Incremental Topological Flipping Works for Regular Triangulations*

H. Edelsbrunner and N. R. Shah
Dept. Comput. Sci., Univ. Illinois, Urbana, IL 61801, USA.

bstract

set of n weighted points in general position in $\Re^d$ de-
nes a unique regular triangulation. This paper proves
at if the points are added one by one then flipping in
topological order will succeed in constructing this tri-
ngulation. If, in addition, the points are added in a
ndom sequence and the history of the flips is used for
cating the next point, then the algorithm takes expected
me at most $O(n \log n + n^{\lceil d/2 \rceil})$. The second term is of
e same order of magnitude as the maximum number
f possible simplices.

Introduction

elaunay triangulations, and their dual Voronoi diagrams, play an important role in a variety of different isciplines of science (see e.g. the survey of Aurenhammer [2]). The computational aspects of Delaunay triangulations have been studied in the area of geometric lgorithms [6, 18], and a large number of different construction algorithms have been produced. This paper onsiders the class of regular triangulations which includes the Delaunay triangulations [17]. A finite point et in $\Re^d$ defines a unique Delaunay triangulation, but here are many regular triangulations of the set. A nique regular triangulation is implied if each point is ssigned a real number as its weight. If all weights are he same then the regular triangulation is the Delaunay triangulation of the set.

NO way !!!

*Research of both authors was supported by the National Science Foundation under grant CCR-8921421 and under the Alan T. Waterman award, grant CCR-9118874. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the view of the National Science Foundation.

8th Annual Computational Geometry, 6/92, Berlin, Germany
©1992 ACM 89791-518-6/92/0006/0043 .... $1.50

Several algorithms proposed for Delaunay triangulations are based on the notion of a local transformation henceforth referred to as a flip. Historically the first such algorithm can be attributed to Lawson [14], see also [15]. Given a finite point set in the real plane, $\Re^2$, the algorithm first constructs an arbitrary triangulation of the set. This triangulation is then gradually altered through a sequence of edge-flips until the Delaunay triangulation is obtained. The generalization of this method to $\Re^3$ has difficulties, and Joe [12] demonstrates that it is indeed incorrect if the flips are applied to an arbitrary initial triangulation. In a different paper, Joe [13] shows that if a single point, p, is added to the Delaunay triangulation of a set S in $\Re^3$ then many different sequences of flips will succeed in constructing the Delaunay triangulation of $S \cup \{p\}$. This can be used as the basis of an incremental algorithm. Rajan [20] considers Delaunay triangulations in arbitrary dimensions, $\Re^d$, and argues that a single point can always be added by a sequence of flips. However, he needs a priority queue to find the appropriate sequence, which takes logarithmic time per flip. On a different front, Guibas, Knuth and Sharir [11] study the complexity of the incremental algorithm in $\Re^2$ when the points are added in a random sequence. While $\Theta(n^2)$ edge-flips are required in the worst case, they prove that under a random insertion sequence the expected number of flips is only $O(n)$. They also provide an elegant and in the expected sense efficient technique for the initial location of the point to be added. This step is a sore point of all incremental methods.

This paper unifies and extends the algorithmic results of Joe [13], Rajan [20], and Guibas, Knuth, Sharir [11]. In particular, we show that many different sequences of flips can be used to add a single point to a regular triangulation in $\Re^d$. This eliminates the need for a priority queue that sorts flips. The priority queue is replaced by a stack that computes a topological ordering of the flips in constant amortized time per flip. We use this result to generalize the incremental method of [11] to regular triangulations in $\Re^d$. The resulting algorithm runs in expected time $O(n \log n + n^{\lceil d/2 \rceil})$. The expectation is computed over all input sequences of the same

n points and is thus independent of the point distribution. However, the actual expected time, which could be much less, does depend on the distribution. Without assumption on the distribution, we cannot expect any better time bound because there are sets of n points in $\Re^d$ with regular triangulations that consist of $\Theta(n^{\lceil d/2 \rceil})$ simplices.

Outline. Section 2 defines regular triangulations and introduces related terminology. Section 3 explains the relationship between regular triangulations in $\Re^d$ and convex hulls in $\Re^{d+1}$. This relationship provides a sometimes enlightening alternative view of all concepts and techniques discussed in this paper. Section 4 discusses the anatomy of flipping in $\Re^d$. A counterexample to a non-incremental method that attempts to construct regular triangulations in $\Re^2$ by flipping is presented in section 5. A minimalist data structure for storing triangulations is described in section 6. The incremental algorithm is given in section 7, and its correctness is proved in section 8. Section 9 analyzes the algorithm under the assumption of a random insertion sequence and derives the complexity result mentioned earlier.

2 Regular Triangulations

Triangulations. We begin by defining the notion of a triangulation used in this paper. For $0 \leq k \leq d$, the convex hull of a set T of $k+1$ affinely independent points is a *k-simplex*, denoted by Δ_T. A collection of simplices, $\mathcal{C}$, is a *simplicial cell complex* if it satisfies the following two conditions.

(i) If $\Delta_T \in \mathcal{C}$ then $\Delta_U \in \mathcal{C}$ for all $U \subseteq T$.

(ii) If $\Delta_T, \Delta_{T'} \in \mathcal{C}$ then $\Delta_T \cap \Delta_{T'} = \Delta_{T \cap T'}$.

Condition (ii) implies that the intersection of any two simplices in the cell complex is either empty or a face of both. If it is a face then condition (i) implies that it also belongs to the cell complex. The *underlying space* of $\mathcal{C}$ is the pointwise union of its simplices. Let S be a finite point set in $\Re^d$. Usually, a *triangulation* of S is defined as a simplicial cell complex so that S is the set of 0-simplices (vertices) and the underlying space of the complex is the convex hull of S. It will be convenient to relax the first condition and to only require that the set of vertices is a subset of S. Notice that the second condition implies that all extreme points of S are vertices of every triangulation of S.

Power distance and power diagrams. Again, let S be a finite set of points in $\Re^d$, and assign a real valued weight w_p to each point $p \in S$. For each p, define $\pi_p : \Re^d \to \Re$ so that

$$\pi_p(x) = |xp|^2 - w_p,$$

where $|xp|$ is the Euclidean distance between points $x = (x_1, x_2, \ldots, x_d)$ and $p = (p_1, p_2, \ldots, p_d)$, given by $|xp| = \sqrt{\Sigma_{i=1}^d (x_i - p_i)^2}$. We call $\pi_p(x)$ the *power distance* of x from p. It is easy to see that for points $p, q \in S$, the locus of points $x \in \Re^d$ with $\pi_p(x) = \pi_q(x)$ is the hyperplane

$$\chi_{p,q} : 2\sum_{i=1}^{d} x_i(q_i - p_i) + \sum_{i=1}^{d}(p_i^2 - q_i^2) - w_p + w_q = 0.$$

We call $\chi_{p,q}$ the *chordale* of the weighted points p and q.

Sometimes it is convenient to interpret a point p with weight w_p as a sphere centered at p and with radius $\sqrt{w_p}$. If w_p is negative we obtain a sphere with imaginary radius.

Let $H_{p,q}$ denote the half-space of points $x \in \Re^d$ for which $\pi_p(x) \leq \pi_q(x)$. For each $p \in S$, define its *power cell* as

$$P(p) = \bigcap_{q \in S - \{p\}} H_{p,q}.$$

Observe that $P(p)$ is a convex polyhedron, the intersection of the interiors of any two power cells is empty, and the union of all power cells $P(p)$, $p \in S$, covers $\Re^d$. The collection of power cells and their faces defines the cell complex $\mathcal{P}(S)$ in $\Re^d$, known as the *power diagram* of S, see e.g. [1].

Orthogonal centers. For the remainder of the paper, we assume that the weighted points of S are in general position. This involves no loss of generality since we can use the method in [9] to computationally simulate this assumption. General position, in this context, means that for every $d+1$ weighted points in S, there is a unique unweighted point $x \in \Re^d$ with the same power distance from all $d+1$ points, and for every $d+2$ weighted points of S, there is no such point. Two weighted points p and z are said to be *orthogonal* if

$$|pz|^2 = w_p + w_z.$$

Note that this is equivalent to $\pi_p(z) = w_z$ and $\pi_z(p) = w_p$. A subset T of $d+1$ (weighted) points of S defines a unique d-simplex $\Delta = \Delta_T = \text{conv}(T)$. There is a unique weighted point $z = z(\Delta)$ that is orthogonal to all weighted points $p \in T$. We call z the *orthogonal center* of Δ. If the weights of all $p \in T$ are zero then the sphere with center z and radius $\sqrt{w_z}$ is the circumsphere of Δ.

Local and global regularity. Observe that $\pi_z(p) = w_p$ for all $p \in T$. Call Δ *(globally) regular* if $\pi_z(q) > w_q$ for all $q \in S - T$. Clearly, if Δ is regular then z is a vertex of $\mathcal{P}(S)$, the power diagram of S. The regular d-simplices, together with their faces, define a simplicial cell complex known as the *regular triangulation* of S,

How do you find w_z? and z?

How do you know the existence???? Yes, Nice!

Easy, is where several chordals intersect!

a line-shelling of its faces, see [3]. Computing the line-shelling is closely related to computing the ordering of the flips for Delaunay triangulations mentioned in [20]. The result of this paper can also be interpreted as finding a "topological line-shelling".

4 Flipping in d Dimensions

Definition and classification of flips. Consider a set T of $d+2$ points in $\Re^d$. According to Lawson [16], there are exactly two ways to triangulate T. Indeed, the two ways correspond to the two sides (lower and upper) of the $(d+1)$-simplex that is the convex hull of the corresponding lifted points in $\Re^{d+1}$. Because of Radon's theorem (below) and because the $(d+1)$-simplex exhausts all $d+2$ d-simplices as facets, there can be no other triangulation of T. A *flip* is the operation that replaces one triangulation of T with the other.

In $\Re^2$, we distinguish two cases depending on whether the tetrahedron of the lifted points in $\Re^3$ projects to a triangle or a quadrilateral, see Figure 4.1. A 4-simplex

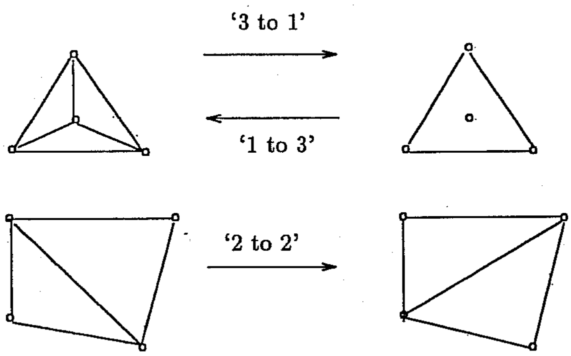

Figure 4.1: There are three types of flips in $\Re^2$, and we denote a flip by the number of triangles before and after the flip. So the flips are of type '1 to 3', '2 to 2', and '3 to 1'. The first type introduces a new point, and the last type removes a point. The last type of flip is not needed for Delaunay triangulations because no point is redundant, and so no point has to be removed from the triangulation.

in $\Re^4$ projects to a single or a double tetrahedron (the convex hull of four or five points) in $\Re^3$. Flips in $\Re^3$ are classified accordingly, see Figure 4.2.

Given $d+2$ weighted points in $\Re^d$, one of the two triangulations is the regular triangulation of the points, the other is not regular. In the construction of $\mathcal{R}(S)$, flips are applied in this directional sense, replacing a non-regular triangulation of $d+2$ points by the regular one.

Flippability. Let $\Delta = \Delta_U$ be a $(d-1)$-simplex of an arbitrary triangulation $\mathcal{T}$ of S, and let $\Delta' = \Delta_{U\cup\{a\}}$ and $\Delta'' = \Delta_{U\cup\{b\}}$ be the two incident d-simplices. The

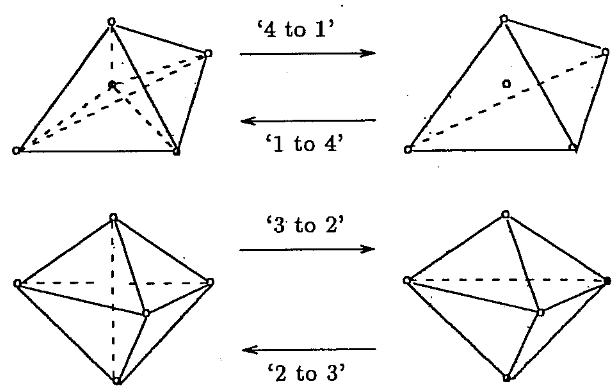

Figure 4.2: The flips in $\Re^3$ can be classified as '4 to 1', '3 to 2', '2 to 3', and '1 to 4'.

induced subcomplex of $T = U \cup \{a, b\}$ consists of all simplices in $\mathcal{T}$ spanned by points in T. Clearly, Δ, Δ' and Δ'' belong to the induced subcomplex of T. We call T (and Δ) *flippable* if conv(T) is the underlying space of the induced subcomplex of T.

Assume that if Δ is given then Δ' and Δ'', and therefore T, can be computed in constant time. This requires that d be a constant. Consider the d $(d-2)$-simplices of Δ. Call such a $(d-2)$-simplex *convex* if there is a hyperplane that contains it and Δ' and Δ'' lie on the same side of this hyperplane; otherwise, call the $(d-2)$-simplex *reflex*. The underlying space of the induced subcomplex of T is equal to conv(T) iff all reflex $(d-2)$-simplices of Δ have degree 3, that is, each is incident to exactly three $(d-1)$-simplices. These are Δ and one other $(d-1)$-simplex each of Δ' and Δ''. Thus, given Δ, it is possible to test in constant time whether or not it is flippable. Recall, however, that our algorithm would attempt to flip Δ only if it is flippable and it is locally non-regular.

A convex geometry theorem. The above discussion is closely related to a classical result in convex geometry known as Radon's theorem [19].

Theorem 4.1 Let T be a set of $d+2$ points in $\Re^d$. Then there exists a partition $T = U \dot\cup V$ so that $\text{conv}(U) \cap \text{conv}(V) \neq \emptyset$.

Every subset of $d+1$ points of T either contains U as a subset, or V, but not both. In the former case, the d-simplex spanned by the subset belongs to the triangulation of T that contains conv(U) as a simplex. In the latter case, it belongs to the other triangulation that contains conv(V).

Assume that $|U| \leq |V|$. Then $|U| \leq \frac{d+2}{2}$ and Δ_U is a k-simplex with $k = |U|-1 \leq \frac{d}{2}$. When T is flipped then Δ_U can only belong to one of the two triangulations of T, the one before or the one after the flip. This implies mentally, that is, points are added one at a time. It is convenient to first construct an artificial d-simplex, $\Delta_{S_0} = \text{conv}(S_0)$, with $S_0 = \{p_{-d}, \ldots, p_0\}$, so that S is contained in it. We should also require that every d-simplex of $\mathcal{R}(S)$ is also a d-simplex of the regular triangulation of $S \cup S_0$. The $d+1$ artificial points can be conveniently chosen at infinity. For example, set $w_{p_i} = 0$, and

$$p_{ij} = \begin{cases} 0 & \text{if } -i > j \\ +\infty & \text{if } -i = j \\ -\infty & \text{if } -i < j \end{cases},$$

where p_{ij} denotes the j-th coordinate of p_i, for $-d \leq i \leq 0$. The symbol '∞' is a place-holder for a large enough number, and this is the easiest way to think of the artificial points and their effect on the computations. The particular choice of points guarantees that $\mathcal{R}(S)$ is a subcomplex of $\mathcal{R}(S_0 \cup S)$. In fact, $\mathcal{R}(S)$ consists of all simplices of $\mathcal{R}(S_0 \cup S)$ that are not incident to any point of S_0.

Global algorithm. Define $S_i = \{p_{-d}, p_{-d+1}, \ldots, p_i\}$. We proceed as follows. Given $\mathcal{R}(S_{i-1})$, let $\Delta = \Delta_T$ be the d-simplex that contains p_i. If, even after adding p_i, Δ is still regular then $\mathcal{R}(S_i) = \mathcal{R}(S_{i-1})$. Otherwise, flip $T \cup \{p_i\}$. This is a flip of type '1 to $d+1$'. Continue flipping locally non-regular $(d-1)$-simplices until none remain. The resulting triangulation is $\mathcal{R}(S_i)$.

We need some more terminology. A $(d-1)$-simplex Δ_U of a triangulation belongs to the *link* of vertex p_i if $\Delta_{U \cup \{p_i\}}$ is a d-simplex of the triangulation. The $(d-1)$-simplices of the link of p_i are called *link facets*. In the algorithm given below only locally non-regular link facets are flipped.

```
1 Construct R(S_0) = Δ_{S_0};
2 for i := 1 to n do
3   locate the d-simplex Δ_T in R(S_{i-1}) that
    contains p_i;
4   if R(T ∪ {p_i}) ≠ Δ_T then
5     flip T ∪ {p_i};
6     while there exist locally non-regular link facets do
7       find a locally non-regular link facet Δ
        that is flippable;
8       flip Δ
      endwhile
    endif
  endfor.
```

In section 8 we will argue that it is indeed sufficient to restrict our attention to link facets when we search for a remaining non-regular $(d-1)$-simplex in step 7. The details of the while loop (steps 6, 7, 8) and the point location operation (step 3) are explained below. As we will see, the implementation of steps 3 and 4 is slightly different than indicated above.

Finding and flipping link facets. We now describe a way to efficiently implement steps 6 and 7. A stack of link facets is maintained. Each time a link facet Δ is flipped, all new link facets are pushed onto the stack. The search for a link facet that is locally non-regular and also flippable begins at the top of the stack. If the topmost link facet is not flippable or it is locally regular or it is not part of the current triangulation then it is simply popped from the stack. In the first case, it could be that this link facet becomes flippable later as the result of some changes in its neighborhood. If this happens then a neighboring link facet will be added whose flip implies the flip of the popped link facet. Consider the case where the link facet Δ is no longer in the current triangulation. Δ is stored in the stack as a pair of pointers to the two d-simplices incident to it. Both d-simplices are no longer part of the current triangulation. To handle this case, whenever a flip removes a d-simplex, it is marked. If the two d-simplices incident to a link facet are marked, it is discarded. In fact, the d-simplices destroyed by flips are maintained in a structure called the *history dag*, see below. Each flip adds at most d facets to the stack. This implies that the total time required by the while loop is proportional to the number of flips performed.

Point location. The method we use to implement step 3 is a generalization of the two-dimensional technique of [11]. The history of performed flips is used as an aid in the search. More specifically, as points are added and flips are carried out, we maintain the collection of discarded d-simplices in a directed acyclic graph, called the *history dag*.

The history dag has a unique root, which is the d-simplex Δ_{S_0}. At any moment, the d-simplices of the current triangulation are the sinks of the dag. Recall that a flip replaces some k d-simplices of the current triangulation with some other $d+2-k$ d-simplices. Before the flip, the k d-simplices are sinks of the dag. Performing the flip means to add the $d+2-k$ new d-simplices as successors to the k old d-simplices. Thus, the k sinks become inner nodes, and $d+2-k$ new sinks are added to the dag.

The search with a point p_i proceeds as follows. Starting at the root of the history dag, we follow the path of d-simplices that contain p_i. Before proceeding from a d-simplex Δ to the next one, we check whether $w_{p_i} < \pi_z(p_i)$, where $z = z(\Delta)$. If it is, then the search terminates because this implies that p_i is redundant in S_i and therefore also in S.

8 Correctness

The algorithm of section 7 could fail for two reasons. First, if all link facets are locally regular although there

Let L' be the subset of d-simplices in L that are incident to locally non-regular link facets. By assumption, $L' \neq \emptyset$. For each $\Delta \in L'$ consider $f(\Delta) = f_{p_i}(\Delta)$ and let $\Delta_{\min} = \Delta_U$ be the d-simplex that minimizes f. We prove below that $T = U \cup \{p_i\}$ is flippable.

By choice, $f(\Delta_{\min}) \leq f(\Delta)$ for all $\Delta \in L'$. All $\Delta \in L - L'$ are incident to locally regular link facets. Therefore, $w_{p_i} < \pi_z(p_i)$, where $z = z(\Delta)$. This implies

$$f(\Delta_{\min}) < w_{p_i} < \pi_z(p_i) = f(\Delta).$$

In other words, $\Delta_{\min}$ minimizes f over all $\Delta \in L$. Consider a half-line, r, emanating from p_i that intersects a link facet in its relative interior. Before intersecting any other d-simplex outside star(p_i), r intersects d-simplices in L. By Lemma 2.2 f increases along the sequence of d-simplices intersecting r. Thus, if r intersects $\Delta_{\min}$ then it cannot intersect any other d-simplex outside star(p_i) before $\Delta_{\min}$. This implies that the subcomplex induced by $T = U \cup \{p_i\}$ has underlying space equal to conv(T). In other words, T is flippable.

9 Randomized Analysis

If the points of S are added in a random sequence we can show that the expected running time of our algorithm is $O(n \log n + n^{\lceil d/2 \rceil})$. The analysis follows the same pattern as in [11]. We begin with a brief worst-case analysis of the number of flips performed.

Maximum number of flips. The $d+2$ points involved in a flip define $d+2$ d-simplices, each occurring either in the triangulation of the $d+2$ points before the flip or the one after the flip. So one of the two triangulations has $k \geq \frac{d+2}{2}$ d-simplices. The k d-simplices intersect in a $(d-k+1)$-simplex, with $d-k+1 \leq \frac{d}{2}$. Set $\delta = \lfloor \frac{d}{2} \rfloor$. This implies that each flip deletes at least one δ-simplex or adds at least one. As mentioned in section 8, a simplex is added and deleted at most once, so the number of flips cannot exceed the total number of δ-simplices defined by n points. A δ-simplex is spanned by $\delta+1$ points, so n points span $\binom{n}{\delta+1}$ δ-simplices. It follows that the maximum number of flips needed for a regular triangulation of n points in $\Re^d$ is at most $2\binom{n}{\delta+1} = O(n^{\lceil (d+1)/2 \rceil})$. This should be compared with the result of Lemma 9.2 below.

The analysis of the running time under the assumption of a random input sequence requires some additional definitions.

Terminology and k-set bounds. Consider an arbitrary subset T of $d+1$ points of S and let $\Delta = \Delta_T$ be the d-simplex defined by T. Let $z = z(\Delta)$ be the orthogonal center of Δ, and define

$$\Sigma(\Delta) = \{p \in S \mid \pi_z(p) < w_p\}.$$

Note that $\Sigma(\Delta) \cap T = \emptyset$, and that $\Sigma(\Delta) = \emptyset$ iff Δ is a d-simplex of $\mathcal{R}(S)$. Call $\sigma(\Delta) = |\Sigma(\Delta)|$ the *width* of Δ.

The analysis is based on bounds for the number of d-simplices with a fixed width k. It will be necessary to also consider d-simplices incident to points of S_0. For each subset $\Omega \subseteq S_0$ and for each $0 \leq k \leq n$, write Y_k^Ω for the collection of subsets $T \subseteq S_n$, $|T| = d+1$, for which $T \cap S_0 = \Omega$ and $\sigma(\Delta_T) = k$. Furthermore, define $Y_{\leq j}^\Omega = \bigcup_{k=0}^{j} Y_k^\Omega$.

For non-empty Ω, the sets Y_k^Ω are somewhat more natural if we consider the lifted set $S_n^+ = \{p^+ \in \Re^{d+1} \mid p \in S_n\}$. As explained in section 3, the orthogonal center of $\Delta = \Delta_T$, $T \subseteq S_n$, corresponds to the hyperplane that contains the points of T^+. The constraint that a hyperplane contain a point with some arbitrarily large or arbitrarily small coordinates (symbolized by $+\infty$ or $-\infty$) really means the hyperplane must contain a certain direction. Recall that $T \cap S_0 = \Omega$ and that T^+ contains $\omega = |\Omega|$ points with arbitrarily large or small coordinates. So the hyperplane spanned by T^+ must contain ω directions, or equivalently, it must be normal to an ℓ-flat, where $\ell = d+1-\omega$. This ℓ-flat can be viewed as an embedding of $\Re^\ell$ in $\Re^{d+1}$.

The maximum cardinalities of the sets Y_k^Ω relate to the maximum number of k-sets of a collection of points in $\Re^\ell$. A *k-set* of a finite point set $A \subseteq \Re^\ell$ is a subset $B \subseteq A$ of size k for which there is a half-space H in $\Re^\ell$ with $B = A \cap H$. Write $g_k^{(\ell)}(A)$ for the number of k-sets of A and define $g_{\leq j}^{(\ell)}(A) = \sum_{k=1}^{j} g_k^{(\ell)}(A)$. The result on k-sets that is most relevant to our analysis is

$$g_{\leq j}^{(\ell)}(A) = O(n^{\lfloor \frac{\ell}{2} \rfloor} j^{\lceil \frac{\ell}{2} \rceil}),$$

where $n = |A|$. The proof of this bound in [4] assumes that ℓ is a constant and j is asymptotically less than n. If j is proportional to n then the bound is trivial. Alternatively, this bound can be obtained by a straightforward extension of the relevant calculations in [11].

The connection between the sets Y_k^Ω and the concept of a k-set is based on the lifting map explained in section 3. Consider a set $T \in Y_k^\Omega$. So $|T| = d+1$, $T \cap S_0 = \Omega$, and for $\Delta = \Delta_T$ we have $\sigma(\Delta) = |\Sigma(\Delta)| = k$. Let h_Δ be the hyperplane in $\Re^{d+1}$ spanned by the points in T^+. The property of the lifting map discussed right before Lemma 3.1 implies that $\Sigma(\Delta)^+ = S^+ \cap H$ for one of the two open half-spaces H bounded by h_Δ. Thus, $\Sigma(\Delta)^+$ is a k-set of S^+. Furthermore, if $\Omega \neq \emptyset, S_0$, then h_Δ is normal to an ℓ-flat F_Ω, where $\ell = d+1-|\Omega|$. For a point $p \in S$ let p_Ω be the orthogonal projection of p^+ into F_Ω. Extend this definition to sets, so that for example $S_\Omega = \{p_\Omega \mid p \in S\}$. With these definitions, $\Sigma(\Delta)_\Omega$ is a k-set of S_Ω. So we can use the above bound on the number of k-sets and obtain the following result.

and the construction of so-called alpha shapes [8, 10]. Indeed, the main motivation for studying the problems solved in this paper is our intention to implement weighted and unweighted alpha shapes in dimensions beyond $\Re^3$. It would be interesting to conduct an experimental study comparing the algorithm of this paper and its main contenders for constructing d-dimensional regular triangulations. These are probably the randomized algorithm of Clarkson and Shor [4] and the output-sensitive algorithm of Seidel [22].

Acknowledgement. We thank Raimund Seidel for sharing an observation that significantly simplified an earlier version of the correctness proof in section 8. Independent from us, Barry Joe found a proof that incremental topological flipping works for computing Delaunay triangulations in d-dimensional space. His proof is similar to ours, although it covers only the case of equal weights for all points.

References

[1] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.* **16** (1987), 78–96.

[2] F. Aurenhammer. Voronoi diagrams – a survey of a fundamental geometric data structure. *ACM Comput. Surveys* **23** (1991), 345–406.

[3] H. Bruggesser and P. Mani. Shellable decompositions of cells and spheres. *Math. Scand.* **29** (1971), 197–205.

[4] K. Clarkson and P. Shor. Applications of random sampling in computational geometry. *Discrete Comput. Geom.* **4** (1989), 387–421.

[5] B. N. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk* **7** (1934), 793–800.

[6] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, Heidelberg, Germany, 1987.

[7] H. Edelsbrunner. An acyclicity theorem for cell complexes in d dimensions. *Combinatorica* **10** (1990), 251–260.

[8] H. Edelsbrunner, D. G. Kirkpatrick and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **IT-29** (1983), 551–559.

[9] H. Edelsbrunner and E. P. Mücke. Simulation of Simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graphics* **9** (1990), 66–104. !!!

[10] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. Manuscript, Dept. Comput. Sci., Univ. Illinois at Urbana-Champaign, 1992.

[11] L. J. Guibas, D. E. Knuth and M. Sharir. Randomized incremental construction of Delaunay and Voronoi diagrams. *In* "Proc. 17th Ann. Internat. Coll. Automata, Lang., Progr. 1990", Lecture Notes in Comput. Sci. **443**, Springer-Verlag, 414–431.

[12] B. Joe. Three dimensional triangulations from local transformations. *SIAM J. Sci. Stat. Comput.* **10** (1989), 718–741.

[13] B. Joe. Construction of three-dimensional Delaunay triangulations using local transformations. *Computer Aided Geometric Design* **8** (1991), 123–142.

[14] C. L. Lawson. Generation of a triangular grid with applications to contour plotting. Memo 299, Jet Propulsion Laboratory, Pasadena, California, 1972.

[15] C. L. Lawson. Software for C^1 surface interpolation. *Mathematical Software III*, edited by J. Rice, Academic Press, New York, 1977, 161–194.

[16] C. L. Lawson. Properties of n-dimensional triangulations. *Computer Aided Geometric Design* **3** (1986), 231–246.

[17] C. Lee. Regular triangulations of convex polytopes. *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, edited by P. Gritzmann and B. Sturmfels, Amer. Math. Soc., Providence, RI, 1991, 443–456.

[18] F. P. Preparata and M. I. Shamos. *Computational Geometry – an Introduction.* Springer-Verlag, New York, 1985.

[19] J. Radon. Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten. *Math. Ann.* **83** (1921), 113–115.

[20] V. T. Rajan. Optimality of the Delaunay triangulations in $\Re^d$. *In* "Proc. 7th Ann. Sympos. Comput. Geom. 1991", 357–363.

[21] E. Schönhardt. Über die Zerlegung von Dreieckspolyedern in Tetraeder. *Math. Ann.* **98** (1928), 309–312.

[22] R. Seidel. Constructing higher-dimensional convex hulls at logarithmic cost per face. *In* "Proc. 18th Ann. ACM Sympos. Theory Comput. 1986", 403–413.

[23] G. F. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *J. Reine Angew. Math.* **133** (1907), 97–178.