i\}_{i=0}^{l-1}$ such that for $i = 1, 2, \ldots, l-1$, b_i is the circumcenter of a triangle in which b_{i-1} is the more recently committed point of an edge opposite an angle less than $\hat{\kappa}$. The point b_0 may be an input point or segment midpoint.

For $i = 0, 1, \ldots, l-2$, let a_i be the *other* endpoint of the short edge of which b_i is the more recently committed endpoint. In the case where a_0, b_0 are both input points, they are committed simultaneously; we imagine a total order on input points which determines the tie. Both a_0, b_0 may be midpoints on distinct, nondisjoint input segments. In this case we assume that the triangle with circumcenter b_1 was removed by a (QUALITY′) operation because of a small angle opposite a_0, b_0. In particular this means that we assume the angle subtended by the input segments containing a_0, b_0 is at least $\pi/3$ in this case.

When talking about such sequences, for $i = 1, 2, \ldots, l-1$, let $\tilde{r}_i$ be the circumradius of the triangle associated with b_i. Note that $\tilde{r}_i = |b_i - b_{i-1}| = |b_i - a_{i-1}|$, and that $|a_i - b_i| \geq \tilde{r}_i$. We let $\tilde{r}_0 = |b_0 - a_0|$, *i.e.*, the length of the first short edge.

Note that for a circumcenter sequence, $\{b_i\}_{i=0}^{l-1}$, the points $b_1, b_2, \ldots, b_{l-2}$ are circumcenters which have been committed, b_{l-1} is a circumcenter, though it may be rejected, and b_0 may be any type of point. If b is a triangle circumcenter, there is always a circumcenter sequence ending with b, although it may be a trivial sequence of two elements. Any circumcenter sequence whose first element, b_0, is a triangle circumcenter may be extended to a maximal sequence whose first element is either a segment midpoint or an input point.

The following geometric lemma is the key result which allows us to make the $\arcsin 2^{-7/6}$ output guarantee. It essentially states that only circumcenter sequences longer than a certain length can "turn" around a 180° feature.

Lemma 5.3.2. *Let S_1, S_2 be two segments with disjoint interiors on a common line, L. Assume that $|S_2| \leq |S_1|$, i.e., S_2 is no longer than S_1. Let b_0 be the midpoint of S_1, and let a_0 be some other point. Let $\{b_i\}_{i=1}^{l-1}$ be a circumcenter sequence such that b_{l-1} is inside the diametral circle of S_2, and such that b_1 is the circumcenter of a triangle with edge (a_0, b_0) opposite an angle smaller than $\hat{\kappa}$. Then $l \geq 4$.*

Note that unlike in the regular terminology of circumcenter sequences, this lemma makes no assumptions about which of a_0, b_0 was committed first. This is why we have chosen to index the circumcenter sequence from $i = 1$ instead of the usual $i = 0$.

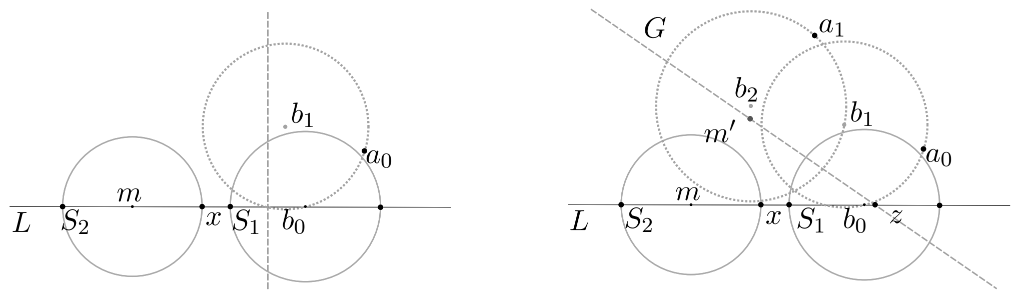

(a) b_1 does not encroach S_2. (b) b_2 does not encroach S_2.

Figure 19: The head of a circumcenter sequence is shown; the point b_1 must be to the right of the bisector of b_0 and x, and so it cannot encroach S_2, which is on the other side of this bisector, as shown in (a). In (b) the bisector of b_1 and the point x is shown. Since b_2 cannot be closer to x than to b_1, and since the diametral circle of S_2 is on the opposite side the bisector, b_2 cannot encroach S_2. In this case, a_0 is shown to be outside the diametral circle of S_1. This is not a necessary hypothesis for this lemma.

Proof. The basic argument is sketched in Figure 19. The point b_1 is the circumcenter of a triangle whose circumcircle does not contain the point x, which is the endpoint of S_1 closer to S_2. However, this circumcircle has b_0 on it, so b_1 must be in the closed halfspace defined by the bisector of x and b_0 and which does not contain x, as shown in Figure 19(a). Thus b_1 cannot be in the diametral circle of S_2, which is in the open halfspace on the other side of this bisector.

Now let G be the bisector of the points b_1 and x. The point b_2 is the center of a circle which does not contain x, but has b_1 on its boundary, since b_1 is one of the vertices of the triangle which b_2 is added to remove. Thus b_2 must be either on the line G, or in the open halfspace defined by G that is closer to the point b_1. In Figure 19(b), this is the halfspace to the upper right of G.

It then suffices to show that the closure of the diametral ball of S_2 is contained in the other open halfspace defined by G, and thus b_2 cannot encroach S_2.

Let z be the intersection of L and G; take m to be the midpoint of S_2, and m' is its projection onto G. Let x' be the projection of x onto G. Let y be the projection of b_1 onto L. See Figure 20. The point x is clearly between m and z, otherwise x would be in the halfspace closer to b_1 than to x, a contradiction. Thus $|m - z| = |m - x| + |x - z|$.

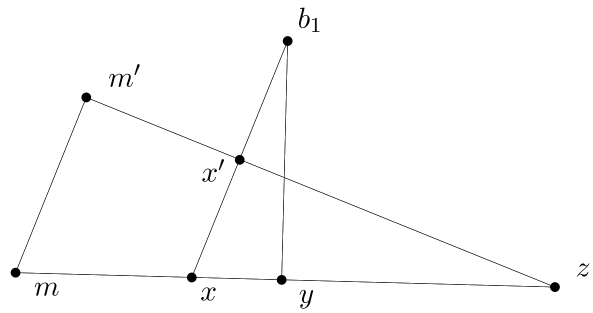

Figure 20: The geometric heart of the argument is shown, with three congruent triangles, $\Delta mm'z, \Delta xx'z, \Delta xyb_1$.

By congruency of the three triangles of Figure 20,

$$\frac{|m-m'|}{|m-z|} = \frac{|x-x'|}{|x-z|} = \frac{|x-y|}{|x-b_1|}.$$

Let $r = \frac{|S_2|}{2} \le \frac{|S_1|}{2}$, by assumption. Since S_1, S_2 have disjoint interiors, $|m-x| \ge r$. Then $|m-z| \ge r + |x-z|$, so

$$\begin{aligned} |m-m'| &= \frac{|x-x'|\,|m-z|}{|x-z|}, \\ &\ge \frac{|x-x'|\,(r+|x-z|)}{|x-z|}, \\ &\ge \frac{|x-x'|}{|x-z|}r + |x-x'| = \frac{|x-y|}{|x-b_1|}r + |x-x'|. \end{aligned}$$

As noted above, b_1 is to the right of the bisector of x and b_0, so $|x-y| \ge \frac{|x-b_0|}{2} = \frac{|S_1|}{4} \ge \frac{r}{2}$. Note also that $|x-b_1| = 2|x-x'|$. Then

$$|m-m'| \ge \frac{r^2}{4|x-x'|} + |x-x'|.$$

The right hand side is minimized when $|x-x'| = \frac{r}{2}$, where the right hand side has value r. Note, however, that $|x-x'| \ge \frac{\tilde{r}_1}{2} \ge \frac{1}{2\sin\hat{\kappa}}\frac{|S_1|}{4} > \frac{r}{2}$, so the right hand side will be strictly larger than r.

That is, $|m-m'| > r$, and thus the distance from m to G, which is $|m-m'|$, is greater than the radius of the diametral circle of S_2. Then the closed diametral circle of S_2 is contained in the open halfspace opposite b_1, as desired. □

This lemma allows us to prove a better output angle for the Delaunay Refinement Algorithm. Previous proofs required $2\sin\hat{\kappa} \le \frac{1}{\sqrt{2}}$; by the lemma, the following proof only requires that $(2\sin\hat{\kappa})^3 \le \frac{1}{\sqrt{2}}$. A better output angle could be guaranteed if the lemma could be improved; this would have to be via some alternation of the algorithm, as the example of Figure 21 shows the lemma cannot be extended in the naïve setting. While it is conceivable some alternation could be made to increase the number of "steps" around input segments that meet at 180°, such a method could not prevent the "creep" of small subsegments along an input segment. Shewchuk notes this plaintively: "[a] small feature could . . . propagate along the whole length of [a] segment." [43, page 65]

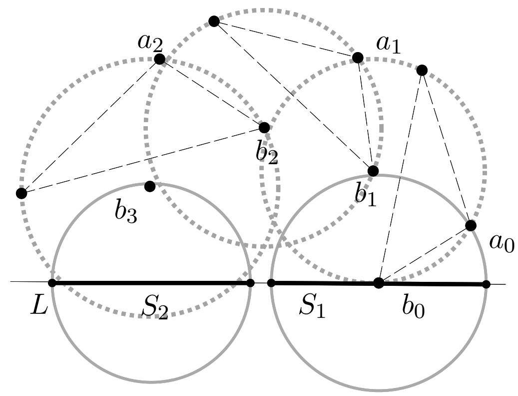

Figure 21: A circumcenter sequence, $\{b_i\}_{i=0}^{3}$, is displayed, which shows that Lemma 5.3.2 cannot be extended. The segments S_1, S_2 are shown, with their diametral circles. The points b_1, b_2, b_3 are circumcenters of triangles (shown) with an angle smaller than $\pi/6$. The point b_3 encroaches S_2.

Since $\hat{\kappa} < \pi/6$, we can establish a geometric series which gives the following lemma and its corollary. The corollary describes how a segment midpoint which is not caused by a midpoint encroaching the segment is caused by some other midpoint or input point.

Lemma 5.3.3. *Suppose $\{b_i\}_{i=0}^{l-1}$ is a circumcenter sequence. For $i > 0$, let $\tilde{r}_i$ be the circumradius associated with b_i. Then for $i = 1, 2, \ldots, l-1$,*

- $\tilde{r}_{i-1} < 2\tilde{r}_i \sin\hat{\kappa}$ *and therefore* $\tilde{r}_i < (2\sin\hat{\kappa})^{l-1-i}\tilde{r}_{l-1}$, *and*
- $|b_{l-1} - b_i| < \frac{\tilde{r}_{l-1}}{1-2\sin\hat{\kappa}}$, *and* $|b_{l-1} - a_i| < \frac{\tilde{r}_{l-1}}{1-2\sin\hat{\kappa}}$.

Proof. By definition, b_i is the circumcenter of a triangle of radius $\tilde{r}_i$, which has a short edge no shorter than $\tilde{r}_{i-1}$ opposite an angle less than $\hat{\kappa}$. By the sine rule, then $2\tilde{r}_i \sin\hat{\kappa} > \tilde{r}_{i-1}$.

Using this repeatedly gives $\tilde{r}_i < (2\sin\hat{\kappa})^{l-i}\tilde{r}_{l-1}$. Since $2\sin\hat{\kappa} < 1$, we may bound the distance from b_i to b_{l-1} by the geometric series, as follows:

$$\begin{aligned}
|b_{l-1} - b_i| &\leq |b_{l-1} - b_{l-2}| + |b_{l-2} - b_{l-3}| + \ldots + |b_{i+1} - b_i|, \\
&\leq \tilde{r}_{l-1} + \tilde{r}_{l-2} + \ldots + \tilde{r}_{i+1}, \\
&< \tilde{r}_{l-1} + (2\sin\hat{\kappa})\tilde{r}_{l-1} + \ldots + (2\sin\hat{\kappa})^{l-i-2}\tilde{r}_{l-1}, \\
&< \frac{1}{1-2\sin\hat{\kappa}}\tilde{r}_{l-1}.
\end{aligned}$$

The bound for $|b_{l-1} - a_i|$ follows since $|b_{i+1} - a_i| = |b_{i+1} - b_i| = \tilde{r}_{i+1}$, and the above analysis suffices. □

Corollary 5.3.4. *Suppose that segment s_p with midpoint p and radius r was split, but the segment was not encroached by a committed point. Then there is some maximal circumcenter sequence $\{b_i\}_{i=0}^{l-1}$ such that b_{l-1} "yielded" to p, causing it to be committed. Moreover, $\tilde{r}_i < (2\sin\hat{\kappa})^{l-1-i}\sqrt{2}r_p$, $|p - b_i| \leq \eta r_p$, and $|p - a_i| \leq \eta r_p$, for $i = 0, 1, \ldots, l-1$, with $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$.*

Proof. As in the classical setting (see Figure 10), since b_{l-1} was the center of an empty circumcircle, but encroached s_p, then $\tilde{r}_{l-1} \leq \sqrt{2}r_p$. Using the lemma gives the desired bound on $\tilde{r}_i$. By the lemma, and since $\hat{\kappa} < \pi/6$, $\tilde{r}_i \leq \tilde{r}_{l-1}$. Then

$$|p - b_i| \leq |p - b_{l-1}| + |b_{l-1} - b_i| \leq r_p + \frac{\tilde{r}_{l-1}}{1 - 2\sin\hat{\kappa}} \leq \left(1 + \frac{\sqrt{2}}{1 - 2\sin\hat{\kappa}}\right) r_p = \eta r_p.$$

The bound on $|p - a_i|$ follows, *mutatis mutandis*, as above. ☐

5.4 Good Grading

Since the algorithms under consideration split segments at midpoints, we are interested in the ratio of the lengths of adjoining segments, modulo powers of two. The following assumption on the input can be satisfied by augmenting the input with at most $2|\mathcal{S}|$ points, splitting some or all of the line segments in the input which share an endpoint so they have acceptable lengths–see Algorithm 2 and Algorithm 4 of Chapter 4.

Assumption 5.4.1. In addition to those of Assumption 2.2.1 we make the following assumption:

(a) If S_1, S_2 are two adjoining input segments that meet at angle other than π, then they have the same length modulo a power of two, that is $\frac{|S_1|}{|S_2|} = 2^k$ for some integer k.

Note that input which satisfy Assumption 5.4.1 also satisfy Assumption 3.3.3. Since the Delaunay Refinement Algorithm and the adaptive variant differ only on the test for skinny angles, we claim that Lemma 3.3.6 applies to input satisfying Assumption 5.4.1 and for the adaptive algorithm.

The following lemma uses the length assumption and Lemma 5.3.2 to obtain a similar result when a segment is split due to a triangle circumcenter. In this case we show an alternative: either the local feature size of the midpoint is bounded, or there is some other segment to "blame" for the split, with the radii non-dwindling.

Lemma 5.4.2. *Let s_p be a subsegment of midpoint p and radius r_p. Suppose that s_p was not encroached by a committed point, rather a circumcenter b_{l-1} was proposed to be committed, but rejected in favor of splitting s_p. Let $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$. Then either* *(a)* $\operatorname{lfs}(p) \leq (\sqrt{2} + \eta)r_p$, *or*

(b) there is a segment s_q *with committed midpoint* q, *and radius* r_q *such that*

(i) the input segments containing s_p, s_q *are nondisjoint,*

(ii) $r_q \leq r_p$,

(iii) $|p - q| \leq \eta r_p$, *and*

(iv) if s_p, s_q *are on the same input segment then* $2r_q \leq r_p$; *if they are on distinct input segments sharing input point* x, *then* $|x - p| < \frac{\eta}{\sin\theta^*} r_p$.

Proof. For convenience, we say that a point q "provokes" a point p, if q is committed before p, and p is the midpoint of a segment, s, which is encroached by q.

Let $\{b_i\}_{i=0}^{l-1}$ be a maximal cirumcenter sequence ending with the circumcenter b_{l-1} which caused p to be committed. Consider the identity of b_0:

- If b_0 is an input point, then by definition so is a_0, and so $\operatorname{lfs}(p) \leq |p - b_0| \vee |p - a_0|$. By Corollary 5.3.4, these are both bounded above by ηr_p, so $\operatorname{lfs}(p) < \sqrt{2} + \eta r_p$.
- If b_0 is a midpoint on an input segment disjoint from the one containing p, then by definition of local feature size and using Corollary 5.3.4, $\operatorname{lfs}(p) \leq |p - b_0| \leq \eta r_p$, which suffices.
- Suppose that b_0 is a midpoint on an input segment nondisjoint to the one containing p. Furthermore suppose that a_0 did not provoke b_0. Let s_q be the segment associated with the midpoint $q = b_0$, and let r_q be its radius. By assumption $r_q \leq |a_0 - b_0| = \tilde{r}_0$. We will show the second alternative for this choice of q. By Corollary 5.3.4, $|p - q| = |p - b_0| \leq \eta r_p$.

  If s_p, s_q are on the same input segment, then by Assumption 5.4.1, r_p/r_q is a power of two. We will try to use Lemma 5.3.2, with $S_1 = s_q, S_2 = s_p$. The lemma does not apply if $|S_1| < |S_2|$, but this inequality actually states that $2r_q < 2r_p$, which would imply that $2r_q \leq r_p$. So assuming otherwise, by use of the lemma, $l \geq 4$, and thus since $\hat{\kappa} \leq \arcsin 2^{-7/6}$, using Corollary 5.3.4

  $$r_q < \sqrt{2}(2\sin\hat{\kappa})^3 r_p \leq \sqrt{2}\left(\frac{1}{\sqrt[6]{2}}\right)^3 r_p = r_p.$$

  Thus we have a contradiction, so it must be that $2r_q \leq r_p$.

  If s_p, s_q are on *distinct* input segments that meet at some angle other than π, by item (a) of Assumption 5.4.1, r_p/r_q is a power of two, and so if $r_p < r_q$, then $r_p \leq \frac{r_q}{2}$. But by Corollary 5.3.4, $r_q < (2\sin\hat{\kappa})^{l-1}\sqrt{2}r_p \leq \sqrt{2}r_p$. Thus we would have the contradiction $r_p \leq \frac{r_p}{\sqrt{2}}$, so it must be that $r_q \leq r_p$.

  If, on the other hand, s_p, s_q are on input segments that meet at angle π, then we will again try to use Lemma 5.3.2, with $S_1 = s_q, S_2 = s_p$. The lemma does not apply if

$|S_1| < |S_2|$, *i.e.*, if $2r_q < 2r_p$, which suffices. So assuming otherwise, as above because $\hat{\kappa} \leq \arcsin 2^{-7/6}$, we have the contradiction $r_q < r_p$.

- Suppose that b_0 is a midpoint on an input segment nondisjoint to the one containing p, and a_0 *did* provoke b_0, *i.e.*, $r_q \geq |a_0 - b_0| = \tilde{r}_0$, where again r_q is the radius associated with b_0.

  If a_0 was (on) an input feature disjoint from the one containing b_0, then by definition $\operatorname{lfs}(p) \leq |p - b_0| \vee |p - a_0|$. By Corollary 5.3.4, these are both bounded above by ηr_p, so $\operatorname{lfs}(p) < (\sqrt{2} + \eta) r_p$.

  The point a_0 could not have been a circumcenter, because it provokes b_0. So the only alternative is that it is a midpoint on an input segment nondisjoint from the one containing b_0.[1]

  Since this a circumcenter sequence, by assumption, ϕ, the angle between the segments containing a_0, b_0 is at least $\pi/3$. Let s_q be the segment associated with $q = a_0$, of radius r_q. By Lemma 3.3.6, $\tilde{r}_0 = |a_0 - b_0| \geq 2r_q \sin \frac{\phi}{2} \geq r_q$, since $\phi \geq \pi/3$. Consider the subcases:

  – The input segments containing a_0, p are disjoint. Then by definition, $\operatorname{lfs}(p) \leq |p - a_0| \leq \eta r_p$, by Corollary 5.3.4, as above, which suffices.

  – The input segments containing a_0, p are nondisjoint, moreover the input point x is shared by all three input segments in consideration. We will mirror the argument from above.

    Let s_q be the segment associated with the midpoint $q = a_0$, and let r_q be its radius. By Corollary 5.3.4, $|p - q| = |p - a_0| \leq \eta r_p$. We have already shown that $r_q \leq \tilde{r}_0$, and thus $r_q \leq (2 \sin \hat{\kappa})^{l-1} \sqrt{2} r_p$.

    If s_p, s_q are on the same input segment, as above, using $\hat{\kappa} \leq \arcsin 2^{-7/6}$, and Lemma 5.3.2 gives $2r_q \leq r_p$.

    If s_p, s_q are on distinct input segments that meet at some angle other than π, by the same argument as above, since r_p, r_q have essential ratio 1, it must be that $r_q \leq r_p$.

    If, on the other hand, s_p, s_q are on input segments that meet at angle π, then as above, using Lemma 5.3.2, $r_q < r_p$.

  – The input segments containing a_0, p are nondisjoint, and no input point is shared by the three input segments. Let x be the input point shared by the input segments containing p, b_0. Let y be the input point shared by the input segments containing a_0, b_0. Then $\operatorname{lfs}(p) \leq |p - y| \leq |p - a_0| + |a_0 - y|$. By Lemma 3.1.4,

[1] Note there are now three input segments in consideration, namely the ones containing p, a_0, b_0. It could be the case that p, b_0 are on the same input segment, or even p and a_0 are.

$|a_0 - y| \leq |a_0 - b_0| 2\cos\phi \leq \tilde{r}_0 \leq \sqrt{2}r_p$. Thus using Corollary 5.3.4 yet again, $\mathrm{lfs}(p) \leq \eta r_p + \sqrt{2}r_p$, as desired.

Suppose that alternative (b) holds and that s_p, s_q are on distinct segments sharing input point x. We have already shown that $|p - q| \leq \eta r_p$. But $|p - q|$ is bounded below by the distance from p to the segment containing q, which is $|x - p| \sin((\theta \wedge \pi/2)) \geq |x - p| \sin\theta^*$. Thus $|x - p| < \frac{\eta}{\sin\theta^*} r_p$. □

We now collapse the two cases of midpoint addition which were considered in the previous lemma and Lemma 3.3.6, namely encroached and nonencroached, to get a single local feature size estimate. The proof proceeds by constructing a sequence of midpoints on input segments sharing a common endpoint, then uses the two lemmata and Lemma 5.3.2 to find a grading estimate for segment midpoints. Note that unlike previous good-grading proofs, the grading of midpoints is proved independently of that for circumcenters.

Lemma 5.4.3 (Midpoint Local Feature). *Suppose that the input to the Adaptive Delaunay Refinement Algorithm conforms to Assumption 5.4.1. Then there is a constant, μ, depending on θ^* and $\hat{\kappa}$ such that if p is the midpoint of a segment, s, of radius r that is committed by the algorithm, then* $\mathrm{lfs}(p) \leq \mu r$.

Moreover, $\mu = 4\eta\left(1 + \frac{1}{\sin\theta^*}\right)$ suffices, where $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$, as above.

Proof. We construct a *hybrid sequence*, actually two sequences, $\{(P_i, R_i)\}_{i=0}^{1-l}, \{X_i\}_{i=1}^{l-1}$, where each P_i will be a midpoint of a segment of radius R_i, and P_i and P_{i-1} will be on input features sharing input point X_i. We may optionally define an X_0.

We will use Lemma 3.3.6 and Lemma 5.4.2 to establish the sequence. The sequence is constructed backwards, so for convenience we pretend that we know how long it will be, *i.e.*, we know l, so that we can set $P_{l-1} = p, R_{l-1} = r$. For convenience, let $S_{l-1} = s$. The sequence is constructed backwards to make it analogous with encroachment sequences. Thus we will claim that when P_i is committed, P_{i-1} has already been committed, and is somehow "responsible" for P_i being committed.

Construct the sequence as follows: given segment S_i with committed midpoint P_i, and radius R_i, consider why P_i was committed:

- Suppose that P_i was committed because an input point or a point on a disjoint input feature encroached S_i. Then $\mathrm{lfs}(P_i) \leq R_i$. Let P_i be the first midpoint in the sequence, *i.e.*, $i = 0$ because l was chosen magically.
- On the other hand, suppose S_i was encroached by a point, q, on a nondisjoint input feature. Let the segments share input point x. Since the input conforms to Assumption 5.4.1 we may apply Lemma 3.3.6, which asserts that $R_i \geq r_q$, the radius associated

with q. Since q encroaches S_i we have $|P_i - q| \leq R_i \leq \eta R_i$. Moreover by Claim 3.1.2, $|x - P_i| < \frac{1}{\sin\theta^*} R_i < \frac{\eta}{\sin\theta^*} R_i$.

If $i \neq l-1$ and $x \neq X_{l-1}$ then let $P_0 = P_i$ be the first point in the sequence, and let $X_0 = x$. Otherwise let $X_i = x = X_{l-1}$, let $P_{i-1} = q$, let $R_{i-1} = r_q$, and let S_{i-1} be the parent segment of q.

- If S_i was not encroached by any point, then by Lemma 5.4.2, then either $\operatorname{lfs}(P_i) \leq (\sqrt{2} + \eta) R_i$, in which case let P_i be the first element of the sequence, *i.e.*, let $i = 0$; or there is some midpoint, q, of some subsegment s_q, with useful properties. The lemma asserts that S_i, s_q are on nondisjoint segments. There are two alternatives:
  - If they are on the same input segment, let P_i be the first midpoint of the sequence, *i.e.*, let $i = 0$. Note that in this case, by the lemma, $R_0 \geq 2r_q$, and $|P_0 - q| \leq \eta R_0$.
  - Otherwise, let x be the single input point shared by the two input segments.
    If $i \neq l-1$ and $x \neq X_{l-1}$, then let P_i be the first midpoint in the sequence, and let $X_0 = x, P_0 = q$. By the lemma $|X_0 - P_0| < \frac{\eta}{\sin\theta^*} R_0$. Note also that (X_0, X_{l-1}) is an input segment.
    Otherwise, let $X_i = x = X_{l-1}$, let $P_{i-1} = q$, let $S_{i-1} = s_q$, let R_{i-1} be the radius of S_{i-1}. The lemma asserts that $R_i \geq R_{i-1}$, $|P_i - P_{i-1}| \leq \eta R_i$, and $|X_i - P_i| < \frac{\eta}{\sin\theta^*} R_i$.

We can claim the following facts about the hybrid sequence, noting their similarity to properties proven about encroachment sequences in Chapter 3:

(a) $R_0 \leq R_1 \leq \ldots R_{l-1}$;

(b) $|P_i - P_{i-1}| \leq \eta R_i$;

(c) $|X_i - P_i| \leq \frac{\eta}{\sin\theta^*} R_i$;

(d) $X_i = X_{l-1}$ for $i = 1, 2, \ldots, l-2$.

(e) Either

  (i) $\operatorname{lfs}(P_0) \leq \max\left\{1, \sqrt{2} + \eta\right\} R_0 = (\sqrt{2} + \eta) R_0$, or

  (ii) there is some $X_0 \neq X_{l-1}$ such that (X_0, X_{l-1}) is an input segment and such that $|P_0 - X_0| \leq \frac{\eta}{\sin\theta^*} R_0$, or

  (iii) there is some midpoint q of radius r_q such that $R_0 \geq 2r_q$, and $|P_0 - q| \leq \eta R_0$.

These facts will be enough to establish the lemma. First note that if $l = 1$, it must be that s was encroached by a point on a nondisjoint input feature, so it suffices to take $\mu \geq 1$. So assume otherwise.

We bound $|P_{l-1} - P_0|$; our analysis will be quite an overestimate for the case where $l = 2$, but will suffice. By the triangle inequality, and item (c), item (a), item (d), and item (b) of above,

$$\begin{aligned}
|P_{l-1} - P_0| &\leq |P_{l-1} - X_{l-1}| + |P_1 - X_1| + |P_1 - P_0|, \\
&\leq \frac{\eta}{\sin\theta^*} R_{l-1} + \frac{\eta}{\sin\theta^*} R_1 + \eta R_1, \\
&\leq \left(\eta + \frac{2\eta}{\sin\theta^*}\right) R_{l-1}.
\end{aligned}$$

We consider the sequence head, *i.e.*, item (e).

- If the first alternative holds, then by the Lipschitz condition,

$$\begin{aligned}
\mathrm{lfs}\,(P_{l-1}) &\leq |P_{l-1} - P_0| + \mathrm{lfs}\,(P_0), \\
&\leq \left(\eta + \frac{2\eta}{\sin\theta^*}\right) R_{l-1} + (\sqrt{2} + \eta) R_0, \\
&\leq \left(\sqrt{2} + 2\eta + \frac{2\eta}{\sin\theta^*}\right) R_{l-1},
\end{aligned}$$

so it suffices to take

$$\boxed{\sqrt{2} + 2\eta + \frac{2\eta}{\sin\theta^*} \leq \mu.}$$

- If the second alternative holds, then $X_0 \neq X_{l-1}$. So by definition of local feature size,

$$\begin{aligned}
\mathrm{lfs}\,(P_{l-1}) &\leq |P_{l-1} - X_{l-1}| \vee |P_{l-1} - X_0|, \\
&\leq \frac{\eta}{\sin\theta^*} R_{l-1} \vee (|P_{l-1} - P_0| + |P_0 - X_0|), \\
&\leq \frac{\eta}{\sin\theta^*} R_{l-1} \vee \left(\left[\eta + \frac{2\eta}{\sin\theta^*}\right] R_{l-1} + \frac{\eta}{\sin\theta^*} R_0\right), \\
&\leq \frac{\eta}{\sin\theta^*} R_{l-1} \vee \left(\eta + \frac{3\eta}{\sin\theta^*}\right) R_{l-1}, \\
&\leq \left(\eta + \frac{3\eta}{\sin\theta^*}\right) R_{l-1},
\end{aligned}$$

so it suffices to take

$$\boxed{\eta + \frac{3\eta}{\sin\theta^*} \leq \mu.}$$

- If the third alternative holds, use this lemma inductively on q to find that $\mathrm{lfs}\,(q) \leq \mu r_q \leq \mu \frac{R_0}{2} \leq \mu \frac{R_{l-1}}{2}$. Then by the Lipschitz condition,

$$\begin{aligned}
\mathrm{lfs}\,(P_{l-1}) &\leq |P_{l-1} - P_0| + |P_0 - q| + \mathrm{lfs}\,(q), \\
&\leq \left(\eta + \frac{2\eta}{\sin\theta^*}\right) R_{l-1} + \eta R_0 + \frac{\mu}{2} R_{l-1}, \\
&\leq \left(2\eta + \frac{2\eta}{\sin\theta^*} + \frac{\mu}{2}\right) R_{l-1},
\end{aligned}$$

so it suffices to take

$$\boxed{4\eta + \frac{4\eta}{\sin\theta^*} \leq \mu.}$$

Simple analysis shows that $\mu = 4\eta\left(1 + \frac{1}{\sin\theta^*}\right)$ suffices to satisfy the boxed constraints. $\square$

The following lemma is then immediate.

Theorem 5.4.4 (Adaptive Good Grading). *Suppose that the input to the Adaptive Delaunay Refinement Algorithm conforms to Assumption 5.4.1. Let μ be the constant depending on $\theta^*, \hat{\kappa}$ from Lemma 5.4.3. Then there is a positive constant C such that when the algorithm, operating with a output angle parameter $\hat{\kappa} \leq \arcsin 2^{-7/6}$, commits or attempts to commit the point p then if q is any previously committed point then*

- *If p is the midpoint of a segment encroached by q, which is a midpoint on an adjoining input segment then*

$$\operatorname{lfs}(p) \leq (1 + \frac{\mu}{2\sin\frac{\theta}{2}})\,|p-q| \leq (1 + \frac{\mu}{2\sin\frac{\theta^*}{2}})\,|p-q|\,,$$

  where θ is the angle subtended by the two segments.

- *If p is a midpoint, and either q did not encroach the parent segment of p or is not a midpoint on an adjoining input segment, then*

$$\operatorname{lfs}(p) \leq \mu\,|p-q|\,.$$

- *If p is the circumcenter of a triangle of circumradius r, then*

$$\operatorname{lfs}(p) \leq Cr.$$

Moreover, $C = 1 + 2\sin\hat{\kappa}(1+\mu)$ suffices.

Proof. We determine sufficient conditions on the constant C. Again, for convenience, we say that a point q "provokes" a point p, if q is committed before p, and p is the midpoint of a segment, s, which is encroached by q.

- Suppose p is the midpoint of subsegment which is encroached by q which is a midpoint on a nondisjoint input feature. Let θ be the angle between the two input segments. Let r_q be the radius associated with q. By Lemma 3.3.6, $|p-q| \geq 2r_q \sin\frac{\theta}{2}$. By Lemma 5.4.3, $\operatorname{lfs}(q) \leq \mu r_q$. Then using the Lipschitz condition,

$$\operatorname{lfs}(p) \leq |p-q| + \operatorname{lfs}(q) \leq |p-q| + \mu r_q \leq \left(1 + \frac{\mu}{2\sin\frac{\theta}{2}}\right)|p-q|\,,$$

  as desired.

- Suppose p is the midpoint of a subsegment. Suppose the subsegment is encroached by q which is not such a midpoint. Then q must be on a disjoint input feature and so $\operatorname{lfs}(p) \leq |p-q| \leq \mu\,|p-q|\,.$

If q does not encroach the subsegment, then $|p-q|$ is at least the radius of the subsegment. r. By Lemma 5.4.3, $\mathrm{lfs}\,(p) \leq \mu r \leq \mu\,|p-q|$.

- If p is a circumcenter of a skinny triangle of circumradius r, then let a, b be the vertices of the shortest edge, and θ the angle opposite this edge. By assumption $\theta < \hat{\kappa}$. Note that by the sine rule, $|a-b| = 2r\sin\theta < 2r\sin\hat{\kappa}$. If a, b are both input points, then they are disjoint and by definition $\mathrm{lfs}\,(p) \leq r$, so it suffices to take $1 \leq C$. Otherwise let b be the most recently committed of the two points. We consider the possible identities of b, a:
  - Suppose b is a midpoint of a subsegment. If a is a midpoint on a nondisjoint input segment, with the two segments subtending angle θ, by definition of the (QUALITY′) operation, we know $\theta \geq \pi/3$, thus $2\sin\frac{\theta}{2} \geq 1$. Using the lemma inductively, since a witnesses b, it must be the case that $\mathrm{lfs}\,(b) \leq (1+\mu)\,|a-b|$.
  - If a was not such a midpoint or did not provoke b, then $\mathrm{lfs}\,(b) \leq \mu\,|a-b|$.
  - If b was a circumcenter, then since it was committed after a, its associated circumradius bounds $|a-b|$ from below, so $\mathrm{lfs}\,(b) \leq C\,|a-b|$.

  Then using the Lipschitz condition,

$$\mathrm{lfs}\,(p) \leq r + \mathrm{lfs}\,(b) \leq r + \max\{1+\mu, C\}\,|a-b| \leq (1 + 2\sin\hat{\kappa}\max\{1+\mu, C\})\,r.$$

  And it suffices to ensure that

$$1 + 2\sin\hat{\kappa}\max\{1+\mu, C\} \leq C.$$

In all it suffices to take $C = \max\left\{\frac{1}{1-2\sin\hat{\kappa}}, 1 + 2\sin\hat{\kappa}(1+\mu)\right\}$. Because $\mu \geq \frac{4\eta}{\sin\theta^*} \geq \frac{4\sqrt{2}}{\sin\theta^*\,(1-2\sin\hat{\kappa})}$, we will have $C = 1 + 2\sin\hat{\kappa}(1+\mu)$. □

Note that for $\hat{\kappa} \leq \arcsin 2^{-7/6}$, simple calculation shows that $\eta < 14$, and we may bound μ:

$$\mu \leq 4\eta\left(1 + \frac{1}{\sin\theta^*}\right) \leq 56\left(1 + \frac{1}{\sin\theta^*}\right) = \mathcal{O}\left(\frac{1}{\sin\theta^*}\right).$$

Thus the grading constants of Theorem 5.4.4 are finite for a given input for any acceptable $\hat{\kappa}$. This is in marked contrast to the constants normally found in grading guarantees, which are unbounded as the output angle approaches a limit value. This also addresses the discrepancy between the classical grading proofs, which predict that point density is unbounded as $\hat{\kappa}$ approaches $\arcsin\frac{1}{2\sqrt{2}}$, and the experimental observation that Ruppert's algorithm generally converges for all κ less than $\pi/6$. The fact that η is finite for all $\hat{\kappa} < \pi/6$ is a promising sign that the output angle guarantee may be improved to match the empirical evidence.

5.4.1 Asymptotic Grading Optimality

Lemma 5.4.3 tells us that when a segment s with radius r and midpoint p is split, then

$$\operatorname{lfs}(p) \leq \mu r = \mathcal{O}\left(\frac{1}{\theta^*}\right) r.$$

Can we do better? We argue that we cannot do better asymptotically, by presenting a counterexample.

Given n large, let $\theta = \pi/n$, and let the input consist of the following points: o, the origin, $v_i = (\cos i\theta, \sin i\theta)$, for $i = 0, 1, \ldots, n$, and a point $p = \left(\frac{1}{2} + \epsilon, -\epsilon\right)$. We let the segments of the input be the segments (o, v_i) for $i = 0, 1, \ldots, n$. Such an input is demonstrated in Figure 22(a) for $n = 34$.

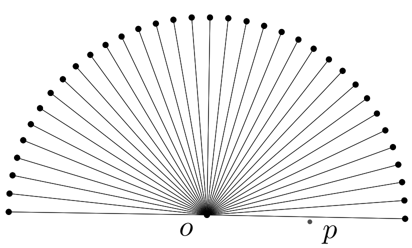

(a) The counterexample lower bound for $n = 34$.

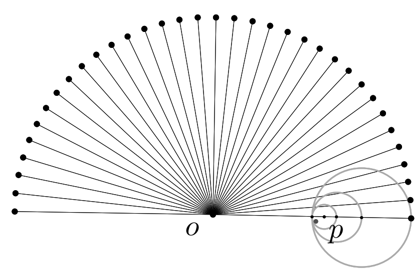

(b) The point p causes splits on (o, v_0).

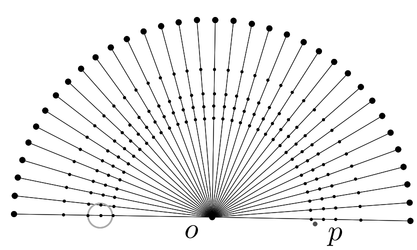

(c) The splits spread to (o, v_n).

Figure 22: The asymptotic lower bound for μ is illustrated. In (a), the input is shown. A small feature near $(0.5, 0)$ causes a number of splits on the segment (o, v_0), as shown in (b). Because θ is small, the splits propagate around onto (o, v_n), as shown in (c), where a very small segment is seen to be encroached. The midpoint of this small segment has local feature size on the order of $\frac{1}{2}$, but the radius of the small segment is on the order of θ.

Employing the (Conformality) rule, the input segment (o, v_0) will be repeatedly bisected into subsegments, adding midpoints of the form $\left(0.5 + 2^{-k}, 0\right)$ for $k = 2, 3, 4, \ldots, m,$

where m is the smallest number such that $2^{-m} < \sqrt{2}\epsilon$. This is shown in Figure 22(b). By making ϵ small, m can be made arbitrarily large.

A ghastly calculation reveals that a point $\left(0.5 + 2^{-k}, 0\right)$ will encroach on a segment with endpoints $\frac{1}{2}\left(\cos\theta, \sin\theta\right)$ and $\left(\frac{1}{2} + 2^{1-k}\right)\left(\cos\theta, \sin\theta\right)$, if

$$\cos\theta > 1 - \frac{1}{2}\left(\frac{2^{-k}}{\frac{1}{2} + 2^{-k}}\right)^2.$$

This is approximately equivalent to the condition $\theta \leq 2^{0.5-k}$. If this segment has length 2^{1-k}, its radius is 2^{-k}. Thus a subsegment will be produced on (o, v_1) of length $\Theta\left(\theta\right)$.

The midpoints will percolate through the (o, v_i), as the argument above can be used to show that the midpoints on (o, v_{i-1}) will encroach on subsegments on (o, v_i). This creates the situation shown in Figure 22(c). A segment of length approximately 2^{-k} exists on (o, v_n), but the local feature size of the midpoint is around $\frac{1}{2}$. If the radius of this segment is r, and m is its midpoint, then

$$\mathrm{lfs}\left(m\right) = \Omega\left(\frac{1}{\theta}\right) r,$$

which establishes asymptotic optimality of μ.

Note, however, this is only a proof that the *analysis* was not too sloppy; the algorithm may still be suboptimal. For example, a smarter algorithm might split the segment (o, v_0) not at the midpoint but a point closer to p. This could result in a single midpoint sufficing to ensure no subsegment on (o, v_0) is encroached. The split would be repeated on each (o, v_i), but no small subsegment would be created, as shown in Figure 23.

This example illustrates the importance of finding a better way to ensure conformality of the input segments. To find such a technique it might be enlightening to study the open problem of Conforming Delaunay Triangulations, though it is not clear how the current best solution to the problem, found by Edelsbrunner and Tan, could be applied to the meshing problem [20].

5.5 Output Quality

Before addressing optimality, we consider the output quality. Recall that the algorithm may leave behind angles smaller than the parameter $\hat{\kappa}$. We will show that small output angles are not too much smaller than a nearby small input angle. The following simple geometric claim gives the output quality guarantee; the idea is to use it with facts about midpoints, the definition of (Quality′), and the Delaunay property to get the bound on output angles.

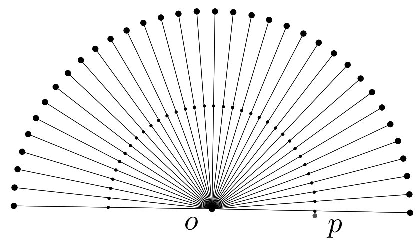

Figure 23: A smarter algorithm might handle the input of Figure 22 better. If the algorithm split (o, v_0) at a point closer to p, and these midpoints percolated to each input segment, the algorithm could guarantee each segment was represented by subsegemnts which were not encroached, and no subsegment was too much smaller than the local feature size of its midpoint, independent of θ^*.

Lemma 5.5.1. *Let x, s, q be three distinct noncollinear points. Let p be a point on the open line segment from x to s. Suppose that $|p - s| \leq |x - p| \leq |x - q|$. Let $\theta = \angle pxq$, and $\phi = \angle psq$. Then*

$$\phi \geq \arctan\left(\frac{\sin\theta}{2 - \cos\theta}\right).$$

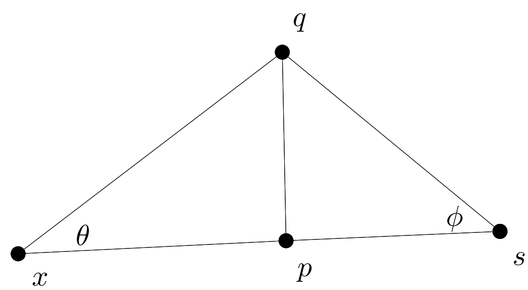

Figure 24: When $|p - s| \leq |x - p| \leq |x - q|$, then $\phi \geq \arctan\left(\frac{\sin\theta}{2-\cos\theta}\right)$.

Proof. The hypothesis is illustrated in Figure 24. By the sine rule,

$$\frac{|x - s|}{\sin\angle xqs} = \frac{|x - q|}{\sin\phi},$$

so $\sin\phi = \frac{|x-q|\sin\angle xqs}{|x-s|}$. Since $|p - s| \leq |x - p| \leq |x - q|$, then

$$\sin\phi \geq \frac{\sin\angle xqs}{2}.$$

Clearly $\angle xqs + \theta + \phi = \pi$, since these are angles of a triangle. So $\sin xqs = \sin(\pi - \theta - \phi) =$

$\sin(\theta + \phi)$. Thus

$$\begin{aligned} \sin\phi &\geq \frac{\sin(\theta+\phi)}{2} = \frac{\sin\theta\cos\phi + \cos\theta\sin\phi}{2} \\ 2\sin\phi - \cos\theta\sin\phi &\geq \sin\theta\cos\phi \\ \sin\phi &\geq \frac{\sin\theta\cos\phi}{2-\cos\theta} \end{aligned}$$

If ϕ is obtuse, then the result holds, as the arctangent is restricted to $(-\pi/2, \pi/2)$. Otherwise ϕ is acute, so $\cos\phi$ is nonnegative, giving $\tan\phi \geq \frac{\sin\theta}{2-\cos\theta}$, which suffices. □

The following claim is a simple consequence of Thales' theorem.

Claim 5.5.2 *(Edge-Apex Rule).* Given a triangle Δpqr in the Delaunay Triangulation of a set of points, $\mathcal{P}$, with L the line through p, q, then $\angle prq \geq \angle pr'q$ for every $r' \in \mathcal{P}$ that is on the same side of L as p, with equality only holding in the case of degeneracy.

We can now state the output guarantee.

Lemma 5.5.3. *Suppose the Adaptive Delaunay Refinement Algorithm terminates for a given input. Let Δpqr be a triangle in the output triangulation. Then either*

(a) The angle $\angle prq > \hat{\kappa}$, or

(b) the points p and q are midpoints on adjoining input segments which meet at angle $\theta < \pi/3$ and

$$\angle prq \geq \arctan\left(\frac{\sin\theta}{2-\cos\theta}\right).$$

Consequently no angle in the output mesh is smaller than $\min\left\{\hat{\kappa}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$.

Proof. Supposing that $\angle prq \leq \hat{\kappa}$, by the definition of the Adaptive Delaunay Refinement Algorithm, it must be that p, q are midpoints on an adjoining input segment, meeting at an angle, θ, less than $\pi/3$. Let x be the input point common to these segments. Without loss of generality, assume that $|x - p| \leq |x - q|$. The midpoint p is the endpoint of two subsegments of this input segment; let the one farther from x be (p, s). By Claim 3.1.1, $|p - s| \leq |p - x|$. Then by Lemma 5.5.1, $\angle psq \geq \arctan\left(\frac{\sin\theta}{2-\cos\theta}\right)$. Letting L be the line through p, q, consider the location of r:

- Suppose r is on the same side of L as x. By Claim 5.5.2, $\angle prq \geq \angle pxq = \theta > \arctan\left(\frac{\sin\theta}{2-\cos\theta}\right)$.
- If r is on the same side of L as s, by Claim 5.5.2, $\angle prq \geq \angle psq \geq \arctan\left(\frac{\sin\theta}{2-\cos\theta}\right)$.

□

The following corollary gives an *upper* bound on output angles that depends on the output angle parameter, $\hat{\kappa}$, but not on the minimum output angle. Given $\hat{\kappa} = \arcsin 2^{-7/6} \approx 26.45°$, it guarantees no output angle is bigger than about $\pi - 2\arcsin\frac{\sqrt{3}-1}{2} \approx 137.1°$.

Corollary 5.5.4. *If Δpqr is a triangle in the output triangulation produced by the Adaptive Delaunay Refinement Algorithm, then*

$$\angle pqr \leq \max\left\{\pi - 2\hat{\kappa}, \pi - 2\arcsin\frac{\sqrt{3}-1}{2}\right\}.$$

Proof. Without loss of generality, assume that $\angle prq$ is the smallest angle of triangle Δpqr. We first prove that

$$\angle pqr \leq \min_{\kappa \leq \hat{\kappa}}\left[(\pi - 2\kappa) \vee \frac{2}{3}(\pi + \arcsin 2\sin\kappa - \kappa)\right].$$

Pick $\kappa \leq \hat{\kappa}$. Considering the two alternatives of Lemma 5.5.3: in the first case $\angle prq \geq \hat{\kappa} \geq \kappa$, and thus, since it is the smallest angle of the triangle, then $\angle pqr \leq \pi - 2\hat{\kappa} \leq \pi - 2\kappa$; So suppose the second case holds, *i.e.*, that p, q are midpoints on adjoining input segments which meet at angle $\theta < \pi/3$, and $\angle prq \geq \arctan\left(\frac{\sin\theta}{2-\cos\theta}\right)$. By trigonometry, if $\theta \geq \arcsin 2\sin\kappa - \kappa$, then $\arctan\left(\frac{\sin\theta}{2-\cos\theta}\right) \geq \kappa$, in which case, again, $\angle pqr \leq \pi - 2\kappa$. So assume otherwise. We will show that $\angle pqr \leq \frac{2}{3}(\pi + \theta)$.

Let p be on input segment S_1, let q be on S_2; let x be the input point shared by S_1,S_2. Assume that r is not on the same side of the line segment (p, q) as x. In the case where r *is* on the same side of the segment as x, an argument similar to that which follows can show that $\angle pqr \leq \frac{2}{3}(\pi - \theta)$.

Because the output of the mesh respects the input segments, it must be the case that $\angle pqr$ is smaller than the angle subtended by (p, q) and S_2. That is, r is "between" S_1 and S_2, as shown in Figure 25(a). This is an external angle of triangle Δxpq at q, thus has magnitude $\theta + \angle xpq$. Then if $\angle xpq \leq \pi/2$, we can bound

$$\angle pqr \leq \theta + \angle xpq \leq \theta + \pi/2 = \frac{2}{3}\theta + \frac{1}{3}\theta + \pi/2 \leq \frac{2}{3}\theta + \frac{\pi}{9} + \pi/2 = \frac{2}{3}\theta + \frac{11}{18}\pi,$$

because $\theta < \pi/3$, which suffices. So assume $\angle xpq$ is obtuse.

Let s be the point on the line containing S_1 such that $\angle pqs = \pi/2$; because we have assumed $\angle xpq$ is obtuse, there is such a point and it is on the same side of (p, q) as r. Let C_1 be the circumcircle of p, q, s; it has center O on the line through S_1. Because line segments in the output are not encroached it must be the case that there is a vertex of the mesh on the line segment (p, s), since $\angle pqs = \pi/2$. It could be the case that this point is s itself. By Claim 5.5.2 it must be the case that r is inside or on C_1, as otherwise, by Thales' theorem, Δpqr would not have the Delaunay property.

Since $\angle prq$ is the smallest angle of the triangle, then (p, q) is the shortest edge of the triangle. Then if C_2 is the circle centered at q of radius $|p - q|$, it must be the case that r is not inside C_2. The point p is a point of intersection of C_1, C_2; let t be the other. See Figure 25(b).

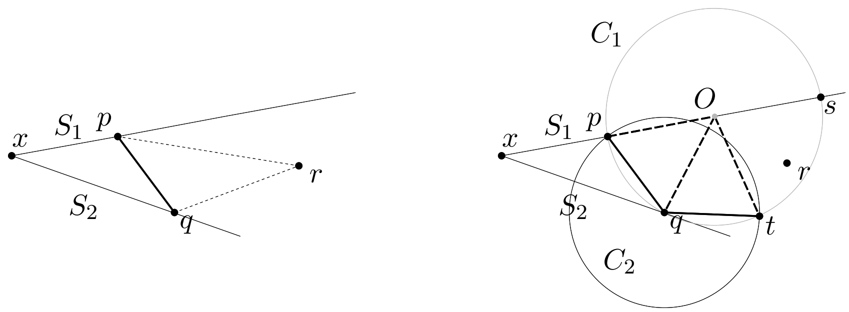

(a) Input (b) With Circles

Figure 25: The proof of Corollary 5.5.4 is shown, for the case where r is opposite (p, q) from x. The point r must be between S_1,S_2, so $\angle pqr$ is smaller than the angle subtended by (p, q) and S_2. The point r must be inside the circle C_1, as otherwise some point on S_1 subtends a larger angle to (p, q) than r does, violating the Delaunay property of the output. By assumption, (p, q) is the shortest edge of the triangle Δpqr so r is outside the circle C_2 . Thus $\angle pqr \leq \angle pqt$. These two bounds together give $\angle pqr \leq \frac{2}{3}(\pi + \theta)$.

Then $\angle pqr \leq \angle pqt$. Looking at the two congruent isosceles triangles of Figure 25(b), it is clear that $\angle pqt$ is twice the external angle of Δxpq at p, that is $\angle pqt = 2(\pi - \angle xpq)$. We have already seen that $\angle pqr \leq \theta + \angle xpq$, thus

$$\angle pqr \leq \min \{2(\pi - \angle xpq), \theta + \angle xpq\} .$$

We have assumed that $\pi/2 \leq \angle xpq \leq \pi - \theta$. Over this range the terms cross at $\angle xpq = \frac{2\pi - \theta}{3}$, and thus $\angle pqr \leq \frac{2}{3}(\pi + \theta)$, as desired.

Now we note that $\pi - 2\kappa$ is decreasing with increasing κ, while $\frac{2}{3}(\pi + \arcsin 2 \sin \kappa - \kappa)$ is increasing. A calculation shows that they cross when $\kappa = \arcsin \frac{\sqrt{3}-1}{2}$. Thus if $\hat{\kappa} \leq \arcsin \frac{\sqrt{3}-1}{2} \approx 21.47°$, then $\angle pqr$ is smaller than $\pi - 2\hat{\kappa}$. If $\hat{\kappa} \geq \arcsin \frac{\sqrt{3}-1}{2}$, then $\angle pqr$ is no larger than $\pi - 2 \arcsin \frac{\sqrt{3}-1}{2} \approx 137.1°$. □

In Figure 26, the guaranteeable minimum output angle is plotted versus θ^*. The bound achievable by Shewchuk's Terminator algorithm, $\arcsin \left[\sin\left(\frac{\theta^*}{2}\right)/\sqrt{2}\right]$, is also plotted [43]. As was noted in Chapter 1, the minimum angle bound for the Terminator is tight. Clearly the Adaptive Delaunay Refinement Algorithm offers significant improvement in minimum angle bounds over the Terminator.

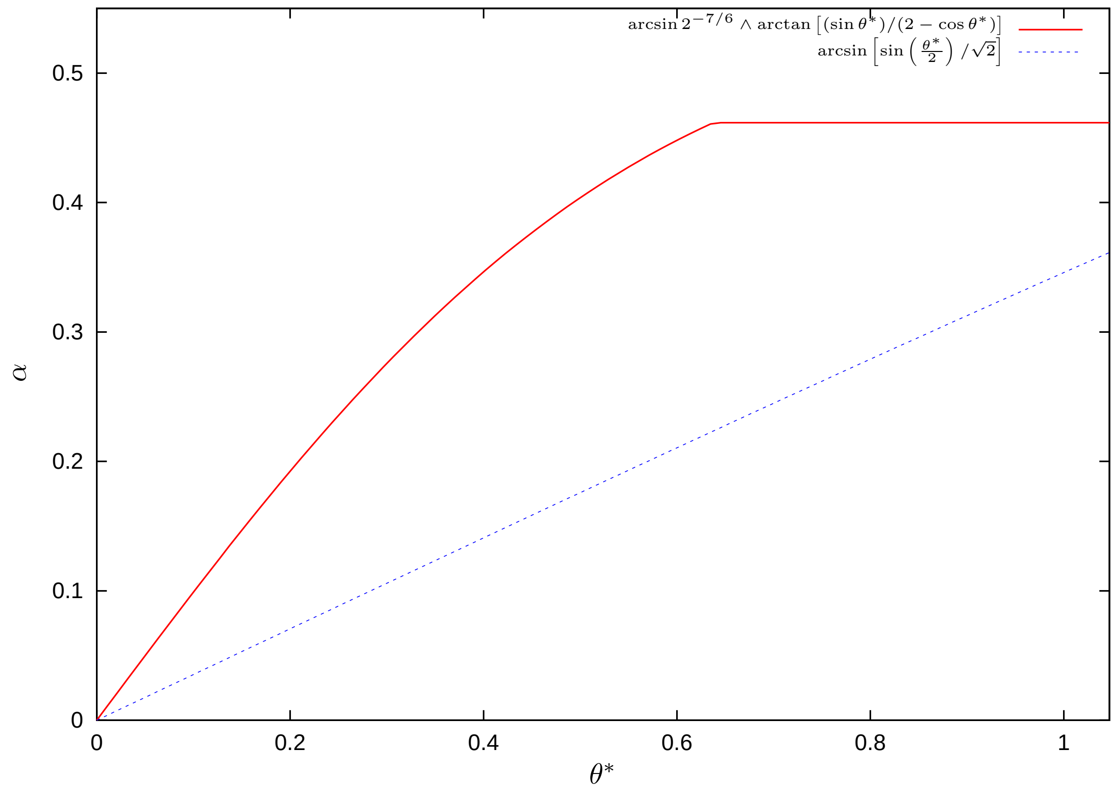

Figure 26: The maximum minimum output angle guaranteeable by the Adaptive Delaunay Refinement Algorithm, *i.e.*, $\arcsin 2^{-7/6} \wedge \arctan\left[(\sin\theta^*)/(2-\cos\theta^*)\right]$, is plotted versus minimum input angle, θ^*. The previous best known bound, $\arcsin\left[\sin\left(\frac{\theta^*}{2}\right)/\sqrt{2}\right]$, is also shown.

5.6 Termination and Optimality

We could rely on the work of Corollary 2.5.7 to make our termination and optimality claims; however, the form of Theorem 5.4.4 indicates that we can make a slightly improved optimality claim. The key idea is that segment midpoints may be near other segment midpoints because of a small angle in the input, but they should be relatively farther away from other segment midpoints on the same input segment. Thus we make our optimality claim by packing disks about segment midpoints that are within input segments. To make this claim, we refer to Chapter 7, where we alter Mitchell's results to allow this one-dimensional packing.

Corollary 5.6.1. *Suppose an input,* $(\mathcal{P}, \mathcal{S})$, *that conforms to Assumption 5.4.1 is fed to the Adaptive Delaunay Refinement Algorithm, using output angle parameter,* $\hat{\kappa} \leq \arcsin 2^{-7/6}$. *Then the algorithm terminates, outputting the Delaunay Triangulation of the set of points* $\mathcal{P}'$, *and with no angle in the triangulation less than* $\alpha = \min\left\{\hat{\kappa}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$, *and with no angle larger than* $\max\left\{\pi - 2\hat{\kappa}, \pi - 2\arcsin\frac{\sqrt{3}-1}{2}\right\}$.

Let $\mathcal{P}_m$ be the Steiner midpoints added by the algorithm, and let $\mathcal{P}_t$ be the Steiner circumcenters. Then we can make the following bounds:

$$\begin{aligned} |\mathcal{P}_m| &\leq 1.02\,(\mu+1)\int_{\mathcal{S}}\frac{1}{\mathrm{lfs}\,(x)}\mathrm{d}x = \mathcal{O}\left(\frac{1}{\theta^*}\right)\int_{\mathcal{S}}\frac{1}{\mathrm{lfs}\,(x)}\mathrm{d}x, \qquad (7)\\ |\mathcal{P}_t| &\leq \frac{1}{\pi}\,(2\mu+3)^2\int_{\Omega}\frac{1}{\mathrm{lfs}^2\,(x)}\mathrm{d}x = \mathcal{O}\left(\left(\frac{1}{\theta^*}\right)^2\right)\int_{\Omega}\frac{1}{\mathrm{lfs}^2\,(x)}\mathrm{d}x. \qquad (8)\end{aligned}$$

Then, if there is any *triangulation on a set of points $\mathcal{P}''$ that conforms to the original input and has minimum angle α then*

$$\left|\mathcal{P}'\right| = \mathcal{O}\left(\alpha^{-3}\right)\left|\mathcal{P}''\right|. \qquad (9)$$

Note that without making separate packing arguments, we could have only been able to claim that $|\mathcal{P}'| = \mathcal{O}\left(\alpha^{-5}\right)|\mathcal{P}''|$, which represents a more serious asymptotic loss of optimality.

Proof. By Theorem 5.4.4, termination follows as in Corollary 2.5.7. We now pack the points separately:

1. **Circumcenters:** The argument here is nearly identical to that in Corollary 2.5.7. First we recall that

$$\begin{aligned}\eta &= 1+\frac{\sqrt{2}}{1-2\sin\hat{\kappa}} \geq 1+\sqrt{2}, \text{ and}\\ \mu &= 4\eta\left(1+\frac{1}{\sin\theta^*}\right) \geq 4\left(1+\sqrt{2}\right)\left(1+\frac{2}{\sqrt{3}}\right) > 20.\end{aligned}$$

But since $\hat{\kappa} < \arcsin 2^{-7/6}$, then $C = 1+2\sin\hat{\kappa}\,(1+\mu) < 1.9+0.9\mu < 2.0+0.9\mu < \mu$. Now, for each circumcenter, p, consider a ball of radius $r_p = \frac{\mathrm{lfs}\,(p)}{2(\mu+1)}$ about p. We claim that the balls are disjoint, *and* no input point or segment midpoint is inside any such ball. If p, p' are two circumcenters, then using the Lipschitz condition and Theorem 5.4.4, we have $\mathrm{lfs}\,(p) \leq (C+1)\,|p-p'| < (\mu+1)\,|p-p'|$, and thus $r_p < \frac{|p-p'|}{2}$. Similarly for $r_{p'}$. If q is an input point, it is committed before p and so by the theorem $\mathrm{lfs}\,(p) \leq C\,|p-q|$, which implies that r_p is smaller than half $|p-q|$. Now consider the case where q is a segment midpoint. If q was committed before p, then the argument of above suffices. If q was committed after p, then we claim that $|p-q|$ is larger than r_q, the radius associated with q, since otherwise p would not have been committed but would have yielded to some segment midpoint. Thus $\mathrm{lfs}\,(q) \leq \mu\,|p-q|$, so $\mathrm{lfs}\,(p) \leq (\mu+1)\,|p-q|$, which shows that r_p is smaller than half $|p-q|$.

Now we can use the argument of Corollary 2.5.7, namely that $r_p < \mathrm{d}_1(p)$ for each circumcenter p. This implies that all of the ball about p is inside Ω, because p is not on $\partial\Omega$. Thus

$$\int_\Omega \frac{1}{\mathrm{lfs}^2(x)}\mathrm{d}x \geq \sum_{p\in\mathcal{P}_t}\int_{B_p}\frac{1}{\mathrm{lfs}^2(x)}\mathrm{d}x \geq \frac{\pi}{(2(\mu+1)+1)^2}\left|\mathcal{P}_t\right|,$$

follows as in Corollary 2.5.7, establishing equation 8.

2. **Midpoints:** Now we pack a one-dimensional disc about each segment midpoint. If p is a segment midpoint, let $r_p = \frac{\mathrm{lfs}(p)}{2(\mu+1)}$, and let B_p be the one-dimensional interval about p of radius r_p running along the input segment containing p. By Theorem 5.4.4 and the Lipschitz property, the B_p do not intersect one another. Thus we have

$$\int_{\mathcal{S}}\frac{1}{\mathrm{lfs}(z)}\mathrm{d}z \geq \sum_{p\in\mathcal{P}_m}\int_{B_p}\frac{1}{\mathrm{lfs}(z)}\mathrm{d}z \geq \left|\mathcal{P}_m\right| 2\ln\left|\frac{r_p+\mathrm{lfs}(p)}{\mathrm{lfs}(p)}\right| = \left|\mathcal{P}_m\right| 2\ln\left|1+\frac{1}{2(\mu+1)}\right|.$$

Then, using the fact that $\mu > 20$, it can be shown that

$$\left|\mathcal{P}_m\right| \leq \frac{1}{2\ln\left|1+\frac{1}{2(\mu+1)}\right|}\int_{\mathcal{S}}\frac{1}{\mathrm{lfs}(z)}\mathrm{d}z < 1.02(\mu+1)\int_{\mathcal{S}}\frac{1}{\mathrm{lfs}(z)}\mathrm{d}z,$$

establishing equation 7.

By Theorem 11 of Mitchell [31],

$$\int_\Omega \frac{1}{\mathrm{lfs}^2(x)}\mathrm{d}x \leq \left(\frac{21.5}{\alpha}+11.9\right)\left|\mathcal{P}''\right| = \mathcal{O}\left(\frac{1}{\alpha}\right)\left|\mathcal{P}''\right|.$$

Thus

$$\left|\mathcal{P}_t\right| \leq \frac{(2(\mu+1)+1)^2}{\pi}\int_\Omega\frac{1}{\mathrm{lfs}^2(x)}\mathrm{d}x = \mathcal{O}\left(\frac{\mu^2}{\alpha}\right)\left|\mathcal{P}''\right| = \mathcal{O}\left(\frac{1}{\theta^{*2}\alpha}\right)\left|\mathcal{P}''\right| = \mathcal{O}\left(\frac{1}{\alpha^3}\right)\left|\mathcal{P}''\right|.$$

By Corollary 7.2.3 of Chapter 7,

$$\int_{\mathcal{S}}\frac{1}{\mathrm{lfs}(x)}\mathrm{d}x = \mathcal{O}\left(\frac{1}{\alpha}\log\frac{1}{\alpha}\right).$$

Thus

$$\left|\mathcal{P}_m\right| = \mathcal{O}\left(\frac{\mu}{\alpha}\log\frac{1}{\alpha}\right)\left|\mathcal{P}''\right| = \mathcal{O}\left(\frac{1}{\alpha^3}\right)\left|\mathcal{P}''\right|.$$

$\square$

5.7 Augmented Input

In this chapter we have been assuming that input conforms to Assumption 5.4.1. Algorithm 2 and Algorithm 4 of Chapter 4 accept arbitrary input conforming to Assumption 2.2.1 and create augmented input which conform to Assumption 5.4.1. Moreover, Theorem 4.1.2

and Theorem 4.2.2 assert that the resultant decrease in local feature size is $\Omega(\sin\theta^*)$. Paired with Corollary 5.6.1, this would appear to add a factor of $\mathcal{O}\left(\frac{1}{\sin^2\theta^*}\right) = \mathcal{O}\left(\frac{1}{\alpha^2}\right)$ to the optimality constant. We here argue that we can do better, *i.e.*, that the increase in the optimality constant is related to γ and not θ^*.

As in Chapter 4, we let $\text{lfs}(z)$ be the local feature size with respect to the given input, and $\text{lfs}'(z)$ be with respect to an input which has been augmented by a γ-Bounded Reduction Augmenter or a γ-Feature Size Augmenter.

Now we reprove Lemma 5.4.3 with respect to $\text{lfs}(\cdot)$ for the case where the algorithm is given the augmented input. The proof is only marginally trickier than before.

Lemma 5.7.1 (Augmented Midpoint Local Feature). *Suppose an input is fed to a γ-Bounded Reduction Augmenter, which produces an input that conforms to Assumption 5.4.1. Suppose that the augmented input is given as input to the Adaptive Delaunay Refinement Algorithm. Then there is a constant, μ, depending on θ^* and $\hat{\kappa}$ such that if p is the midpoint of a segment, s, of radius r that is committed by the algorithm, then* $\text{lfs}(p) \leq \mu r$.

Moreover, $\mu = (2\gamma+1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right)$ *suffices, where, as above,* $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$.

Proof. As in the proof of Lemma 5.4.3, we construct a hybrid sequence. The sequence is constructed exactly as in that previous proof, substituting the words "augmented input" for "input," and referring to $\text{lfs}'(\cdot)$ instead of $\text{lfs}(\cdot)$. We recall that by Theorem 4.1.2, that $\text{lfs}'(x) \geq \frac{\sin\theta^*}{2\gamma-1}\text{lfs}(x)$ for all x.

We skip then to the following facts about the hybrid sequence, which follow as in the proof of Lemma 5.4.3:

(a) $R_0 \leq R_1 \leq \ldots R_{l-1}$;

(b) $|P_i - P_{i-1}| \leq \eta R_i$;

(c) $|X_i - P_i| \leq \frac{\eta}{\sin\theta^*} R_i$;

(d) $X_i = X_{l-1}$ for $i = 1, 2, \ldots, l-2$.

(e) Either

  (i) $\text{lfs}'(P_0) \leq \max\left\{1, \sqrt{2}+\eta\right\} R_0 = (\sqrt{2}+\eta)R_0$, so $\text{lfs}(P_0) \leq \frac{(2\gamma-1)(\sqrt{2}+\eta)}{\sin\theta^*}R_0$, or

  (ii) there is some point of the augmented input $X_0 \neq X_{l-1}$ such that $|P_0 - X_0| \leq \frac{\eta}{\sin\theta^*}R_0$; moreover, (X_0, X_{l-1}) is a segment of the augmented input; or

  (iii) there is some midpoint q of radius r_q such that $R_0 \geq 2r_q$, and $|P_0 - q| \leq \eta R_0$.

These facts will be enough to establish the lemma. First note that if $l = 1$, it must be that s was encroached by a point on a nondisjoint feature of the augmented input, so then $\text{lfs}'(p) \leq r$, so $\text{lfs}(p) \leq \frac{2\gamma-1}{\sin\theta^*}r$, so it suffices to take

$$\boxed{\frac{2\gamma-1}{\sin\theta^*} \leq \mu.}$$

So assume $l \neq 1$.

We bound $|P_{l-1} - P_0|$; our analysis will be quite an overestimate for the case where $l = 2$, but will suffice. By the triangle inequality, and item (c), item (a), item (d), and item (b) of above,

$$\begin{aligned} |P_{l-1} - P_0| &\leq |P_{l-1} - X_{l-1}| + |P_1 - X_1| + |P_1 - P_0|, \\ &\leq \frac{\eta}{\sin\theta^*} R_{l-1} + \frac{\eta}{\sin\theta^*} R_1 + \eta R_1, \\ &\leq \left(\eta + \frac{2\eta}{\sin\theta^*}\right) R_{l-1}. \end{aligned}$$

We consider the sequence head, *i.e.*, item (e).

- If the first alternative holds, then by the Lipschitz condition,

$$\begin{aligned} \operatorname{lfs}(P_{l-1}) &\leq |P_{l-1} - P_0| + \operatorname{lfs}(P_0), \\ &\leq \left(\eta + \frac{2\eta}{\sin\theta^*}\right) R_{l-1} + \frac{(2\gamma-1)(\sqrt{2}+\eta)}{\sin\theta^*} R_0, \\ &\leq \left(\eta + \frac{2\eta + (2\gamma-1)(\sqrt{2}+\eta)}{\sin\theta^*}\right) R_{l-1}, \end{aligned}$$

  so it suffices to take

$$\boxed{\eta + \frac{2\eta + (2\gamma-1)(\sqrt{2}+\eta)}{\sin\theta^*} \leq \mu.}$$

- If the second alternative holds, then $X_0 \neq X_{l-1}$, and (X_0, X_{l-1}) is a segment of the augmented input. Since a γ-Bounded Reduction Augmenter is a γ-Feature Size Augmenter, then the local feature size of points on this segment are bounded by the segment length, *i.e.*, $\operatorname{lfs}(X_0) \leq \gamma |X_0 - X_{l-1}|$.
  Let $R^* = (|P_{l-1} - X_{l-1}| \vee |P_{l-1} - X_0|)$. In the proof of Lemma 5.4.3, it was shown that $R^* \leq \left(\eta + \frac{3\eta}{\sin\theta^*}\right) R_{l-1}$. Since the circle centered at P_{l-1} of radius R^* contains both X_0 and X_{l-1}, then $|X_0 - X_{l-1}| \leq 2R^*$. By the Lipschitz condition,

$$\begin{aligned} \operatorname{lfs}(P_{l-1}) &\leq |P_{l-1} - X_0| + \operatorname{lfs}(X_0), \\ &\leq R^* + \gamma |X_0 - X_{l-1}|, \\ &\leq R^* + 2\gamma R^*, \\ &\leq (2\gamma+1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right) R_{l-1}. \end{aligned}$$

  So it suffices to take

$$\boxed{(2\gamma+1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right) \leq \mu.}$$

- If the third alternative holds, then as in the proof of Lemma 5.4.3, use this lemma inductively on q to find that $\mathrm{lfs}\,(q) \leq \mu r_q \leq \mu \frac{R_0}{2} \leq \mu \frac{R_{l-1}}{2}$. Then by the Lipschitz condition,

$$\begin{aligned} \mathrm{lfs}\,(P_{l-1}) &\leq |P_{l-1} - P_0| + |P_0 - q| + \mathrm{lfs}\,(q)\,, \\ &\leq \left(\eta + \frac{2\eta}{\sin\theta^*}\right) R_{l-1} + \eta R_0 + \frac{\mu}{2} R_{l-1}, \\ &\leq \left(2\eta + \frac{2\eta}{\sin\theta^*} + \frac{\mu}{2}\right) R_{l-1}, \end{aligned}$$

so it suffices to take

$$\boxed{4\eta + \frac{4\eta}{\sin\theta^*} \leq \mu.}$$

Simple analysis, and using $\gamma \geq 2$ shows that $\mu = (2\gamma+1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right)$ suffices to satisfy the boxed constraints. ☐

We actually get the same bounds when using a γ-Feature Size Augmenter, assuming $\gamma \geq 3$. Given the space, we state the lemma:

Lemma 5.7.2 (Augmented Midpoint Local Feature). *Suppose an input is fed to a γ-Feature Size Augmenter, with $\gamma \geq 3$, and which produces an input that conforms to Assumption 5.4.1. Suppose that the augmented input is given as input to the Adaptive Delaunay Refinement Algorithm. Then there is a constant, μ, depending on θ^* and $\hat{\kappa}$ such that if p is the midpoint of a segment, s, of radius r that is committed by the algorithm, then* $\mathrm{lfs}\,(p) \leq \mu r$.

Moreover, $\mu = (2\gamma+1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right)$ *suffices, where, as above,* $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$.

Proof (Sketch). The proof is exactly as that of Lemma 5.7.1. We note that at one point in that proof we use the fact that a γ-Bounded Reduction Augmenter is a γ-Feature Size Augmenter. The constraints which μ must satisfy are given as

$$\begin{aligned} \frac{2\gamma+3}{\sin\theta^*} &\leq \mu, \\ \eta + \frac{2\eta + (2\gamma+3)(\sqrt{2}+\eta)}{\sin\theta^*} &\leq \mu, \\ (2\gamma+1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right) &\leq \mu, \\ 4\eta + \frac{4\eta}{\sin\theta^*} &\leq \mu. \end{aligned}$$

Only the first two differ from the proof of Lemma 5.7.1, and these reflect the slight loss of bound on $\mathrm{lfs}'\,(\cdot)$ when comparing a Bounded Reduction Augmenter to a Feature Size Augmenter. (*cf.* Theorem 4.1.2 and Theorem 4.2.2.)

Simple analysis using $\gamma \geq 3$ shows that $\mu = (2\gamma+1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right)$ suffices to satisfy these constraints. $\square$

We can reprove Theorem 5.4.4 using the new μ, and get the following termination and optimality result, which follows as Corollary 5.6.1.

Corollary 5.7.3. *Suppose an input,* $(\mathcal{P}, \mathcal{S})$, *that conforms to Assumption 2.2.1 is fed to a* γ*-Bounded Reduction Augmenter or a* γ*-Feature Size Augmenter (with* $\gamma \geq 3$*), whose output is fed as input to the Adaptive Delaunay Refinement Algorithm, using an output angle parameter* $\hat{\kappa} \leq \arcsin 2^{-7/6}$.

Then the algorithm terminates, outputting the Delaunay Triangulation of the set of points $\mathcal{P}'$*, and with no angle in the triangulation less than* $\alpha = \min\left\{\hat{\kappa}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$, *and no angle is larger than* $\max\left\{\pi - 2\hat{\kappa}, \pi - 2\arcsin\frac{\sqrt{3}-1}{2}\right\}$. *Moreover, if there is* any *triangulation on a set of points* $\mathcal{P}''$ *that conforms to the original input and has minimum angle* α *then*

$$\left|\mathcal{P}'\right| = \mathcal{O}\left(\alpha^{-3}\right)\left|\mathcal{P}''\right|.$$

5.8 How Good is "Optimal?"

We here consider the size of the optimality constant of Corollary 5.7.3. Suppose we have an input with a given lower bound on minimum input angle, θ^*. Suppose this input is fed to the γ-Feature Size Augmenter given as Algorithm 4, with $\gamma = \frac{3+\sqrt{13}}{2}$. Finally suppose the output of the augmenter is fed to the Adaptive Delaunay Refinement Algorithm, with $\hat{\kappa} = \arcsin 2^{-7/6}$. Let $\alpha = \min\left\{\arcsin 2^{-7/6}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$. Let $\mathcal{P}'$ be the set of output points; let $\mathcal{P}''$ be the set of points of a triangulation which conforms to the given input and has no output angle smaller than α. We let $\zeta(\theta^*)$ be the bound on $\frac{|\mathcal{P}'|}{|\mathcal{P}''|}$ guaranteed by Corollary 5.7.3.

Then we have the unwieldy expression:

$$\begin{aligned}\zeta(\theta^*) = 1 \quad &+ \quad 1.02(\mu+1)6\left(\ln\frac{3}{2} + h_\alpha\pi + \sqrt{h_\alpha^2+1}\left[\ln(h_\alpha + 1/h_\alpha) - \ln(2\sin\frac{\ln 3}{8})\right]\right) \\ &+ \quad \frac{1}{\pi}(2\mu+3)^2\left(\frac{21.5}{\alpha} + 11.9\right),\end{aligned}$$

where $h_\alpha = \frac{1}{\alpha}\ln 2\cos\alpha$, and $\mu = (2\gamma+1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right)$, and $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$. This is only a collection of all the terms relevant to Corollary 5.7.3 (see Chapter 7). The graph of ζ versus θ^* is of some interest; it is shown in Figure 27.

To show the improvements made by use of linear packings, and to justify Chapter 7, we also plot the naïve optimality constant which would be had from packing all Steiner points

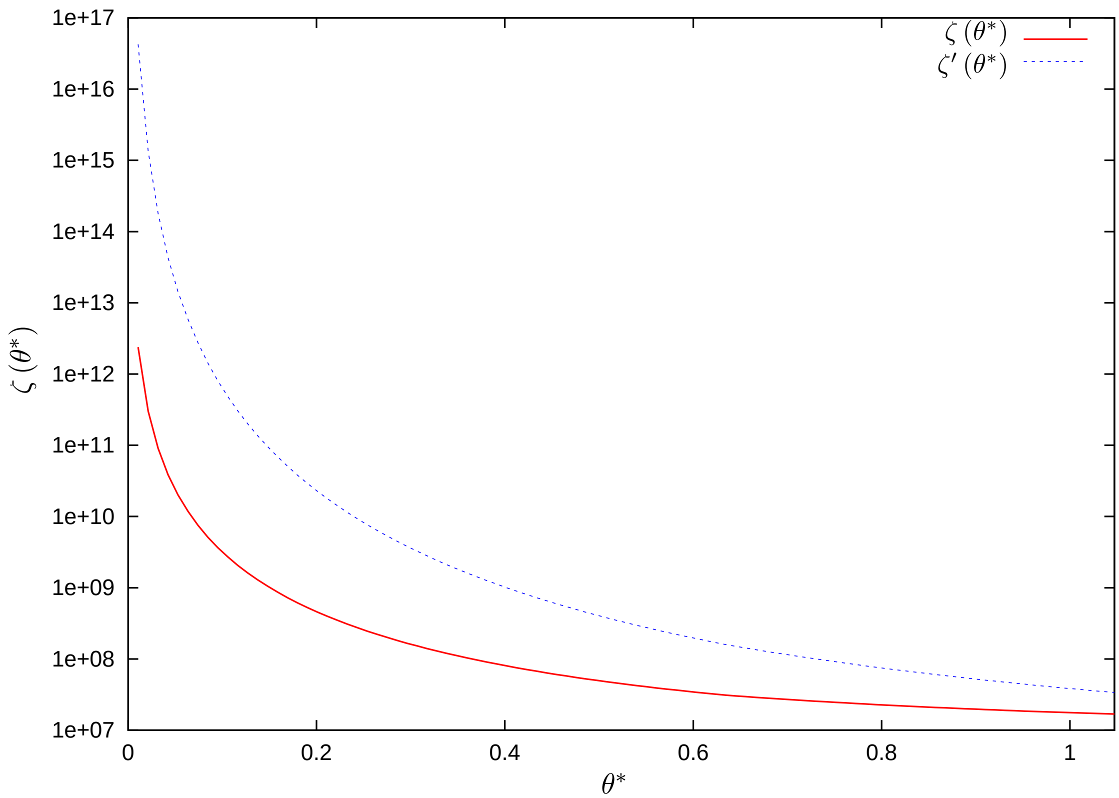

Figure 27: The optimality constant $\zeta(\theta^*)$ is plotted versus θ^*. The cardinality of a mesh constructed by the Adaptive Delaunay Refinement Algorithm is compared against any mesh conforming to the input which has minimum angle $\alpha = \min\left\{\arcsin 2^{-7/6}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$. The naïve optimality constant $\zeta'(\theta^*)$ is also plotted.

together and using Mitchell's work. We claim this constant is given by

$$\zeta'(\theta^*) = 1 + \frac{2}{\pi}\left(\frac{\mu}{\sin\frac{\theta^*}{2}} + 5\right)^2 \left(\frac{21.5}{\alpha} + 11.9\right).$$

The improvements made by using the anisotropic packings is about one order of magnitude for modest θ^*, but can be several orders of magnitude for smaller θ^*. However, $\zeta(\theta^*)$ is still too large to be of any practical use, taking on values larger than 1×10^7.

In Figure 28 the two factors which comprise $\zeta(\theta^*)$ are plotted, *i.e.*,

$$\begin{aligned}
\zeta_m(\theta^*) &= 1.02(\mu+1)\,6\left(\ln\frac{3}{2} + h_\alpha\pi + \sqrt{h_\alpha^2+1}\left[\ln(h_\alpha + 1/h_\alpha) - \ln(2\sin\frac{\ln 3}{8})\right]\right),\\
\zeta_t(\theta^*) &= \frac{1}{\pi}(2\mu+3)^2\left(\frac{21.5}{\alpha} + 11.9\right).
\end{aligned}$$

The factor $\zeta_t(\theta^*)$, which is the optimality of the set of Steiner circumcenters, is clearly the dominating factor. Thus any attempt at improving the optimality guarantee needs to focus on this factor. This could be achieved by proving a better grading constant for circumcenters, improving Mitchell's work, or employing a smarter packing argument. Note, however, that ζ_m, the optimality of the set of Steiner midpoints, is also fairly large, taking on values at least 1×10^4; improvements of this bound are also welcome.

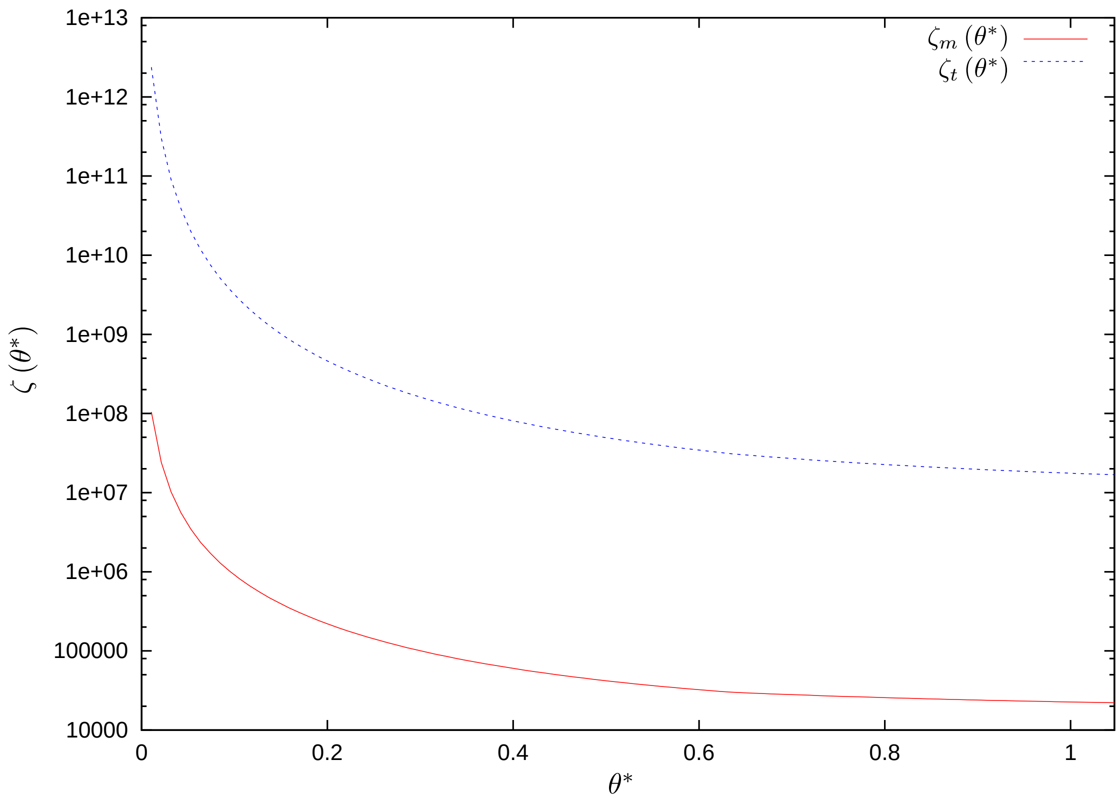

Figure 28: The two optimality constants which comprise $\zeta(\theta^*)$ are shown. The factor $\zeta_m(\theta^*)$ is the optimality associated with the Steiner midpoints, while $\zeta_t(\theta^*)$ is that associated with the Steiner circumcenters. The latter factor is clearly dominating.

Before attempting to improve the analysis of the algorithm one should consider how poorly the algorithm may perform. Gary Miller (personal communication) provided the following example which demonstrates that the optimality constant is $\Omega\left(\frac{1}{\theta^*}\right)$. The input consists of the two segments $(x, y), (x, z)$, with $\theta^* = \angle yxz$ small, as shown in Figure 29(a). If the two segments do not conform to Assumption 5.4.1, some augmenting procedure will be used to put them into the requisite form. If one of the algorithms of Chapter 4 or splitting on concentric circular shells is used as the augmenter, the result will be at least one augmenting point on each of the input segments, and some augmented input segments which are "far" from x, and with small values of $\mathrm{lfs}'(\cdot)$. The result, as shown in Figure 29(b), is two segments of length on the order of $\frac{1}{2}$, with a distance of approximately $\sin(\theta^*)$ between them. The algorithm will fill this region with triangles having no angle smaller than $\hat{\kappa}$. There can be $\Omega\left(\frac{1}{\theta^*}\right)$ of these triangles, as shown in Figure 29(c). An adversary's mesher, however, might

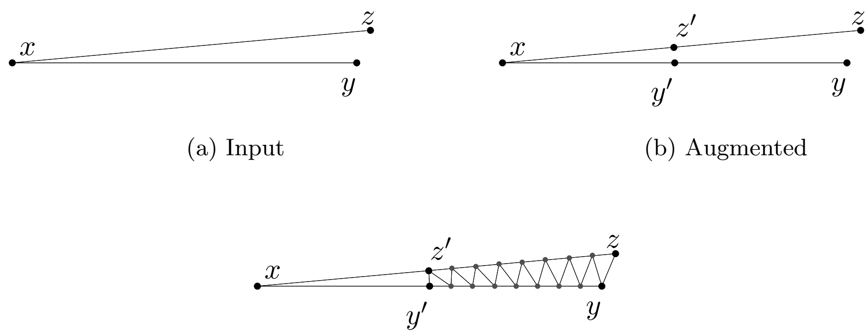

(a) Input

(b) Augmented

(c) Final

Figure 29: An input showing the limitations of the algorithm is shown in (a). Because the segments do not conform to Assumption 5.4.1, they will be split by an augmenting procedure, as in (b). In the augmented input the segments $(y, y'), (z, z')$ are disjoint, and the algorithm will fill in the area between them with $\Omega\left(\frac{1}{\theta^*}\right)$ triangles, as shown in (c). An adversary's algorithm might merely return the single triangle Δyxz.

just return the mesh of the single triangle Δyxz, which ultimately has the same (or better!) minimum angle as the mesh returned by the Adaptive Delaunay Refinement Algorithm.

As was the case for the example of Subsection 5.4.1, the problem is how the algorithm deals with conformality. There is no obvious fix.

CHAPTER VI

VARIATIONS

"How absolute the knave is! We must speak by the card, or equivocation will undo us."
–Hamlet

In this chapter, two variations of the main exposition are considered, namely incorporating an augmenter with the algorithm, and Miller's variant, which comes with a timing analysis [28]. We prove only good grading for each of these, as the angle bounds and optimality guarantees follow exactly as in Chapter 5. Either or both of these variants could have been incorporated into the description of the Adaptive Delaunay Refinement Algorithm to form a more general algorithm for analysis, making this chapter unnecessary. Unfortunately, such an exposition could be painfully complicated, as this chapter illustrates. It is likely that, having understood the analysis up to this point, even the most fervid reader would be satisfied to merely skim this chapter.

6.1 Delaunay Refinement with an Augmenter

The analysis of Section 5.7 applies to input which has been augmented by a Bounded Reduction Augmenter or Feature Size Augmenter, then given to the Adaptive Delaunay Refinement Algorithm. In practice, however, using an augmenter as a preprocessor can result in a fairly large number of Steiner Points. Ruppert's idea of splitting on concentric circular shells shows that the augmenting procedure can be performed on an as-needed basis. In Figure 30, two meshes of Lake Superior are shown. In Figure 30(a), Algorithm 2 has been used as a preprocessor to make the input conform to Assumption 5.4.1. Ruppert's strategy was incorporated into the Adaptive Delaunay Refinement Algorithm to produce the mesh of Figure 30(b), which has far fewer Steiner Points.

In this section we show that an augmenter can be incorporated into the Adaptive Delaunay Refinement Algorithm without sacrificing good grading. Although the grading and optimality guarantees will be no better than those of Chapter 4 for the preprocessing augmenter, the observed practical performance is superior, pardon the pun, as illustrated in Figure 30.

We consider a variation of the Adaptive Delaunay Refinement Algorithm which accepts arbitrary input, and adaptively performs "off-center" splits as needed. This strategy avoids the possibility of adding unnecessary augmenting points during the preprocessing step. This

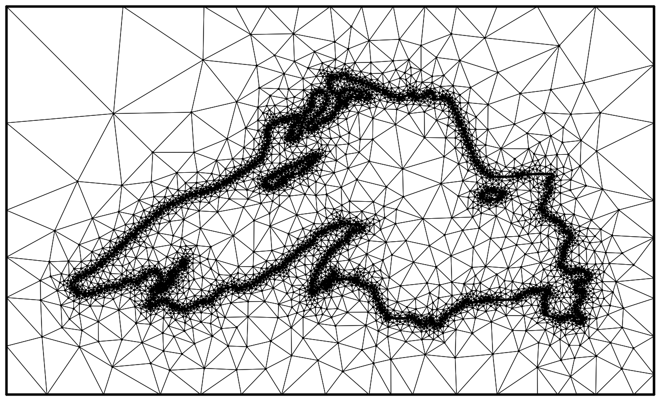

(a) Preprocessed by Algorithm 2

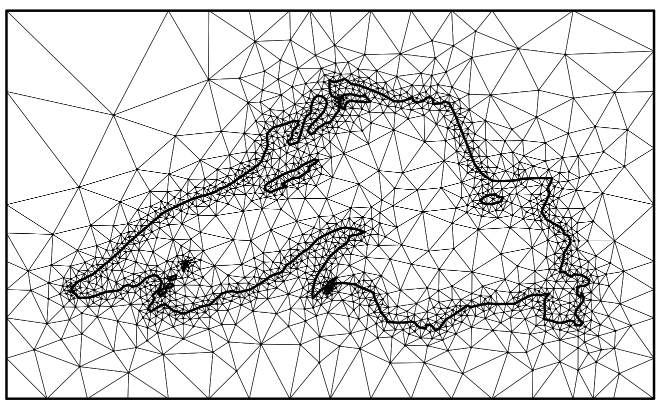

(b) Using Splitting on Concentric Circular Shells

Figure 30: In (a) a mesh of Lake Superior is shown. The mesh was generated by feeding the input to Algorithm 2 to make it conform to Assumption 5.4.1, then to the Adaptive Delaunay Refinement Algorithm. The mesh has 6007 vertices. The mesh in (b) is generated by the Adaptive Delaunay Refinement Algorithm using splitting on concentric circular shells. The mesh has only 1750 vertices. The input consists of 522 vertices and 522 segments, and has minimum angle $\theta^* \approx 15.02°$. Both meshes have minimum output angle of $14.07° \approx \arctan\left[(\sin\theta^*)/(2-\cos\theta^*)\right]$, though all but three triangles in each mesh have minimum angle at least $\arcsin 2^{-7/6}$. Edges of the input are shown in bold.

analysis will apply, for example, to Ruppert's strategy of splitting on concentric circular shells [39].

Again we will assume that the algorithm maintains a set of points and a set of segments, initialized as the input. For the adaptive quality test, we assume the algorithm maintains, for every midpoint, a pointer to the two input points which are endpoints of the input segment containing the midpoint.

For the sake of the *analysis only* we imagine a PSLG being maintained during the lifetime of the algorithm. This PSLG is initialized as the input, and is occasionally modified by an augmenting procedure during the course of the algorithm. We call this PSLG the "imaginary input." We let $\mathrm{lfs}'(x)$ be the local feature size with respect to the imaginary input, but take caution that this is time dependant.

We generalize the idea of splitting a segment, where, as usual, we reserve the word "segment" to mean a segment in the set maintained by the algorithm. We redefine what we mean by splitting a segment. When we say the algorithm "splits" a segment s we mean the following: If s happens to be in the imaginary input, the algorithm can decide to break s in two subsegments not necessarily at the midpoint, but at an "off-center point." Alternatively, the segment s could be split regularly, at its midpoint. If s is not in the imaginary input, it must be split at its midpoint.

Note that we allow an off-center point to actually be the midpoint of a segment. This allows augmenters like splitting on concentric circular shells, which first splits an input segment at its midpoint. The algorithm has to somehow know the difference between midpoints and off-center points.

To achieve offcenter splits, the algorithm must also somehow know that a segment is still in the imaginary input, because the imaginary input need not actually be maintained by the algorithm. In the case of concentric shell splitting, for example, the use of a few flags on segments suffices to achieve this. The algorithm must also know that an off-center point is special, *i.e.*, that it is not to be treated as a midpoint, even if it happens to be the midpoint of a segment, with regard to the adaptive quality test.

When a current segment is split by an off-center point it is also split, in our minds, in the imaginary input, and the off-center point is added to the imaginary input. When an imaginary input segment is split by a midpoint, it should not be so split in the imaginary input, rather it becomes a "finished," a term we define shortly.

We define the radius of a segment split by an off-center point as the length of the shorter subsegment.

We assume that there is some γ such that at any time during the course of the algorithm, the imaginary input could be the output of a γ-Feature Size Augmenter which was fed the original input. This means that, by Theorem 4.2.2, $\mathrm{lfs}\,(x) \le \frac{2\gamma+3}{\sin\theta^*}\mathrm{lfs}'\,(x)$ for all x.

We also assume that the augmenter is somehow working towards putting the imaginary input into the form of Assumption 5.4.1. To define this precisely, we say that a segment of the imaginary input is "finished" if either it is not present in the set of segments maintained by the algorithm, or it cannot be broken by an off-center point without violating the γ-augmenting condition. Then Assumption 5.4.1 is adapted as the following assumption.

Assumption 6.1.1. We make the following assumption on the algorithm:

(a) If S_1, S_2 are two adjoining imaginary input segments which are both finished, and which meet at angle other than π, then they have the same length modulo a power of two, that is $\frac{|S_1|}{|S_2|} = 2^k$ for some integer k.

The algorithm then has two major operations, which are nearly the same as the Adaptive Delaunay Refinement Algorithm, with some variation to reflect that off-center points need to be viewed as input.

(CONFORMALITY) If s is a current segment, and there is a committed point that encroaches s, then split s.

(QUALITY) If a, b, c are committed points, the circumcircle of the triangle Δabc contains no committed point, $\angle acb < \hat{\kappa}$, the circumcenter, p, of the triangle is inside Ω and either (i) both a, b are midpoints (not off-center points) on distinct nondisjoint *imaginary* input segments, sharing input endpoint x, and $\angle axb > \pi/3$, or (ii) a, b are not midpoints on adjoining *imaginary* input segments, then attempt to commit p. If, however, the point p encroaches any current segment, then do not commit to point p, rather in this case split one, some, or all of the current segments which are encroached by p.

6.1.1 Good Grading

It should be clear why this variant should work: off-center splits can be blamed on the augmenting process (with factor γ), while in the "endgame," the algorithm should act exactly like the Adaptive Delaunay Refinement Algorithm, so the analysis of Chapter 5 suffices. This is exactly how good grading is proven.

We start with a basic fact of the algorithm:

Claim 6.1.2. If p is an off-center point that is committed, and r is the "radius" associated with p, *i.e.*, the length of the shorter of two subsegments created by p, then $\mathrm{lfs}\,(p) \le \gamma r$.

If p is the first midpoint added to an imaginary input segment, thereby "finishing" the segment, and r is the associated radius, then $\mathrm{lfs}\,(p) \le 2\gamma r$.

Proof. The first follows since the off-center points are added as part of a γ-Feature Size Augmenter, and p is a member of both of its subsegments.

The second follows since p must be the center of a segment which has the property that the local feature size at every point on it is no greater than γ times the length of the whole segment, which is $2r$. □

The difficulty with proving good grading is that the proofs of Section 5.4 rely heavily on Assumption 5.4.1, so separate arguments are required when the algorithm splits a subsegment on an imaginary input segment that is not finished. These will use the assumption that off-center splits can be seen as part of a γ-Feature Size Augmenter augmenter. Unlike in Section 5.7, we cannot use Lemma 5.4.2 with $\mathrm{lfs}'(\cdot)$ replacing $\mathrm{lfs}(\cdot)$, since there is a problem with timestamp. So we reprove the lemma:

Lemma 6.1.3. *Suppose that a circumcenter b_{l-1} was proposed to be committed, but rejected in favor of splitting a subsegment s_p. Suppose s_p is split by midpoint or off-center point p, and that the smaller of the two subsegments has length r_p. Let $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$. Then either*

1. $\mathrm{lfs}(p) \le \frac{2\gamma+3}{\sin\theta^*}\eta r_p$, *or*
2. *there is a segment s_q with committed midpoint q, and radius r_q such that*
   (a) *the imaginary input segments containing s_p, s_q are nondisjoint,*
   (b) $r_q \le r_p$,
   (c) $|p - q| \le \eta r_p$, *and*
   (d) *if s_p, s_q are on the same imaginary input segment then $2r_q \le r_p$; if they are on distinct imaginary input segments sharing imaginary input point x, then $|x - p| < \frac{\eta}{\sin\theta^*} r_p$.*

Proof. We can immediately assume that p is a midpoint, as otherwise, by Claim 6.1.2, $\mathrm{lfs}(p) \le \gamma r_p$, which suffices.

For convenience, we say that a point q "provokes" a point p, if q is committed before p, and p is the midpoint of a segment, s, which is encroached by q.

Let $\{b_i\}_{i=0}^{l-1}$ be a maximal cirumcenter sequence ending with the circumcenter b_{l-1} which caused p to be committed. Recalling facts about circumcenter sequences, $\sqrt{2}r_p \ge \tilde{r}_{l-1} > \tilde{r}_0$. Now consider the identity of b_0:

- If b_0 is an input point, then by definition so is a_0, and so $\mathrm{lfs}(p) \le |p - b_0| \vee |p - a_0|$. By Corollary 5.3.4, these are both bounded above by ηr_p, so $\mathrm{lfs}(p) < \eta r_p$, which suffices.
- If b_0 is a midpoint or off-center point on a real input segment disjoint from the one containing p, then by definition of local feature size and using Corollary 5.3.4, $\mathrm{lfs}(p) \le |p - b_0| \le \eta r_p$, which suffices.

- If b_0 is a midpoint or off-center point on an *imaginary* input segment disjoint from the one containing p at the time p is considered for commission, then $\mathrm{lfs}'(p) \leq |p - b_0| \leq \eta r_p$. Using Theorem 4.2.2, then $\mathrm{lfs}(p) \leq \frac{2\gamma+3}{\sin\theta^*}\eta r_p$.
- Suppose that b_0 is a midpoint or off-center point on an *imaginary* input segment nondisjoint to the one containing p, at the time p is considered for commission. Furthermore suppose that a_0 did not provoke b_0. Let $q = b_0$, let r_q be the radius associated with b_0. By the assumption that a_0 did not provoke b_0, we have $\tilde{r}_0 = |a_0 - b_0| \geq r_q$. If b_0 is an off-center point, then by Claim 6.1.2, $\mathrm{lfs}(b_0) \leq \gamma r_q$. In this case we bound:

  $$\mathrm{lfs}(p) \leq |p - b_0| + \mathrm{lfs}(b_0) \leq \eta r_p + \gamma r_q \leq \eta r_p + \gamma \tilde{r}_{l-1} \leq \left(\eta + \sqrt{2}\gamma\right) r_p,$$

  which suffices.

  So we assume b_0 is a midpoint. In this case the imaginary input segments containing p, b_0 are both finished. Careful inspection of the case where a_0 does not provoke b_0 in Lemma 5.4.2 show that the analysis of that proof apply here. This is because adjoining finished segments satisfy Assumption 5.4.1 locally.
- Suppose that b_0 is a midpoint or off-center point on an input segment nondisjoint to the one containing p, and a_0 *did* provoke b_0. We first consider the identity of a_0 :
  - If a_0 is (on) a real input feature disjoint from the one containing b_0, then by definition $\mathrm{lfs}(p) \leq |p - b_0| \vee |p - a_0|$. By Corollary 5.3.4, these are both bounded above by ηr_p, so $\mathrm{lfs}(p) < \eta r_p$, which suffices.
  - If a_0 is (on) an imaginary input feature disjoint from the one containing b_0, where we are considering the imaginary input at the time that p is committed. As above $\mathrm{lfs}'(p) < \eta r_p$, so $\mathrm{lfs}(p) < \frac{2\gamma+3}{\sin\theta^*}\eta r_p$, by Theorem 4.2.2.
  - It cannot be the case that a_0 is a circumcenter, as it would not have been committed before b_0 if it provoked b_0.
  - Suppose a_0 is a midpoint or off-center point which is on an imaginary input feature nondisjoint from the one containing b_0, where we are again considering the imaginary input at the time p is committed.

    Let $q = a_0$, let r_q be the radius associated with a_0. By definition of circumcenter sequences, the input segments opposite a_0, b_0 subtends angle ϕ at least $\pi/3$. Then by Lemma 3.3.6, $\tilde{r}_0 = |a_0 - b_0| \geq 2r_q \sin\frac{\phi}{2} \geq r_q$.

    If a_0 is an off-center point, then by definition of our algorithm $\mathrm{lfs}(a_0) \leq \gamma r_q$. We bound

    $$\mathrm{lfs}(p) \leq |p - a_0| + \mathrm{lfs}(a_0) \leq \eta r_p + \gamma r_q \leq \eta r_p + \gamma \tilde{r}_{l-1} \leq \left(\eta + \sqrt{2}\gamma\right) r_p,$$

    which suffices.

So we assume a_0 is a midpoint, thus the imaginary input segment containing it is finished. What can we say about b_0? If b_0 is an off-center point, then it is an imaginary input feature disjoint from the imaginary input segment containing a_0. Then by definition $\operatorname{lfs}'(a_0) \leq |a_0 - b_0| = \tilde{r}_0 \leq \tilde{r}_{l-1} \leq \sqrt{2}r_p$. Using Theorem 4.2.2 and the Lipschitz condition we bound

$$\operatorname{lfs}(p) \leq |p - a_0| + \frac{2\gamma + 3}{\sin\theta^*}\operatorname{lfs}'(a_0) \leq \left(\eta + \frac{2\gamma + 3}{\sin\theta^*}\sqrt{2}\right) r_p.$$

Some extra work is required to show that this suffices to bound $\operatorname{lfs}(p)$ as desired; this relies on $\gamma \geq 3$.

Then we can assume that b_0 is also a midpoint, as is a_0, and p. Careful inspection of the proof of Lemma 5.4.2 shows that the analysis there suffices in this case, again using the fact that because the imaginary input segments in question are finished, they conform to Assumption 5.4.1 locally.

□

We now can return to hybrid sequences to bound the local feature size on a midpoint or off-center splits.

Lemma 6.1.4 (Adaptively Augmented Midpoint Local Feature). *There is a constant, μ, depending on γ, θ^* and $\hat{\kappa}$ such that if p is the midpoint or off-center point of a segment, s, of radius r that is committed by the algorithm, then* $\operatorname{lfs}(p) \leq \mu r$.

Moreover, $\mu = (2\gamma + 1)\left(\eta + \frac{3\eta}{\sin\theta^*}\right)$ suffices, where, as above, $\eta = 1 + \frac{\sqrt{2}}{1 - 2\sin\hat{\kappa}}$.

Proof. First, if p is an off-center point, then Claim 6.1.2 shows that $\operatorname{lfs}(p) \leq \gamma r$, which suffices. So assume p is a midpoint.

As in the proof of Lemma 5.4.3, we construct a hybrid sequence. Because of the complications involved with off-center splits and imaginary input, we define the hybrid sequence explicitly.

Recall we construct two sequences, $\{(P_i, R_i)\}_{i=0}^{1-l}, \{X_i\}_{i=1}^{l-1}$, where each P_i will be a midpoint of a segment of radius R_i, and P_i and P_{i-1} will be on input features sharing input point X_i. We may optionally define an X_0. The point P_0 may be an off-center point, but the remaining P_i will be midpoints.

We will use Lemma 3.3.6 and Lemma 6.1.3 to establish the sequence. The sequence is constructed backwards, so for convenience we pretend that we know how long it will be, *i.e.*, we know l, so that we can set $P_{l-1} = p, R_{l-1} = r$. For convenience, let $S_{l-1} = s$. The sequence is constructed backwards to make it analogous with encroachment sequences. Thus we will claim that when P_i is committed, P_{i-1} has already been committed, and is somehow "responsible" for P_i being committed.

Construct the sequence as follows: given segment S_i with committed midpoint P_i, and radius R_i, consider why P_i was committed:

- Suppose that P_i is an off-center point. Then we have chosen l magically so that $i = 0$. In this case by Claim 6.1.2, $\text{lfs}(P_0) \leq \gamma R_0$.
- Suppose that P_i was a midpoint committed because an input point or a point on a disjoint feature of the real input encroached S_i. Then we let $i = 0$ again and $\text{lfs}(P_0) \leq R_0$, by definition.
- Suppose that P_i was a midpoint committed because an off-center point or a point on a disjoint feature of the imaginary input encroached S_i. Then we let $i = 0$ again and $\text{lfs}'(P_0) \leq R_0$, so by Theorem 4.2.2, we have $\text{lfs}(P_0) \leq \frac{2\gamma+3}{\sin\theta^*} R_0$.
- Suppose that P_i was a midpoint committed because a midpoint, q, on a nondisjoint feature of the imaginary input encroached S_i. Let the input segments in question share input point x. Because both P_i and q are midpoints, the imaginary input segments are finished and conform to Assumption 5.4.1 locally, so we may apply Lemma 3.3.6, which asserts that $R_i \geq r_q$, the radius associated with q. Since q encroaches S_i we have $|P_i - q| \leq R_i \leq \eta R_i$. Moreover by Claim 3.1.2, $|x - P_i| < \frac{1}{\sin\theta^*} R_i < \frac{\eta}{\sin\theta^*} R_i$.
  If $i \neq l-1$ and $x \neq X_{l-1}$ then let $P_0 = P_i$ be the first point in the sequence, and let $X_0 = x$. Otherwise let $X_i = x = X_{l-1}$, let $P_{i-1} = q$, let $R_{i-1} = r_q$.
- If s was not encroached by any point, then by Lemma 6.1.3, then either $\text{lfs}(P_i) \leq \frac{2\gamma+3}{\sin\theta^*}\eta R_i$, in which case let P_i be the first element of the sequence; or there is some midpoint, q, of some subsegment s_q, with useful properties. The lemma asserts that S_i, s_q are on nondisjoint segments. If they are on the same input segment, let P_i be the first midpoint of the sequence. Note that in this case, by the lemma, $R_0 \geq 2r_q$, and $|P_0 - q| \leq \eta R_0$. Otherwise, let x be the single input point shared by the two input segments. If $i \neq l-1$ and $x \neq X_{l-1}$, then let P_i be the first midpoint in the sequence, and let $X_0 = x, P_0 = q$. By the lemma $|X_0 - P_0| < \frac{\eta}{\sin\theta^*} R_0$.
  Otherwise, let $X_i = x = X_{l-1}$, let $P_{i-1} = q$, let R_{i-1} be the radius associated with q. The lemma asserts that $R_i \geq R_{i-1}$, $|P_i - P_{i-1}| \leq \eta R_i$, and $|X_i - P_i| < \frac{\eta}{\sin\theta^*} R_i$. Let S_{i-1} be the subsegment of which P_{i-1} is midpoint.

We can claim the following facts about the hybrid sequence:

(1) $R_0 \leq R_1 \leq \ldots R_{l-1}$;

(2) $|P_i - P_{i-1}| \leq \eta R_i$;

(3) $|X_i - P_i| \leq \frac{\eta}{\sin\theta^*} R_i$;

(4) $X_i = X_{l-1}$ for $i = 1, 2, \ldots, l-2$.

(5) Either

  (a) $\text{lfs}(P_0) \leq \frac{(2\gamma+3)}{\sin\theta^*}\eta R_0$, or

(b) there is some point of the augmented input $X_0 \neq X_{l-1}$ such that $|P_0 - X_0| \leq \frac{\eta}{\sin\theta^*} R_0$; moreover, (X_0, X_{l-1}) is a segment of the augmented input; or

(c) there is some midpoint q of radius r_q such that $R_0 \geq 2r_q$, and $|P_0 - q| \leq \eta R_0$.

As in the proof of Lemma 5.7.2, these conditions suffice to give the conclusion of the lemma. □

Theorem 6.1.5 (Adaptive Good Grading). *Let μ be the constant depending on $\gamma, \theta^*, \hat{\kappa}$ from Lemma 6.1.4. Then there is a positive constant C such that when the algorithm, operating with a output angle parameter $\hat{\kappa} \leq \arcsin 2^{-7/6}$, commits or attempts to commit the point p then if q is any previously committed point then*

- *If p is an off-center point, and q is any kind of previously committed point, then* $\operatorname{lfs}(p) \leq \mu |p - q|$.
- *If p is the midpoint of a segment encroached by q, which is a midpoint or off-center point on an adjoining input segment then*

$$\operatorname{lfs}(p) \leq (1 + \frac{\mu}{2\sin\frac{\theta}{2}})\,|p - q| \leq (1 + \frac{\mu}{2\sin\frac{\theta^*}{2}})\,|p - q|\,,$$

  where θ is the angle subtended by the two segments.
- *If p is a midpoint, and either q did not encroach the parent segment of p or is not a midpoint on an adjoining input segment, then*

$$\operatorname{lfs}(p) \leq \mu\,|p - q|\,.$$

- *If p is the circumcenter of a triangle of circumradius r, then*

$$\operatorname{lfs}(p) \leq Cr.$$

Moreover, $C = 1 + 2\sin\hat{\kappa}(1 + \mu)$ suffices.

Proof. We determine sufficient conditions on the constant C. Again, for convenience, we say that a point q "provokes" a point p, if q is committed before p, and p is the midpoint of a segment, s, which is encroached by q.

- Suppose p is an off-center point, with associated radius r. By Lemma 6.1.4, $\operatorname{lfs}(p) \leq \mu r$. If q did not provoke p, then $r \leq |p - q|$, which suffices. If q *did* provoke p, then q could not have been a circumcenter. The points p, q cannot be on the same input segments. Since p is an off-center point it must be a feature disjoint from the one containing q in the imaginary input. In this case $\operatorname{lfs}'(p) \leq |p - q|$. By Theorem 4.2.2 this gives $\operatorname{lfs}(p) \leq \frac{2\gamma+3}{\sin\theta^*}|p - q|$. By definition of μ then $\operatorname{lfs}(p) \leq \mu\,|p - q|$.

The remaining cases for the identity of p follow as in the proof of Theorem 5.4.4, but rely, of course, on Lemma 6.1.4 instead of Lemma 5.4.3. □

Corollary 6.1.6. *Suppose an input* $(\mathcal{P}, \mathcal{S})$ *that conforms to Assumption 2.2.1 is fed to the Adaptive Delaunay Refinement Algorithm with an adptive* γ*-Feature Size Augmenter, using an output angle parameter* $\hat{\kappa} \leq \arcsin 2^{-7/6}$.

Then the algorithm terminates, outputting the Delaunay Triangulation of the set of points $\mathcal{P}'$*, and with no angle in the triangulation less than* $\alpha = \min\left\{\hat{\kappa}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$*, and no angle is larger than* $\max\left\{\pi - 2\hat{\kappa}, \pi - 2\arcsin\frac{\sqrt{3}-1}{2}\right\}$*. Moreover, if there is* any *triangulation on a set of points* $\mathcal{P}''$ *that conforms to the original input and has minimum angle* α *then*

$$\left|\mathcal{P}'\right| = \mathcal{O}\left(\alpha^{-3}\right)\left|\mathcal{P}''\right|.$$

6.2 A Runtime Analyzable Algorithm

Miller analyzed the runtime of a variant of the Delaunay Refinement Algorithm. The changes made to the algorithm are designed to allow such an analysis [28]. Here we prove that the algorithm does indeed terminate with good grading and bounded angle.

The main variation is that to avoid unnecessary work, the algorithm doesn't so readily "yield" when a circumcenter is being committed. The second change is that the algorithm maintains a Constrained Delaunay Triangulation during its lifetime, as opposed to most implementations of the Delaunay Refinement Algorithm, which maintain a Delaunay Triangulation. To facilitate explanation of the first change, we say that a triangle Δabc whose circumcenter is inside the diametral circle of segment (p, q) "yields" to the segment if (a) the midpoint of (p, q) is inside the circumcircle of Δabc, or (b) the circumcenter of Δabc is inside the diametral circle of one of the two subsegments of (a, b). In this case, we also say that the circumcenter yields to the segments, or to their midpoints. In Figure 31 the yield cases are illustrated.

Again we will assume that the algorithm maintains a set of points and a set of segments, initialized as the input. For the adaptive quality test, we assume the algorithm maintains, for every midpoint, a pointer to the two input points which are endpoints of the input segment containing the midpoint.

With respect to the given input, we say that two points are "visible" to each other if both points are on an input segment, or if the open line segment between the two points does not intersect any segment of the input. A triangle Δabc is "Constrained Delaunay" if there is no committed point inside the circumcircle of the triangle which is visible to any of a, b, c. Two simplices are visible to one another if every point in one is visible to every point in the other.

The algorithm has two major operations:

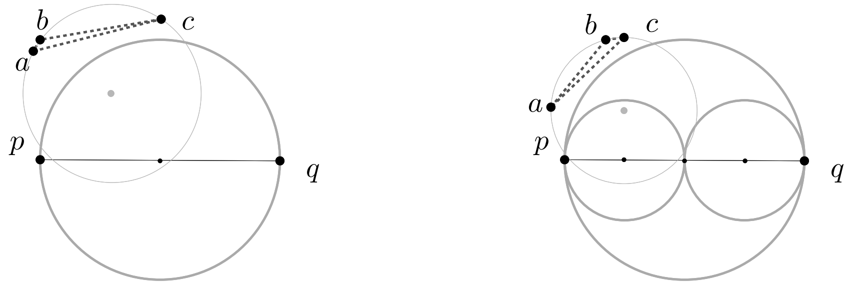

(a) Midpoint in Circumcircle (b) Circumcenter Encroaches Subsegment

Figure 31: The two "yield" conditions are shown. In (a), the midpoint of segment (p, q) is inside the diametral circle of triangle Δabc. In (b), the circumcenter of Δabc encroaches the diametral circle of one of the two subsegments of (p, q).

(CONFORMALITY′) If s is a current segment, and there is a committed point that encroaches s, and the committed point is visible to s, then split s.

(QUALITY') If a, b, c are committed points, the triangle Δabc is Constrained Delaunay, $\angle acb < \hat{\kappa}$, the circumcenter, p, of the triangle is inside Ω and either (i) both a, b are midpoints on distinct nondisjoint input segments, sharing input endpoint x, and $\angle axb > \pi/3$, or (ii) a, b are not midpoints on adjoining input segments, then attempt to commit p. If, however, the point p encroaches some current segments, $\{S_i\}_{i=0}^{l}$ all of which are visible to Δabc, then if Δabc yields to any of these segments, then do not commit to point p, rather in this case split one, some or all of the segments of $\{S_i\}_{i=0}^{l}$. If however, Δabc does not yield to any of these segments, then commit to p, and then split every segment of $\{S_i\}_{i=0}^{l}$.

Since we've exchanged the (CONFORMALITY) operation for the (CONFORMALITY′) operation, we may have sacrificed the guarantee that the output is truly Delaunay. A simple argument due to Edelsbrunner and Tan [20, Lemma 4.4] shows that the output mesh is indeed conforming: Take the segment with nonempty diametral circle which has no *visible* point encroaching it, and has the global minimum number of segments of the input blocking visibility from the segment to the encroaching point over all such segments. Then select one of the current segments blocking visibility. Then either there is a visible point encroaching this segment or this segment violates the minimality assumption of our choice of original segment.

Since any triangle which is Delaunay is Constrained Delaunay, if the algorithm terminates, there is no triangle in the final mesh with smallest angle less than $\hat{\kappa}$, except those "across" from small input angles. That is, we are claiming that the minimum and maximum angle guarantees of Section 5.5 apply.

6.2.1 Good Grading

We start with a claim that shows this algorithm is not too far different from the Adaptive Delaunay Refinement Algorithm.

Claim 6.2.1. After the application of each major operation, if (p, q) is a current segment, then there is no committed circumcenter visible to (p, q) which is inside its diametral circle.

Proof. We proceed by induction. At the beginning of the algorithm, this is obviously true. If a (CONFORMALITY′) operation is performed, since the diametral circles nest, the property is maintained. Suppose a (QUALITY') operation is performed. If the circumcenter is not committed then the property is maintained. Suppose to the contrary that the circumcenter *is* committed. By the definition of the (QUALITY') operation, if the committed circumcenter would encroach on a segment, and the circumcenter is visible to the segment, then the triangle does not yield to the segment. Thus in particular, the circumcenter is not inside the diametral circles of the subsegments of this segment. Since the segment is split, it is no longer current, and its subsegments are not encroached by the committed circumcenter. □

We assume that the input conforms to Assumption 5.4.1. This can be enforced with an augmenting preprocessor, or the augmenter can be integrated with the algorithm. We follow the analysis of Section 5.4, sketching some of the proofs where a complete exposition is counterproductive. Many of the proofs of that section go through with perhaps only notational change, largely because they make claims about radii and local feature size, and don't make lower bound claims on distances between committed points. All of the work on circumcenter sequences can be reproven without any changes.

Lemma 5.4.2 goes through with only a notational change:

Lemma 6.2.2. *Let s_p be a subsegment of midpoint p and radius r_p. Suppose that s_p was not committed during a (CONFORMALITY′) operation, rather it was split during a (QUALITY') operation attempting to commit circumcenter b_{l-1}. Let $\eta = 1 + \frac{\sqrt{2}}{1-2\sin\hat{\kappa}}$. Then either*

1. $\operatorname{lfs}(p) \le (\sqrt{2} + \eta) r_p$, *or*
2. *there is a segment s_q with committed midpoint q, and radius r_q such that*
   (a) *the input segments containing s_p, s_q are nondisjoint,*
   (b) $r_q \le r_p$,
   (c) $|p - q| \le \eta r_p$, *and*
   (d) *if s_p, s_q are on the same input segment then $2r_q \le r_p$; if they are on distinct input segments sharing input point x, then $|x - p| < \frac{\eta}{\sin\theta^*} r_p$.*

This leaves only the good grading proof. The only change comes in the case where a segment midpoint is being committed during a (Quality') operation in which the circumcenter does not yield.

Theorem 6.2.3 (Adaptive Good Grading). *Suppose that the input to the analyzable algorithm conforms to Assumption 5.4.1. Let μ be the constant depending on $\theta^*, \hat{\kappa}$ from Lemma 5.4.3. Then there is a positive constant C such that when the algorithm, operating with a output angle parameter $\hat{\kappa} \leq \arcsin 2^{-7/6}$, commits or attempts to commit the point p then if q is any previously committed point* that is visible to p *then*

- *If p is the midpoint of a segment encroached by q, which is a midpoint on an adjoining input segment then*

$$\operatorname{lfs}(p) \leq (1 + \frac{\mu}{2\sin\frac{\theta}{2}})\,|p-q| \leq (1 + \frac{\mu}{2\sin\frac{\theta^*}{2}})\,|p-q|\,,$$

  where θ is the angle subtended by the two segments.

- *If p is a midpoint, and either q did not encroach the parent segment of p or is not a midpoint on an adjoining input segment, then*

$$\operatorname{lfs}(p) \leq \mu\,|p-q|\,.$$

- *If p is the circumcenter of a triangle of circumradius r, then*

$$\operatorname{lfs}(p) \leq Cr.$$

Moreover, $C = 1 + 2\sin\hat{\kappa}(1+\mu)$ suffices.

Proof. We determine sufficient conditions on the constant C. Again, for convenience, we say that a point q "provokes" a point p, if q is committed before p, and p is the midpoint of a segment, s, which is encroached by q.

- Suppose p is the midpoint of subsegment which is encroached by q which is a midpoint on a nondisjoint input feature. Let θ be the angle between the two input segments. Let r_q be the radius associated with q. By Lemma 3.3.6, $|p-q| \geq 2r_q\sin\frac{\theta}{2}$. By Lemma 5.4.3, $\operatorname{lfs}(q) \leq \mu r_q$. Then using the Lipschitz condition,

$$\operatorname{lfs}(p) \leq |p-q| + \operatorname{lfs}(q) \leq |p-q| + \mu r_q \leq \left(1 + \frac{\mu}{2\sin\frac{\theta}{2}}\right)|p-q|\,,$$

  as desired.

- Suppose p is the midpoint of a subsegment, and suppose this subsegment is encroached by q which is not a midpoint on a nondisjoint input feature. If q is an input point or on a nondisjoint input feature, then by definition $\operatorname{lfs}(p) \leq |p-q| \leq \mu\,|p-q|\,.$

Suppose that the segment is encroached by q, and q is a circumcenter which is committed before p is committed. Since the triangle of q does not yield to the segment, it must be the case that $r_q \leq |p-q|$, where r_q is the cirumradius of the triangle associated with q. Using this theorem inductively, $\operatorname{lfs}(q) \leq C r_q$. Using the Lipschitz condition, $\operatorname{lfs}(p) \leq |p-q| + \operatorname{lfs}(q) \leq (1+C)|p-q|$. Then it suffices that

$$\boxed{1 + C \leq \mu.}$$

If q does not encroach the subsegment, then $|p-q|$ is at least the radius of the subsegment. r. By Lemma 5.4.3, $\operatorname{lfs}(p) \leq \mu r \leq \mu |p-q|$.

- If p is a circumcenter of a skinny triangle of circumradius r, then let a, b be the vertices of the shortest edge, and θ the angle opposite this edge. By assumption $\theta < \hat{\kappa}$. Note that by the sine rule, $|a-b| = 2r\sin\theta < 2r\sin\hat{\kappa}$. If a, b are both input points, then they are disjoint and by definition $\operatorname{lfs}(p) \leq r$, so it suffices to take $1 \leq C$. Otherwise let b be the most recently committed of the two points. We consider the possible identities of b, a:
  - Suppose b is a midpoint of a subsegment. If a is a midpoint on a nondisjoint input segment, with the two segments subtending angle θ, by definition of the (QUALITY′) operation, we know $\theta \geq \pi/3$, thus $2\sin\frac{\theta}{2} \geq 1$. Using the lemma inductively, since a witnesses b, it must be the case that $\operatorname{lfs}(b) \leq (1+\mu)|a-b|$.
  - If a was not such a midpoint or did not provoke b, then $\operatorname{lfs}(b) \leq \mu |a-b|$.
  - If b was a circumcenter, then since it was committed after a, its associated circumradius bounds $|a-b|$ from below, so $\operatorname{lfs}(b) \leq C|a-b|$.

  Then using the Lipschitz condition,

$$\operatorname{lfs}(p) \leq r + \operatorname{lfs}(b) \leq r + \max\{1+\mu, C\}|a-b| \leq (1 + 2\sin\hat{\kappa}\max\{1+\mu, C\})\, r.$$

  And it suffices to ensure that

$$\boxed{1 + 2\sin\hat{\kappa}\max\{1+\mu, C\} \leq C.}$$

Some work is required to show that both boxed constraints are satisfied by our choice of C. This is not too difficult, given the definition of μ,η, and because $\hat{\kappa} \leq \arcsin 2^{-7/6}$. □

Corollary 6.2.4. *Suppose an input* $(\mathcal{P}, \mathcal{S})$ *that conforms to Assumption 5.4.1 is fed to the analyzable variant of the Adaptive Delaunay Refinement Algorithm, using an output angle parameter* $\hat{\kappa} \leq \arcsin 2^{-7/6}$.

Then the algorithm terminates, outputting the Delaunay Triangulation of the set of points $\mathcal{P}'$, *and with no angle in the triangulation less than* $\alpha = \min\left\{\hat{\kappa}, \arctan\left(\frac{\sin\theta^*}{2-\cos\theta^*}\right)\right\}$, *and no angle is larger than* $\max\left\{\pi - 2\hat{\kappa}, \pi - 2\arcsin\frac{\sqrt{3}-1}{2}\right\}$. *Moreover, if there is* any *triangulation on a set of points* $\mathcal{P}''$ *that conforms to the original input and has minimum angle* α *then*

$$\left|\mathcal{P}'\right| = \mathcal{O}\left(\alpha^{-3}\right)\left|\mathcal{P}''\right|.$$

CHAPTER VII

OPTIMALITY

"The supreme misfortune is when theory outstrips performance." –Leonardo Da Vinci

By Da Vinci's standards, we are rather fortunate, since in practice the Delaunay Refinement Algorithm, and its variants, far outperform our theoretical guarantees, which all rely on worst-case estimates. We attempt to slightly improve one aspect of the standard optimality proof.

The goal of this chapter is to show that if $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ is a planar triangulation with minimum angle α, then

$$\int_E \frac{1}{\mathrm{lfs}_{\mathcal{T}}(z)} \mathrm{d}z = \mathcal{O}\left(\frac{1}{\alpha} \log\left(\frac{1}{\alpha}\right)\right) |\mathcal{V}|, \tag{10}$$

where $\mathrm{lfs}_{\mathcal{T}}(z)$ is the distance from z to the second nearest disjoint neighbor among $(\mathcal{V}, \mathcal{E})$, and $E = \cup_{e \in \mathcal{E}} e$. We also present a lower bound to show that this result cannot be improved asymptotically.

This result is similar to that of Mitchell, who proved that

$$\int_\Omega \frac{1}{\mathrm{lfs}^2_{\mathcal{T}}(z)} \mathrm{d}z = \mathcal{O}\left(\frac{1}{\alpha}\right) |\mathcal{V}|,$$

where Ω is the region of the plane covered by the triangulation $\mathcal{T}$ [31]. The new bound, however, can be applied in situations where point density is anisotropic, and skewed along certain predefined edges. For example, the Adaptive Delaunay Refinement Algorithm produces meshes with edges that can be comparatively short "across" small input angles, while being more modest along input segments (*cf.* Theorem 5.4.4).

Throughout this chapter we will assume that $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ is a triangulation with minimum angle α. Moreover, we assume that the triangulation is maximal, *i.e.*, that no edge may be added to $\mathcal{E}$ without crossing an existing edge.[1] The following lemma is due to Mitchell.

Lemma 7.0.5 (Mitchell [31]). *Let $e, f \in \mathcal{E}$ share a common endpoint in a triangulation with minimum angle α. Then*

$$\frac{|e|}{|f|} \leq (2 \cos \alpha)^{\frac{\angle ef}{\alpha}}.$$

[1] Mitchell drops this requirement when the input domain is a polygon with polygonal holes. In this case the holes represent regions where the output mesh is not maximal.

7.1 The Whirl

The bound of Mitchell's lemma involves a parametric shape of some interest. We define it precisely as follows

Definition 7.1.1 (α-whirl). Given points p, q and minimum angle α, we define the α-*whirl* centered at p to be the set

$$\mathcal{W}(p,q) = \left\{ x \;\middle|\; |p-x| = |p-q| \left(2\cos\alpha\right)^{\frac{-\angle qpx}{\alpha}} \right\},$$

where we assume $0 \leq \angle qpx \leq \pi$. Thus the α-whirl can be imagined as a parametric curve defined in polar coordinates as

$$\left(\theta, |p-q|\, g_\alpha(|\theta|)\right),$$

where p is the origin of the coordinate system, q is on the positive x-axis, $\theta \in (-\pi, \pi]$, and $g_\alpha(\theta) = (2\cos\alpha)^{\frac{-\theta}{\alpha}} = e^{-h_\alpha \theta}$, where $h_\alpha = \frac{1}{\alpha} \ln 2\cos\alpha$. An α-whirl is shown in Figure 32.

The α-whirl is two pieces of an equiangular spiral of angle $\arctan \frac{1}{h_\alpha}$, thus if z is a point on the curve, then segment (p, z) forms angle $\arctan \frac{1}{h_\alpha}$ with the tangent to the α-whirl at z.

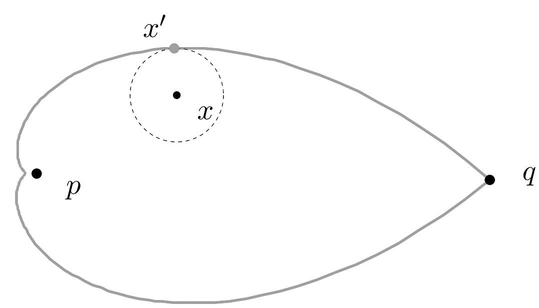

Figure 32: The α-whirl, $\mathcal{W}(p,q)$, with $\alpha \approx 27.6°$ is shown. For a given point x inside the closed curve, we wish to find the nearest point on the α-whirl, here marked as x'. We display a circle with center at x to show that this x' actually is the closest point on the α-whirl to this x. We wish to find x' in terms of the angle $\angle qpx$ and the distance $|p-x|$. It is clear that either (x, x') is normal to the α-whirl at x' or x' is the point collinear with p, q, at the "pinch point" of the α-whirl. There may be two points on the upper portion of the α-whirl whose normals pass through x, though one will be a local maximum to distance, the other a local minimum. The whirl becomes "flatter" as α becomes small.

Now we consider the question: what is the distance from an α-whirl to a point inside of it? The answer will be used in establishing upper and lower bounds on the integral $\int_E \frac{1}{\mathrm{lfs}_{\mathcal{T}}(z)} \mathrm{d}z$.

We consider the α-whirl, $\mathcal{W}(p,q)$. We adjust the units so that, for convenience, $|p-q| = 1$. We let z be the point $(\phi, \lambda g_\alpha(\phi))$, for $\phi \in [0,\pi]$, $\lambda \in [0,1]$. For this choice of ϕ, λ, we only need to consider the distance to the "upper" half of the α-whirl, so we can restrict θ to be in $[0, \pi]$. We consider the *squared* distance from z to the point $(\theta, g_\alpha(\theta))$, which is

$$\begin{aligned} f_{\lambda,\phi}(\theta) &= |(\phi, \lambda g_\alpha(\phi)) - (\theta, g_\alpha(\theta))|^2 \\ &= (\lambda g_\alpha(\phi)\cos\phi - g_\alpha(\theta)\cos\theta)^2 + (\lambda g_\alpha(\phi)\sin\phi - g_\alpha(\theta)\sin\theta)^2 \\ &= \lambda^2 g_\alpha(2\phi) + g_\alpha(2\theta) - 2\lambda g_\alpha(\theta+\phi)\cos(\theta-\phi). \end{aligned} \tag{11}$$

We have and will continue to use ubiquitously the fact that $g_\alpha(\cdot)$ is an exponential function, and thus $g_\alpha(a)g_\alpha(b) = g_\alpha(a+b)$.

Since $g_\alpha(\theta)$ is an exponential function, we have $\frac{dg_\alpha(\theta)}{d\theta} = -h_\alpha g_\alpha(\theta)$. We take the derivative of $f_{\lambda,\phi}(\theta)$:

$$\frac{df_{\lambda,\phi}(\theta)}{d\theta} = -2h_\alpha g_\alpha(2\theta) + 2\lambda g_\alpha(\theta+\phi)\left[\sin(\theta-\phi) + h_\alpha\cos(\theta-\phi)\right]. \tag{12}$$

Note that this derivative is continuous. To minimize $f_{\lambda,\phi}(\theta)$ we look for zeroes of its derivative.

Letting $\psi = \theta - \phi$, we have $\psi \in [-\phi, \pi - \phi] \subset [-\pi, \pi]$. We set the derivative to zero and find a solution, assuming ψ is in this range.

$$\begin{aligned} 0 &= -2h_\alpha g_\alpha(2\theta) + 2\lambda g_\alpha(\theta+\phi)\left[\sin(\theta-\phi) + h_\alpha\cos(\theta-\phi)\right] \\ \lambda g_\alpha(\phi)\left[\sin(\theta-\phi) + h_\alpha\cos(\theta-\phi)\right] &= h_\alpha g_\alpha(\theta) \\ \lambda &= \frac{h_\alpha g_\alpha(\theta-\phi)}{\sin(\theta-\phi) + h_\alpha\cos(\theta-\phi)} = \frac{h_\alpha g_\alpha(\psi)}{\sin\psi + h_\alpha\cos\psi}. \end{aligned}$$

If $\lambda = 0$, this last equation has no solution, as the numerator on the far right hand side is strictly positive. The denominator is positive only on $(-\arctan h_\alpha, \pi - \arctan h_\alpha)$.

Letting $j_\alpha(\psi) = \frac{h_\alpha g_\alpha(\psi)}{\sin\psi + h_\alpha\cos\psi}$, then $\frac{df_{\lambda,\phi}(\theta)}{d\theta} = 0$ if and only if $\lambda = j_\alpha(\theta - \phi)$. So we wish to find those λ for which there is $\psi \in (-\arctan h_\alpha, \pi - \arctan h_\alpha) \cap [-\phi, \pi - \phi]$ such that $\lambda = j_\alpha(\psi)$. To examine the behaviour of $j_\alpha(\cdot)$, we take its derivative. By the quotient rule and some algebra, we have

$$\frac{dj_\alpha(\psi)}{d\psi} = \frac{-h_\alpha g_\alpha(\psi)(h_\alpha^2+1)\cos\psi}{(\sin\psi + h_\alpha\cos\psi)^2}. \tag{13}$$

This derivative has the same sign as $-\cos\psi$. Since we take $\psi \in (-\arctan h_\alpha, \pi - \arctan h_\alpha)$, this derivative can only have a zero at $\pi/2$, and is negative if $\psi < \pi/2$, positive if $\psi > \pi/2$. As ψ approaches $-\arctan h_\alpha$ or $\pi - \arctan h_\alpha$, $j_\alpha(\psi)$ goes to infinity. Depending on λ, and ϕ there may be zero, one, or two roots to the equation $\lambda = j_\alpha(\psi) = j_\alpha(\theta - \phi)$ since we restrict $\psi \in (-\arctan h_\alpha, \pi - \arctan h_\alpha) \cap [-\phi, \pi - \phi]$. If there is only one root, it is less than $\pi/2$. If there are two, one is less than $\pi/2$, the other is greater.

We can then determine the sign of the derivative of $f_{\lambda,\phi}(\theta)$. Consider the left endpoint of the domain:

$$\begin{aligned} \left.\frac{df_{\lambda,\phi}(\theta)}{d\theta}\right|_{\theta=0} &= -2h_\alpha + 2\lambda g_\alpha(\phi)\left[\sin(-\phi) + h_\alpha\cos(-\phi)\right] \\ &= 2h_\alpha\left[\lambda g_\alpha(\phi)\cos\phi - 1\right] - 2\lambda g_\alpha(\phi)\sin\phi. \end{aligned}$$

Because h_α is positive and λ and $g_\alpha(\phi)$ are in $[0,1]$, the first term is non-positive. Because $\phi \in [0,\pi]$, the second term is nonpositive, so the derivative is nonpositive, and zero only if $\lambda = 1, \phi = 0$. In this latter case, the point $(\phi, \lambda g_\alpha(\phi))$ is actually on the α-whirl, so the distance to the curve is zero.

Assuming otherwise, the derivative of $f_{\lambda,\phi}(\theta)$ is negative at the left end of the domain, and has either zero, one or two roots (*i.e.*, the valid solutions to $\lambda = j_\alpha(\theta - \phi)$.) If there are zero roots, then the derivative stays negative and $f_{\lambda,\phi}(\theta)$ is minimized at $\theta = \pi$. If there is a single root, say θ_1, and it is of multiplicity one (*i.e.*, $\theta_1 \neq \pi/2$), then it must be a minimum, so $f_{\lambda,\phi}(\theta)$ is minimized at θ_1. If there are two roots, with θ_1 being the one smaller than $\pi/2$, then $f_{\lambda,\phi}(\theta)$ is minimized either at θ_1, or at π.

Since $j_\alpha(\psi)$ is minimized at $\pi/2$, then the derivative of $f_{\lambda,\phi}(\theta)$ clearly has no roots if $\lambda < j_\alpha(\pi/2) = h_\alpha g_\alpha(\pi/2)$; the derivative has a double root when $\lambda = j_\alpha(\pi/2)$. So assume $j_\alpha(\pi/2) < \lambda \leq 1$. Let $\psi_1 < \pi/2 < \psi_2$ be the roots to $\lambda = j_\alpha(\psi)$. These roots may or may not be in $[-\phi, \pi - \phi]$. Only if ϕ is sufficiently small will both roots be in this range. Otherwise, zero or one of these roots is in the range, and $f_{\lambda,\phi}(\theta)$ will be minimized at $\theta = \pi \wedge (\phi + \psi_1)$.

If both roots are in the range, then $\pi/2 < \psi_2 < \pi - \phi$, so then

$$\begin{aligned} f_{\lambda,\phi}(\pi) &= \lambda^2 g_\alpha(2\phi) + g_\alpha(2\pi) - 2\lambda g_\alpha(\pi+\phi)\cos(\pi-\phi) \\ &\geq \lambda^2 g_\alpha(2\phi), \end{aligned} \tag{14}$$

because $\lambda, g_\alpha(\cdot)$ are positive and $\cos(\pi - \phi)$ is negative. However

$$\begin{aligned} f_{\lambda,\phi}(\phi+\psi_1) &= \lambda^2 g_\alpha(2\phi) + g_\alpha(2\phi)g_\alpha(2\psi_1) - 2\lambda g_\alpha(2\phi)g_\alpha(\psi_1)\cos\psi_1 \\ &= \lambda^2 g_\alpha(2\phi) + g_\alpha(2\phi)g_\alpha(2\psi_1) - \frac{2h_\alpha g_\alpha(\psi_1)g_\alpha(2\phi)g_\alpha(\psi_1)\cos(\psi_1)}{\sin\psi_1 + h_\alpha\cos\psi_1} \\ &= \lambda^2 g_\alpha(2\phi) + g_\alpha(2\phi)g_\alpha(2\psi_1)\frac{\sin\psi_1 - h_\alpha\cos\psi_1}{\sin\psi_1 + h_\alpha\cos\psi_1}. \end{aligned} \tag{15}$$

Given that ψ_1 is a root to $\lambda = j_\alpha(\psi)$, the denominator of the fraction in equation 15 is positive. However, the numerator may be nonpositive; this will hold if $\psi_1 \leq \arctan h_\alpha$. Thus if $\phi < \pi/2$, and

$$\lambda \geq j_\alpha(\arctan h_\alpha) = \frac{h_\alpha g_\alpha(\arctan h_\alpha)}{\sin\arctan h_\alpha + h_\alpha\cos\arctan h_\alpha} = g_\alpha(\arctan h_\alpha)\frac{\sqrt{h_\alpha^2+1}}{2},$$

$\lambda \leq j_\alpha(\pi/2)$	$j_\alpha(\pi/2) < \lambda \leq j_\alpha(\arctan h_\alpha)$	$j_\alpha(\arctan h_\alpha) < \lambda \leq 1$
π	π or $\phi + \psi_1$	$\phi + \psi_1$

Table 1: The θ which minimizes $f_{\lambda,\phi}(\theta)$ for the given values of λ, ϕ is shown. We assume that $\phi \leq \pi/2$. For $\lambda > j_\alpha(\pi/2)$, ψ_1 is the root of $\lambda = j_\alpha(\psi)$ which is less than $\pi/2$. In the second column, there is some ambiguity as to which of $\pi, \phi + \psi_1$ minimizes $f_{\lambda,\phi}(\theta)$, due to limitations of the analysis. These results will be used in the establishment of both upper and lower bounds on the integral of interest. We note that $j_\alpha(\pi/2) = h_\alpha g_\alpha(\pi/2)$, and $j_\alpha(\arctan h_\alpha) = g_\alpha(\arctan h_\alpha)\frac{\sqrt{h_\alpha^2+1}}{2}$.

then $f_{\lambda,\phi}(\theta)$ is minimized at $\theta = \phi + \psi_1$.

We collect these results in Table 1, which shows the value of θ which minimizes $f_{\lambda,\phi}(\theta)$ given various values of λ, ϕ, with $\phi \leq \pi/2$. In this table, ψ_1 is the root of $\lambda = j_\alpha(\psi)$ less than $\pi/2$ if there is such a root. If there is such a root then, starting from equation 15,

$$\begin{aligned}
f_{\lambda,\phi}(\phi + \psi_1) &= \lambda^2 g_\alpha(2\phi) + g_\alpha(2\phi) g_\alpha(2\psi_1) \frac{\sin\psi_1 - h_\alpha \cos\psi_1}{\sin\psi_1 + h_\alpha\cos\psi_1}. \\
&= \frac{g_\alpha(2\psi_1) h_\alpha^2 g_\alpha(2\phi)}{(\sin\psi_1 + h_\alpha\cos\psi_1)^2} + g_\alpha(2\phi) g_\alpha(2\psi_1)\frac{\sin\psi_1 - h_\alpha\cos\psi_1}{\sin\psi_1 + h_\alpha\cos\psi_1}. \\
&= g_\alpha(2\phi) g_\alpha(2\psi_1) \frac{h_\alpha^2 + \sin^2\psi_1 - h_\alpha^2\cos^2\psi_1}{(\sin\psi_1 + h_\alpha\cos\psi_1)^2} \\
&= \frac{g_\alpha(2\phi) g_\alpha(2\psi_1)(h_\alpha^2+1)\sin^2\psi_1}{(\sin\psi_1 + h_\alpha\cos\psi_1)^2}. \qquad (16)
\end{aligned}$$

We will be interested, at times, in the "inverse" of the function $j_\alpha(\cdot)$, that is $j_\alpha^{-1}(\lambda)$. By inverse, we mean the unique $0 \leq \psi_1 \leq \pi/2$ such that $\lambda = j_\alpha(\psi_1)$. The following claim gives a lower bound sufficient for our analysis. It assumes that $h_\alpha \neq 0$, as otherwise, $j_\alpha(\psi)$ is identically zero.

Claim 7.1.2. Suppose $h_\alpha \neq 0$, and let $j_\alpha(\pi/2) \leq \lambda \leq 1$; then

$$j_\alpha^{-1}(\lambda) \geq \frac{-\ln\lambda}{h_\alpha + \frac{1}{h_\alpha}}. \qquad (17)$$

Proof. First we compare $\frac{\sin\psi + h_\alpha\cos\psi}{h_\alpha}$ to $e^{\frac{\psi}{h_\alpha}}$. Both functions take the value 1 when $\psi = 0$. The first derivative (with respect to ψ) of both functions takes value $\frac{1}{h_\alpha}$ at $\psi = 0$. The latter has positive derivative everywhere, the former has a derivative which experiences sign change. The second derivative of the former is nonpositive for $\psi \in [0, \pi/2]$, while the latter function has positive curvature. Thus the first derivative of $\frac{\sin\psi + h_\alpha\cos\psi}{h_\alpha}$ is less than that of $e^{\frac{\psi}{h_\alpha}}$ in this range. And thus the former function is less than the latter in this range.

Taking the reciprocals of these functions gives

$$e^{\frac{-\psi}{h_\alpha}} \leq \frac{h_\alpha}{\sin\psi + h_\alpha\cos\psi} = j_\alpha(\psi) e^{h_\alpha\psi}.$$

Thus $e^{-\psi(h_\alpha+\frac{1}{h_\alpha})} \leq j_\alpha(\psi)$. Note that $e^{-\psi(h_\alpha+\frac{1}{h_\alpha})}$ takes the same value as $j_\alpha(\psi)$ when $\psi = 0$, (namely 1), and bounds it below. Moreover it is monotone decreasing. Thus if $\lambda = e^{-\psi^*(h_\alpha+\frac{1}{h_\alpha})}$, and $0 \leq \psi_1 \leq \pi/2$ is a root to $\lambda = j_\alpha(\psi)$, then $\psi^* \leq \psi_1$. $\square$

We consider the maximum value of $j_\alpha(\pi/2)$ for $\alpha \in [0, \pi/3]$. We know that

$$j_\alpha(\pi/2) = h_\alpha g_\alpha(\pi/2) = h_\alpha e^{\frac{-\pi}{2}h_\alpha}.$$

Since h_α can take all values in $[0, \infty)$, it suffices to maximize $xe^{\frac{-\pi}{2}x}$ on $[0, \infty)$. Simple calculus shows this is maximized when $x = \frac{2}{\pi}$, and there has value $\frac{2}{\pi}e^{-1} < \frac{1}{4}$.

We also consider the maximum value of $j_\alpha(\arctan h_\alpha)$ for appropriate α. Again, h_α can take any nonnegative value, so it suffices to maximize $e^{-x \arctan x}\frac{\sqrt{x^2+1}}{2}$ for positive x. Calculus shows this function is monotone, and decreasing on $[0, \infty)$. Thus it has maximum value at 0, where it has value $\frac{1}{2}$.

7.2 An Upper Bound

We now focus on using the results concerning the α-whirl to find an upper bound for the integral $\int_E \frac{1}{\operatorname{lfs}_{\mathcal{T}}(z)}\mathrm{d}z$. The analysis is fairly involved, although none of it would be beyond an advanced student of calculus. We present an overview here to guide the reader: First we show that if x is a point on a segment (p, q) of a mesh, and x is closer to p than q (or x is the midpoint of the segment), then there is some point q' on the segment such that the distance from x to the α-whirl $\mathcal{W}(p, q')$ is a *lower* bound on $\operatorname{lfs}(x)$. This is done simply in Lemma 7.2.1 by showing that no edge or point of the mesh disjoint from (p, q) or the point p is inside this α-whirl. Then in Theorem 7.2.2, the distance from a point to a α-whirl is used to bound the integral.

Lemma 7.2.1. *Let $e = (p, q)$ be an edge of a triangulation with minimum angle α. Let q' be the point on (p, q) such that $|p - q'| = \frac{\sqrt{3}}{2}|p - q|$. Then there is no edge or vertex of the triangulation which is disjoint from e or from p and contained inside the closed curve $\mathcal{W}(p, q')$.*

Proof. Let $\{q_i\}_{i=0}^{n-1}$ be the set of vertices of the triangulation such that (p, q_i) is an edge of the triangulation; moreover, assume $q_0 = q$, and the vertices q_i are ordered counterclockwise around p. All the triangles in the triangulation which have p as a corner are of the form Δpq_iq_{i+1} for some $0 \leq i < n$, where q_n is read to be q_0.

Then it suffices to show that if Δpq_iq_{i+1} is a triangle of the triangulation then no point of (q_i, q_{i+1}) is inside $\mathcal{W}(p, q')$. Let θ be the counterclockwise angle $\angle q_0pq_i$; without loss of generality we assume this is less than the counterclockwise angle $\angle q_0pq_{i+1}$, and that $\theta \in [0, \pi)$, as otherwise we look from the other side of the plane.

Let z be a point on (q_i, q_{i+1}). Again, without loss of generality we can assume that $\angle q_0 pz \in [\theta, \pi]$, as otherwise we look from the other side of the plane. It suffices to show that $|p - z| \geq \frac{\sqrt{3}}{2} |p - q_0| \, g_\alpha(\angle q_0 pz)$.

Let $\phi = \angle q_i pz$. By assumption, $0 \leq \phi \leq \pi - 2\alpha$. Let $\psi = \angle pq_i q_{i+1}$. Using the sine rule on the triangle $\Delta pq_i z$, we have

$$\frac{|p - z|}{\sin \psi} = \frac{|p - q_i|}{\sin(\phi + \psi)}.$$

By Lemma 7.0.5, $|p - q_i| \geq |p - q_0| \, g_\alpha(\theta)$, and so

$$|p - z| \geq \frac{|p - q_0| \, g_\alpha(\theta) \sin \psi}{\sin(\phi + \psi)}.$$

Since $\angle q_0 pz = \theta + \phi$, it suffices to prove that $\frac{\sin \psi}{\sin(\phi+\psi)} \geq \frac{\sqrt{3}}{2} g_\alpha(\phi)$. So we attempt to minimize $k_\psi(\phi) = \frac{\sin \psi}{\sin(\phi+\psi)} g_\alpha(-\phi)$.

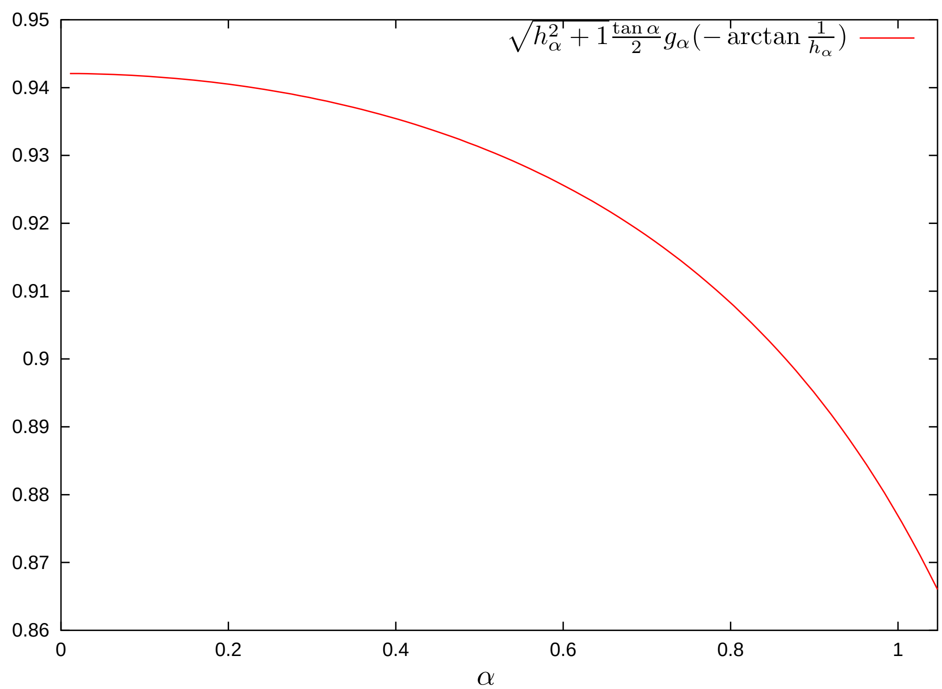

Figure 33: The function $\sqrt{h_\alpha^2 + 1}\frac{\tan\alpha}{2} g_\alpha(-\arctan\frac{1}{h_\alpha})$ takes its minimal value of $\frac{\sqrt{3}}{2}$ at $\alpha = \pi/3$.

We take the derivative, which is

$$\frac{dk_\psi(\phi)}{d\phi} = g_\alpha(-\phi) \sin \psi \frac{h_\alpha \sin(\phi + \psi) - \cos(\phi + \psi)}{\sin^2(\phi + \psi)}.$$

The derivative is zero when $\tan \phi + \psi = \frac{1}{h_\alpha}$. Note that ϕ, ψ are angles of $\Delta pq_i z$ so $\phi < \pi - \psi$, and thus there is at most one extreme point for the function $k_\psi(\phi)$. If it happens to be the case that $\arctan \frac{1}{h_\alpha} \le \psi$, then there is no extreme point, but in this case it is simple to show that $\left.\frac{dk_\psi(\phi)}{d\phi}\right|_{\phi=0}$ is positive, so the function has minimal value $k_\psi(0) = 1$.

Otherwise $\psi < \arctan \frac{1}{h_\alpha}$, and there is an extreme point which is a minimum. The value of $k_\psi(\phi)$ at this minimum is

$$k_\psi\left(\arctan \frac{1}{h_\alpha} - \psi\right) = \sqrt{h_\alpha^2 + 1}\sin\psi g_\alpha(\psi) g_\alpha(-\arctan \frac{1}{h_\alpha}).$$

We note that $\sin\psi g_\alpha(\psi)$ is increasing on $\left[0, \arctan \frac{1}{h_\alpha}\right)$, so we may assume ψ takes its minimum value, α. Thus $k_\psi(\phi)$ is bounded below by $\sqrt{h_\alpha^2+1}\frac{\tan\alpha}{2} g_\alpha(-\arctan\frac{1}{h_\alpha})$. We claim that this function takes a minimum value of $\frac{\sqrt{3}}{2}$ when $\alpha = \pi/3$, which establishes the lemma.

To verify this claim, most readers will be satisfied to look at the graph of the function versus α, as shown in Figure 33. The skeptic (or masochist), however, may insist on a purely analytic demonstration of the monotonicity of this function. The analysis is fairly difficult and not easily forthcoming, so we omit it; the interested reader is encouraged to verify that the derivative of this function is indeed negative. ☐

The idea of Lemma 7.2.1 is illustrated in Figure 34 for two different triangulations. The lemma becomes tight for triangulations which have many angles of size α, as shown in Figure 34(b), otherwise it can be a severe overestimate, as shown in Figure 34(a).

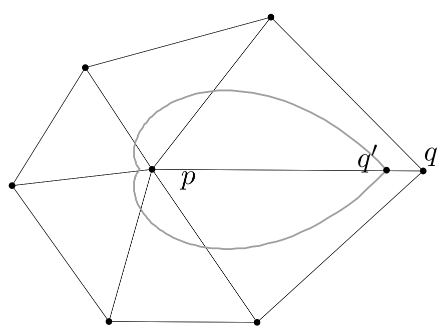

(a) "Normal" Triangulation

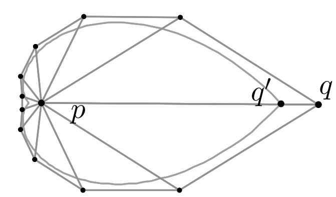

(b) "Pathological" Triangulation

Figure 34: Lemma 7.2.1 is illustrated. In (a), a "normal" triangulation is shown, with the α-whirl, $\mathcal{W}(p, q')$, where $|p - q'| = \frac{\sqrt{3}}{2}|p - q|$. In (b), a more pathological triangulation is shown; all but one triangle are isosceles with two angles equal to α, which is around 32° in this figure.

Theorem 7.2.2. *Let $e = (p, q)$ be an edge of a triangulation with minimum angle α. Let m be the midpoint of the edge. Let $\mathrm{lfs}_{\mathcal{T}}(z)$ be the local feature size with respect to the* *triangulation. Then*

$$\int_p^m \frac{1}{\operatorname{lfs}_{\mathfrak{T}}(z)} \mathrm{d}z \le \ln\frac{3}{2} + h_\alpha \pi + \sqrt{h_\alpha^2 + 1}\left[\ln(h_\alpha + 1/h_\alpha) - \ln(2\sin\frac{\ln 3}{8})\right], \qquad (18)$$

and

$$\int_p^m \frac{1}{\operatorname{lfs}_{\mathfrak{T}}(z)} \mathrm{d}z \le \ln\frac{3}{2} + h_\alpha \pi + \sqrt{h_\alpha^2 + 1}\left[\frac{h_\alpha^2 + 1}{2\sin\frac{\ln 3}{4}} - h_\alpha\right]. \qquad (19)$$

Note that inequality 18 is a tighter approximation when α is small, while inequality 19 is more appropriate for larger α. The tradeoff point of these two approximations occurs when α is approximately $37.59°$.

Proof. Let q' be the point on segment (p, q) such that $|p - q'| = \frac{\sqrt{3}}{2}|p - q|$. By Lemma 7.2.1, if Γ is the α-whirl, $\mathcal{W}(p, q')$, there is no edge or vertex of the triangulation which is disjoint from the edge e or the point p, and which intersects the region interior to the closed curve Γ. Thus if z is a point on the line segment (p, m), then $\operatorname{lfs}_{\mathfrak{T}}(z)$ is at least the distance from z to Γ.

Let $\lambda = \frac{2}{\sqrt{3}}\frac{|p-z|}{|p-q|}$. Then

$$I =_{\mathrm{df}} \int_p^m \frac{1}{\operatorname{lfs}_{\mathfrak{T}}(z)} \mathrm{d}z \le \frac{\sqrt{3}|p-q|}{2}\int_0^{\frac{1}{\sqrt{3}}} \frac{1}{|p-q'|\sqrt{f_{\lambda,0}(mn(\lambda))}} \mathrm{d}\lambda,$$

where $mn(\lambda)$ is that θ which minimizes $f_{\lambda,0}(\theta)$, which is the squared distance from the point $(0, \lambda)$, (in polar coordinates) to the unit-length α-whirl centered on the origin and running along the x-axis.

We now use the results tabulated in Table 1 to find the values of $mn(\lambda)$. Since there is some ambiguity in the second column, we "double-up" the difference. That is if $l = a \wedge b$, for positive quantities a, b, then $\frac{1}{l} \le \frac{1}{a} + \frac{1}{b}$. Thus

$$I \le \int_0^{j_\alpha(\arctan h_\alpha)} \frac{1}{\sqrt{f_{\lambda,0}(\pi)}} \mathrm{d}\lambda + \int_{j_\alpha(\pi/2)}^{\frac{1}{\sqrt{3}}} \frac{1}{\sqrt{f_{\lambda,0}(\psi_1)}} \mathrm{d}\lambda,$$

where, again, ψ_1 is the root of $\lambda = j_\alpha(\psi)$ which is less than $\pi/2$. We have shown in the previous section that $j_\alpha(\pi/2) < \frac{1}{4} < \frac{1}{\sqrt{3}}$, thus the second integral does not have inverted limits. We have also noted that $j_\alpha(\arctan h_\alpha) \le \frac{1}{2} < \frac{1}{\sqrt{3}}$, so we have not made an overestimate in the first integral.

We substitute the value of $f_{\lambda,0}(\pi)$ in the first integral, and make a change of variables to the second; every $\lambda \in \left[j_\alpha(\pi/2), \frac{1}{\sqrt{3}}\right]$ has some associated root ψ_1, given by $\lambda = j_\alpha(\psi_1)$.

By equation 13, we can find $\mathrm{d}\lambda$ in terms of $\mathrm{d}\psi_1$, and ψ_1. Thus, using equation 16,

$$\begin{aligned} I &\leq \int_0^{j_\alpha(\arctan h_\alpha)} \frac{1}{\lambda+g_\alpha(\pi)}\mathrm{d}\lambda \\ &\quad + \int_{\pi/2}^{j_\alpha^{-1}(\frac{1}{\sqrt{3}})} \frac{(\sin\psi_1+h_\alpha\cos\psi_1)}{g_\alpha(\psi_1)\sqrt{h_\alpha^2+1}\sin\psi_1} \frac{-h_\alpha g_\alpha(\psi_1)(h_\alpha^2+1)\cos\psi_1}{(\sin\psi_1+h_\alpha\cos\psi_1)^2}\mathrm{d}\psi_1, \\ &= \ln\left|\frac{j_\alpha(\arctan h_\alpha)+g_\alpha(\pi)}{g_\alpha(\pi)}\right| - h_\alpha\sqrt{h_\alpha^2+1}\int_{\pi/2}^{j_\alpha^{-1}(\frac{1}{\sqrt{3}})} \frac{\cos\psi_1}{\sin\psi_1(\sin\psi_1+h_\alpha\cos\psi_1)}\mathrm{d}\psi_1, \\ &\leq \ln\left|\frac{3}{2g_\alpha(\pi)}\right| + h_\alpha\sqrt{h_\alpha^2+1}\int_{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}^{\pi/2} \frac{\cos\psi_1}{\sin\psi_1(\sin\psi_1+h_\alpha\cos\psi_1)}\mathrm{d}\psi_1, \end{aligned}$$

where in the last line we have used the fact that $j_\alpha(\arctan h_\alpha) \leq \frac{1}{2}$, and $g_\alpha(\pi) \leq 1$. Since $\ln g_\alpha(\pi) = -h_\alpha\pi$, the first term on the right hand side above becomes $\ln\frac{3}{2} + h_\alpha\pi$, which is the common term of inequality 18 and inequality 19. The analysis for the two inequalities diverges here, so we consider them separately.

- For inequality 18, it suffices to prove that

$$J =_{\mathrm{df}} h_\alpha \int_{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}^{\pi/2} \frac{\cos\psi_1}{\sin\psi_1(\sin\psi_1+h_\alpha\cos\psi_1)}\mathrm{d}\psi_1 \leq \left[\ln(h_\alpha+1/h_\alpha) - \ln(2\sin\frac{\ln 3}{8})\right].$$

  The integral splits conveniently as

$$\begin{aligned} J &= h_\alpha\int_{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}^{\pi/2} \frac{1}{h_\alpha\sin\psi_1} - \frac{1}{h_\alpha(\sin\psi_1+h_\alpha\cos\psi_1)}\mathrm{d}\psi_1, \\ &\leq \int_{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}^{\pi/2} \frac{1}{\sin\psi_1}\mathrm{d}\psi_1, \\ &= \ln\tan\frac{\psi_1}{2}\bigg|_{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}^{\pi/2} = \left[\ln\left(\tan\frac{\pi}{4}\right) - \ln\left(\tan\frac{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}{2}\right)\right], \\ &\leq \ln\left(\cot\frac{-\ln\frac{1}{\sqrt{3}}}{2(h_\alpha+\frac{1}{h_\alpha})}\right), \\ &\leq \ln\left(\csc\frac{\ln 3}{4(h_\alpha+\frac{1}{h_\alpha})}\right). \end{aligned}$$

  In the second line we have thrown away the negative term in the integral; we flipped the sign of the term by inverting tangent to cotangent, then used the lower bound of Claim 7.1.2, *i.e.*, $j_\alpha^{-1}(\lambda) \geq \frac{-\ln\lambda}{h_\alpha+\frac{1}{h_\alpha}}$. We then bounded the cosine and manipulated the inner log to get the last line.

  We now bound the cosecant. For admissible values of α, h_α may assume any nonnegative value. Then $\frac{\ln 3}{4(h_\alpha+1/h_\alpha)}$ can assume any value in $\left(0, \frac{\ln 3}{8}\right]$. In this range, $\sin x$ is bounded from below by $(\sin\frac{\ln 3}{8})x/(\frac{\ln 3}{8})$. Thus the cosecant is bounded above:

$$\csc\frac{\ln 3}{4(h_\alpha+1/h_\alpha)} \leq \frac{\ln 3}{8\sin\frac{\ln 3}{8}}\frac{1}{(\ln 3)/(4(h_\alpha+1/h_\alpha))} = \frac{h_\alpha+1/h_\alpha}{2\sin\frac{\ln 3}{8}}.$$

This establishes the inequality.

- For inequality 19, it suffices to prove that

$$J =_{\mathrm{df}} h_\alpha \int_{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}^{\pi/2} \frac{\cos\psi_1}{\sin\psi_1(\sin\psi_1 + h_\alpha\cos\psi_1)}\mathrm{d}\psi_1 \leq \left[\frac{h_\alpha^2+1}{2\sin\frac{\ln 3}{4}} - h_\alpha\right].$$

Since h_α is nonnegative, and $\cos\psi_1$ is positive in the given range, we can bound $\frac{1}{\sin\psi_1 + h_\alpha\cos\psi_1} \leq \frac{1}{\sin\psi_1}$, to get

$$\begin{aligned} J &\leq h_\alpha \int_{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}^{\pi/2} \frac{\cos\psi_1}{\sin^2\psi_1}\mathrm{d}\psi_1, \\ &= h_\alpha\left(-\csc\psi_1\Big|_{j_\alpha^{-1}(\frac{1}{\sqrt{3}})}^{\pi/2}\right) = h_\alpha\left(\csc j_\alpha^{-1}(\frac{1}{\sqrt{3}}) - 1\right), \\ &\leq h_\alpha\left(\csc\frac{\ln 3}{2(h_\alpha + \frac{1}{h_\alpha})} - 1\right), \\ &\leq h_\alpha\left(\frac{h_\alpha + 1/h_\alpha}{2\sin\frac{\ln 3}{4}} - 1\right). \end{aligned}$$

As above we have used the lower bound on $j_\alpha^{-1}(\cdot)$ and bounded the cosecant, establishing the desired inequality.

□

Corollary 7.2.3. *Let $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ be a triangulation with minimum angle α, where $\alpha \leq \pi/6$. Let $\mathrm{lfs}_{\mathcal{T}}(z)$ be the local feature size of a point z with respect to the triangulation, and let $E = \cup_{e\in\mathcal{E}} e$. Then*

$$\int_E \frac{1}{\mathrm{lfs}_{\mathcal{T}}(z)}\mathrm{d}z = \mathcal{O}\left(\frac{1}{\alpha}\log\left(\frac{1}{\alpha}\right)\right)|\mathcal{V}|.$$

Proof. Let $e = (p, q)$ be an edge of the triangulation. By the theorem

$$\int_p^q \frac{1}{\mathrm{lfs}_{\mathcal{T}}(z)}\mathrm{d}z < 2\left[\ln\frac{3}{2} + h_\alpha\pi + \sqrt{h_\alpha^2+1}\left(\ln(h_\alpha + \frac{1}{h_\alpha}) - \ln 2\sin\frac{\ln 3}{8}\right)\right].$$

Since $\alpha \in (0, \pi/6]$, $h_\alpha > 1$, and $h_\alpha = \mathcal{O}\left(\frac{1}{\alpha}\right)$. Then $\int_p^q \frac{1}{\mathrm{lfs}_{\mathcal{T}}(z)}\mathrm{d}z = \mathcal{O}\left(\frac{1}{\alpha}\ln\frac{1}{\alpha}\right)$. Thus

$$\int_E \frac{1}{\mathrm{lfs}_{\mathcal{T}}(z)}\mathrm{d}z = \mathcal{O}\left(\frac{1}{\alpha}\log\left(\frac{1}{\alpha}\right)\right)|\mathcal{E}|.$$

By Euler's formula $|\mathcal{E}| \leq 3|\mathcal{V}|$, which suffices. □

7.3 A Lower Bound

We now show that the upper bound of the previous section is optimal by demonstrating an example where the integral is large. The example consists of a triangulation where the local feature size is bounded from above by the distance to an α-whirl.

Definition 7.3.1. Let $n \geq 4$ be some large even integer, and let $\alpha = \pi/n$. Let $\mathcal{V}_n = \{p\} \cup \{q_i\}_{i=0}^{n}$, where p is the origin of a polar coordinate system, and q_i is the point, in polar coordinates, $(\alpha i, g_\alpha(\alpha i))$. Let $\mathcal{E}_n = \{(p, q_i)\}_{i=0}^{i=n} \cup \{(q_i, q_{i+1})\}_{i=0}^{n-1}$.

Then the *spiral mesh* on n is the triangulation $\mathcal{T}_n = (\mathcal{V}_n, \mathcal{E}_n)$. The spiral mesh is a triangulation with minimum angle α. We illustrate such a triangulation, for a small n, in Figure 35.

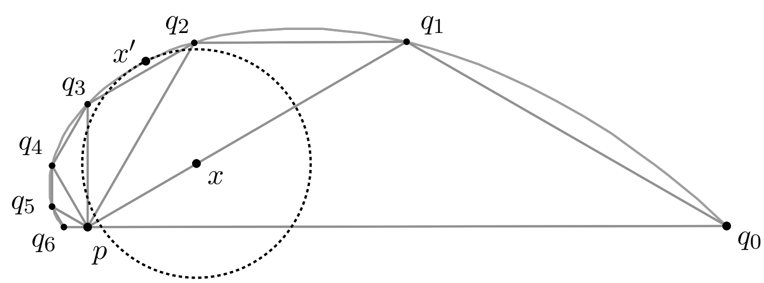

Figure 35: The spiral mesh on $n = 6$ is shown. We use small n for visualization purposes, as the triangulation becomes flattened along the segment (p, q_0) for large n. We also show the top half of the α-whirl, $\mathcal{W}(p, q_0)$, for $\alpha = \pi/6$. For the point x shown on the segment (p, q_1), the distance from x to $\mathcal{W}(p, q_0)$ is greater than the local feature size of x with respect to the mesh, which is the distance from x to (q_1, q_2). This holds when $\angle xpx' \geq \alpha$, where x' is the closest point to x on $\mathcal{W}(p, q_0)$.

Given an edge (p, q_i) of a spiral mesh, for many points on the edge, the local feature size of the point with respect to the mesh is approximately the distance from the point to the α-whirl, $\mathcal{W}(p, q_0)$. This is shown in Figure 35, and holds for any x on an edge such that $\angle xpx' \geq \alpha$, where x' is the point on $\mathcal{W}(p, q_0)$ closest to x. Thus the following lower bound proof uses much of the same technology as the upper bound proof.

Theorem 7.3.2. *Let $\epsilon > 0$ be given. Then there is some $N \geq 4$ such that for $n \geq N$, if $\operatorname{lfs}_{\mathcal{T}}(z)$ is the local feature size with respect to the spiral mesh on n, and $0 \leq i \leq n/2$, then*

$$\int_p^{q_i} \frac{1}{\operatorname{lfs}_{\mathcal{T}}(z)} \mathrm{d}z \geq (1 - \epsilon) \ln 2 \left[\frac{1}{\alpha} \ln \frac{1}{\alpha}\right],$$

where $\alpha = \pi/n$.

Proof. We determine sufficient conditions on N. Assume that $n \geq N \geq 4$. Let $\phi = \alpha i$. Let z be a point on the segment (p, q_i) such that $|p - z| = \lambda g_\alpha(\phi)$. We claim that $\operatorname{lfs}_{\mathcal{T}}(z) \leq \sqrt{f_{\lambda,\phi}(\theta)}$ for any $\theta \in [\phi + \alpha, \pi]$. This is the case because the line segment from z to the point, in polar coordinates, $(\theta, g_\alpha(\phi))$, will cut through a line segment of the form (q_j, q_{j+1}) for $i < j < n$ if $\alpha \leq \theta$. This line segment is disjoint from (p, q_i), giving the upper bound on local feature size.

We make a change of variables from z to λ, then use this upper bound. We let $\vec{v_i}$ be the vector from p to q_i.

$$I = \int_p^{q_i} \frac{1}{\mathrm{lfs}_{\mathfrak{J}}(z)} \mathrm{d}z = g_\alpha(\phi) \int_0^1 \frac{1}{\mathrm{lfs}_{\mathfrak{J}}(p + \lambda \vec{v_i})} \mathrm{d}\lambda \geq g_\alpha(\phi) \int_{j_\alpha(\pi/2)}^{j_\alpha(\alpha)} \frac{1}{\sqrt{f_{\lambda,\phi}(mn(\lambda))}} \mathrm{d}\lambda,$$

where, again, $mn(\lambda)$ is the point on the α-whirl closest to $(\phi, \lambda g_\alpha(\phi))$. We have made the domain smaller, which suffices as the integrand is positive, to ensure that the closest point is on the α-whirl for an angle $\theta \geq \phi + \alpha$. We've also truncated the lower end of the domain for simplicity of calculation.

Over the given range of λ, $f_{\lambda,\phi}(mn(\lambda)) \leq f_{\lambda,\phi}(\phi + \psi_1)$, where ψ_1 is the root to $\lambda = j_\alpha(\psi)$ which is less than $\pi/2$. Here the ambiguity in the second column of Table 1 is no matter to us, since we need only an upper bound on $f_{\lambda,\phi}(mn(\lambda))$. Note also that since $i \leq n/2$, that $\phi \leq \pi/2$, and thus that the results of Table 1 really do apply.

We now make a change of variables again, to ψ_1. By equation 13, we can find $\mathrm{d}\lambda$ in terms of $\mathrm{d}\psi_1$, and ψ_1. Thus, using equation 16,

$$\begin{aligned}
I &\geq g_\alpha(\phi) \int_{\pi/2}^{\alpha} \frac{(\sin\psi_1 + h_\alpha \cos\psi_1)}{g_\alpha(\phi) g_\alpha(\psi_1)\sqrt{h_\alpha^2+1}\sin\psi_1} \frac{-h_\alpha g_\alpha(\psi_1)(h_\alpha^2+1)\cos\psi_1}{(\sin\psi_1 + h_\alpha\cos\psi_1)^2} \mathrm{d}\psi_1, \\
&= h_\alpha\sqrt{h_\alpha^2+1} \int_\alpha^{\pi/2} \frac{1}{h_\alpha \sin\psi_1} - \frac{1}{h_\alpha(\sin\psi_1 + h_\alpha\cos\psi_1)} \mathrm{d}\psi_1, \\
&= \sqrt{h_\alpha^2+1}\ \ln\tan\frac{\psi_1}{2}\Big|_\alpha^{\pi/2} - \sqrt{h_\alpha^2+1} \int_\alpha^{\pi/2} \frac{1}{(\sin\psi_1 + h_\alpha\cos\psi_1)} \mathrm{d}\psi_1.
\end{aligned}$$

We now find a simple lower bound on $\sin x + h_\alpha \cos x$ on $(0, \pi/2]$. Since this function is concave (down) on this domain, we can bound it from below by the linear which interpolates its endpoints. This is the line $y = \frac{1-h_\alpha}{\pi/2}x + h_\alpha$. Thus we can bound

$$\begin{aligned}
I &\geq -\sqrt{h_\alpha^2+1}\ln\tan\frac{\alpha}{2} - \sqrt{h_\alpha^2+1} \int_\alpha^{\pi/2} \frac{1}{h_\alpha + \frac{1-h_\alpha}{\pi/2}\psi_1} \mathrm{d}\psi_1, \\
&= \sqrt{h_\alpha^2+1}\ln\cot\frac{\alpha}{2} - \sqrt{h_\alpha^2+1}\frac{\pi/2}{1-h_\alpha}\ \ln\left(h_\alpha + \frac{1-h_\alpha}{\pi/2}\psi_1\right)\Bigg|_\alpha^{\pi/2}, \\
&\geq h_\alpha \ln\cot\frac{\alpha}{2} + \sqrt{h_\alpha^2+1}\frac{\pi/2}{1-h_\alpha}\ln\left(h_\alpha + \frac{1-h_\alpha}{\pi/2}\alpha\right) \qquad (20)
\end{aligned}$$

We examine the first term of inequality 20. Since $n \geq N \geq 4$, we have $\cos\frac{\alpha}{2} > 0.923$.

We then bound the sine of x from above by x to get the following bound:

$$\begin{aligned} h_\alpha \ln \cot \frac{\alpha}{2} &= h_\alpha \ln \cos \frac{\alpha}{2} - h_\alpha \ln \sin \frac{\alpha}{2}, \\ &\geq h_\alpha \ln 0.923 - h_\alpha \ln \frac{\alpha}{2}, \\ &= h_\alpha \ln 0.923 + h_\alpha \ln 2 + h_\alpha \ln \frac{1}{\alpha}, \\ &= h_\alpha \ln 1.846 + h_\alpha \ln \frac{1}{\alpha} \geq h_\alpha \ln \frac{1}{\alpha}. \end{aligned} \tag{21}$$

Now we suppose that $N \geq 10$. This insures that $h_\alpha \geq 2$, and the second term from inequality 20 has a leading negative factor (its denominator). Thus we bound the log part from above.

$$\begin{aligned} \ln\left(h_\alpha + \frac{1-h_\alpha}{\pi/2}\alpha\right) &= \ln\left(h_\alpha + \frac{2}{n} - \frac{2}{\pi}\ln 2\cos\alpha\right), \\ &\leq \ln\left(h_\alpha + \frac{2}{n}\right) \leq \ln\left(h_\alpha + \frac{h_\alpha}{n}\right) \leq \ln\left(h_\alpha + \frac{h_\alpha}{10}\right), \\ &= \ln\left(\frac{11}{10}\ln 2\cos\alpha\right) + \ln\frac{1}{\alpha}, \\ &\leq \ln 0.763 + \ln\frac{1}{\alpha} \leq \ln\frac{1}{\alpha}. \end{aligned} \tag{22}$$

We now bound the leading part of the second term. Clearly $\sqrt{h_\alpha^2+1} \leq h_\alpha + 1$. Since we supposed that $h_\alpha \geq 2$, then

$$\frac{\sqrt{h_\alpha^2+1}}{1-h_\alpha} \geq \frac{h_\alpha+1}{1-h_\alpha} = -1 + \frac{2}{1-h_\alpha}.$$

Given the bound on h_α, the right hand side is greater than -3. Thus

$$\frac{\sqrt{h_\alpha^2+1}}{1-h_\alpha}\pi/2 \geq -5.$$

Combining this with inequalities (20)-(22) gives

$$I \quad \geq \quad (h_\alpha - 5)\ln\frac{1}{\alpha} \tag{23}$$

Now we assume that

$$N \geq \max\left\{\frac{10\pi}{\epsilon\ln 2}, \frac{\pi}{\arccos(2^{-\frac{\epsilon}{2}})}\right\}.$$

The first bound on N insures that $5\alpha \leq 5\pi/N \leq \frac{\epsilon\ln 2}{2}$. The second insures that $\ln\cos\alpha \geq -\frac{\epsilon\ln 2}{2}$. Then

$$h_\alpha - 5 = [\ln 2\cos\alpha - 5\alpha]\frac{1}{\alpha} = [\ln 2 + \ln\cos\alpha - 5\alpha]\frac{1}{\alpha} \geq \left[\ln 2 - \frac{\epsilon\ln 2}{2} - \frac{\epsilon\ln 2}{2}\right]\frac{1}{\alpha},$$

and thus $h_\alpha - 5 \geq [(1-\epsilon)\ln 2]\frac{1}{\alpha}$. Combining with inequality 23 gives the theorem.

We note that the chosen N is not optimal, but suffices for a rough analysis. Thus for example, for all $\epsilon < 1$, we require $N > 46$, but a less general analysis can show that, for example, if $N \geq 30$, then $I \geq \frac{2}{5}\left[\frac{1}{\alpha}\ln\frac{1}{\alpha}\right]$. $\square$

Corollary 7.3.3. *For sufficiently small α, there exist triangulations, $(\mathcal{V}, \mathcal{E})$ with minimum angle α such that*

$$\int_E \frac{1}{\operatorname{lfs}_{\mathcal{T}}(z)} \mathrm{d}z = \Omega\left(\frac{1}{\alpha}\log\left(\frac{1}{\alpha}\right)\right)|\mathcal{V}|,$$

where $\operatorname{lfs}_{\mathcal{T}}(z)$ is the local feature size of point z with respect to the triangulation, and $E = \cup_{e\in\mathcal{E}} e$.

Proof. We apply the theorem for $\epsilon = \frac{1}{2}$, getting an appropriate N. We let $\alpha = \pi/N$, then consider the spiral mesh on N. We can bound the integral by the value on the $\frac{N}{2}+1$ spokes for which the theorem applies:

$$\int_E \frac{1}{\operatorname{lfs}_{\mathcal{T}}(z)} \mathrm{d}z \geq \sum_{i=0}^{N/2} \int_p^{q_i} \frac{1}{\operatorname{lfs}_{\mathcal{T}}(z)} \mathrm{d}z \geq (1+N/2)\frac{\ln 2}{2}\frac{1}{\alpha}\ln\frac{1}{\alpha} = \frac{\ln 2}{4}\left(\frac{1}{\alpha}\ln\frac{1}{\alpha}\right)|\mathcal{V}|.$$

□

7.4 *Using the Upper Bound*

The upper bound was used in Corollary 5.6.1 to get asymptotic improvements in the optimality constants of the Adaptive Delaunay Refinement Algorithm. As discussed in Section 5.8, the improvements from using this upper bound are tangible, but the optimality constants are ultimately too large to be of any real value.

We briefly consider the improvements in optimality that can be stated for Ruppert's Algorithm under the non-acute input condition. Ruppert proved [39] the following grading constants for the case where input meet at non-acute angles:

$$C_T = \frac{1+2\sin\alpha}{1-2\sqrt{2}\sin\alpha}, \quad C_S = \frac{1+\sqrt{2}}{1-2\sqrt{2}\sin\alpha},$$

where $\alpha < \arcsin\frac{1}{2\sqrt{2}}$ is the output angle parameter, and minimum angle in the output mesh. Using Mitchell's work, the optimality constant for Ruppert's Algorithm is

$$\zeta'(\alpha) = \frac{2}{\pi}(2C_S+3)^2\left(\frac{21.5}{\alpha}+11.9\right),$$

The improved optimality constant is

$$\zeta(\alpha) = 1 + \zeta_m(\alpha) + \zeta_t(\alpha),$$

where

$$\begin{aligned}
\zeta_m(\alpha) &= \frac{6}{2\log\left|1+\frac{1}{2(C_S+1)}\right|}\left(\ln\frac{3}{2} + h_\alpha\pi + \sqrt{h_\alpha^2+1}\left[\ln(h_\alpha + 1/h_\alpha) - \ln(2\sin\frac{\ln 3}{8})\right]\right),\\
\zeta_t(\alpha) &= \frac{1}{\pi}(2C_T+3)^2\left(\frac{21.5}{\alpha}+11.9\right).
\end{aligned}$$

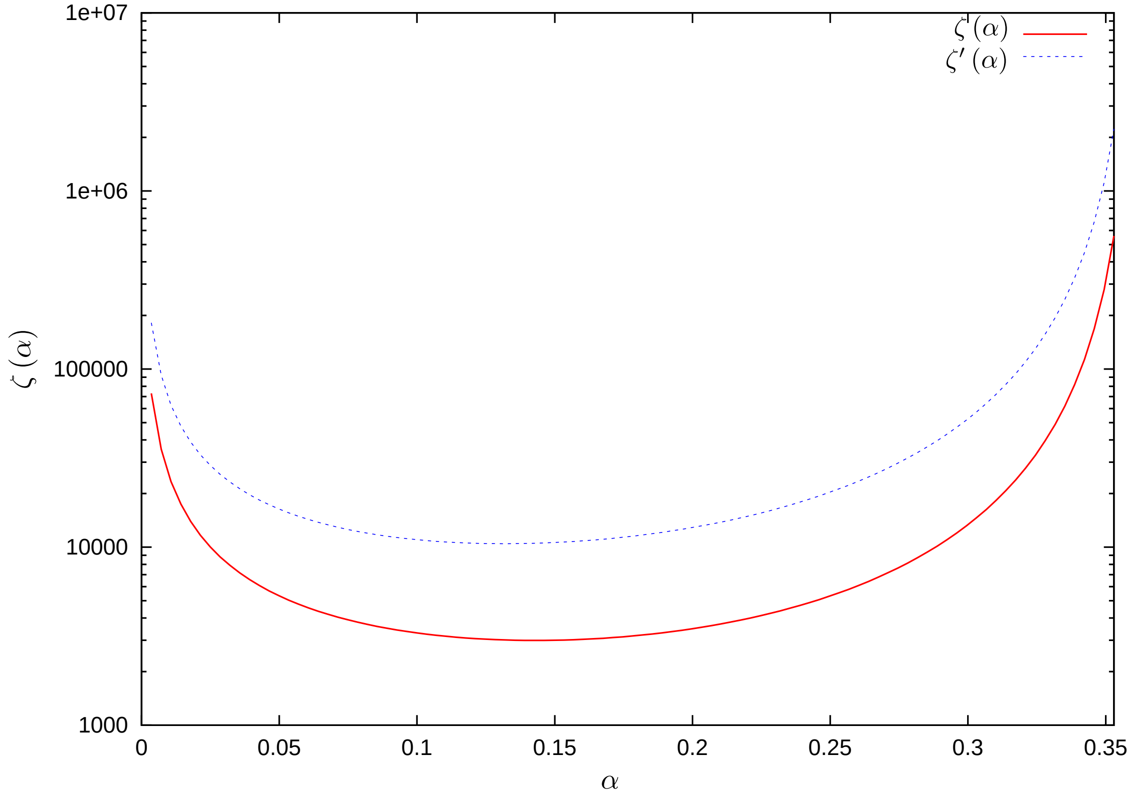

Figure 36: The improvements in the optimality constant for Ruppert's algorithm are shown. The improved constant is $\zeta(\alpha)$. The improvement is approximately an order of magnitude.

The constants $\zeta(\alpha), \zeta'(\alpha)$ are plotted in Figure 36. The results of this chapter only yield improvements of one order of magnitude or less for moderate values of α.

In Figure 37, the separate grading constants with respect to midpoints and circumcenters, *i.e.*, $\zeta_m(\alpha)$ and $\zeta_t(\alpha)$, are plotted. For α between 8° and 16°, note that $\zeta_m(\alpha)$ takes values less than 1000. Considering that this optimality constant relies on a number of worst-case assumptions, a value as small as 1000 is nearly practical.

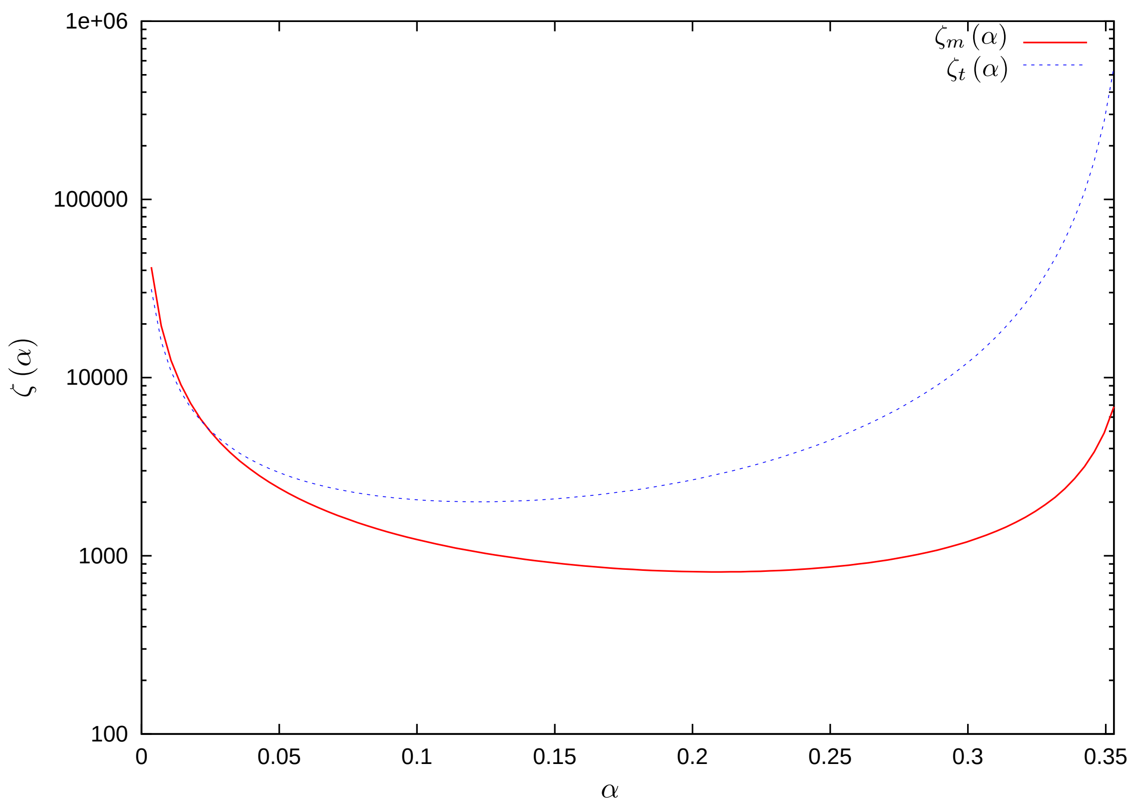

Figure 37: The two constants associated with the optimality of the Steiner circumcenters, and Steiner midpoints, respectively $\zeta_t(\alpha)$ and $\zeta_m(\alpha)$ are plotted versus α, for Ruppert's Algorithm. Note that $\zeta_m(\alpha)$ takes values as small as 1000 for some α.

REFERENCES

[1] Th. Apel, M. Berzins, P. K. Jimack, G. Kunert, A. Plaks, I. Tsukerman and M. Walkley. Mesh shape and anisotropic elements: theory and practice. In *The mathematics of finite elements and applications, X, MAFELAP 1999 (Uxbridge)*, pages 367–376. Elsevier, Oxford, 2000. URL `citeseer.nj.nec.com/373015.html`.

[2] I. Babuška and A. K. Aziz. On the angle condition in the finite element method. *SIAM Journal on Numerical Analysis*, 13(2):214–226, 1976.

[3] Brenda S. Baker, Eric Grosse and Conor S. Rafferty. Nonobtuse triangulation of polygons. *Discrete Comput. Geom.*, 3(2):147–168, 1988. ISSN 0179-5376.

[4] Eric B. Becker, Graham F. Carey and J. Tinsley Oden. *Finite elements. Vol. I*. The Texas Finite Element Series, I. Prentice Hall Inc., Englewood Cliffs, NJ, 1981. ISBN 0-13-317057-8. An introduction.

[5] M. Bern, S. Mitchell and J. Ruppert. Linear-size nonobtuse triangulation of polygons. *Discrete Comput. Geom.*, 14(4):411–428, 1995. ISSN 0179-5376. ACM Symposium on Computational Geometry (Stony Brook, NY, 1994).

[6] Marshall Bern and David Eppstein. Mesh generation and optimal triangulation. In Ding Zhu Du and Frank Hwang, editors, *Computing in Euclidean geometry*, volume 1 of *Lecture Notes Series on Computing*, pages 23–90. World Scientific Publishing Co. Inc., River Edge, NJ, 1992. ISBN 981-02-0966-5.

[7] Marshall Bern, David Eppstein and John Gilbert. Provably good mesh generation. *J. Comput. System Sci.*, 48(3):384–409, 1994. ISSN 0022-0000. 31st Annual Symposium on Foundations of Computer Science (FOCS) (St. Louis, MO, 1990).

[8] Charles Boivin and Carl F. Ollivier-Gooch. Guaranteed-quality triangular mesh generation for domains with curved boundaries. *International Journal for Numerical Methods in Engineering*, 55(10):1185–1213, 2002.

[9] A. Bowyer. Computing Dirichlet tessellations. *Comput. J.*, 24(2):162–166, 1981. ISSN 0010-4620.

[10] Siu-Wing Cheng, Tamal K. Dey, Herbert Edelsbrunner, Michael A. Facello and Shang-Hua Teng. Sliver exudation. *J. ACM*, 47(5):883–904, 2000. ISSN 0004-5411.

[11] L. Paul Chew. Constrained Delaunay triangulations. *Algorithmica*, 4(1):97–108, 1989. ISSN 0178-4617. Computational geometry (Waterloo, ON, 1987).

[12] L. Paul Chew. Guaranteed-quality triangular meshes. *CS 89-983*, Computer Science Department, Cornell University, 1989.

[13] L. Paul Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proceedings of the Ninth Annual Symposium on Computational Geometry (San Diego, California, 1993)*, pages 274–280. ACM, New York, 1993.

[14] H. S. M. Coxeter and S. L. Greitzer. *Geometry Revisited*. The Mathematical Association of America, Washington, DC, 1967. ISBN 0-88385-600-X.

[15] Mark de Berg, Marc van Kreveld, Mark Overmars and Otfried Schwarzkopf. *Computational geometry*. Springer-Verlag, Berlin, revised edition, 2000. ISBN 3-540-65620-0. Algorithms and applications.

[16] Tamal K. Dey, Kōkichi Sugihara and Chandrajit L. Bajaj. Delaunay triangulations in three dimensions with finite precision arithmetic. *Comput. Aided Geom. Design*, 9(6):457–470, 1992. ISSN 0167-8396.

[17] H. Edelsbrunner, X. Li, Miller, G. Stathopoulos, D. A. Talmor, S. Teng, A Ungor and N. J. Walkington. Smoothing and cleaning up slivers. In *ACM Symposium on Theory of Computing*, pages 273–277. 2000. URL `citeseer.nj.nec.com/edelsbrunner00smoothing.html`.

[18] Herbert Edelsbrunner. *Geometry and topology for mesh generation*, volume 7 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2001. ISBN 0-521-79309-2.

[19] Herbert Edelsbrunner and Ernst Peter Mücke. Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. In *Proceedings of the Fourth Annual Symposium on Computational Geometry (Urbana, IL, 1988)*, pages 118–133. ACM, New York, 1988.

[20] Herbert Edelsbrunner and Tiow Seng Tan. An upper bound for conforming Delaunay triangulations. *Discrete Comput. Geom.*, 10(2):197–213, 1993. ISSN 0179-5376.

[21] David A. Field. Qualitative measures for initial meshes. *International Journal For Numerical Methods In Engineering*, 47:887–906, 2000.

[22] Steven Fortune. A sweepline algorithm for Voronoĭ diagrams. *Algorithmica*, 2(2):153–174, 1987. ISSN 0178-4617.

[23] Stephen Guattery, Gary L. Miller and Noel Walkington. Estimating interpolation error: a combinatorial approach. In *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 1999)*, pages 406–413. ACM, New York, 1999.

[24] Claes Johnson. *Numerical solution of partial differential equations by the finite element method*. Cambridge University Press, Cambridge, 1987. ISBN 0-521-34514-6; 0-521-34758-0.

[25] Clark Kimberling. Triangle centers and central triangles. *Congr. Numer.*, 129:xxvi+295, 1998. ISSN 0384-9864.

[26] D. T. Lee and A. K. Lin. Generalized Delaunay triangulation for planar graphs. *Discrete Comput. Geom.*, 1(3):201–217, 1986. ISSN 0179-5376.

[27] G. L. Miller, D. Talmor, S. Teng and N. J. Walkington. A Delaunay based numerical method for three dimensions: generation, formulation and partition. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*, pages 683–692. ACM Press, 1995.

[28] Gary L. Miller. A timing analysis of a Delaunay refinement mesh generation algorithm, December 2002. In preparation.

[29] Gary L. Miller, Steven E. Pav and Noel J. Walkington. Fully incremental 3d Delaunay mesh generation. In *Proceedings of the 11th International Meshing Roundtable*, pages 75–86. Sandia National Laboratory, September 2002.

[30] Gary L. Miller, Steven E. Pav and Noel J. Walkington. An incremental Delaunay meshing algorithm. *Technical Report 02-CNA-011*, Center for Nonlinear Analysis, Carnegie Mellon University, 2002. URL `http://www.math.cmu.edu/cna`.

[31] Scott A. Mitchell. Cardinality bounds for triangulations with bounded minimum angle. In *Sixth Canadian Conference on Computational Geometry*, pages 326–331. 1994.

[32] Scott A. Mitchell and Stephen A. Vavasis. Quality mesh generation in higher dimensions. *SIAM J. Comput.*, 29(4):1334–1370 (electronic), 2000. ISSN 1095-7111.

[33] David M. Mount and Alan Saalfeld. Globally-equiangular triangulations of co-circular points in $O(n \log n)$ time. In *Proceedings of the Fourth Annual Symposium on Computational Geometry (Urbana, IL, 1988)*, pages 143–152. ACM, New York, 1988.

[34] Ernst P. Mücke. A robust implementation for three-dimensional Delaunay triangulations. *Internat. J. Comput. Geom. Appl.*, 8(2):255–276, 1998. ISSN 0218-1959.

[35] Aleksandar Nanevski, Guy Blelloch and Robert Harper. Automatic generation of staged geometric predicates. In *International Conference on Functional Programming*, pages 217–228. Florence, Italy, September 2001.

[36] J. T. Oden and J. N. Reddy. *An introduction to the mathematical theory of finite elements*. Wiley-Interscience [John Wiley & Sons], New York, 1976. Pure and Applied Mathematics.

[37] Atsuyuki Okabe, Barry Boots, Kokichi Sugihara and Sung Nok Chiu. *Spatial tessellations: concepts and applications of Voronoi diagrams*. John Wiley & Sons Ltd., Chichester, second edition, 2000. ISBN 0-471-98635-6. With a foreword by D. G. Kendall.

[38] Carl F. Ollivier-Gooch and Charles Boivin. Guaranteed-quality simplicial mesh generation with cell size and grading control. *Engineering with Computers*, 17(3):269–286, 2001.

[39] Jim Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18(3):548–585, 1995. ISSN 0196-6774. Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Austin, TX, 1993).

[40] Michael Ian Shamos and Dan Hoey. Closest-point problems. In *16th Annual Symposium on Foundations of Computer Science (Berkeley, Calif., 1975)*, pages 151–162. IEEE Computer Society, Long Beach, Calif., 1975.

[41] Jonathan Richard Shewchuk. *Delaunay Refinement Mesh Generation*. Ph.D. thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 1997. Available as Technical Report CMU-CS-97-137.

[42] Jonathan Richard Shewchuk. Constrained Delaunay tetrahedralizations and probably good boundary recovery. In *Proceedings of the Eleventh International Meshing Roundtable (Ithaca, New York)*, pages 193–204. Sandia National Labs, September 2002.

[43] Jonathan Richard Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Comput. Geom.*, 22(1-3):21–74, 2002. ISSN 0925-7721. 16th ACM Symposium on Computational Geometry (Hong Kong, 2000).

[44] Jonathan Richard Shewchuk. What is a good linear element? interpolation, conditioning, and quality measures. In *Proceedings of the Eleventh International Meshing Roundtable (Ithaca, New York)*, pages 115–126. Sandia National Labs, September 2002.

[45] K. Sugihara, M. Iri, H. Inagaki and T. Imai. Topology-oriented implementation—an approach to robust geometric algorithms. *Algorithmica*, 27(1):5–20, 2000. ISSN 0178-4617. Implementation of geometric algorithms.

[46] Monique Teillaud. *Towards dynamic randomized algorithms in computational geometry*. Springer-Verlag, Berlin, 1993. ISBN 3-540-57503-0.

[47] D. F. Watson. Computing the n-dimensional Delaunay tessellation with application to Voronoĭ polytopes. *Comput. J.*, 24(2):167–172, 1981. ISSN 0010-4620.

[48] Chee-Keng Yap. A geometric consistency theorem for a symbolic perturbation scheme. In *Proceedings of the Fourth Annual Symposium on Computational Geometry (Urbana, IL, 1988)*, pages 134–142. ACM, New York, 1988.