Smoothing and Cleaning up Slivers

Herbert Edelsbrunner,[*] Xiang-Yang Li,[†] Gary Miller,[‡] Andreas Stathopoulos,[§]
Dafna Talmor,[¶] Shang-Hua Teng,[||] Alper Üngör[†] and Noel Walkington[‡]

ABSTRACT

A sliver is a tetrahedron whose four vertices lie close to a plane and whose perpendicular projection to that plane is a convex quadrilateral with no short edge. Slivers are both undesirable and ubiquitous in 3-dimensional Delaunay triangulations. Even when the point-set is well-spaced, slivers may result. This paper shows that such a point set permits a small perturbation whose Delaunay triangulation contains no slivers. It also gives deterministic algorithms that compute the perturbation of n points in time $\mathrm{O}(n \log n)$ with one processor and in time $\mathrm{O}(\log n)$ with $\mathrm{O}(n)$ processors.

Keywords. Mesh generation, computational geometry, tetrahedral meshes, Delaunay triangulations, slivers, mesh smoothing, mesh clean-up.

1. INTRODUCTION

This paper presents a smoothing and clean-up algorithm for 3-dimensional Delaunay triangulations that removes all slivers. A necessary assumption of the algorithm is that the input triangles and tetrahedra have a bounded circumradius to shortest edge length ratio.

Mesh generation. A *mesh* of a geometric domain is a decomposition into primitive pieces called *elements*. The generation of meshes has a wide range of applications in physical simulation, computer graphics, computer vision, geographic information systems, and elsewhere. In this paper we consider meshes that decompose domains in $\mathbb{R}^3$ into tetrahedra.

The size and shape of the triangles and tetrahedra in the mesh are important because they influence the convergence and stability of numerical algorithms such as the finite element method [4; 13]. Given a constant $\varrho_0 > 0$, we say a triangle or tetrahedron has the *Ratio Property* $[\varrho_0]$ if its circumradius divided by the length of its shortest edge does not exceed ϱ_0. If the constant is either not important or clear from the context we will drop ϱ_0 from the notation. A mesh has the *Ratio Property* if all its triangles and tetrahedra have it, and a point set S has the *Ratio Property* if its Delaunay triangulation has it.

The most common tetrahedral meshes are Delaunay triangulations. Given a finite set $S \subseteq \mathbb{R}^3$ in general position, the *Delaunay triangulation* is the unique triangulation $\mathrm{Del}\, S$ that contains a tetrahedron iff all points other than the vertices of the tetrahedron lie outside the circumsphere of the tetrahedron. Fast algorithms for constructing Delaunay triangulations can be found in textbooks of computational geometry, for example [5]. Algorithms that maintain Delaunay triangulations under point insertions until the mesh has the Ratio Property $[\varrho_0]$ are given in [10; 11; 12].

Previous work. The ubiquity of slivers in 3-dimensional Delaunay triangulations has been documented in Cavendish, Field and Frey [1]. Eliminating slivers seems difficult though, and most meshing algorithms based on Delaunay triangulations can generate the Ratio Property but cannot guarantee the absence of slivers. Talmor [14] notes that even well-spaced points do not prevent slivers.

The first positive result on slivers was an algorithm by Chew [3] that eliminates slivers by adding new points to generate a uniformly dense mesh. In a recent breakthrough, Cheng *et al.* [2] showed how to assign weights to the points so that the weighted Delaunay triangulation is free of slivers without adding new points. Similar to the method in this paper, their algorithm assumes that the input data has the Ratio Property $[\varrho_0]$.

The work in this paper is related to techniques that improve the quality of a mesh in a post-processing operation. The recently adopted terminology in mesh generation refers to such an operation as *smoothing* if only the embedding of the mesh is affected, and as *clean-up* if the mesh is altered through local changes in the connectivity. Among the most popular smoothing techniques are Laplacian [6] and optimization-based smoothing [7]. Clean-up operations are usually coupled with smoothing techniques [8].

[*] Department of Computer Science, Duke University, Durham, NC 27708, and Raindrop Geomagic, Research Triangle Park, NC 27709.

[†] Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801.

[‡] Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213.

[§] Department of Computer Science, College of William and Mary, Williamsburg, VA 23187.

[¶] LMS-CADSI, Suite 104, 3150 Almaden Expwy, San Jose, CA 95118.

[||] Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, and Akamai Technologies, Cambridge, MA 02139.

Results. The main result of this paper is a mesh smoothing and clean-up algorithm that eliminates slivers without adding new points. Most slivers disappear after a mild perturbation of the point set. The remaining slivers migrate to the boundary where they can be peeled off or can be treated with boundary enforcement heuristics. As in [2] we consider periodic sets of the form $S = S_0 + \mathbb{Z}^3$, where $S_0 \subseteq [0,1)^3$ is a set of n points in the half-open unit cube. The algorithm assumes the input Delaunay triangulation has the Ratio Property $[\varrho_0]$ for a constant $\varrho_0 > 0$. The algorithm achieves the improvement in time $\mathrm{O}(n \log n)$ with a single processor and in time $\mathrm{O}(\log n)$ with $\mathrm{O}(n \log n)$ processors. The mathematical derivations needed to prove the algorithm correct owe their inspiration to the sliver exudation work of Cheng *et al.* [2]. The algorithm in this paper and its proof provide theoretical backing for the experimental observation that mesh smoothing combined with clean-up gives better results than either technique by itself [8].

Outline. Section 2 introduces perturbations and proves basic geometric results. Section 3 prepares the proof of the Sliver Theorem, which is presented in Section 4. Section 5 presents the algorithm that removes slivers by perturbing the points. Section 6 concludes the paper.

2. PRELIMINARIES

This section introduces perturbations and shows that many geometric properties are preserved after a mild perturbation.

Linear relations. We adopt the notation of [2] and call two real quantities X and Y *linearly related*, denoted as $X \sim Y$, if there are constants $c, C > 0$ with $cX \leq Y \leq CX$.

For a sample application of this concept consider a tetrahedron $pqrs$. Denote the volume of $pqrs$ by V and the length of the shortest edge by L. We use, as in [2], $\sigma = \sigma(pqrs) = \frac{V}{L^3}$ as a measure of shape quality of $pqrs$. Many types of tetrahedra can have small value of σ, but only slivers can simultaneously have the Ratio Property and a small value of σ. It is useful to relate this measure of shape quality with a distance to radius ratio defined for $pqrs$. Let D be the distance of point p to the plane passing through qrs, and let Y be the circumradius of qrs. Figure 1 illustrates these definitions. As proved in [2], we have $c_1 \frac{D}{Y} \leq \sigma \leq C_1 \frac{D}{Y}$ for $c_1 = 1/96\varrho_0^3$

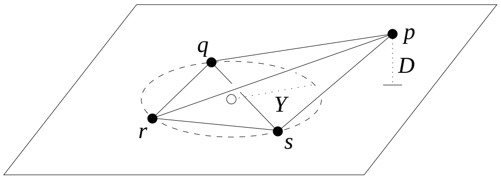

Figure 1: Y is the radius of the circle passing through q, r, s and D is the distance of p from the plane of that circle.

and $C_1 = \pi\varrho_0^3/3$. We state this result for later reference.

QUALITY LEMMA. If $pqrs$ has the Ratio Property then $\frac{D}{Y} \sim \sigma$.

Observe that the ratio $\frac{D}{Y}$ depends on the ordering of the vertices, and more specifically on which vertex is ordered first. The Quality Lemma implies that all four possible ratios are linearly related to σ and therefore to each other.

Perturbation. Let S be a set of n points in $\mathbb{R}^3$. A *perturbation* is a function $' : S \to \mathbb{R}^3$. We denote the image of S by S' and the image of a point $p \in S$ by p'. The effect of a perturbation on S is quantified using a function $N = N_S : \mathbb{R}^3 \to \mathbb{R}$ that maps every point $x \in \mathbb{R}^3$ to the distance from x to the second closest point in S. For example, if $x \in S$ then the smallest distance is zero and $N(x) = N_S(x)$ is the distance to the nearest neighbor. Note that N is continuous and Lipschitz with constant one, that is, $|N(x) - N(y)| \leq \|x - y\|$. The *impact* of the perturbation is the smallest real number φ such that $\|p - p'\| \leq \varphi N(p)$ for every $p \in S$. It follows that the sphere of radius $\varphi N(p)$ around p contains p'. In this paper we only consider *mild* perturbations with impact at most some constant $\varphi_0 < \frac{1}{2}$. For mild perturbations, the spheres around the points are pairwise disjoint.

The distances between two points before and after a mild perturbation are linearly related. More precisely, $c_2\|p - q\| \leq \|p' - q'\| \leq C_2\|p - q\|$ with $c_2 = 1 - 2\varphi_0$ and $C_2 = 1 + 2\varphi_0$. It is easy to prove these bounds as we may assume that the two points are the nearest neighbors of each other, that is, $N(p) = N(q) = \|p - q\|$. We state this result for later reference.

DISTANCE LEMMA. If S' is the image of a mild perturbation of S then $\|p - q\| \sim \|p' - q'\|$ for all $p, q \in S$.

Consider the function $N' = N_{S'}$ that maps every $x \in \mathbb{R}^3$ to the distance from x to the second closest point in S'. For mild perturbations, N and N' are linearly related.

N LEMMA. If S' is the image of a mild perturbation of S then $N'(x) \sim N(x)$ for every $x \in \mathbb{R}^3$.

PROOF. We establish $c_3 N(x) \leq N'(x) \leq C_3 N(x)$ for constants $c_3 = 1 - 2\varphi_0$ and $C_3 = 1 + 2\varphi_0$. By definition of mildness,

$$\|x - p'\| \quad \geq \quad \|x - p\| - \varphi_0 N(p)$$

for each $p \in S$. Using the Lipschitz property we get $N(p) \leq N(x) + \|x - p\|$ and therefore

$$\|x - p'\| \quad \geq \quad (1 - \varphi_0)\|x - p\| - \varphi_0 N(x).$$

We have $\|x - p\| \geq N(x)$ except if p is closest to x, and therefore $\|x - p'\| \geq (1 - 2\varphi_0)N(x)$ for all points p other than the closest. This establishes the lower bound. A symmetric argument proves the claimed upper bound. $\square$

3. RATIO, LENGTH, DEGREE

This section recalls a result of Talmor [14] and uses it to extend a result of Miller *et al.* [11] from Delaunay triangulations to complexes formed by taking unions of Delaunay triangulations.

Gap Property. It is convenient to denote a sphere with center $z \in \mathbb{R}^3$ and radius $Z \in \mathbb{R}$ by the pair (z, Z). We call a sphere *empty* of points in S if it encloses none of the points of S. In other words, every point of S either lies on or outside the sphere. Given a constant γ_0, we say S has the *Gap Property* $[\gamma_0]$ if every empty sphere has radius at most $\gamma_0 N(x)$, where x is any point on the sphere. We prove that perturbations with small impact preserve the Gap Property. Specifically, we define $\varphi_1 = 1/(4\gamma_0 + 2)$ and restrict ourselves to *very mild* perturbations whose impact is $\varphi \leq \varphi_1$.

GAP PRESERVATION LEMMA. If S has the Gap Property $[\gamma_0]$ and S' is the image of a very mild perturbation then there is a constant γ_0' depending only on γ_0 such that S' has the Gap Property $[\gamma_0']$.

PROOF. Let x be a point and (z, Z) a sphere passing through x that is empty of points in S'. We prove below that $Z \leq \gamma_0' N'(x)$, for $\gamma_0' = 2\gamma_0 + 2 + 1/2\gamma_0$. Figure 2 illustrates the argument which follows.

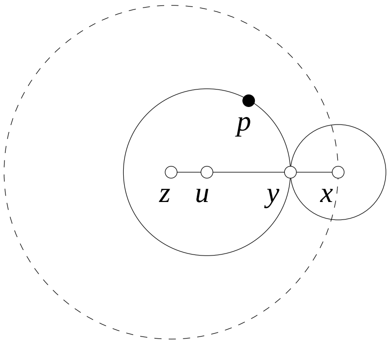

Figure 2: The solid spheres are empty of points in S and the dashed sphere is empty of points in S'.

We may assume that $Z > N(x)$. Then the line segment from x to z intersects the sphere $(x, N(x))$ in a point, which we denote as y. Let $\vec{xz}$ denote the half-line that starts at x and passes through z, and let the point $u \in \vec{xz}$ be furthest from x with the property that the sphere $(u, U = \|u - y\|)$ is empty of points in S. Since S has the Gap Property $[\gamma_0]$, we have $U \leq \gamma_0 N(y)$, and because of the Lipschitz property we have $\gamma_0 N(y) \leq 2\gamma_0 N(x)$. Let $p \in S$ be a point on the sphere (u, U). Since p and y both lie on this sphere we have $\|p - y\| \leq 2U \leq 4\gamma_0 N(x)$, and therefore

$$\begin{aligned} N(p) &\leq N(x) + \|p - x\| \\ &\leq 2N(x) + \|p - y\| \\ &\leq (4\gamma_0 + 2)N(x) \\ &= N(x)/\varphi_1. \end{aligned}$$

The distance between p and its perturbed image p' is therefore at most $N(x)$. But since p' lies on or outside (z, Z) by assumption, the radius of the sphere around u is $U \geq \|z - y\|$. This implies $Z \leq U + N(x)$, and by the N Lemma,

$$\begin{aligned} Z &\leq (2\gamma_0 + 1)N(x) \\ &\leq \frac{2\gamma_0 + 1}{1 - 2\varphi_1} N'(x) \\ &\leq (2\gamma_0 + 2 + \frac{1}{2\gamma_0})N'(x), \end{aligned}$$

which proves the claim. $\Box$

Ratio Property. Talmor proves that if S has the Ratio Property $[\varrho_0]$ then there is a constant γ_0 depending only on ϱ_0 such that S has the Gap Property $[\gamma_0]$. The bound on the constant γ_0 is rather pessimistic but it exists. Symmetrically, Talmor proves that if S' has the Gap Property $[\gamma_0']$ then there is a constant ϱ_0' depending only on γ_0' such that S' has the Ratio Property $[\varrho_0']$. In short, a finite set has the Ratio Property iff it has the Gap Property. Together with the Gap Preservation Lemma this implies that very mild perturbations also preserve the Ratio Property.

RATIO PRESERVATION LEMMA. If S has the Ratio Property $[\varrho_0]$ and S' is the image of a very mild perturbation then there is a constant ϱ_0' depending only on ϱ_0 such that S' has the Ratio Property $[\varrho_0']$.

Union of triangulations. Next we consider a simplicial complex K obtained by mapping the Delaunay triangulation of very mild perturbations back to S. Specifically, a tetrahedron $pqrs$ belongs to K if there is a very mild perturbation such that $p'q'r's'$ is a tetrahedron in the Delaunay triangulation of S'. The complex K furthermore contains all triangles, edges, vertices of these tetrahedra.
Miller *et al.* [11] prove that if S has the Ratio Property $[\varrho_0]$ then the length of edges sharing a common vertex is bounded from above by some constant ν_0 depending only on ϱ_0: $\|p - q\|/\|p - u\| \leq \nu_0$ for all edges pq, pu in $\operatorname{Del} S$. We show that the same is true for the edges in K, but for a possibly different constant ν_1.

LENGTH VARIATION LEMMA. If S has the Ratio Property $[\varrho_0]$ then there exists a constant ν_1 depending only on ϱ_0 and φ_1 such that $\|p - q\|/\|p - u\| \leq \nu_1$ for all edges pq, pu in K.

PROOF. We prove the claim for $\nu_1 = \nu_0' \cdot \frac{1+2\varphi_1}{1-2\varphi_1}$, where ν_0' is the constant bound on the length variation for a point set S' that has the Ratio Property $[\varrho_0']$.
We simplify the argument by assuming the distance between p and its nearest neighbor in S is $N(p) = 1$. The Distance Lemma implies $1 - 2\varphi_1 \leq N'(p') \leq 1 + 2\varphi_1$. $N'(p')$ is also the length of the shortest edge with endpoint p' in $\operatorname{Del} S'$. The Ratio Preservation Lemma implies that S' has the Ratio Property $[\varrho_0']$, so there exists a constant upper bound ν_0' on the length variation around p'. The length of the longest edge with endpoint p' in $\operatorname{Del} S'$ is $\|p' - v'\| \leq \nu_0'(1 + 2\varphi_1)$. The length of the corresponding edge in $\operatorname{Del} S$ is $\|p - v\| \leq \nu_0' \cdot \frac{1+2\varphi_1}{1-2\varphi_1}$, which proves the claimed upper bound. $\Box$
A volume argument can now be used to prove an upper bound on the degree of vertices in K.

DEGREE LEMMA. If S has the Ratio Property $[\varrho_0]$ then there exists a constant δ_1 such that every vertex belongs to at most δ_1 edges in K.

Specifically, [2; 11] prove that the vertex degree is bounded from above by $\delta_1 = (2\nu_1^2 + 1)^3$. As a consequence of the Degree Lemma, each vertex belongs to at most a constant number of tetrahedra in K.

4. SLIVER THEOREM

Sections 2 and 3 provide the technical prerequisites for the proof of the Sliver Theorem II, which motivates the algorithm for sliver removal presented in Section 5. We begin by studying conditions under which a perturbation creates slivers.

Circles and Tori. Let (z, Z) be the circumsphere of a tetrahedron $pqrs$ and let (y, Y) be the circumcircle of the triangle qrs. If $pqrs$ has the Ratio Property $[\varrho_0]$ then $Z/2\varrho_0 \leq Y \leq Z$. Call $pqrs$ a *sliver* if $\sigma(pqrs) = \frac{V}{L^3} < \sigma_0$, where $\sigma_0 > 0$ is a constant that we specify later. We now prove that if $pqrs$ is a sliver, then the distance P from p to the closest point on the circle (y, Y) cannot be large.

TORUS LEMMA. If $pqrs$ is a sliver and has the Ratio Property $[\varrho_0]$ then there exists a constant C_4 such that $P \leq C_4 \sigma_0 Z$.

PROOF. We prove the claim for $C_4 = \frac{c+1}{c \cdot c_1}$, where $c = 1 - \sqrt{1 - 1/4\varrho_0^2}$, which is a positive constant less than 1. Let H be the distance from z to the plane of qrs, and let α be the angle between qy and qz, as shown in Figure 3. Then $Y = Z \cos\alpha$, which implies $\cos\alpha = \frac{Y}{Z} \geq \frac{1}{2\varrho_0}$. We also have $H = Z \sin\alpha \leq Z(1-c)$. Next consider p and its closest

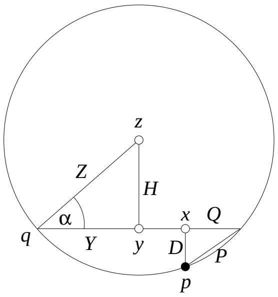

Figure 3: The circumsphere (z, Z) of $pqrs$ and the circumcircle (y, Y) of qrs projected along a direction in the plane of qrs.

point x on the plane of qrs. By the Quality Lemma, the distance from p to x is $D \leq \sigma_0 Y / c_1$. Let Q bet the distance from x to its closest point on the circle (y, Y). We have $\frac{D}{Q} \geq \frac{Z-H}{Y} \geq \frac{Z-H}{Z} \geq c$. The distance from p to its closest point on (y, Y) is therefore

$$P \;\leq\; D + Q \;\leq\; D(1 + \frac{1}{c}),$$

which proves the claim. □

Forbidden regions. Keeping q, r, s fixed, the *forbidden region* F_{qrs} is the set of points p such that $pqrs$ is a sliver and has the Ratio Property $[\varrho_0']$. By the Torus Lemma, F_{qrs} is contained in the solid torus of points at distance at most P from the circle (y, Y), as illustrated in Figure 4. The volume of that torus is the perimeter of the circle times the area of the sweeping disk, which is $2\pi Y \cdot \pi P^2$.

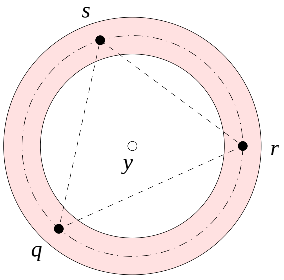

Figure 4: Every triangle qrs that forms a tetrahedron with p defines a forbidden region inside a torus of points around the circumcircle of qrs.

FORBIDDEN VOLUME LEMMA. If $pqrs$ has the Ratio Property $[\varrho_0]$ then there is a constant C_5 that depends only on ϱ_0 such that the volume of the forbidden region F_{qrs} is at most $C_5 N^3(p)$.

PROOF. We prove the claim for $C_5 = 2\pi^2 C_4^2 \sigma_0^2 \varrho_0^3 \nu_1^3$. The Torus Lemma implies that the volume of F_{qrs} is

$$\begin{aligned} V &\leq 2\pi Y \cdot \pi P^2 \\ &\leq 2\pi^2 C_4^2 \sigma_0^2 Z^3. \end{aligned}$$

Let L be the length of the shortest among the edges pq, pr, ps. The Length Variation Lemma implies $L \leq \nu_1 N(p)$, and the Ratio Property implies $Z \leq \varrho_0 L \leq \varrho_0 \nu_1 N(p)$. The claimed upper bound on the volume follows. □

Existence. It remains to put all pieces together and prove the existence of a very mild perturbation that removes all slivers. The number II differentiates the following result from the theorem with the same name in [2].

SLIVER THEOREM II. If S has the Ratio Property $[\varrho_0]$ then there is a constant $\sigma_0 > 0$ and a very mild perturbation with $\sigma \geq \sigma_0$ for all tetrahedra in $\mathrm{Del}\, S'$.

PROOF. We establish the result for the constant

$$\sigma_0 \;=\; \sqrt{\frac{2}{3\pi C_4^2} (\frac{\varphi_1}{\varrho_0 \nu_1 \delta_1})^3}.$$

We determine the image of each point in some fixed but arbitrary sequence. While considering $p \in S$, we keep the locations of all other points fixed. The image p' of p is chosen inside the sphere with center p and radius $\varphi_1 N(p)$ and outside the forbidden regions F_{qrs} for all triangles qrs that form a tetrahedron with p in K. To prove that such a point exists, we show that the inside of the sphere is not covered by the forbidden regions, as suggested by Figure 5. This is

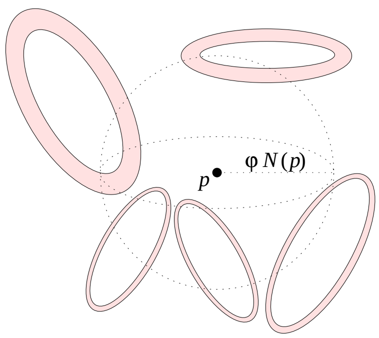

Figure 5: The forbidden regions of p do not cover the entire inside of the ball of possible images p'.

done by comparing the volume of that inside with the sum of volumes of the forbidden regions. By the Forbidden Volume Lemma, each triangle qrs defines a forbidden region with volume at most $C_5 N^3(p)$. The Degree Lemma implies that there are fewer than δ_1^3 triangles qrs. The volume covered by the forbidden regions is therefore less than $\delta_1^3 C_5 N^3(p)$. With σ_0 as defined above, this is less than the volume of the inside of the sphere, which is $\frac{4\pi}{3} \cdot \varphi_1^3 N^3(p)$. □

5. ALGORITHM

The proof of the Sliver Theorem II is constructive and implies an algorithm that computes the perturbation whose existence is asserted by the theorem. That algorithm assumes the set $S = S_0 + \mathbb{Z}^3$ has the Ratio Property $[\varrho_0]$. It consists of three steps:

1. computing the union complex K of S,
2. perturbing the points in S_0 to get $S' = S_0' + \mathbb{Z}^3$,
3. constructing the Delaunay triangulation of S'.

Since S contains infinitely many points, we can of course not compute K explicitly, but we can compute a piece that determines the rest by translation along integer vectors. The same applies to computing the Delaunay triangulation of S'. Computing K is difficult, but we can efficiently construct a complex K' that contains K as a subcomplex. For every point $p \in S$, the new complex K' contains all edges, triangles, and tetrahedra that connect p with all points $q \in S$ at distance at most $\nu_1 N(p)$ from p. The Length Variation Lemma implies $K \subseteq K'$. The volume argument proving the Degree Lemma still applies and proves a constant upper bound on the maximum vertex degree in K'. Each tetrahedron sharing p defines a forbidden region, which we represent by the containing torus described in the Torus Lemma. Since there is only a constant number of tori, we can find in constant time a point p' inside the sphere with center p and radius $\varphi_1 N(p)$ and outside all tori.

Using the divide-and-conquer algorithm in [9], the complex K' for n points can be computed in time $O(n \log n)$ with one processor and in time $O(\frac{n \log n}{k})$ with $k < n$ processors. The single processor version implies an $O(n \log n)$ time implementation of the entire algorithm. For the parallel version, we also need to perturb the points in parallel. This can be done by coloring the vertices of K' with a constant number of colors such that no two adjacent vertices receive the same color. A whole color class of points is then perturbed concurrently, which leads to a parallel implementation of the algorithm that runs in time $O(\frac{n \log n}{k})$ using $k < n$ processors.

The algorithm sketched above is probably too complicated to have any practical relevance. To simplify, we would have to eliminate Step 1, as in the weight assignment algorithm of [2]. There is however a subtle difference between increasing the weight of a point p and moving p to a new location p': in the former case all new simplices share p while in the latter case the motion may create tetrahedra connecting vertices all different from p. Moving p might introduce a sliver connecting other points, and if these other points are already perturbed images of original points, then we need to backtrack to repair the damage, which takes time. We leave the formulation of an efficient deterministic algorithm that avoids the construction of K as an open question.

6. DISCUSSION

Inspired by the sliver exudation method in [2], this paper shows that slivers in a 3-dimensional Delaunay triangulation with the Ratio Property can be removed by a small perturbation of the points. One serious drawback of the current result are the pessimistic estimates of the constants used to express it. These constants are either too large or too small for the results to have any practical significance. The most serious deterioration of the constant estimates happens in the transition from the Ratio Property to the degree bound in the Degree Lemma. It might be worthwhile to design algorithms that permit a direct analysis of vertex degrees. An example of such an algorithm is the biting method of Li *et al.* [10], which produces a mesh from a sphere covering. Alternatively, we may gain insights by experimenting with an implementation of the algorithm and measuring its average and worst-case performance. Another drawback is that the result holds only sufficiently far from the domain boundary, if any. It would be interesting to extend the method to include possibly constrained perturbations of points on the boundary of a meshed domain.

7. REFERENCES

[1] J. C. Cavendish, D. A. Field and W. H. Frey. An approach to automatic three-dimensional finite element mesh generation. *Internat. J. Numer. Methods Engrg.* **21** (1985), 329–347.

[2] S.-W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello and S.-H. Teng. Sliver exudation. *J. Assoc. Comput. Mach.*, to appear.

[3] L. P. Chew. Guaranteed-quality Delaunay meshing in 3D. Short paper in "Proc. 13th Ann. Sympos. Comput. Geom., 1997", 391–393.

[4] P. G. Ciarlet. *The Finite Element Method for Elliptic Problems.* North-Holland, 1978.

[5] M. de Berg, M. van Kreveld, M. Overmars and O. Schwarzkopf. *Computational Geometry. Algorithms and Applications.* Springer-Verlag, Berlin, Germany, 1997.

[6] D. A. Field. Laplacian smoothing and Delaunay triangulations. *Comm. Appl. Numer. Meth.* **4** (1988), 709–712.

[7] L. Freitag, M. Jones and P. Plassmann. An efficient parallel algorithm for mesh smoothing. *In* "Proc. 4th Internat. Meshing Roundtable, 1995", 47–58.

[8] L. Freitag and C. Ollivier-Gooch. Tetrahedral mesh improvement using swapping and smoothing. *Internat. J. Numer. Methods Engin.* **40** (1997), 3979–4002.

[9] A. M. Frieze, G. L. Miller and S.-H. Teng. Separator based divide and conquer in computational geometry. *In* "Proc. 5th ACM Sympos. Parallel Algor. Arch., 1992", 420–430.

[10] X.-Y. Li, S.-H. Teng and A. Üngör. Biting: advancing front meets sphere packing. *Internat. J. Numer. Methods Engin.*, to appear.

[11] G. L. Miller, D. Talmor, S.-H. Teng and N. Walkington. A Delaunay based numerical method for three dimensions: generation, formulation and partition. *In* "Proc. 27th Ann. ACM Sympos. Theory Comput. 1995", 683–692.

[12] J.R. Shewchuck. Tetrahedral mesh generation by Delaunay refinement. *In* "Proc. 14th Ann. Sympos. Comput. Geom. 1998", 86–95.

[13] G. Strang and G. J. Fix. *An Analysis of the Finite Element Method.* Prentice Hall, Englewood Cliffs, New Jersey, 1973.

[14] D. Talmor. Well-spaced points for numerical methods. Report CMU-CS-97-164, Dept. Comput. Sci., Carnegie-Mellon Univ., Pittsburgh, Penn., 1997.