Size Complexity of Volume Meshes vs. Surface Meshes *

Benoît Hudson
Toyota Technological Institute, Chicago
bhudson@tti-c.org

Gary L. Miller Todd Phillips Don Sheehy
Carnegie Mellon University
{glmiller,tp517,dsheehy}@cs.cmu.edu

July 3, 2008. (Under Submission)

Abstract

Typical volume meshes in three dimensions are designed to conform to an underlying two-dimensional surface mesh, with volume mesh element size growing larger away from the surface. The surface mesh may be uniformly spaced or highly graded, and may have fine resolution due to extrinsic mesh size concerns. When we desire that such a volume mesh have good aspect ratio, we require that some space-filling *scaffold* vertices be inserted off the surface. We analyze the number of scaffold vertices in a setting that encompasses many existing volume meshing algorithms. We show that under simple preconditions, the number of scaffold vertices will be linear in the number of surface vertices.

1 Introduction

Given a surface mesh, many scientific computing and graphics applications will want to produce a volume mesh. Conversely, to build a surface mesh from another description of an input geometry, one might temporarily build a point location structure such as an oct-tree, which is a volume mesh. A natural question arises: can we relate the size of the surface mesh to the size of the volume mesh? A volume mesh will obviously have more vertices than the corresponding surface mesh, but in most settings, the spacing between vertices should grow quickly away from the surface. Since the density of the volume mesh is driven only by the surface, it is intuitive that the surface vertices should dominate in quantity. Our main result is to show that given a surface mesh in a well-proportioned domain, *the total number of vertices in the volume is linear in the number of vertices on the surface.* We will make this statement specific later as the Scaffold Theorem (Theorem 3.1).

This result has immediate and important ramifications concerning the asymptotic work and space of a large host of existing meshing and surface reconstruction algorithms. For example, in volume meshing, the user may specify a closed surface and ask for its interior to be meshed. Typical algorithms enclose the surface in a *bounding box* that contains the closed surface, incrementally add points until the surface is recovered and the volume mesh has good quality, then strip away the exterior volume vertices (see Figure 1). The surface and interior vertices are then returned to the user. This approach is widespread and is used for many two-, three-, and higher-dimensional meshing algorithms [BEG94, ABE98, CDE+00, ELM+00, She98, MV00, Üng04, MPW02, HMP06, CDR07]. The work and space complexity of these algorithms is output-sensitive and depends on the number of exterior vertices, *even though these vertices are transient.* Our new analysis is the first to control this exterior work. Since we show that the number of transient vertices is bounded by

*This work was supported in part by the National Science Foundation under grants ACI 0086093, CCR-0085982 and CCR-0122581.

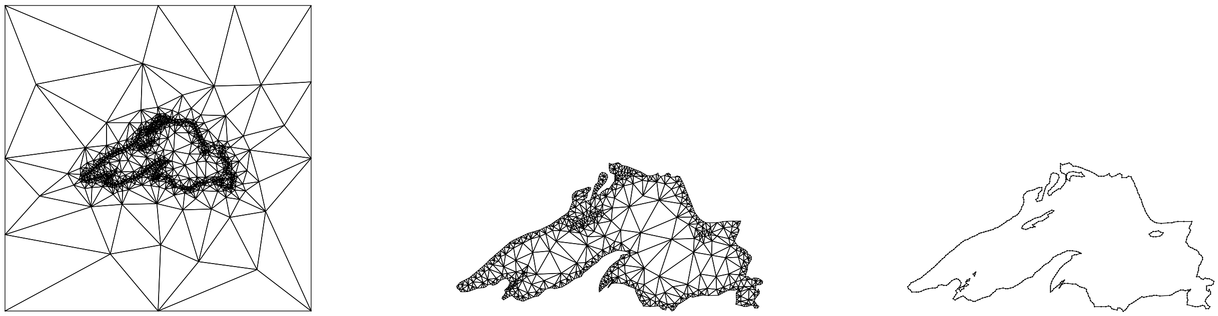

Figure 1: Incremental mesh refinement algorithms first generate a mesh over a bounding box (left), then remove the scaffold vertices and elements (center). Some applications may be interested in only the surface mesh (right). The (one-dimensional) Lake Superior surface shown mesh has 530 surface vertices. The volume mesh shown has 1072 total volume vertices; 258 interior and 284 exterior. We offer the first theoretical analysis of the costs of this scaffolding.

the surface vertices, this for the first time implies that these algorithms run output-sensitively with respect to the true user-desired output.

In the rest of this work we make our results precise. A good deal of care is taken to ensure the generality of these results, so that this analysis may be applied to many of the existing meshing algorithms with theoretical guarantees. Our proofs are in two parts. In the first part, we prove that if a good-quality volume mesh respects a surface, the volume vertices outnumber the surface vertices by only a constant. Our definition of respecting a surface is much looser than that of most prior work: the Voronoi cells of the surface vertices must cover the surface, but there is no topological requirement. In addition, our definition of a surface is extremely loose; it need not be manifold, or even connected. Additionally, our surface need not be $d-1$ dimensional: for instance, it could be a curve in 3D. Our only requirements are that the surface have a bounded number of connected components, and that each connected component of the surface have diameter within a constant factor of the diameter of the bounding domain.

In the second part, we show how this result relates to standard concepts from mesh refinement and surface reconstruction. In particular, we show that our result proves that a volume mesh of an ϵ-net of a surface is only a constant factor larger than the surface. We also show that many prior quality mesh refinement algorithms are susceptible to our analysis. This implies that they still run in the time (and memory usage) bounds they claim, even when the volume actually meshed is larger than what the user asked to mesh.

2 Preliminary Geometric Definitions

In this paper, we assume there exists a surface $\mathcal{S}$ embedded in $\mathbb{R}^d$. It need not be connected. Let D be the minimum diameter of any connected component of $\mathcal{S}$, where the diameter is the maximum Euclidean distance between two points in the component. We require that the diameter of all other components, and the diameter of $\mathcal{S}$, be in $\Theta(D)$. Around $\mathcal{S}$ there is a compact and connected *domain* Ω with $\mathcal{S} \subset \Omega \subset \mathbb{R}^d$. Typically, Ω will be a box or a hypercube. The diameter of Ω must be in $\Theta(D)$.

Let Γ_d denote the volume of the unit ball in $\mathbb{R}^d$. For $x \in \mathbb{R}^d$ and $r \in \mathbb{R}$, let $\mathcal{B}(x, r)$ be the open ball centered at x or radius r (whose volume is given by $\Gamma_d\, r^d$.)

Suppose we have a set of points (vertices) $M \subset \Omega$. A vertex-set M induces a *local feature size* function $f_M : \Omega \to \mathbb{R}$. At a point $x \in \Omega$, the local feature size is the distance from x to the second-nearest vertex. We frequently use the fact that f_M is 1-Lipschitz: that is, $f_M(x) \leq f_M(y) + |x - y|$ for all x and y in $\mathbb{R}^d$ (this is easily verified by the triangle inequality. At a vertex $v \in M$, the local feature size coincides with the distance

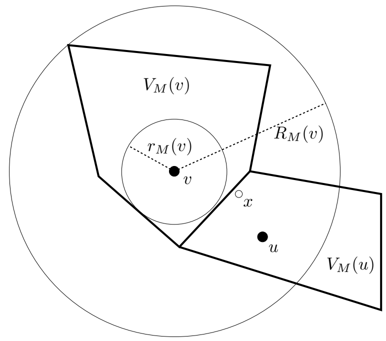

Figure 2: The Voronoi cells of two vertices u and v in a vertex-set M (not pictured). The radii of the inner-ball and outer-ball of v are labeled. The point x is 0.9-medial.

to the nearest neighbour, which we denote $NN_M(v)$.

Given the vertex-set M, we denote by $V_M(v)$ the closed Voronoi cell of v: those points in $\mathbb{R}^d$ for which no vertex in M is closer than is v. We identify two natural balls with v: the *inner-ball* $b_M(v)$ is the largest ball centered at v that is contained within $V_M(v)$, while the *outer-ball* $B_M(v)$ is the smallest ball centered at v that contains all of $V_M(v) \cap \Omega$. We denote by $r_M(v)$ and $R_M(v)$ their respective radii (i.e. $B_M(v) = \mathcal{B}(v, R_M(v))$.) See Figure 2.

2.1 Well-spaced, Well-paced, and Medial Points

If there is a constant τ for which every v in a vertex-set M, has the property $R_M(v) \leq \tau r_M(v)$, then we say that v is *τ-well-spaced*. Loosely, this implies that the $V_M(v) \cap \Omega$ is roundly shaped and that v is roughly centered. Figure 3(right) shows a set of well-spaced points; contrast this with the vertices of Figure 4(left).

The problem of scaffolding a surface mesh with a volume mesh is that of finding a minimal well-spaced superset (volume mesh) M of a vertex-set (surface mesh) N.

We will make use of a theorem from [MPS08]. First, we introduce some relevant definitions. The boundaries of the Voronoi cells of each vertex in M form the medial axis of M. Miller *et al* [MPS08] generalize this and say that a point x is *θ-medial* with respect to M if it lies near the medial axis, in the sense that $NN_M(x) \geq \theta f_M(x)$. Notice that whenever we had a point x to a set M, it will decrease the feature-size f_M in the vicinity of x. A key observation is that adding a θ-medial point will only decrease the feature-size by a constant amount.

Given an arbitrary vertex-set N, and an ordered set of vertices $E \equiv \langle v_1, \ldots, v_k \rangle$, we say that E is a θ-well-paced extension of N if v_1 is a θ-medial point of N, and each v_i is a θ-medial point of $N \cup \{v_1 \ldots v_{i-1}\}$. Informally, the name arises from the fact that the local feature size shrinks only slowly after each insertion.

Well-paced extensions are *not* well-spaced in general, but they have useful similarities to surface meshes. We now state for completeness the Well-Pacing Theorem ([MPS08], Corollary 3).

Theorem 2.1 *(***Well-Pacing Theorem***)*
There is a constant $C_{2.1}$, such that if N is a well-paced extension of a well-spaced set, then there exists a well-spaced superset $M \supset N$, with $|M| \leq C_{2.1}|N|$.

(**For Review Purposes:** [MPS08] will appear in August. Appendix section C contains relevant proof details.)

3 Scaffold Theorem

3.1 Overview

Our main result is the Scaffold Theorem 3.1, showing that given a volume mesh M with underlying surface mesh N, $|M|$ is bounded above by a constant times the size of $|N|$. (Informally, $|M| \lesssim |N|$.) Section 3.2 defines the formal setting in where Theorem applies.

The easiest proof would be to show that N can be written as a well-paced extension of a well-spaced set, then we could simply apply the Well-Pacing Theorem 2.1 to directly bound $|M|$. This is difficult, so instead we define the concept of a spacing-equivalent surface mesh N', and show that one exists (Theorem 3.6). We can then construct a scaffolding volume mesh M' for N'. The proof then proceeds as three straightforward lemmas (3.3, 3.4, and 3.5). First, Lemma 3.3 shows that $M \lesssim M'$. This follows because the surface meshes N and N' are roughly equivalent, so then their associated scaffolding volume meshes are also roughly equivalent. Next, Lemma 3.4 applies Theorem 2.1 to show that $M' \lesssim N'$. Lastly, Lemma 3.5 shows that $N' \lesssim N$ and follows simply by the spacing-equivalence of N'. Putting these together yields $M \lesssim M' \lesssim N' \lesssim N$. We now proceed with a formal proof.

3.2 Definitions: Scaffold Mesh and Seeded Surface Mesh

Suppose we are given a domain Ω as in Section 2. Further suppose we are given a "surface" $\mathcal{S} \subset \Omega$. We require only that $\mathcal{S}$ is a closed subset. Suppose we have a finite vertex-set $M \subset \Omega$. Define the surface vertices as $N = M \cap \mathcal{S}$. To capture the notion that the sizing of this point set is solely due to the surface, we wish to relate the Voronoi cells of M to the feature-size due to N. Formally, the input must satisfy the following two conditions:

$$\exists\, \alpha^+ \in (0, \infty),\ \forall\, m \in M,\ R_M(m) \le \alpha^+ f_N(m) \tag{1}$$

$$\exists\, \alpha^- \in (0, 1),\ \forall\, m \in M,\ r_M(m) \ge \alpha^- f_N(m) \tag{2}$$

Note this implicitly implies that M is a well-spaced set of vertices, with $R(m)/r(m) \le \alpha^+/\alpha^-$. When M and N satisfy conditions (1) and (2), we will say that M is an **α-scaffolding volume mesh** of the set N in Ω. We further require that the surface vertices N cover $\mathcal{S}$ in the following sense:

$$\mathcal{S} \subset \bigcup_{n \in N} V_M(n) \tag{3}$$

Lastly, we require that the volume being filled is well-proportioned to the underlying surface. We accomplish this by requiring the existence of a **seed** $N_0 \subset N$, The seed N_0 must contain two vertices in each connected component of $\mathcal{S}$ and must be a well-spaced set of points. Formally:

$$\text{For each connected component } \mathcal{T} \subset \mathcal{S},\ \exists p, q \in N_0 \cap \mathcal{T},\ \text{ s.t. } p \neq q \tag{4}$$

$$N_0 \text{ is } \rho\text{-well-spaced, i.e.} \exists \rho \in (1, \infty) \forall n \in N_0,\ R_{N_0}(n) \le \rho r_{N_0}(n) \tag{5}$$

If a set N with scaffolding volume mesh M meets conditions (3)-(5), then we say that N is a **seeded surface mesh** of $\mathcal{S}$ in Ω. Figures 3 and 4 illustrate these definitions.

3.3 Definitions: Spacing-Equivalent Sets

Suppose $N \subset \mathcal{S}$ is a vertex-set. Consider a vertex-set $N' \subset \mathcal{S}$ that satisfies the following two conditions:

$$\exists\, \beta^+ \in (0, \infty),\ \forall\, s \in \mathcal{S},\ f_{N'}(s) \le \beta^+ f_N(s) \tag{6}$$

$$\exists\, \beta^- \in (0, \beta^+),\ \forall\, s \in \mathcal{S},\ f_{N'}(s) \ge \beta^- f_N(s) \tag{7}$$

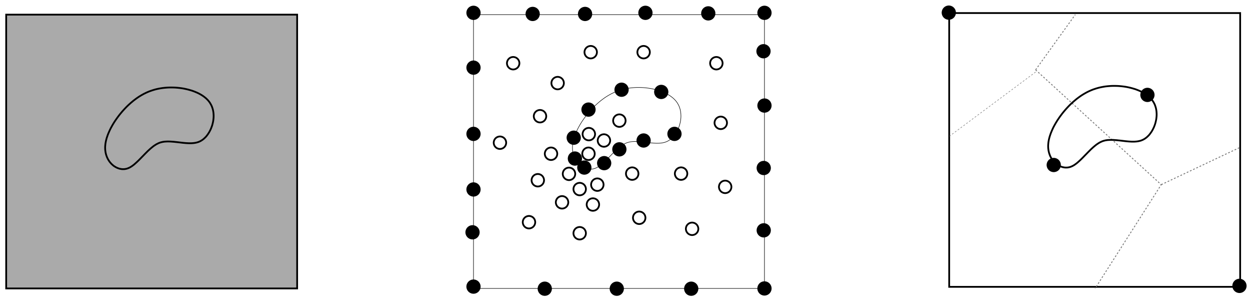

Figure 3: The definitions of Section 3.2 illustrated abstractly. **Left,** a surface $\mathcal{S}$ is composed of *both* black shapes, with the domain Ω shaded. **Center,** vertices form a volume mesh M of Ω. The subset of surface vertices N are shown in black, with the volume vertices in white. Note how the density of the volume vertices is driven only by the density of surfaces vertices. **Right,** a possible seed N_0, containing at least two points from each component of the surface $\mathcal{S}$. Notice the four points are well-spaced and have quality Voronoi cells (shown in dashed lines).

We then say that such an N' is **β-spacing-equivalent** to N on $\mathcal{S}$. Theorem 3.6 will show that a seeded surface mesh N (with seed N_0) always has a β-spacing-equivalent set N', with the additional property that N' is also a $(1/3)$-well-paced extension of the seed N_0. For any vertex-set N', it is possible to construct (using off the shelf algorithms [Rup95]) a well-spaced superset of vertices M' that is a γ-scaffolding volume mesh, for some γ determined by the algorithm selected. This mesh M' need not be aware of the surface at all, since we will *not* require that N' be a seeded surface mesh.

3.4 Scaffold Theorem Proof

Theorem 3.1 *(**Scaffold Theorem**) Suppose M is an α-scaffolding volume mesh of N, and that N is a seeded surface mesh. Then there exists a constant $C_{3.1}$ depending only on α^+ and α^- such that:*

$$|M| \leq C_{3.1}|N|$$

The main theorem will follow immediately from the existence of a spacing-equivalent surface mesh (Theorem 3.6) and then composing Lemmas 3.3, 3.4, and 3.5. The proof begins with a small technical lemma with the goal of extending spacing-equivalence to the entire domain Ω, since it is only guaranteed on $\mathcal{S}$.

Lemma 3.2 *Suppose N is a surface mesh as in Section 3.2 and N' is a β-spacing-equivalent vertex-set, then:*

$$\forall\, x \in \Omega,\ f_{N'}(x) \leq (1 + 2\beta^+) f_N(x)$$

Proof: Let $x \in \Omega$. Let $s \in \mathcal{S}$ be a point on $\mathcal{S}$ with minimum Euclidean distance to x (s exists by the closure of $\mathcal{S}$.) Because $N \subset \mathcal{S}$, we have $f_N(x) \geq |x - s|$. Combining this fact with the Lipschitz conditions and condition (6), we find:

$$f_{N'}(x) \leq f_{N'}(s) + |x - s| \leq \beta^+ f_N(s) + |x - s| \leq \beta^+(f_N(x) + |x - s|) + |x - s| \leq (1 + 2\beta^+) f_N(x) \tag{8}$$

■

Using the previous lemma, we can now show that an induced scaffold M' of a spacing-equivalent set N' is at least as large (up to a constant) as the scaffold M of the original set N. (A bound in the other direction is also true with differing constants, but we will not need it.)

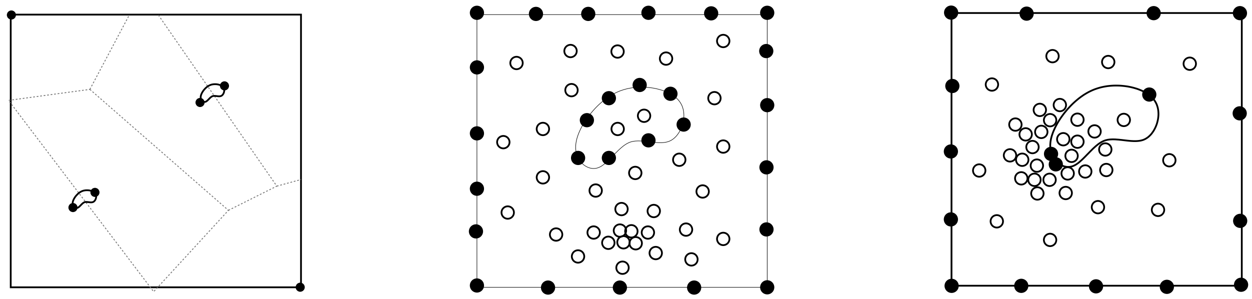

Figure 4: Examples of violations of the conditions in Section 3.2. **Left,** when the surface $\mathcal{S}$ has disproportionately small components, it will be too costly to fill the volume in a way that resolves these small surface features. Note that no seed can exist in this example. An attempted seed is shown, but as the surface components grow relatively small, there is no way to fit two points on each component in a way that is well-spaced. **Center,** this volume mesh is not a scaffold mesh, because the sizing is not driven by the surface. The sink in the lower-center could contain arbitrarily many volume vertices. Note how this violates equation (2). **Right,** If the surface vertices do not cover $\mathcal{S}$ as in equation 3, then there could be a small surface feature requiring many volume vertices to resolve without having to add any more surface vertices. Note that a seed still exists in this example.

Lemma 3.3 *Suppose M is an α-scaffolding volume mesh of N as in Section 3.2, and suppose M' is a γ-scaffolding volume mesh for a β-spacing-equivalent surface mesh N', then:*

$$|M| \leq C_{3.3}|M'| \qquad \textit{where} \qquad C_{3.3} = \left(\frac{2(1+2\beta^+)\gamma^+}{\alpha^-\gamma^-}\right)^d$$

Proof: The proof is a packing argument showing that if we partition the vertices M into the Voronoi cells $V_{M'}$, no cell receives more than a constant-sized portion of M. There are disjoint empty balls around the vertices of M with radii that are lower bounded by f_N which lower bounds $f_{N'}$ (by Lemma 3.2). But $f_{N'}$ also upper bounds the size of the Voronoi cells of M' because it is a γ-scaffolding volume mesh of N'. A detailed proof is given in the appendix. ■

Lemma 3.4 *Suppose M' is a γ-scaffolding volume mesh of a N', and that N' is a well-paced extension of a well-spaced seed, then:*

$$|M'| \leq C_{3.4}|N'|$$

Proof: This is a direct application of the Well-Pacing Theorem 2.1. The constant $C_{3.4}$ will be inversely proportional to γ^-. ■

Lemma 3.5 *Suppose N' is β-spacing-equivalent to a seeded surface mesh N with α-scaffolding volume mesh M, then we have:*

$$|N'| \leq C_{3.5}|N| \qquad \textit{where} \qquad C_{3.5} = \left(\frac{4\alpha^+}{\alpha^-\beta^-}\right)^d$$

Proof: Lemma 3.5 follows from a packing argument and the proof is similar Lemma 3.3. See Appendix for details. ■

3.5 Existence of Spacing-Equivalent Sets

Theorem 3.6 *(**Spacing-Equivalent Existence**) Suppose N is a seeded surface mesh of $\mathcal{S}$ in Ω (with seed N_0) as specified by Section 3.2. Then there exists a β-spacing-equivalent N' that is a $(1/3)$-well-paced extension of N_0, for any $\beta^+ \in (0, 1]$ and $\beta^- \in (0, \beta^+/17)$.*

Proof:

Let N, N_0, $\mathcal{S}$, β^+, and β^- be given satisfying the hypothesis of the theorem. Proof is by construction according to the following algorithm:
Begin with the seed N_0 and a counter i, initially $i = 0$.
While $\exists x \in \mathcal{S}$ such that $f_{N_i}(x) > \beta^+ f_N(x)$, let $p \in N_i$ such that $x \in V_{N'}(p)$:

1. If $x \notin b_{N_i}(p)$. Set $N_{i+1} = \{x\} \cup N_i$.
2. If $x \in b_{N_i}(p)$. Let $y \in \partial b_{N_i}(p) \cap \mathcal{S}$. (Note y exists by 4.) Set $N_{i+1} = \{y\} \cup N_i$.
3. Increment i.

The finiteness of N guarantees this will terminate at some index I. Take $N' = N_I$.

First, we claim that any points added to N_i is $(1/3)$-medial wrt N_i. Consider a point x on any iteration i of the loop. Let $q \in N_i$ be the nearest neighbor of p. If x is inserted in Step 1.:

$$f_{N_i}(x) \le |x - p| + f_{N_i}(p) = |x - p| + 2r_{N_i}(p) \le 3|x - p| = 3NN_{N_i}(x)$$

so x is $(1/3)$-medial. If y is inserted in Step 2., we have:

$$f_{N_i}(y) \le |y - p| + f_{N_i}(p) = 3|y - p| = 3NN_{N_i}(y)$$

so y is $(1/3)$-medial.

Thus N' is a well-paced extension of the well-spaced N_0.

It remains to show that N' is β-spacing-equivalent to N, namely that conditions (6) and (7) hold.

Clearly, the upper bound condition (6) holds by construction. Either there are no more violations of (6) or $N' = \mathcal{S}$, but the latter would clearly violate finiteness of N. The proof of condition (7) is more subtle but follows similar arguments from previous work [Rup95]. We will use the following lemma to show condition (7).

Lemma 3.7 *Suppose $x \in N'$ is inserted at iteration i (i.e. $x \in N_{i+1} - N_i$), then:*

$$NN_{N_i}(x) > \frac{\beta^+}{7} f_N(x)$$

Note that this Lemma would imply the immediate corollary:

Corollary 3.8

$$\forall x \in N',\ f_{N'}(x) > \frac{\beta^+}{8} f_N(x)$$

Proof of corollary: Let $x \in N'$ (iteration i) be given and take $y \in N'$ (iteration j) s.t. $|x - y| = NN_{N'}(x)$. If $i > j$, then we simply apply the lemma directly:

$$NN_{N'}(x) = NN_{N_i}(x) > \frac{\beta^+}{7} f_N(x) > \frac{\beta^+}{8} f_N(x)$$

If $i < j$, then we use Lipschitz of f_N and apply the lemma to y:

$$f_N(x) \le f_N(y) + |x - y| < \frac{7}{\beta^+} NN_{N_j}(y) + |x - y| \le \frac{7}{\beta^+}|x - y| + |x - y| < \frac{8}{\beta^+}|x - y| = \frac{8}{\beta^+} NN_{N'}(x)$$

We return to proving Lemma 3.7. Proof is by induction. Note that the corollary always holds inductively whenever the lemma does. For $x \in N_0$, the lemma clearly holds since $N_0 \subset N$. Let x be a point added in Case (1) at iteration i. The lemma follows immediately since x was selected as a point where f_N was large and x was (1/3)-medial:

$$NN_{N_i}(x) \geq \frac{1}{3} f_{N_i}(x) > \frac{\beta^+}{3} f_N(x) > \frac{\beta^+}{7} f_N(x)$$

.

Let y be a point added in Case (2) at iteration i relative to some $x \in \mathcal{S}$ and $p \in N_i$. Let q be the nearest neighbor of p in N_i. We will use the Lipschitz condition on f_N and f_{N_i}, as well as applying the corollary as an inductive hypothesis:

$$f_N(y) \leq f_N(x) + |x - y| < \frac{1}{\beta^+} f_{N_i}(x) + |x - y| \tag{9}$$

$$\leq \frac{1}{\beta^+} f_{N_i}(y) + \frac{1}{\beta^+}|x - y| + |x - y| = \frac{1}{\beta^+} f_{N_i}(y) + (1 + \frac{1}{\beta^+})|x - y| \tag{10}$$

$$\leq \frac{1}{\beta^+}|y - q| + (1 + \frac{1}{\beta^+})2r_{N_i}(p) \tag{11}$$

$$\leq \frac{1}{\beta^+} 3r_{N_i}(p) + (1 + \frac{1}{\beta^+})2r_{N_i}(p) = (2 + \frac{5}{\beta^+})NN_{N_i}(y) \leq \frac{7}{\beta^+} NN_{N_i}(y) \tag{12}$$

Having proved the lemma and corollary, we recall that condition (7) must hold for any $x \in \mathcal{S}$. Let $x \in \mathcal{S}$ be given and take $v \in N'$ such that $x \in V_{N'}(v)$. then. We first notice that $f_{N'}(x) \geq r_{N'}(p)$. We use Lipschitz property of f_N and apply the corollary at v:

$$f_N(x) \leq f_N(v) + |x - v| < \frac{8}{\beta^+} f_{N'}(v) + |x - v| = \frac{16}{\beta^+} r_{N'}(v) + |x - v| \leq \frac{16}{\beta^+} f_{N'}(x) + f_{N'}(x) \leq \frac{17}{\beta^+} f_{N'}(x) < \frac{1}{\beta^-} f_{N'}(x)$$

■

4 Algorithms

Our result assumes that the surface $\mathcal{S}$, the volume mesh M, the surface mesh N, and the seed N_0 were all given. Ideally, we should not need to know so much, and instead we would have an algorithm to fill in the unknowns. There are many mesh refinement algorithms in the literature that need only know either $\mathcal{S}$ or N. Provided the seed exists—it need not be known—said mesh refinement algorithms will produce an output that matches the requirements of the Scaffold Theorem 3.1: $|M| \in \Theta(|N|)$. The surprising conclusion is that in terms of runtime and output size, when the ambient dimension is bounded, it is asymptotically free to mesh a volume rather than meshing only a surface—again, provided the surface has a seed.

4.1 Meshing a surface sample

The simplest application is to take as input a set of points N that all lie on a manifold surface (for example, the famous Stanford Bunny model), and construct from it the volume mesh M. This is a useful endeavor if we are to animate the model. To generate the volume mesh, we use a Voronoi (or Delaunay) refinement algorithm. The volume mesher first wraps the points of N into an appropriate bounding box, of diameter only a constant factor larger than the diameter of N. It initializes M with N, then finds a vertex v with $R_M(v) \geq \tau r_M(v)$, and identifies some point p that is in the Voronoi cell of v, but far from it: $|vp| \leq |up|$ for all $u \in M$, but $|vp| \geq \tau r_M(v)$. The algorithm then adds p to M, and continues this process until M is τ-well-spaced. A large number of algorithms implement this process (*e.g.* [Rup95, She98, HPU05, HMP06]).

Our theorem requires that the surface is covered by the Voronoi cells of just the surface vertices—that is, no point of $\mathcal{S}$ lies in the Voronoi cell of a vertex in $M \backslash N$. Under certain assumptions on N, we can prove this holds. We require that there be some ϵ such that for all x, there is a vertex $v \in N$ such that $|vx| \leq \epsilon$; but for all $u \in N$, all other vertices $v \in N$ lie at distance $|uv| \geq \epsilon/2$. In other words, N is an ϵ-net of $\mathcal{S}$.

Lemma 4.1 *Consider a point $x \in \mathcal{S}$ whose nearest neighbour in N is v. If the volume mesh M is computed with $\tau > 4$, then x remains in the Voronoi cell $V_M(v)$.*

Proof: For any vertex u created during refinement, there is some u' that created u: when u was inserted, its nearest neighbour was u', and $|uu'| \geq \tau r_M(u')$. In other words, created vertices have nearest neighbour larger than the distance between the closest pair of points in N. The closest pair must be at least $\epsilon/2$ from each other, by assumption, so any $u \in M \backslash N$ has $r_M(u) \geq \tau\epsilon/4$. Return now to consider x. For a contradiction, we assume that the nearest neighbour of x in M is a created vertex u. Then $|ux| < |vx|$. Given that $|vx| \leq \epsilon$, we know that $|uv| \leq 2\epsilon$. Clearly, $r_M(u)$ can be no larger than half the distance to any vertex: $r_M(u) \leq |uv|/2 \leq \epsilon$. Remembering the previous bound, we conclude that $\epsilon \geq \tau\epsilon/4$, or equivalently, $\tau \leq 4$, a contradiction. ■

As a corollary, this means that every point $x \in \mathcal{S}$ is in the Voronoi cell of some vertex in N, and therefore Theorem 3.1 holds. Then $|M| \in O(|N|)$, assuming N is an ϵ-net of $\mathcal{S}$, and that $\tau > 4$. But N is input, so $n = |N|$: the volume mesh contains a number of vertices only linear in the size of the input! We can relax the requirement on τ by remembering that a τ-well-spaced mesh and a τ'-well-spaced mesh have size within a constant factor of each other, where the constant is a function of τ, τ'. This lets us conclude:

Corollary 4.2 *A τ-well-spaced mesh of an ϵ-net has size $O(n)$ for any τ and ϵ.*

4.2 Meshing a surface

In mesh refinement for engineering and scientific applications, the input is typically specified as a piecewise linear complex or a piecewise smooth complex, made up of a collection of vertices, segments or curves, and polygons or smooth surfaces (and so on, in higher dimension). As in the prior subsection, we assume the algorithm first places a box around the input complex, then iteratively inserts vertices. In the face of linear or smooth features, this requires greater care than before although the details are nearly irrelevant to our results here. The mesher continues adding vertices until two conditions are met: that the vertices are well-spaced, and that the Delaunay triangulation "respects" the input complex. In the case of piecewise linear complexes, we say a triangulation respects it if each linear facet appears as the union of a set of Delaunay simplices [Rup95, She98]. In other words, the Voronoi diagram of surface vertices covers the input. For smooth complexes [CDR07, RY07], respecting a surface requires even more than the covering condition.

The analysis of algorithms that mesh complexes typically rely on a global function called the *local feature size*, denoted lfs, which represents how much refinement will be necessary locally. For our purposes, we require that the local feature size be defined on $\mathcal{S}$, and extended to the entire domain via the minimum 1-Lipschitz function: at $x \in \Omega \backslash \mathcal{S}$, $\mathrm{lfs}(x) \equiv \min_{y \in \mathcal{S}} \mathrm{lfs}(y) + |xy|$. This is within a factor of three of Ruppert's more traditional local feature size function define on linear complexes, but extends more easily to smooth complexes. Most mesh refinement algorithms arising from the computational geometry community guarantee that vertices are not too closely packed: at every $v \in M$, $r_M(v) \geq \gamma^- \mathrm{lfs}(v)$. For our purposes here, we also need the algorithm to guarantee that every vertex v must have $R_M(v) \leq \gamma^+ \mathrm{lfs}(v)$, and every point x must have some neighbour v with $|vx| \leq \gamma^+ \mathrm{lfs}(x)$. Parallel mesh refinement algorithms happen to require this for fast parallel runtime [STÜ07, HMP07]. These conditions are extremely reminiscent of conditions (1) and (2), assuming f_N and lfs are related.

Lemma 4.3 *For all $x \in \Omega$, $f_N(x) \in [\delta^- \mathrm{lfs}(x), \delta^+ \mathrm{lfs}(x)]$.*

Proof: For the lower bound on f_N, consider a point x in the domain. It lies in the Voronoi cell of some vertex $v \in M$. Since local feature size is 1-Lipschitz, $\mathrm{lfs}(x) \leq \mathrm{lfs}(v) + |vx|$. By the assumption on the algorithm,

$r_M(v) \geq \gamma^- \operatorname{lfs}(v)$. We also know that the second-nearest vertex to x is at least as far as $\max(r_M(v), |vx|)$. Thus we know $\operatorname{lfs}(x) \leq (1 + 1/\gamma^-) f_M(x)$. Finally, removing vertices can only increase f: $f_M(x) \leq f_N(x)$, which proves that $f_N(x) \geq \frac{\gamma^-}{1+\gamma^-} \operatorname{lfs}(x)$. Then $\delta^- = \frac{\gamma^-}{1+\gamma^-}$.

For the upper bound, first consider a vertex $u \in N$. Because the Voronoi cells of N cover $\mathcal{S}$, we know that u has a neighbour u' on N with $|uu'| \leq 2R_M(u)$. The algorithm guarantees $R_M(u) \leq \gamma^+ \operatorname{lfs}(u)$. Thus $f_N(u) \leq 2\gamma^+ \operatorname{lfs}(u)$. Next, consider a a point $y \in \mathcal{S}$ whose nearest neighbour is u. We know that $f_N(y) \leq f_N(u) + |uy| \leq 2\gamma^+ \operatorname{lfs}(u) + |uy|$. At u, the Lipschitz condition bounds $\operatorname{lfs}(u) \leq \operatorname{lfs}(y) + |uy|$. Finally, u is the nearest vertex to y, but there must be some vertex within distance $\gamma^+ \operatorname{lfs}(y)$, which means $|uy| \leq \gamma^+ \operatorname{lfs}(y)$. Thus $f_N(y) \leq 2\gamma^+ \operatorname{lfs}(y) + 2|uy| \leq 4\gamma^+ \operatorname{lfs}(y)$. We can now consider an arbitrary point x. Let y be the point of $\mathcal{S}$ for which $\operatorname{lfs}(x) = \operatorname{lfs}(y) + |xy|$. By the Lipschitz condition on f_N, we know $f_N(x) \leq f_N(y) + |xy|$. We just proved that $f_N(y) \leq 4\gamma^+ \operatorname{lfs}(y)$. By definition, $\operatorname{lfs}(y) \leq \operatorname{lfs}(x)$, so we know $f_N(y) \leq 4\gamma^+ \operatorname{lfs}(x)$. Equally, $|xy| \leq \operatorname{lfs}(x)$. Therefore, $f_N(x) \leq (1 + 4\gamma^+) \operatorname{lfs}(x)$. Then $\delta^+ = (1 + 4\gamma^+)$. ■

Corollary 4.4 *When presented with an input piecewise smooth or piecewise linear complex for which there exists a seed, a quality mesh refinement algorithms that outputs a mesh of optimal size creates a volume mesh M and a surface mesh N, with $|M| \in O(|N|)$.*

5 Conclusions

Accounting for scaffolding costs is a pressing question in the timing and output-size analysis of many mesh generation algorithms that are used in practice. The Scaffold Theorem shows that these costs are not dominant, as has so often been assumed without proof in prior work. This analysis is made applicable to many algorithms by abstracting the meshing problem to that of simply generating a minimal well-spaced superset of a vertex-set. This ignores many of the topological and geometric intricacies that make meshing algorithms difficult to analyze, while still preserving enough distribution information about the vertices to make meaningful statements on mesh-size.

Reflecting on the analysis, the surface vertices are paramount and the underlying surface itself plays only a small role in controlling the size of the volume mesh, It is then theoretically of interest to simply consider the size of a minimal well-spaced superset M of a vertex-set $N \subset \Omega$. It is well-established that:

$$|M| \in \Theta\left(\int_\Omega \frac{1}{f_N^d}\right)$$

A worst case upper bound on this integral is $O(|N| \log \Delta)$, where Δ is the spread of the domain; the ratio of the diameter of Ω to the closest pair in N. In general, this bears no combinatorial relationship to $|N|$. The Scaffold Theorem provides sufficient conditions (that are highly relevant in practice) for a setting wherein $|M|$ is linear in $|N|$. But these conditions are nowhere near necessary. It is an interesting question whether there exist simple necessary and sufficient conditions that will combinatorially bound $|M|$ when N is given arbitrarily.

References

[ABE98] Nina Amenta, Marshall Bern, and David Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical models and image processing: GMIP*, 60(2), 1998.

[BEG94] Marshall Bern, David Eppstein, and John R. Gilbert. Provably Good Mesh Generation. *Journal of Computer and System Sciences*, 48(3):384–409, June 1994.

[CDE+00] Siu-Wing Cheng, Tamal Krishna Dey, Herbert Edelsbrunner, Michael A. Facello, and Shang-Hua Teng. Sliver Exudation. *Journal of the ACM*, 47(5):883–904, September 2000.

[CDR07] Siu-Wing Cheng, Tamal K. Dey, and Edgar A. Ramos. Delaunay refinement for piecewise smooth complexes. In *SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1096–1105, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics.

[ELM+00] Herbert Edelsbrunner, Xiang-Yang Li, Gary L. Miller, Andreas Stathopoulos, Dafna Talmor, Shang-Hua Teng, Alper Üngör, and Noel Walkington. Smoothing and cleaning up slivers. In *STOC*, pages 273–277, 2000.

[HMP06] Benoît Hudson, Gary Miller, and Todd Phillips. Sparse Voronoi Refinement. In *Proceedings of the 15th International Meshing Roundtable*, pages 339–356, Birmingham, Alabama, 2006. Long version available as Carnegie Mellon University Technical Report CMU-CS-06-132.

[HMP07] Benoît Hudson, Gary L. Miller, and Todd Phillips. Sparse Parallel Delaunay Refinement. In *19th Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pages 339–347, San Diego, June 2007.

[HPU05] Sariel Har-Peled and Alper Üngör. A Time-Optimal Delaunay Refinement Algorithm in Two Dimensions. In *Symposium on Computational Geometry*, 2005.

[Moo95] Doug Moore. The cost of balancing generalized quadtrees. In *SMA '95: Proceedings of the Third Symposium on Solid Modeling and Applications*, pages 305–312, 1995.

[MPS08] Gary L. Miller, Todd Phillips, and Donald Sheehy. Linear-size meshes. In *Canadian Conference on Computational Geometry*, 2008. To appear.

[MPW02] Gary L. Miller, Steven E. Pav, and Noel J. Walkington. Fully Incremental 3D Delaunay Refinement Mesh Generation. In *Eleventh International Meshing Roundtable*, pages 75–86, Ithaca, New York, September 2002. Sandia National Laboratories.

[MV00] Scott A. Mitchell and Stephen A. Vavasis. Quality mesh generation in higher dimensions. *SIAM J. Comput.*, 29(4):1334–1370 (electronic), 2000.

[Rup95] Jim Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18(3):548–585, 1995. Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Austin, TX, 1993).

[RY07] Laurent Rineau and Mariette Yvinec. Meshing 3D domains bounded by piecewise smooth surfaces. In *16th International Meshing Roundtable*, pages 442–460, 2007.

[She98] Jonathan Richard Shewchuk. Tetrahedral Mesh Generation by Delaunay Refinement. In *Proceedings of the Fourteenth Annual Symposium on Computational Geometry*, pages 86–95, Minneapolis, Minnesota, June 1998. Association for Computing Machinery.

[STÜ07] Daniel Spielman, Shang-Hua Teng, and Alper Üngör. Parallel Delaunay refinement: Algorithms and analyses. *IJCGA*, 17:1–30, 2007.

[Üng04] Alper Üngör. Off-centers: A new type of steiner points for computing size-optimal guaranteed-quality delaunay triangulations. In *Proceedings of LATIN*, 2004.

A Proof of Lemma 3.3

Lemma 3.3 Suppose M is an α-scaffolding volume mesh of N as in Section 3.2, and suppose M' is a γ-scaffolding volume mesh for a β-spacing-equivalent surface mesh N', then:

$$|M| \leq C_{3.3}|M'| \qquad \text{where} \qquad C_{3.3} = \left(\frac{2(1+2\beta^+)\gamma^+}{\alpha^-\gamma^-}\right)^d$$

Proof:

To form the packing argument, we will first lower-bound the size of the $b_M(v)$ for $v \in M$, which we know are disjoint. Let $u \in M'$ be given. Define the children $C(u) = M \cap V_{M'}(u)$, those vertices of M that are partitioned into u's Voronoi cell. By condition (2) and Lemma 3.2

$$\forall v \in M,\ r_M(v) \geq \alpha^- f_N(v) \geq \frac{\alpha^-}{1+2\beta^+} f_{N'}(v) \tag{13}$$

Since $N' \subset M'$, $f_{N'} \geq f_{M'}$ everywhere.
Since $v \in V_{M'}(u)$, the definition of $f_{M'}$ implies that $f_{M'}(v) \geq r_{M'}(u)$.
Furthermore, since M is γ-scaffolding, conditions (1) and (2) yield $r_{M'}(u) \geq \frac{\gamma^-}{\gamma^+} R_{M'}(u)$
Combining these three facts into Equation (13) yields:

$$\forall\ v \in\ C(u),\ r_M(v) \geq \frac{\alpha^-\gamma^-}{(1+2\beta^+)\gamma^+} R_{M'}(u) \tag{14}$$

Let $r = \min_{v \in C(u)} r_M(v)$. Clearly $\mathcal{B}(v,r) \subset b_M(v)$, so the set $\{\mathcal{B}(v,r) \mid v \in C(u)\}$ are all disjoint.

We can now pack these disjoint balls into one large ball around u. Since $v \in V_{M'}(u) \subset B_{M'}(u)$, we have $\mathcal{B}(v,r) \subset \mathcal{B}(u, R_{M'}(u)+r)$.

Using $|\cdot|$ to denote the volume of a set, we have:

$$\Gamma_d(R_{M'}(u)+r)^d = |\mathcal{B}(u, R_{M'}(u)+r)| \geq |\bigcup_{v \in C(u)} \mathcal{B}(v,r)| = \sum_{v \in C(u)} |\mathcal{B}(v,r)| = |C(u)|\Gamma_d\, r^d \tag{15}$$

From which it follows that $(R_{M'}(u)+r)^d \geq |C(u)|r^d$.
If $r > R_{M'}(u)$, then the above simply reduces to $|C(u)| \leq 2^d$ by substituting out $R_{M'}(u)$.
If $r \leq R_{M'}(u)$, then we apply (14) and obtain:

$$2^d R_{M'}(u)^d \geq |C(u)|\left(\frac{\alpha^-\gamma^-}{(1+2\beta^+)\gamma^+}\right)^d R_{M'}(u)^d \qquad \text{or more simply:} \qquad |C(u)| \leq \left(\frac{2(1+2\beta^+)\gamma^+}{\alpha^-\gamma^-}\right)^d \tag{16}$$

Note that the second case is always a worse bound on $|C(u)|$, so we have $|C(u)| \leq C_{3.3}$; plugging in to $|M| \leq \sum_{u \in M'} |C(u)|$ will finish the proof. ∎

B Proof of Lemma 3.5

Lemma 3.5 Suppose N' is β-spacing-equivalent to a seeded surface mesh N with α-scaffolding volume mesh M, then we have:

$$|N'| \leq C_{3.5}|N| \qquad \text{where} \qquad C_{3.5} = \left(\frac{4\alpha^+}{\alpha^-\beta^-}\right)^d$$

Proof:

We will use the Voronoi cells of N to cover $\mathcal{S}$. Then we distribute the vertices of N' into these cells, and claim that no cell of N contains more than a constant number of points from N'. Note this packing uses condition (3) to guarantee that $N' = \cup_{n \in N} N' \cap V_M(n)$.

To form the packing argument, we will first lower-bound the size of the $b_{N'}(v)$ for $v \in N'$, which we know are disjoint.

Let $u \in N$. Define the children of $C(u) = N' \cap V_M(u)$, those vertices of N' that are partitioned into u's Voronoi cell.

Let $v \in C(u)$. We note by definition that for $v \in N'$, $r_{N'}(v) = 1/2 f_{N'}(v)$. Utilizing condition (7), we then have:

$$r_{N'}(v) \geq \frac{\beta^-}{2} f_N(v)$$

Since $v \in C(u) \subset V_N(u)$, we have $f_N(v) \geq r_N(u)$, so we are reduced to:

$$r_{N'}(v) \geq \frac{\beta^-}{2} r_N(u)$$

Lastly, conditions (1) and (2) yield $r_N(u) \geq \frac{\alpha^-}{\alpha^+} R_N(u)$. Combining, we find:

$$r_{N'}(v) \geq \frac{\beta^- \alpha^-}{2\alpha^+} R_N(u) \tag{17}$$

Let $r = \min_{v \in C(u)} r_{N'}(v)$. Clearly $\mathcal{B}(v, r) \subset b_{N'}(v)$, and the $\mathcal{B}(v, r)$ are all disjoint for $v \in C(u)$. We can now pack the disjoint balls into one large ball around u. Since $v \in V_N(u) \subset B_N(u)$, we have $\mathcal{B}(v, r) \subset \mathcal{B}(u, R_N(u) + r)$.

Using $| \cdot |$ to denote the volume of a set, we have:

$$\Gamma_d (R_N(u) + r)^d = |\mathcal{B}(u, R_N(u) + r)| \geq \left| \bigcup_{v \in C(u)} \mathcal{B}(v, r) \right| = \sum_{v \in C(u)} |\mathcal{B}(v, r)| = |C(u)| \Gamma_d \, r^d \tag{18}$$

So it follows:

$$(R_N(u) + r)^d \geq |C(u)| r^d \tag{19}$$

If $r > R_N(u)$, then the above simply reduces to $2^d \geq |C(u)|$ by substituting for $R_N(u)$. If $r \leq R_N(u)$, then we apply (17) and obtain:

$$2^d R_N(u)^d \geq |C(u)| \left(\frac{\beta^- \alpha^-}{2\alpha^+} \right)^d R_N(u)^d$$

or simply:

$$|C(u)| \leq \left(\frac{4\alpha^+}{\beta^- \alpha^-} \right)^d$$

Note that the second bound is always worse, so we have $|C(u)| \leq C_{3.5}$, so plugging in to $|N'| \leq \sum_{u \in N} |C(u)|$ will finish the proof. ■

C Proof of Well-Paced Theorem

This is a reproduction of a portion of [MPS08] that will appear in August at CCCG08, provided here for the convenience of interested reviewers.

The cost of going from well-paced to well-spaced

Running a meshing algorithm on a point set P will add Steiner points until the resulting set P' is well-spaced. The *cost of cleaning* a point set P, denoted by $Cost(P)$ is defined as $|P'|$, the size of the well-spaced output. The result of this section generalizes previous work on the linear cost of balancing quad trees [Moo95].

Theorem C.1 *If P is a θ-well-paced extension of Q, then $Cost(Q \cup P) = O(Cost(Q) + |P|)$.*

Proof: The proof will be by induction on $n = |P|$. Let $\text{lfs}^{(i)}$ be the local feature size function induced by $Q \cup \{p_1, \ldots, p_i\}$. Let $\Psi_i = c_1 \int_{x\in\Omega} \frac{1}{\text{lfs}^{(i)}(x)^d} dx$, where c_1 is the constant from the upper bound in Theorem 1. In general, c_1 will depend on the particular meshing algorithm used. Theorem 1 says that $Cost(Q \cup \{p_1, \ldots, p_i\}) \leq \Psi_i$ and $\Psi_0 = O(Cost(Q))$, the base of our induction.

By induction, we assume $\Psi_{n-1} \leq Cost(Q) + c_2(n-1)$ for some constant c_2. It will suffice to show that $\Psi_n - \Psi_{n-1} < c_2$. We can split the Ruppert sizing integral as follows.

$$\Psi_n = c_1 \int\limits_{x\in\Omega} \frac{1}{\text{lfs}^{(n)}(x)^d} dx \tag{20}$$

$$\leq \Psi_{n-1} + c_1 \int\limits_{x\in U} \frac{1}{\text{lfs}^{(n)}(x)^d} - \frac{1}{\text{lfs}^{(n-1)}(x)^d} dx \tag{21}$$

where $U \subseteq \Omega$ is the set of all points for which the local feature size was changed by the insertion of p_n. Let $R = r_{p_n}$. The following two inequalities hold for all $x \in U$, the first is trivial and the second follows from the definition of well-paced points.

$$\overset{(n)}{\text{lfs}}(x) \geq |p_n - x|, \text{ and} \tag{22}$$

$$\overset{(n-1)}{\text{lfs}}(x) \leq |p_n - x| + \frac{R}{\theta}. \tag{23}$$

We use these inequalities to compute the integral above using spherical coordinates. Since the integrand is positive everywhere, we can upper bound the integral by integrating over all of $\mathbb{R}^d$ instead of just U:

$$\Psi_n - \Psi_{n-1} \leq c_1 \int\limits_{x\in U} \frac{1}{(|x|)^d} - \frac{1}{(|x| + \frac{R}{\theta})^d} dV, \tag{24}$$

$$\leq c_1 \int_0^\infty \int\limits_{S_r} \left(\frac{1}{r^d} - \frac{1}{(r + \frac{R}{\theta})^d} \right) dA dr, \tag{25}$$

$$\leq c_1 s_d \int_0^\infty \left(\frac{1}{r^d} - \frac{1}{(r + \frac{R}{\theta})^d} \right) r^{d-1} dr, \tag{26}$$

where S_r is the sphere of radius r and s_d is the surface area of the unit d-sphere. In the ball of radius $\frac{R}{2}$ around p_n the lfs is at least $\frac{R}{2}$, so the contribution of this region is to Ψ_n at most some constant c_3.

$$\Psi_n - \Psi_{n-1} \leq c_3 + c_1 s_d \int_{\frac{R}{2}}^\infty \left(\frac{1}{r^d} - \frac{1}{(r + \frac{R}{\theta})^d} \right) r^{d-1} dr \tag{27}$$

By the change variable $yR/\theta = r$ and simplifying we get:

$$\Psi_n - \Psi_{n-1} \leq c_3 + c_1 s_d \int_{\frac{\theta}{2}}^\infty \left(\frac{(y+1)^d - y^d}{y(y+1)^d} \right) dy \tag{28}$$

$$\leq c_3 + c_1 s_d \sum_{i=0}^{d-1} \binom{d}{i} \int_{\frac{\theta}{2}}^\infty \frac{y^i}{y^{d+1}} dy \tag{29}$$

$$\leq c_3 + c_1 s_d d^2 \binom{d}{d/2} (2/\theta)^d \tag{30}$$

The last inequality follows from the fact that each integral is bounded by $d(2/\theta)^d$. Choosing c_2 larger than this constant completes the proof. ■

One interpretation of this theorem is that the amortized increase in the cost of cleaning a point set is constant if you add a θ-medial point.

Corollary C.2 *If Q is a well-spaced point set and P is a well-paced extension then $Cost(Q \cup P) = O(|Q| + |P|)$.*

Proof: Follows from the above theorem and the linear cost of cleaning points that are already well-spaced. ■