

New Bounds on the Size of Optimal Meshes

Donald R. Sheehy[†]

INRIA Saclay

Abstract

The theory of optimal size meshes gives a method for analyzing the output size (number of simplices) of a Delaunay refinement mesh in terms of the integral of a sizing function over the input domain. The input points define a maximal such sizing function called the feature size. This paper presents a way to bound the feature size integral in terms of an easy to compute property of a suitable ordering of the point set. The key idea is to consider the pacing of an ordered point set, a measure of the rate of change in the feature size as points are added one at a time. In previous work, Miller et al. showed that if an ordered point set has pacing ϕ , then the number of vertices in an optimal mesh will be $O(\phi^d n)$, where d is the input dimension. We give a new analysis of this integral showing that the output size is only $\Theta(n + n \log \phi)$. The new analysis tightens bounds from several previous results and provides matching lower bounds. Moreover, it precisely characterizes inputs that yield outputs of size $O(n)$.

Categories and Subject Descriptors (according to ACM CCS): F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—Geometrical problems and computations

1. Introduction

Mesh generation is the process of decomposing a domain into simple geometric elements (simplices). It is fundamental to applications in physical simulation, graphics, and data analysis (see the recent survey by Shewchuk [She11b]). In a common variant of the problem, the input is a set of n points $P \subset \mathbb{R}^d$ and a bounding domain Ω containing P and the output is a simplicial complex decomposing Ω such that the points of P are among the vertices. Throughout, the *size* or *complexity* of a mesh is the number of simplices in the output.

The *quality* of a mesh is measured by the shapes of the simplices. Definitions of quality vary by application, but one commonly used criterion is that all simplices have an upper bound on the circumradius to shortest edge ratio (see Figure 1) [MTTW95]. In such cases, we say that the output mesh has bounded radius-edge. The Delaunay refinement paradigm of meshing adds extra points called Steiner points to a set of input points so that all of the simplices of the Delaunay triangulation have bounded radius-edge. The dual

view of Delaunay refinement is Voronoi refinement where quality is measured by the aspect of the Voronoi cells (see Figure 2 for examples and the precise definition). Figure 3 shows an example of an input, its Delaunay triangulation, and the quality output mesh.

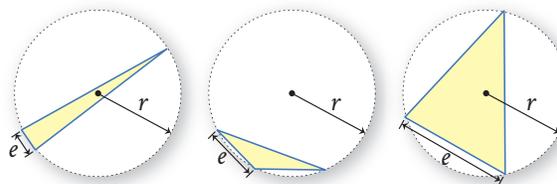


Figure 1: Examples of circumradius to shortest edge ratios.

One consequence of this definition of quality is that the number of vertices and the number of simplices in the output are the same up to constants, i.e. each vertex participates in only a constant number of simplices [MTTW95]. So, bounding the asymptotic mesh size is the same as bounding the number of vertices in the output.

A critical challenge in mesh generation is to balance the competing demands of having both small mesh complexity and high quality elements. To achieve this balance, it is often

[†] This work was partially supported by the National Science Foundation under grant number CCF-1065106, by GIGA grant ANR-09-BLAN-0331-01, and by the European project CG-Learning No. 255827.

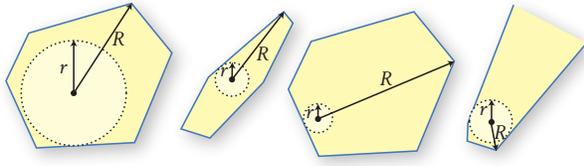


Figure 2: The Voronoi aspect ratio is R/r where R is the distance from the vertex to the farthest vertex of the cell and r is the radius of the largest contained ball centered at the vertex.

required that the output mesh is *graded*, with smaller simplices in dense areas gradually growing into larger simplices in sparse regions. This grading can be seen in Figure 3 with smaller simplices near the curve and larger simplices around the outside. *Optimal meshes* are those that have asymptotically minimum size among all quality meshes of a given input. Delaunay refinement naturally produces graded meshes of asymptotically optimal size up to constants of the form $2^{O(d)}$. The main result of this paper gives new bounds on the size of such optimal meshes. In particular, we show when these meshes have linear size.

1.1. Optimal Size Meshes

A celebrated result of Ruppert [Rup95] showed how to construct quality Delaunay meshes in the plane with an asymptotically optimal number of vertices. The key idea in Ruppert’s analysis is the **local feature size** $\mathbf{f}_P : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ induced by the input set P ; it is defined as

$$\mathbf{f}_P(x) := \min\{r : |P \cap \text{ball}(x, r)| \geq 2\}.$$

Equivalently, \mathbf{f}_P measures the distance to the second-nearest neighbor in P . Note that \mathbf{f}_P is 1-Lipschitz: $\mathbf{f}_P(a) \leq \mathbf{f}_P(b) + \|a - b\|$ for any pair of points $a, b \in \mathbb{R}^d$. We have given the definition of \mathbf{f}_P for point sets, but Ruppert’s work was also concerned with more general inputs including features like edges.

The straightforward generalization of Ruppert’s work to \mathbb{R}^d says that the number of vertices in an optimal mesh of a point set P in a domain $\Omega \subset \mathbb{R}^d$ is bounded (up to constant factors) by the **feature size measure**:

$$\mu_P(\Omega) := \frac{1}{\mathbb{V}_d} \int_{\Omega} \frac{1}{\mathbf{f}_P(x)^d} d\lambda(x),$$

where \mathbb{V}_d denotes the volume of the unit d -ball, and λ is the volume measure on \mathbb{R}^d . Because of the form of the definition, the feature size measure is sometimes called the *feature size integral*. The constant factors in the Ruppert bound may also depend on the desired quality of the output.

The precise constants in the Ruppert bounds depend on details of the particular meshing algorithm used. For example, different refinement schemes prioritize some refine-

ments over others, which can have an impact on the output size. Moreover, the constants depend on the desired output radius-edge ratio. Lastly, the relationship between the feature size measure and the output of a meshing algorithm depends on the *fatness* of the domain to be meshed; a long skinny convex set can have non-negligible measure and yet may contain no vertices of a quality mesh. Thus, some fatness condition on the domain is necessary as will be seen in the Theorem statements below. It is an active area of research to find heuristics that improve these constants in both theory and practice, but these efforts only affect the constant factors in asymptotically tight bounds. This paper deals considers the most general case, analyzing the feature size measure directly rather than attempting to treat only a specific algorithm.

The Ruppert lower bounds are quite general. Any mesh of P in Ω with bounded radius-edge will have size at least $\mu_P(\Omega)$ up to constants. It applies equally to meshing algorithms applying heuristics or post-processes to smooth and improve mesh quality and reduce size. These include variational methods [ACSYD05], Laplacian smoothing, off-center methods [Üng09] and Hodge optimization [MMdGD11] among many others. Again, these methods may improve the constants, but the asymptotic growth is governed by the Ruppert bound.

The Ruppert approach presented here appears in different forms in other work. Other sizing functions were considered in Alliez et al. [ADA07]. Pav also gave anisotropic bounds on mesh size in terms of the smallest angles allowed that relied on integrating a density induced by the sizing field over the domain [Pav04].

1.2. The Search for Meaningful Bounds

The tight bounds on mesh sizing from the Ruppert bounds are nice, but they are not very informative because they depend on the feature size integral rather than n , the number of input points. They do not, for example, tell when the output size will be $O(n^2)$ or $O(n)$ or any other tidy function of n . Moreover, they do not give an easy way to evaluate the amortized change in the output size as a result of adding a single new point.

There have been several previous attempts to give more meaningful bounds on the size of output meshes. In their work on quadtree meshes, Bern et al. [BEG94] prove that the output size is $O(n \log(\text{aspect}(\text{Del}(P))))$, where $\text{aspect}(\text{Del}(P))$ is the maximum aspect ratio among all simplices in the Delaunay triangulation of P . A similar bound comes from considering the gap ratio introduced by Talmor [Tal97]. In both cases, the upper bounds are $O(n \log \Delta_P)$, where Δ_P is the spread of P defined as the ratio of largest to smallest interpoint distances.

There is also a local version of these bounds based on the spread. It can be shown using similar packing methods that

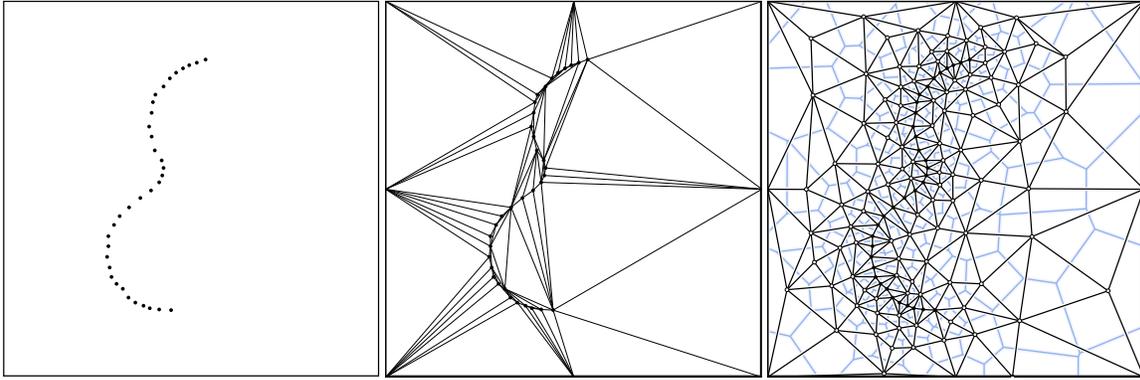


Figure 3: An example of an input point set (left), its Delaunay triangulation (center), and the output mesh with its Voronoi diagram (right).

the optimal output size is at most

$$\sum_{p \in P} \log(\text{aspect}(\text{Vor}(p))).$$

As we will show in Section 2, these bounds can be quite loose.

Later optimal meshing algorithms such as those by Mitchell and Vavasis [MV00] and by Cheng and Dey [CD04] used volume packing arguments similar to those employed by Ruppert. However, in these cases, the algorithms consider higher-dimensional features and thus the optimality proofs must be intertwined with the removal of badly shaped simplices called *slivers*. For point set inputs, existing sliver removal methods such as those by Li and Teng [Li03] or by Cheng et al. [CDE*00] imply that the size of the optimal mesh with slivers is within a constant factor of the size of the optimal mesh without slivers. So, we don't need to consider slivers in our analysis.

The immediate ancestor of the current paper is the work of Miller et al. [MPS08] that introduced the notion of pacing. That paper gave the first bounds on optimal mesh size that could recognize a linear size mesh in the presence of bad aspect ratio input simplices. However, the upper bound was far from tight; it could easily be off by a factor of n . Moreover, the bounds were dependent on a particular (best) permutation of the input points unlike Theorem 1.1 (below), which applies to any permutation starting with the farthest pair.

1.3. Contribution

In this paper, we give a new analysis of the feature size measure that provides tight upper and lower bounds. Our analysis makes it clear exactly when an input will yield an optimal mesh of $O(n)$ size. Moreover, the analysis, gives a tight bound on the influence of a single new point, which may

have implications for future algorithms, particularly in dynamic meshing.

We need a few definitions in order to state the main result. For an ordered point set $P = \{p_1, \dots, p_n\}$, the i th **prefix** of the ordering is defined as $P_i = \{p_1, \dots, p_i\}$. For any point p_i in the ordering ($i \geq 3$), the **pacing** is defined as the ratio ϕ_i of the feature sizes at p_i induced by P_{i-1} and P_i :

$$\phi_i := \frac{f_{P_{i-1}(p_i)}}{f_{P_i(p_i)}}.$$

Equivalently, ϕ_i is the ratio of distances to the second and first nearest neighbors of p_i among its predecessors. The definition implies that $\phi_i \geq 1$ for all i . Define $\phi_1 := \phi_2 := 1$ by convention. The **pacing of the ordering** is the geometric mean of the ϕ_i s and is denoted by $\phi_P := (\prod_{i=1}^n \phi_i)^{1/n}$. The more useful form of this definition is

$$n \log \phi_P = \sum_{i=1}^n \log \phi_i.$$

Theorem 1.1 Let $P = \{p_1, \dots, p_n\}$ be an ordered set of n points in \mathbb{R}^d such that $r = \|p_1 - p_2\|$ is the largest pairwise distance. Let Ω be a subset of \mathbb{R}^d such that $\text{ball}(p_1, 2r) \subseteq \Omega \subseteq \text{ball}(p_1, cr)$, for some constant $c \geq 2$. Then,

$$\mu_P(\Omega) = \Theta(n + n \log \phi_P).$$

In previous work, it was shown that the feature size measure is at most $O(n \hat{\phi}_P^d)$ [MPS08], where $\hat{\phi}_P = \max_i \phi_i$ is a possibly larger notion of pacing. Theorem 1.1 implies a bound of $O(n + n \log \phi_P)$. This tightens the previous results using this method, moving them from theoretically novel to practically useful. Also, as a side effect, it eliminates the need to optimize over permutations in order to express the bound.

The proof of Theorem 1.1 will be broken up into two parts: the upper bound in Theorem 5.3 and the lower bound in Theorems 6.2 and 6.3. There are two theorems for the

lower bound because the $\Omega(n)$ and $\Omega(n \log \phi_P)$ lower bounds are proven separately to handle the case when $\log \phi_P < 1$.

The upper and lower bounds in Theorem 1.1 only assume that the order on the points starts with the farthest pair. Consequently, the same bound holds asymptotically for any such ordering. Section 2.1 gives an example where this condition or a similar one is necessary. The result is stated for ordered point sets because the pacing is a function of the ordering and is not well defined otherwise.

Using the pacing instead of the spread gives a sharper bound, because for all point sets P and all orderings of P , $\Delta_P \geq \phi_P$. Moreover, for simple examples such as the one described below in Section 2.2, the spread can be exponentially larger than the pacing.

Our presentation makes no assumption about the dimension of the inputs. Although most meshing applications deal with 2- or 3-dimensional domains, recent work has shown that higher dimensional meshes are ideal for geometric and topological inference as they provide a nice basis for a space of smooth functions graded according to the density of the input points [HMOS10, She11a]. Functions such as the distance function to the input or the distance to the empirical measure can be used to recover the homology of the underlying space from which the input was sampled [CL06, CCSM11].

2. Some Difficult Inputs for Mesh Generation

In this section, we present two simple examples of input point sets that have caused problems in guaranteed mesh generation. The first example is just two points close together in the middle of a big empty annulus with a sparse bounding box around it. It is a canonical configuration of points leading to superlinear output size in that any superlinear output mesh contains a similar configuration of input points up to scaling and rotation. This bad example inspired many of the looser mesh size analyses described in Section 1.2 such as those based on the spread or the aspect ratio of the input Delaunay triangulation. Such bounds are tight on this bad example.

The second bad example we present is the foil to previous efforts to give a meaningful bound on the output mesh size in terms of n and a simple geometric parameter of the input. The example is a set of exponentially spaced points on a line. As we show below, any of several previous bounds give $O(n^2)$ bounds whereas the true output size is only $O(n)$. The linear output size is easily computed from Theorem 1.1.

2.1. The Two Point Example

Let P contain two points of distance 1 in the center of a bounding box of sidelength Δ composed of 2^d vertices arranged on the corners of a d -cube. This case is illustrated in Figure 4. From the picture, it is clear that the grading of

the refined mesh causes it to have only a constant number of Steiner points in each of the geometrically growing annuli around the points in the center. There are only $O(\log \Delta)$ such annuli and so the total output size is $O(n \log \Delta)$.

This example is simple enough that it is also possible to derive the $O(n \log \Delta)$ bound by evaluating the integral in the definition of the feature size measure. To turn this constant size example into an asymptotic lower bound, it suffices to pack $O(n)$ copies of the bounding box into a box of sidelength $O(n^{1/d} \Delta)$.

The two point example also demonstrates the reason for farthest pair condition on the ordering of the input points. If the two close points in the center are allowed to be first in the permutation, the pacing can be constant. Then, a bound of $O(n \log \phi_P)$ would predict a linear size output, which we have already shown is not the case.

2.2. The Exponentially Spaced Example

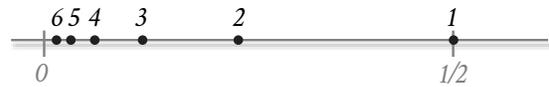


Figure 5: Exponentially spaced points on a line.

The following example will show that previous approaches at giving meaningful optimality bounds are not tight on all instances.

Consider the set of points $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^2$ where $p_i := (2^{-i}, 0)$. Moreover, such a point set only requires $O(\log n)$ bits of floating point precision to represent, so it is reasonable. The spread Δ of P is $2^{\Theta(n)}$. Thus, $n \log \Delta = \Theta(n^2)$. This is not fixed by taking a more local approach. For example, the aspect ratio of the Voronoi cell of p_i is also $2^{\Theta(i)}$, assuming a bounding domain with constant radius. Thus,

$$\sum_{i=1}^n \log(\text{aspect}(\text{Vor}(p_i))) = \sum_{i=1}^n \Theta(i) = \Theta(n^2).$$

A similar result holds if one looks instead at the aspect ratios or radius-edge ratios of the Delaunay triangles. In all cases, the previous methods predict a quadratic size mesh. However, moving p_n to the beginning yields a permutation such that the pacing of p_i is less than 2 for each $i = 1, \dots, n-1$. So, the pacing of this permutation is $\Theta(1)$. Thus, Theorem 1.1 implies that the actual mesh size will only be $\Theta(n)$.

3. The History and Usefulness of Pacing Analysis

We refer to the bounding of the feature size measure in terms of the pacing of an ordering as *pacing analysis*. Pacing anal-

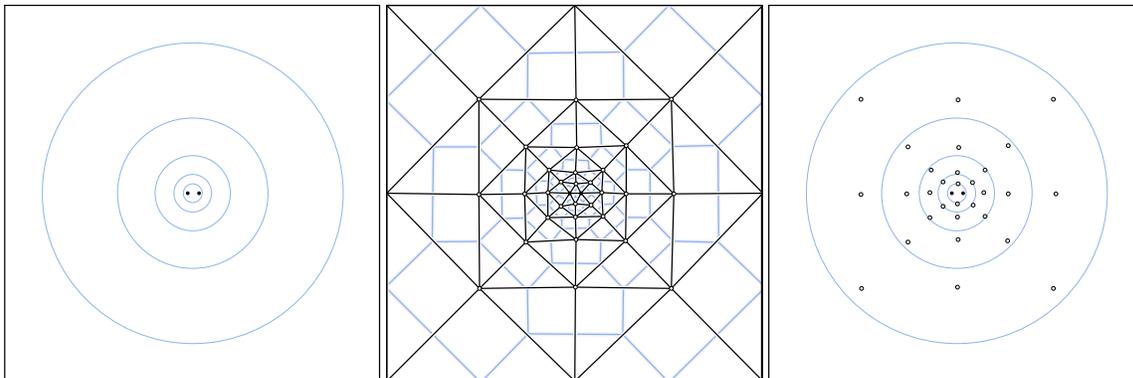


Figure 4: *The Two Point Example (left) and its resulting mesh after refinement (center) including its Voronoi diagram. The refined mesh only has a constant number of vertices in each of the geometrically growing annuli (right).*

ysis has proven useful both in the analysis of meshing algorithms and also in guiding the development of new algorithms. We will give some history of this idea, its origins, and its uses.

Pacing analysis was developed out of a desire to generalize certain properties of quadtree-based meshes to Delaunay meshes. Here, we use “quadtree” as a general term for quadtrees, octtrees, and their higher dimensional analogues. A quadtree may be constructed from a point set by iteratively splitting every box that contains more than one input point into 2^d equal boxes. To *balance* the quadtree, more boxes are split so that no two adjacent boxes have dramatically different sizes. It was observed early on that $O(n)$ new splits sufficed to balance a quadtree with n boxes. This is what is meant by the so-called “linear cost of balancing a quadtree” (see [Moo95] for an extensive treatment of quadtree balancing analysis though the idea is older than that paper).

Balancing quadtrees plays the same role as refinement in Delaunay refinement; it fills out the ambient space so that sizes change gradually, i.e. so that the output is graded. The original motivation for pacing analysis was the search for an analogue of the linear cost of balancing a quadtree that applies to Delaunay refinement. This would require a class of point sets to stand in analogy with the unbalanced quadtrees so that the output of Delaunay refinement on any such point set has only a linear number of vertices. In light of the Rupert bounds, this amounts to a family of n -point sets P such that

$$\mu_P(\Omega) = O(n).$$

There are obvious point sets that satisfy such a criterion. For example, the output of any size-optimal meshing algorithm works, but this is too restrictive to be useful. A much richer class of point sets is generated by the following technique. Starting with any two points, iteratively add a point on the Voronoi diagram of the current point set. Repeat. That is,

each new point is equidistant from its two nearest neighbors. The only requirement is that the first two points must remain the farthest pair throughout. In particular, the vertices of a quadtree may be generated this way. This corresponds to a point set with pacing $\phi_P = 1$.

To generalize this idea further, Miller et al. showed that as long as the distance to the first and second nearest neighbors of each new point differed by at most a constant factor ϕ , then the output size would be $O(\phi^d n)$ [MPS08]. This was used to generalize an idea from Bern, Eppstein and Gilbert [BEG94], who showed that quadtree methods could be used to produce linear size meshes in any dimension by relaxing the usual balancing and aspect ratio conditions. This relaxation was nearly equivalent to the notion of a compressed quadtree, a variant that permits large changes of scale. In this light, the Miller et al. paper showed that compressed quadtrees also have a natural analogue in Delaunay refinement.

The next major use of pacing was in the proof of the Scaffold Theorem of Hudson et al. [HMPS09]. This theorem gives simple conditions on the a sample of points from a surface, under which the surface mesh and the volume mesh will have the same size up to constant factors. This was essential for showing that existing output-sensitive optimal-size meshing algorithms remained output sensitive even if they preprocessed the input by placing it in a bounding box and then later eliminated the extra simplices generated around the outside. This is a very common procedure as it is difficult to identify the inside and outside of a surface before actually building the mesh. The main open problem presented in the conclusion of the original paper on the Scaffold Theorem is whether necessary and sufficient conditions could be found for bounding mesh sizes in terms of the number of vertices. This question is settled by the tight bounds of Theorem 1.1.

Miller et al. extended the linear-size Delaunay meshes of

their previous work to *hierarchical quality meshes* [MPS11]. They showed that these meshes can be constructed with $O(n)$ size in $O(n \log n)$ time. Pacing analysis plays two vital roles in the analysis, both to bound the size of the mesh and also in the analysis of the point location data structure. In fact, the algorithm explicitly computes the pacing of each newly inserted point and behaves differently for well-paced insertions than for others.

In the current work, we present an asymptotically tight version of pacing analysis. This implicitly tightens bounds in each of the preceding results. We also achieve mesh size bounds from pacing analysis in a way that does not depend on the specific ordering given. This means that the change in the final mesh size induced by the insertion of any single point can be approximated up to a constant factor as long as that points is not one of the farthest pair.

4. Basic Definitions

A **measure** μ on \mathbb{R}^d is a function from Σ , the set of (Borel) subsets of \mathbb{R}^d , to \mathbb{R} that is

1. **nonnegative:** $\mu(A) \geq 0$ for all sets $A \in \Sigma$,
2. **countably additive:** $\mu(\bigcup_{i \in \mathbb{N}} A_i) = \sum_{i \in \mathbb{N}} \mu(A_i)$ for any countable collection of disjoint sets $A_i \in \Sigma$, and
3. **has null empty set:** $\mu(\emptyset) = 0$.

An immediate consequence of the nonnegativity and countable additivity properties is that measures are **subadditive** with respect to unions, i.e. $\mu(A \cup B) \leq \mu(A) + \mu(B)$.

Throughout, we use λ to denote the standard volume measure on \mathbb{R}^d . We use σ to denote the surface area measure over the unit $(d-1)$ -sphere. So, letting $B = \text{ball}(c, r)$ with boundary $(d-1)$ -sphere S , $\lambda(B) = \int_{x \in B} d\lambda(x) = r^d \nabla_d$ and $\sigma(S) = dr^{d-1} \nabla_d$.

Given two measures μ and ν , we say that μ is a **submeasure** of ν , denoted $\mu \preceq \nu$, if $\mu(A) \leq \nu(A)$ for all $A \in \Sigma$. We do some basic arithmetic on measures. For example, the sum of two measures is a measure defined as $(\mu + \nu)(A) = \mu(A) + \nu(A)$. If $\mu \preceq \nu$ then $\nu - \mu$ is a measure, where $(\nu - \mu)(A) = \nu(A) - \mu(A)$.

The feature size measure is defined with respect to a density. That is, for a fixed point set P , we can define the mesh density $f = \frac{1}{\nabla_d \mathbf{f}_P}$. This allows us to write $\mu_P = f \cdot \lambda$. In this case, we say that μ_P has density f with respect to the volume measure λ .

For point sets $P \subseteq Q$, $\mathbf{f}_P \geq \mathbf{f}_Q$, because the distance to two points in Q is at most the distance to two points P . This implies that $\mu_P \preceq \mu_Q$. This defines a kind of monotonicity on the feature size measures induced by increasing point sets. Moreover, $P \subseteq Q$ implies that $\mu_Q - \mu_P$ is a measure.

5. Upper Bound

The proof of the upper bound on the feature size measure will follow a simple pattern. First, we prove a bound for inputs consisting of only two points (Lemma 5.1). Then, we bound the change in the measure upon adding a single new point (Lemma 5.2). Finally, we apply this Lemma inductively to get the final bound (Theorem 5.3).

Lemma 5.1 (Just two points) If $P = \{p, q\}$ and $\Omega \subseteq B = \text{ball}(p, \frac{c}{2} \|p - q\|)$ for some constant $c > 4$, then

$$\mu_P(\Omega) \leq 1 + d \ln c.$$

Proof The assumption that $\Omega \subseteq B$ implies $\mu_P(\Omega) \leq \mu_P(B)$. Rewriting $\mu_P(B)$ in polar coordinates yields

$$\mu_P(B) = \frac{1}{\nabla_d} \int_B \frac{d\lambda(x)}{\mathbf{f}_P(x)^d} = \frac{1}{\nabla_d} \int_{s \in S_{d-1}} \int_0^{\frac{c}{2} \|p-q\|} \frac{r^{d-1}}{\mathbf{f}_P(rs)^d} dr d\sigma(s),$$

where S_{d-1} denotes the unit $(d-1)$ -sphere bounding $\text{ball}(0, 1)$, rs denotes the point $s \in \mathbb{R}^d$ with all coordinates scaled by a factor r , and we assume without loss of generality that p is the origin.

For all $x \in \mathbb{R}^d$, $\mathbf{f}_P(x) \geq \max\{\frac{1}{2} \|p - q\|, \|x - p\|\}$. So,

$$\begin{aligned} \mu_P(\Omega) &\leq \frac{1}{\nabla_d} \int_{s \in S_{d-1}} \int_0^{\frac{c}{2} \|p-q\|} \frac{r^{d-1}}{(\max\{\frac{1}{2} \|p - q\|, r\})^d} dr d\sigma(s) \\ &= d \left(\int_0^{\frac{1}{2} \|p-q\|} \frac{r^{d-1} dr}{(\frac{1}{2} \|p - q\|)^d} + \int_{\frac{1}{2} \|p-q\|}^{\frac{c}{2} \|p-q\|} \frac{r^{d-1} dr}{r^d} \right) \\ &= 1 + d \ln c. \end{aligned}$$

□

We now bound the change in the feature size measure induced by the addition of a single point.

Lemma 5.2 (One more point upper bound) Let $P \subset \mathbb{R}^d$ be a finite point set and let $P' = P \cup \{q\}$ for some $q \in \mathbb{R}^d \setminus P$. Let $\phi = \frac{\mathbf{f}_P(q)}{\mathbf{f}_{P'}(q)}$. Then, for any domain $\Omega \subseteq \mathbb{R}^d$,

$$\mu_{P'}(\Omega) - \mu_P(\Omega) \leq 1 + d \ln(3d\phi).$$

Proof Since $P \subseteq P'$, $\nu := \mu_{P'} - \mu_P$ is a measure and thus it will suffice to prove the statement for $\Omega = \mathbb{R}^d$, i.e. that

$$\nu(\mathbb{R}^d) \leq 1 + d \ln(3d\phi).$$

Let U be the subset of \mathbb{R}^d where $\mathbf{f}_P \neq \mathbf{f}_{P'}$. It follows that

$$\nu(\mathbb{R}^d \setminus U) = 0. \tag{1}$$

Let $R := \mathbf{f}_{P'}(q)$; this is the distance from q to the nearest point of P . The definitions of R and ϕ imply that $R\phi = \mathbf{f}_P(q)$. For all points x in the ball $B = \text{ball}(q, \frac{R}{2})$, $\mathbf{f}_{P'}(x) \geq \frac{R}{2}$, so

$$\nu(B) = \frac{1}{\nabla_d} \int_B \left(\frac{1}{\mathbf{f}_{P'}(x)^d} - \frac{1}{\mathbf{f}_P(x)^d} \right) d\lambda(x)$$

$$\leq \frac{1}{\sqrt[d]{d}} \int_B \left(\frac{2}{R}\right)^d d\lambda(x) = 1. \quad (2)$$

For any $x \in U$,

$$\mathbf{f}_P(x) \leq \|x - q\| + R\phi, \quad (3)$$

because \mathbf{f}_P is 1-Lipschitz. Moreover,

$$\mathbf{f}_{P'}(x) \geq \|x - q\|, \quad (4)$$

because q must be one of the two nearest neighbors of x if $\mathbf{f}_P(x) \neq \mathbf{f}_{P'}(x)$. We apply the bounds in (3) and (4) as follows.

$$\begin{aligned} v(U \setminus B) &= \frac{1}{\sqrt[d]{d}} \int_{U \setminus B} \left(\frac{1}{\mathbf{f}_{P'}(x)^d} - \frac{1}{\mathbf{f}_P(x)^d} \right) d\lambda(x) \\ &\leq \frac{1}{\sqrt[d]{d}} \int_{U \setminus B} \left(\frac{1}{\|x - q\|^d} - \frac{1}{(\|x - q\| + R\phi)^d} \right) d\lambda(x) \\ &\leq d \int_{R/2}^{\infty} \left(\frac{1}{r^d} - \frac{1}{(r + R\phi)^d} \right) r^{d-1} dr \\ &< d \ln(3d\phi). \end{aligned} \quad (5)$$

We have extended the integral over all of $\mathbb{R}^d \setminus B$ (the function is nonnegative) and rewrote it in polar coordinates with origin at q . The final inequality follows from a straightforward calculus exercise (the full proof may be found in Lemma A.1 below). The claim now follows from the subadditivity of v and the inequalities (1), (2), and (5) as follows.

$$v(\mathbb{R}^d) \leq v(\mathbb{R}^d \setminus U) + v(U \setminus B) + v(B) < 1 + d \ln(3d\phi).$$

□

Theorem 5.3 (Upper bound) Let $P = \{p_1, \dots, p_n\}$ be an ordered set of points such that $\|p_1 - p_2\| = \text{diameter}(P)$. Let $\Omega \subseteq \text{ball}(p_1, \frac{c}{2}\|p_1 - p_2\|)$ for some constant $c > 4$ be the bounding region. Then,

$$\mu_P(\Omega) < 1 + d \ln c + n + dn \ln(3d\phi_P).$$

Proof Recall that $P_i = \{p_1, \dots, p_i\}$ is the i th prefix. We rewrite the measure as a telescoping sum:

$$\mu_P = \mu_{P_n} = \mu_{P_2} + \sum_{i=3}^n (\mu_{P_i} - \mu_{P_{i-1}}).$$

The bounds from Lemmas 5.1 and 5.2 imply that

$$\mu_{P_i}(\Omega) < 1 + d \ln c + \sum_{i=3}^n (1 + d \ln(3d\phi_i)).$$

The result now follows from the definition of ϕ_P as the geometric mean of the ϕ_i s. □

6. Lower Bound

The proof of the lower bound will be similar to the proof of the upper bound in that we will use the pacing of a single new point to bound the change in the feature size integral.

Lemma 6.1 (One more point lower bound) Let P be a set of at least 2 points and let $P' = P \cup \{q\}$ for some $q \in \mathbb{R}^d$. Let $\Omega \subset \mathbb{R}^d$ be a set containing ball($q, \text{diameter}(P')$). Let $\phi = \frac{\mathbf{f}_P(q)}{\mathbf{f}_{P'}(q)}$. Then,

$$\mu_{P'}(\Omega) - \mu_P(\Omega) \geq \frac{1}{2^d} \left(d \ln \frac{\phi}{3} - 1 \right)$$

Proof The bound is trivial if $\phi \leq 3$, so we may assume that $\phi > 3$. Let $R = \mathbf{f}_{P'}(q)$. Since $P \subset P'$, $\nu := \mu_{P'} - \mu_P$ is a measure. By the subadditivity of ν , it will suffice to prove a lower bound on $\nu(U)$, where $U = \{x : R \leq \|x - q\| \leq \frac{R\phi}{3}\} \subseteq \Omega$. For all $x \in U$,

$$\mathbf{f}_P(x) \geq \frac{2R\phi}{3}, \text{ and}$$

$$\mathbf{f}_{P'}(x) \leq R + \|x - q\| \leq 2\|x - q\|.$$

The lower bound follows because there is at most one point of P in the interior of ball($q, R\phi$). The upper bound follows because $\mathbf{f}_{P'}(x)$ is 1-Lipschitz. We apply these bounds and convert to polar coordinates:

$$\begin{aligned} \nu(U) &= \frac{1}{\sqrt[d]{d}} \int_U \left(\frac{1}{\mathbf{f}_{P'}(x)^d} - \frac{1}{\mathbf{f}_P(x)^d} \right) d\lambda(x) \\ &\geq \frac{1}{\sqrt[d]{d}} \int_U \left(\frac{1}{(2\|x - q\|)^d} - \left(\frac{3}{2R\phi} \right)^d \right) d\lambda(x) \\ &> \left(d \int_R^{\frac{R\phi}{3}} \frac{r^{d-1}}{(2r)^d} dr \right) - \frac{1}{2^d} \\ &= \frac{1}{2^d} \left(\left(d \int_R^{\frac{R\phi}{3}} \frac{dr}{r} \right) - 1 \right) \\ &= \frac{1}{2^d} \left(d \ln \frac{\phi}{3} - 1 \right). \end{aligned}$$

□

When we apply the preceding lemma to a set of n points and use the same telescoping sum argument as in Theorem 5.3, we get the following lower bound.

Theorem 6.2 (Pacing Lower Bound) Let $P = \{p_1, \dots, p_n\}$ be an ordered set of points with pacing ϕ_P . If $\Omega \subset \mathbb{R}^d$ is a set containing ball($p_1, 2\text{diameter}(P)$), then

$$\mu_P(\Omega) > \frac{n}{2^d} \left(d \ln \frac{\phi_P}{3} - 1 \right).$$

The Pacing Lower Bound can be sublinear in n (or negative!) because of the negative term and the possibility that ϕ_P is small. For completeness, we also include a proof of a linear lower bound on $\mu_P(\Omega)$. This could also be deduced directly from the Ruppert lower bounds and the observation that the output mesh must contain all the input points.

Theorem 6.3 (Linear Lower Bound) Let $P = \{p_1, \dots, p_n\}$ be an ordered set of points with pacing ϕ_P . If $\Omega \subset \mathbb{R}^d$ is a set containing ball($p_1, 2\text{diameter}(P)$), then

$$\mu_P(\Omega) \geq \frac{n}{3^d}.$$

Proof For each $p_i \in P$, let $B_i := \text{ball}(p_i, \frac{1}{2}\mathbf{f}_P(p_i))$. Observe that the interiors of the B_i s are pairwise disjoint. For $x \in B_i$, the Lipschitz property of \mathbf{f}_P implies that $\mathbf{f}_P(x) \leq \mathbf{f}_P(p_i) + \|x - p_i\| \leq \frac{3}{2}\mathbf{f}_P(p_i)$. This bound implies that for any $p_i \in P$,

$$\mu_P(B_i) = \frac{1}{\nabla_d} \int_{B_i} \frac{d\lambda(x)}{\mathbf{f}_P(x)^d} \geq \frac{1}{\nabla_d} \int_{B_i} \frac{d\lambda(x)}{(\frac{3}{2}\mathbf{f}_P(p_i))^d} = \frac{1}{3^d}.$$

Let $S := \bigcup_{i=1}^n B_i$. Since $S \subset \Omega$,

$$\mu_P(\Omega) \geq \mu_P(S) = \sum_{i=1}^n \mu_P(B_i) \geq \frac{n}{3^d}.$$

□

7. Some directions for future work

The work of Ruppert on optimal meshing has been extended to feature size functions that also take into account input features beyond just point sets, including piecewise linear or even piecewise smooth complexes [DL09]. One direction for future work is to extend these methods for bounding the feature size measure in these settings as well.

It would also be interesting to extend this approach to anisotropic sizing functions, such as in [BWY11]. In that setting, it is not known how to relax the quality constraints to guarantee a linear size mesh.

Given that the spread has been used extensively to bound the complexity of Voronoi diagrams [Eri01, Eri05] (without Steiner points) and the pacing provides provably tighter results for mesh sizing than the spread, it might be possible to use pacing analysis to prove tighter per-instance bounds on the complexity of Voronoi diagrams.

Moreover, since the proof of Theorem 1.1 bounds the cost of adding a single point, it makes sense to apply these analytic techniques to dynamic meshing problems in which points may be added or deleted. In this case, the pacing of a point to be inserted tells us how to bound the work needed to complete the insertion. Similarly, we can estimate the cost of a deletion. In both cases, we can change the behavior of the algorithm when the number of local changes induced by an insertion or deletion will be exorbitant. This approach was shown to be effective in decreasing the total point location work in the static setting [MPS11].

References

- [ACSYD05] ALLIEZ P., COHEN-STEINER D., YVINEC M., DESBRUN M.: Variational tetrahedral meshing. *ACM Transactions on Graphics* 24, 3 (2005), 617–625.
- [ADA07] ANTANI L., DELAGE C., ALLIEZ P.: Mesh sizing with additively weighted voronoi diagrams. In *Proceedings of the 16th International Meshing Roundtable* (2007), pp. 335–346.
- [BEG94] BERN M. W., EPPSTEIN D., GILBERT J. R.: Provably good mesh generation. *J. Computer & Systems Sciences* 48, 3 (June 1994), 384–409. Special issue for 31st FOCS.
- [BWY11] BOISSONNAT J.-D., WORMSER C., YVINEC M.: *Anisotropic Delaunay Mesh Generation*. Rapport de recherche RR-7712, INRIA, Aug. 2011. URL: <http://hal.inria.fr/inria-00615486/en/>.
- [CCSM11] CHAZAL F., COHEN-STEINER D., MÉRIGOT Q.: Geometric inference for probability measures. *Foundations of Computational Mathematics* 11 (2011), 733–751.
- [CD04] CHENG S.-W., DEY T. K.: Quality meshing with weighted Delaunay refinement. *SICOMP: SIAM Journal on Computing* 33 (2004).
- [CDE*00] CHENG S.-W., DEY T. K., EDELSBRUNNER H., FACELLO M. A., TENG S.-H.: Sliver exudation. *JACM: Journal of the ACM* 47 (2000).
- [CL06] CHAZAL F., LIEUTIER A.: Topology guaranteeing manifold reconstruction using distance function to noisy data. In *Proceedings of the 22nd ACM Symposium on Computational Geometry* (2006).
- [DL09] DEY T. K., LEVINE J. A.: Delaunay meshing of piecewise smooth complexes without expensive predicates. *Algorithms* 2, 4 (2009), 1327–1349.
- [Eri01] ERICKSON J.: Nice point sets can have nasty Delaunay triangulations. In *Proceedings of the 17th ACM Symposium on Computational Geometry* (2001), pp. 96–105. URL: citeseer.ist.psu.edu/erickson01nice.html.
- [Eri05] ERICKSON J.: Dense point sets have sparse delaunay triangulations or “... but not too nasty”. *Discrete & Computational Geometry* 33 (2005), 83–115.
- [HMOS10] HUDSON B., MILLER G. L., OUDOT S. Y., SHEEHY D. R.: Topological inference via meshing. In *Proceedings of the 26th ACM Symposium on Computational Geometry* (2010), pp. 277–286.
- [HMPS09] HUDSON B., MILLER G. L., PHILLIPS T., SHEEHY D. R.: Size complexity of volume meshes vs. surface meshes. In *SODA: ACM-SIAM Symposium on Discrete Algorithms* (2009).
- [Li03] LI X.-Y.: Generating well-shaped d -dimensional Delaunay meshes. *Theor. Comput. Sci.* 296, 1 (2003), 145–165.
- [MMdGD11] MULLEN P., MEMARI P., DE GOES F., DESBRUN M.: HOT: Hodge-optimized triangulations. *ACM Trans. Graph* 30, 4 (2011).
- [Moo95] MOORE D.: The cost of balancing generalized quadtrees. In *SMA '95: Proceedings of the Third Symposium on Solid Modeling and Applications* (1995), pp. 305–312.
- [MPS08] MILLER G. L., PHILLIPS T., SHEEHY D. R.: Linear-size meshes. In *CCCG: Canadian Conference in Computational Geometry* (2008), pp. 175–178.
- [MPS11] MILLER G. L., PHILLIPS T., SHEEHY D. R.: Beating the spread: Time-optimal point meshing. In *SOCG: Proceedings of the 26th ACM Symposium on Computational Geometry* (2011), pp. 321–330.
- [MTTW95] MILLER G. L., TALMOR D., TENG S.-H., WALKINGTON N.: A Delaunay based numerical method for three dimensions: generation, formulation, and partition. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing* (1995), pp. 683–692.
- [MV00] MITCHELL S. A., VAVASIS S. A.: Quality mesh generation in higher dimensions. *SIAM J. Comput.* 29, 4 (2000), 1334–1370 (electronic).
- [Pav04] PAV S. E.: An anisotropic cardinality bound for triangulations. In *CCCG: The Canadian Conference on Computational Geometry* (2004), pp. 95–98. URL: <http://www.cccg.ca/proceedings/2004/34.pdf>.

- [Rup95] RUPPERT J.: A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms* 18, 3 (1995), 548–585.
- [She11a] SHEEHY D. R.: *Mesh Generation and Geometric Persistent Homology*. PhD thesis, Carnegie Mellon University, 2011.
- [She11b] SHEWCHUK J. R.: *Unstructured Mesh Generation*. CRC Press, 2011, ch. 10, pp. 259–298.
- [Tal97] TALMOR D.: *Well-Spaced Points for Numerical Methods*. PhD thesis, Carnegie Mellon University, 1997.
- [Üng09] ÜNGÖR A.: Off-centers: A new type of Steiner points for computing size-optimal quality-guaranteed Delaunay triangulations. *Comput. Geom.* 42, 2 (2009), 109–118.

Appendix A: A bit of calculus

Lemma A.1 Given positive constants $\phi \geq 1$ and R ,

$$\int_{R/2}^{\infty} \left(\frac{1}{r^d} - \frac{1}{(r+R\phi)^d} \right) r^{d-1} dr < \ln(3d\phi).$$

Proof We bound this integral using the change of variables $u = \frac{R\phi}{r} + 1$ as follows.

$$\begin{aligned} \int_{R/2}^{\infty} \left(\frac{1}{r^d} - \frac{1}{(r+R\phi)^d} \right) r^{d-1} dr &= \int_1^{1+2\phi} \left(\frac{u^d - 1}{u^d(u-1)} \right) du \\ &= \sum_{i=0}^{d-1} \int_1^{1+2\phi} u^{i-d} du \\ &< \ln(1+2\phi) + \sum_{i=0}^{d-2} \frac{1}{d-i-1} \\ &< \ln(1+2\phi) + \ln d \\ &\leq \ln(3d\phi). \end{aligned}$$

□