



HAL
open science

On the randomized construction of the Delaunay tree

Jean-Daniel Boissonnat, Monique Teillaud

► **To cite this version:**

Jean-Daniel Boissonnat, Monique Teillaud. On the randomized construction of the Delaunay tree. [Research Report] RR-1140, INRIA. 1989. inria-00075419

HAL Id: inria-00075419

<https://inria.hal.science/inria-00075419v1>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



UNITÉ DE RECHERCHE
INRIA-SOPHIA ANTIPOLIS

Institut National
de Recherche
en Informatique
et en Automatique

Sophia Antipolis
B.P. 109
06561 Valbonne Cedex
France
Tél.: 93 65 77 77

Rapports de Recherche

N°1140

Programme 4
Robotique, Image et Vision

**ON THE RANDOMIZED
CONSTRUCTION
OF THE DELAUNAY TREE**

Jean-Daniel BOISSONNAT
Monique TEILLAUD

Décembre 1989
Révisé Janvier 1991

On the Randomized Construction of the Delaunay Tree

Construction randomisée de l'Arbre de Delaunay

Jean-Daniel Boissonnat[†]

Monique Teillaud[†]



Programme 4 : Robotique, Image et Vision.

Rapport INRIA 1140, Décembre 1989. Révisé en Janvier 1991

A paraître dans Theoretical Computer Science.

Keywords : Computational Geometry, Delaunay Triangulation, Randomized Analysis, On-line algorithms.

[†]Institut National de Recherche en Informatique et Automatique
BP 93 — 06902 Sophia Antipolis Cedex (France)
Phone : +33 93 65 77 60 — Telex : 970 050F — Fax : +33 93 65 76 43
E-mail : boissonn@alcor.inria.fr, teillaud@alcor.inria.fr

This work has been supported in part by the ESPRIT Basic Research Action Nr. 3075 (ALCOM).

Abstract

The Delaunay Tree is a hierarchical data structure that was introduced in [BT86]. It is defined from the Delaunay triangulation and, roughly speaking, represents a triangulation as a hierarchy of balls. It allows a semi-dynamic construction of the Delaunay triangulation of a finite set of n points in any dimension. In this paper, we prove that a randomized construction of the Delaunay Tree (and thus, of the Delaunay triangulation) can be done in $O(n \log n)$ expected time in the plane and in $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ expected time in d -dimensional space. These results are optimal for fixed d . The algorithm is extremely simple and experimental results are given.

Résumé

L'Arbre de Delaunay est une structure hiérarchique qui a été introduite dans [BT86]. L'Arbre de Delaunay est défini à partir de la triangulation de Delaunay et représente une triangulation par une hiérarchie de boules. Il permet une construction semi-dynamique de la triangulation de Delaunay d'un ensemble de n points, en dimension quelconque. Nous prouvons ici qu'une construction randomisée de l'Arbre de Delaunay (et donc de la triangulation de Delaunay) peut être effectuée en temps moyen $O(n \log n)$ dans le plan et $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ en dimension d . Ces résultats sont optimaux pour d fixé. L'algorithme est extrêmement simple et des résultats expérimentaux sont présentés.

1 Introduction

Much attention has been paid, in the Computational Geometry literature, to the Voronoï diagram of a finite set of points and to its geometric dual, the Delaunay triangulation. For general references, we refer the reader to the books by Preparata and Shamos [PS85] and by Edelsbrunner [Ede87]. Aurenhammer recently wrote a very complete survey of this fundamental data structure [Aur91].

In the plane, Shamos [Sha78] established an $\Omega(n \log n)$ lower bound for the time for computing a planar triangulation and provided an algorithm reaching this bound, thus optimal in the worst-case sense. Other optimal algorithms have been proposed by Lee and Schacter [LS80], Preparata and Hong [PH77], Guibas and Stolfi [GS85] and Fortune [For87].

In higher dimensional spaces, Klee [Kle80] has shown that the Delaunay triangulation may have $\Omega\left(\frac{n^{\lceil \frac{d}{2} \rceil}}{\lceil \frac{d}{2} \rceil!}\right)$ simplices. See also [Pas82, Sei82]. It is well known that the problem of computing the Delaunay triangulation of a set of n points in d -space reduces to that of computing the convex hull of a set of n points, obtained by lifting the initial ones to a paraboloid of revolution in $d + 1$ -space. The beneath-beyond algorithm of Seidel [Sei81] (see also [Ede87]) computes the convex hull of a set of n points in d -space in time $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ for $d > 2$. This algorithm is incremental, and the insertion of a new site reduces to an intersection problem in the dual space. It can be used to compute the Delaunay triangulation of n points in d -space in time $O(n^{\lfloor \frac{d}{2} \rfloor + 1})$, which is optimal for odd dimensions.

The previously mentioned algorithms are rather involved, even in the plane, and a much simpler one, proposed by Green and Sibson [GS78] in two dimensions, is often preferred. This algorithm is incremental and can be generalized to higher dimensional spaces as shown by Bowyer [Bow81]. Each point is introduced one after another and the triangulation is updated after each insertion. The worst-case complexity of the algorithm is $O(n^2)$ in the 2-dimensional case and $O\left(n^{\lceil \frac{d}{2} \rceil + 1}\right)$ in the d -dimensional case. In order to improve the algorithm, the authors perform a walk from neighbor to neighbor which yields an estimated complexity of $O(n^{1 + \frac{1}{d}})$ for homogeneous distributions of points. However no precise analysis can be found in the papers and the performances tend to become rather poor with degenerate distributions of points, such as points belonging to a subset of lower dimension. The same remarks hold even when using sophisticated bucketing techniques such as the ones used by Asano et al. [AEH85].

A clear advantage of this algorithm however is that it is quite simple and allows insertions of new points in a dynamic way. This has motivated further investigations. A first idea [BT86] to improve the time complexity of the incremental algorithm consisted of keeping all the incremental versions of the triangulation in a structure called the Delaunay Tree. In order to do that efficiently, the points were introduced in a random way. Other randomized techniques appeared recently [CS89, MMO91, GKS92, Mul89]. A common point to all these randomized algorithms is that no distribution assumptions are made as it is the case, for example, in Dwyer [Dwy91]. Hence the results remain valid for any set of points, provided that the points are inserted at random.

The algorithms in [CS89, MMO91, Mul89] are incremental in the sense that the points are introduced one at a time. But all the points need to be known in advance and maintained in an auxiliary data structure, the so called conflict graph. The algorithms in [BT86, GKS92] do not impose such a restriction and thus, are more “on-line”.

The algorithm used in [GKS92] has not been generalized to higher dimensions. The algorithms in [CS89, Mul89] work in any dimensions. However, as already noticed, they are static.

The algorithm based on the Delaunay Tree can compute the Delaunay triangulation of a set of points in any dimension. In this paper, we prove that a randomized construction of the Delaunay Tree (and, thus, of the Delaunay triangulation) can be done in $O(n \log n)$ expected time in the plane and in $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ expected time in d -dimensional space, for fixed d . These bounds are best possible for fixed d . Indeed, in the plane, computing the Delaunay triangulation of n points lying on a convex curve can be reduced to sorting, which takes $\Omega(n \log n)$ time, even for a randomized algorithm. In d -space, for fixed d , the complexity is identical to the size of the triangulation in the worst-case and, thus, is clearly optimal.

The algorithm is extremely simple and experimental results are given.

2 The Delaunay Tree

2.1 Definition and fundamental results about the Delaunay triangulation

We just recall the definition of the Delaunay triangulation and some general results of importance for the sequel. More details can be found in [Rog64] and [PS85]. Let \mathcal{E}^{\lceil} be a d -dimensional euclidean space and \mathcal{M} a set of n sites M_1, \dots, M_n . We will assume that no subset of $d + 2$ sites are cospherical or lie on a same hyperplane. The *Voronoi diagram* associated to \mathcal{M} is a sequence V_1, \dots, V_n of convex polyhedra covering \mathcal{E}^{\lceil} where V_i consists of all the points of \mathcal{E}^{\lceil} that have M_i as a nearest site in the set \mathcal{M} .

$$V_i = \{P \in \mathcal{E}^{\lceil}, \forall j, \infty \leq | \leq \setminus, \delta(P, M_j) \leq \delta(P, M_i)\}$$

where δ denotes the euclidean distance.

The geometrical dual of the Voronoi diagram, obtained by linking the sites M_i whose Voronoi polyhedra are adjacent, is called the *Delaunay triangulation* of \mathcal{M} . The joins are taken to be straight line segments and we use them as a framework for a simplicial subdivision of space. The Delaunay triangulation is a collection of simplices covering the convex hull of \mathcal{M} . The half-spaces, not containing \mathcal{M} , whose boundary contains a facet of the convex hull of \mathcal{M} , will be considered as *infinite* simplices (they are defined by d sites of \mathcal{M} and a site at infinity). Together with the infinite simplices, the triangulation is a tessellation of the whole space \mathcal{E}^{\lceil} . A fundamental property of the triangulation is related to the balls circumscribing the simplices of the triangulation (the *Delaunay balls*). Notice that the balls circumscribing infinite simplices are half-spaces.

Fundamental property :

No site of \mathcal{M} belongs to the interior of a Delaunay ball.

If a site lies inside the circumscribing ball of a simplex, we say that the site is *in conflict* with the simplex. In these terms, the Delaunay triangulation is the set of simplices without conflict.

2.2 Construction of the Delaunay Tree

The Delaunay Tree is a hierarchical structure based on the incremental procedure of [GS78, Bow81]. During the incremental algorithm, each site is introduced one after another and the triangulation is updated after each insertion. Let M be a site to be introduced in the triangulation. All the simplices in conflict with M can no longer be simplices of the triangulation (and are removed in Green and Sibson's algorithm). The union of these simplices is a simply connected region $R(M)$. Let $F(M)$ denote the set of facets of the boundary of $R(M)$. The new simplices are obtained by linking M to the facets of $F(M)$ (Figure 1).

The Delaunay Tree is constructed in a similar way. But, instead of removing simplices during the different steps of the construction, we store all the simplices which have been constructed as nodes of the Delaunay Tree and at each step we define relationships between simplices of the successive Delaunay triangulations. The aim of this structure is to find $R(M)$ efficiently.

2.2.1 Initialization step

Let d denote the dimension of the space. For the initialization step we choose $d + 1$ sites. They generate one finite simplex and $d + 1$ infinite ones (see Figure 2). These $d + 2$ simplices will be the sons of the root of the tree.

2.2.2 Inserting a new site M

After the insertion of site M , the simplices in conflict with M are called *dead* and M is their *killer*. Observe that a killed simplex may not be incident to a facet of $F(M)$.

Let T be one of the simplices in conflict with M that has a facet F belonging to $F(M)$. We construct the new simplex S having vertex M and facet F . Let N be the simplex sharing facet F with T . Because the triangulation is a Delaunay one, we have the following property (see Figure 3) :

(\mathcal{P}) the circumscribing ball of S is included in the union of the two balls circumscribing T and N

This property is fundamental for the correctness of the algorithm, as will be seen in the sequel.

The newly created simplex S will be called : *son of T* and *stepson of N* through facet F .

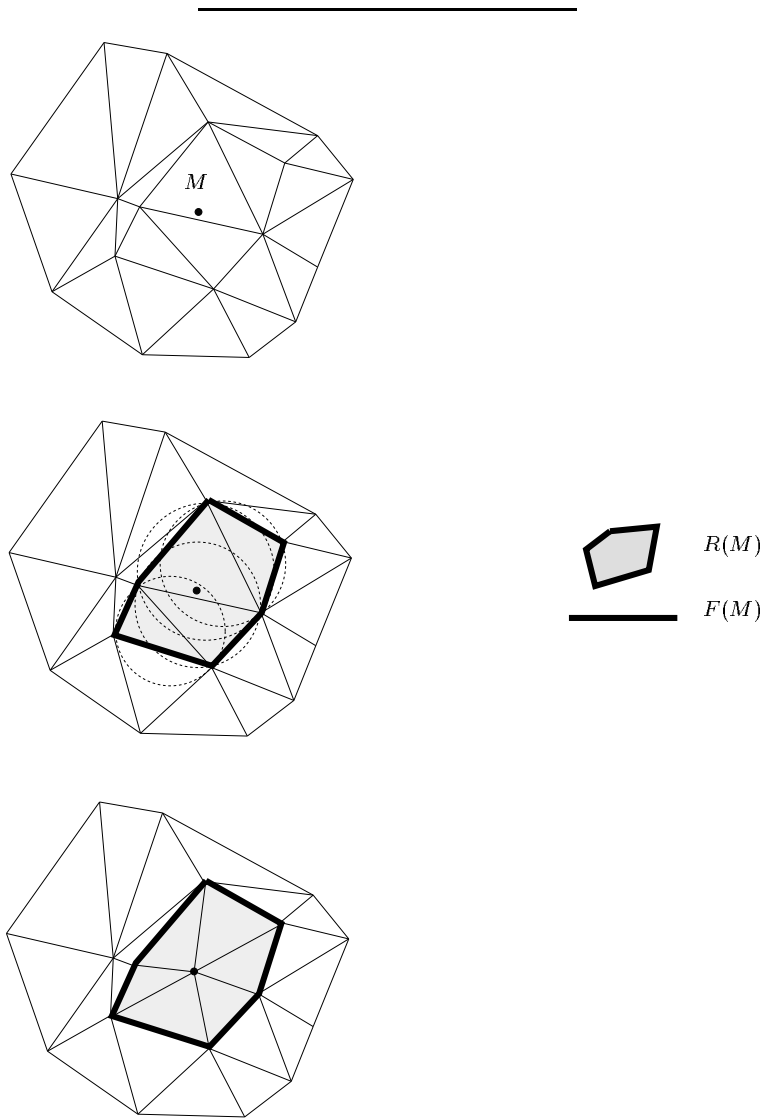


Figure 1: Inserting a new site

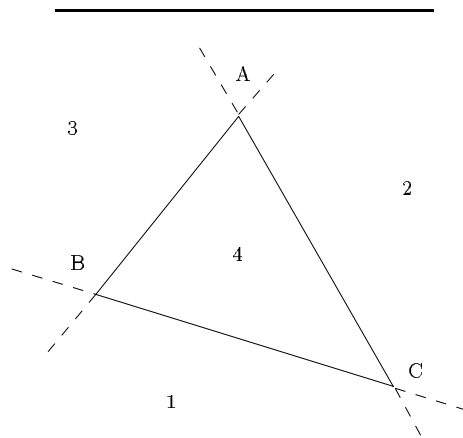


Figure 2: Initialization step

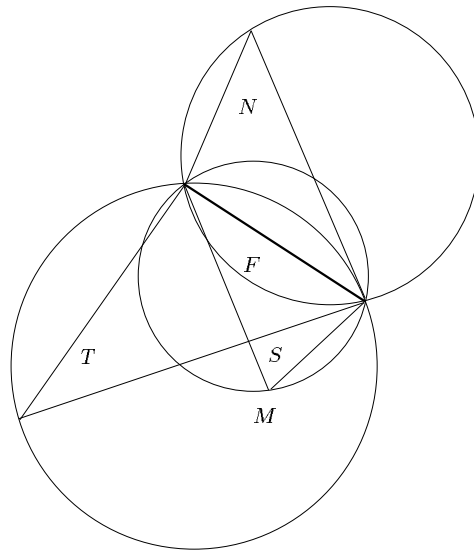


Figure 3: Property (\mathcal{P})

If we now insert a new site M' belonging to the circumscribing ball of S but not to that of N , S will be killed in turn, and its son S' having vertex M' and facet F will be another stepson of N . Thus a node has at most one son and one list of stepsons through each facet, that is : 0 to $d + 1$ sons and 0 to $d + 1$ lists of stepsons.

Remark 2.1 Nevertheless, the total size of the stepson lists in the Delaunay Tree is less than the number of nodes, since each newly created node (the $d + 2$ first ones excepted) has exactly one stepfather. This is true in any dimension.

This hierarchical structure is called a *Delaunay Tree* for short, but it is more exactly a rooted direct acyclic graph. This graph contains a tree : the tree whose links are the links between fathers and sons.

A simplex of the current triangulation is not dead, and so corresponds to a node having no son, but possibly stepsons. We also maintain adjacency relationships between the simplices of the current triangulation. This will be developed in section 2.3.2.

2.3 Using the Delaunay Tree for constructing the Delaunay triangulation

Let M be a site to be introduced in the triangulation. Two steps are performed: first, we locate M in order to find the set $R(M)$ of all the simplices in conflict with M and then we create the new simplices.

2.3.1 Locating M

If M is in conflict with a simplex, we know, by Property (\mathcal{P}), that it is in conflict with its father or its stepfather. So (\mathcal{P}) implies that we will be able to find all the simplices which are killed by M by exploring the Delaunay Tree following Procedure **location** described in Figure 4.

Remark 2.2 In fact, it is also possible to find only one simplex of $R(M)$ by searching the Delaunay Tree, and to deduce the others using neighborhood relationships.

Remark 2.3 let us show that all links from a node to its successive stepsons can be used to locate a new site. In Figure 5, sites 1 to 7 are numbered in insertion order.

The subtree of the Delaunay Tree corresponding to triangles (123), (234), (235), (256) and (237) is shown on the right side of the figure.

If site 8 is now introduced, it is in conflict with (235) and (256), but not with (123). We can locate it using the link from (234) to (235) and the link from (235) to (256). There exists no way to reach (256) from (234) through (237). If we want to locate 9, which is in conflict with (237) but not with (235), we must use the link from (234) to (237). Thus, it is necessary to store both stepsons of (234).

Procedure location(M, T) :

```

if  $T$  has not been visited yet and  $M$  is in conflict with  $T$ 
  for each stepson  $S$  of  $T$  location( $M, S$ ) ;
  for each son  $S$  of  $T$  location( $M, S$ ) ;
  if  $T$  is not dead then
    mark  $T$  killed by  $M$  ;
    add  $T$  to the list  $R(M)$  of the killed simplices
  endif
endif
endif

```

Figure 4: Location of a site in the Delaunay Tree




Figure 5: All stepsons may be useful

2.3.2 Creating the new simplices

We go through the list of the killed simplices. The facets of these simplices which will remain in the new simplices are those of $F(M)$. Let F be such a facet. We create the new simplex $S = (M, F)$ and the two relations between S , its father and its stepfather. The father of S is the killed simplex T containing facet F . The stepfather of S is the neighbour N of T through facet F (recall that we know the adjacency relationships in the current triangulation). Procedure **creation** is described in Figure 6.

```

Procedure creation( $M$ ):
    for each simplex  $T$  killed by  $M$ 
        for each neighbor  $N$  of  $T$  through a facet  $F$ 
            if  $M$  is not in conflict with  $N$ 
                create the simplex  $S$  having vertex  $M$  and facet  $F$  ;
                replace the adjacency relation between  $N$  and  $T$ 
                    by the adjacency relation between  $N$  and  $S$  ;
                create the relations:  $S$  son of  $T$  and  $S$  stepson of  $N$  through facet  $F$ 
            endif
        endfor
    endfor ;
    create the adjacency relationships between the new simplices.

```

Figure 6: Creating the new simplices

In order to create the adjacency relationships between the new simplices, we proceed as follows. Let F be a $(d-1)$ -face of $F(M)$, and f be a $(d-2)$ -face of F , we compute the sequence T_1, T_2, \dots, T_k of killed simplices, such that

- F is a facet of T_1
- T_i is adjacent to T_{i-1} for $i = 2, \dots, k$
- $f \subset T_i$ for $i = 1, \dots, k$
- T_k has a facet F' belonging to $F(M)$, such that $F' \cap F = f$.

We declare that the new simplices MF and MF' are adjacent through facet Mf . By repeating this process for each $(d-1)$ -face F of $F(M)$ and each $(d-2)$ -face f of F , we obtain all the adjacency relationships between the new simplices.

Each killed simplex has at most $\frac{d(d+1)}{2}$ $(d-2)$ -faces in $F(M)$ and thus is examined at most $\frac{d(d+1)}{2}$ times during this process.

3 Analysis of the randomized construction of the Delaunay triangulation

This section proves the main result of this paper, stated in the following theorem:

Theorem 3.1 *The randomized expected time for inserting the n^{th} site in the Delaunay Tree is $O(\log n)$ in two dimensions and $O\left(n^{\lceil \frac{d}{2} \rceil - 1}\right)$ in dimension $d > 2$. The expected storage of the Delaunay Tree is $O\left(n^{\lceil \frac{d}{2} \rceil}\right)^1$ in dimension $d \geq 2$.*

A short proof for the complexity of the whole construction could be given by referring to Clarkson and Shor's theorem for randomized incremental construction [CS89]. Nevertheless we sketch an alternate proof, which provides a non amortized result.

We first introduce some additional notations and definitions. We define the width of a simplex to be the number of sites of \mathcal{M} in conflict with that simplex. $\mathcal{S}_{\parallel}(\mathcal{M})$ will denote the set of simplices of width k and $\mathcal{S}_{\leq \parallel}(\mathcal{M})$ the set of simplices of width at most k . $|\mathcal{A}|$ is the cardinality of a set \mathcal{A} . We define $s_k(n) = \max_{|\mathcal{M}|=n} |\mathcal{S}_{\parallel}(\mathcal{M})|$, and $s_{\leq k}(n) = \max_{|\mathcal{M}|=n} |\mathcal{S}_{\leq \parallel}(\mathcal{M})|$.

We also define a bicycle as a pair of simplices sharing a facet. A site is said to be in conflict with the bicycle, if it is in conflict with one of the two simplices but is not one of their vertices. $\mathcal{B}_{\parallel}(\mathcal{M})$ will denote the set of bicycles of width k and $\mathcal{B}_{\leq \parallel}(\mathcal{M})$ the set of bicycles of width at most k . We define $b_k(n)$ and $b_{\leq k}(n)$ in the same way as for the simplices.

Notice that, when a site M is inserted, each new simplex S is the son of a simplex T and the stepson of a simplex N ; ST and SN (Figure 3) are bicycles without conflict (while NT has one conflict with M). More generally, an edge of the Delaunay Tree corresponds to a bicycle.

The first lemma, due to Clarkson and Shor [CS89, Theorem 3.1], bounds the numbers of simplices and bicycles with width at most k defined by a set of n sites. The proof of this lemma uses the random sampling technique, and the fact that the number of simplices arising in the Delaunay triangulation of n sites is $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ (see for example [Kle80]) and thus, that the number of bicycles of width zero has the same complexity (d is fixed).

Lemma 3.1 *The number of simplices having width at most k is, for $k \geq 2$, at most*

$$s_{\leq k}(n) = O\left(n^{\lceil \frac{d}{2} \rceil} k^{\lfloor \frac{d}{2} \rfloor + 1}\right)$$

The number of bicycles having width at most k is, for $k \geq 2$, at most

$$b_{\leq k}(n) = O\left(n^{\lceil \frac{d}{2} \rceil} k^{\lfloor \frac{d}{2} \rfloor + 2}\right)$$

¹The constants depend on the dimension as a $\frac{(d+5)!}{\lfloor \frac{d}{2} \rfloor!}$ factor for the time complexity and a $\frac{1}{(\lceil \frac{d}{2} \rceil - 1)!}$ factor for the space complexity. We will omit the computations of these constants in this paper.

Remark 3.1 For $k < 2$, we can trivially deduce the following bounds :

$$s_0(n) \leq s_{\leq 1}(n) \leq s_{\leq 2}(n) = O\left(2^{\lfloor \frac{d}{2} \rfloor + 1} n^{\lceil \frac{d}{2} \rceil}\right)$$

and similarly for bicycles.

Remark 3.2 In the two-dimensional case, from [Lee82], $s_k(n)$ and $b_k(n)$ are $O(k(n-k))$.

The sequel analyzes the insertion of the last site M of the set \mathcal{M} of the n already inserted elements.

Analysis of the randomized expected cost of Procedure location

We compute the expected number of nodes $\mu(n)$ visited to locate the n^{th} site M . A simplex S is visited during the location of M if S is the son or the stepson of a simplex T such that M is in conflict with T . So we need here to count the number of edges of the Delaunay Tree traversed during the location of the last site M , which is also the number of bicycles in conflict with M .

Let $B = ST$ be a bicycle of $\mathcal{B}_{\parallel}(\mathcal{M})$. B is in conflict with M if one of the k sites in conflict with B is M , which happens with probability $\frac{k}{n}$. B appears before the insertion of M if the $k-1$ other sites in conflict with B have been inserted after the $d+2$ sites defining B ; S is the son or the stepson of T if the site that created S has been inserted after the $d+1$ vertices of T . This occurs with probability $\frac{(d+1)!(k-1)!}{(k+d+1)!}$. The expected number of edges traversed by Procedure `location` is then obtained in the following way :

$$\begin{aligned} \mu(n) &= \sum_{k=0}^{n-d-2} \sum_{B=ST \in \mathcal{B}_{\parallel}(\mathcal{M})} \frac{k}{n} \text{Prob} \left(\begin{array}{l} S \text{ arises as a son} \\ \text{or a stepson of } T \\ \text{during the construction} \end{array} \right) \\ &= \sum_{k=0}^{n-d-2} \frac{(d+1)! |\mathcal{B}_{\parallel}(\mathcal{M})|}{\backslash(\parallel + \infty) \dots (\parallel + \lceil + \infty)} \end{aligned}$$

Using Lemma 3.1, we deduce that $\mu(n)$ is $O\left(n^{\lceil \frac{d}{2} \rceil - 1}\right)$ in dimension d , and $O(\log n)$ in dimension 2.

Analysis of the randomized expected cost of Procedure creation

We first compute the number of nodes created by the n^{th} site M . A given simplex is created by M if its width is zero after the insertion of M , and if M is one of its $d+1$ vertices. The expected number of created simplices is $\nu(n) = \frac{d+1}{n} |\mathcal{S}_{\lceil}(\mathcal{M})|$, which is bounded by $\frac{d+1}{n} s_0(n) = O\left(n^{\lceil \frac{d}{2} \rceil - 1}\right)$.

The second part of the cost of inserting M is due to the killed simplices, which must be examined $O(d^2)$ times to create the adjacency relationships between the new

simplices. The width of a simplex killed by M is one after the insertion of M , so the expected number of simplices killed by M is $\frac{1}{n}|\mathcal{S}_\infty(\mathcal{M})| \leq \frac{\infty}{\sqrt[n]}|\mathcal{S}_{\leq\infty}(\mathcal{M})|$, which is less than $\frac{1}{n}s_{\leq 1}(n) = O\left(n^{\lceil\frac{d}{2}\rceil-1}\right)$.

The cost of Procedure **creation**, for the insertion of the n^{th} site is thus $O\left(n^{\lceil\frac{d}{2}\rceil-1}\right)$ in dimension d .

Analysis of the expected space used by the Delaunay Tree

The expected number $\eta(n)$ of nodes in the Delaunay Tree is the number of all successive simplices arising in the Delaunay triangulation during the construction, and can be computed by summing the numbers $\nu(i)$ of simplices created by the i^{th} site :

$$\begin{aligned}\eta(n) &= \sum_{i=1}^n \nu(i) \\ &= O\left(\sum_{i=1}^n i^{\lceil\frac{d}{2}\rceil-1}\right) \\ &= O\left(n^{\lceil\frac{d}{2}\rceil}\right)\end{aligned}$$

We already know that the total number of edges in the Delaunay Tree is linear in the number of nodes (Remark 2.1). The expected total storage is thus $O(\eta(n))$.

This achieves the proof of Theorem 3.1.

Remark 3.3 The results stated in Theorem 3.1 are obtained by using worst-case bounds for the size of a Delaunay triangulation ($s_0(n)$ is used to bound $|\mathcal{S}_r(\mathcal{M})|$).

If we assume that the set of sites to be triangulated is uniformly distributed in the space, we know from the results of [Dwy91] that $|\mathcal{S}_r(\mathcal{M})|$ (and thus $|\mathcal{B}_r(\mathcal{M})|$) is $O(n)$. Consequently, the random sampling technique used in Lemma 3.1 provides a linear complexity for the size of $\mathcal{S}_{\leq\parallel}(\mathcal{M})$ and $\mathcal{B}_{\leq\parallel}(\mathcal{M})$ and our algorithm runs in $O(n \log n)$ expected time in any dimension.

4 Experimental results

The algorithm has been implemented in both dimensions 2 and 3. It is to be noted that the algorithm is extremely simple.

Moreover, the numerical computations involved are also very simple : When a simplex is created, we can compute the coordinates of the center of its circumscribing sphere, and its squared radius. This is achieved by first writing the equation of a sphere in d dimensions, then writing that the $d + 1$ sites defining the simplex belong to the sphere, and solving the linear system which results from that, with $d + 1$ equations and

$d + 1$ unknowns, in $O(d^3)$ time.

The center and the squared radius of a simplex is computed once for all and stored in the corresponding node. For testing if a site is in conflict with a simplex, we only have to compute its squared distance to the center of the simplex, which costs $O(d)$, and compare it with the squared radius. The constant in Footnote 1 has been computed according to this method.

The algorithm has run on many examples with different kinds of point distributions, in two and three dimensions.

In each case, several orders of insertion of the sites have been tried to analyse the running time down to the constants. All randomized orders yielded the same results.

Some results are presented on Figures 7, 8 and 9 for the planar case.

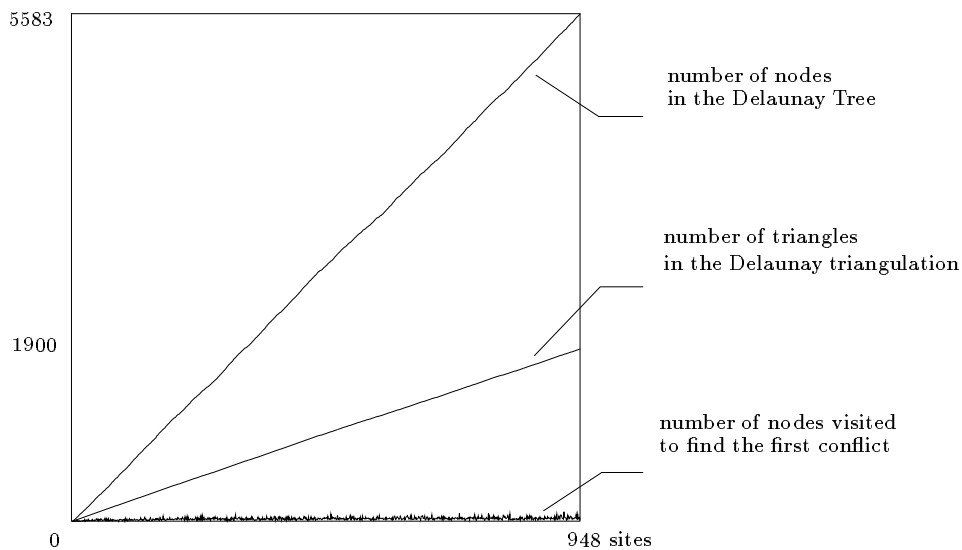


Figure 7: Random sites in the plane

Figures 10, 11 and 12 present results in 3-space. Figures 10 and 11 show results in the case where the size of the final Delaunay triangulation is linear in the number of sites, whereas Figure 12 was obtained for a distribution where the number of tetrahedra was quadratic (the sites are lying on two non-coplanar line segments).

Storage of the Delaunay Tree

The number of nodes in the Delaunay Tree has been studied with respect to the number of inserted sites. This must be compared with the number of simplices in the final triangulation, which is the size of the output.

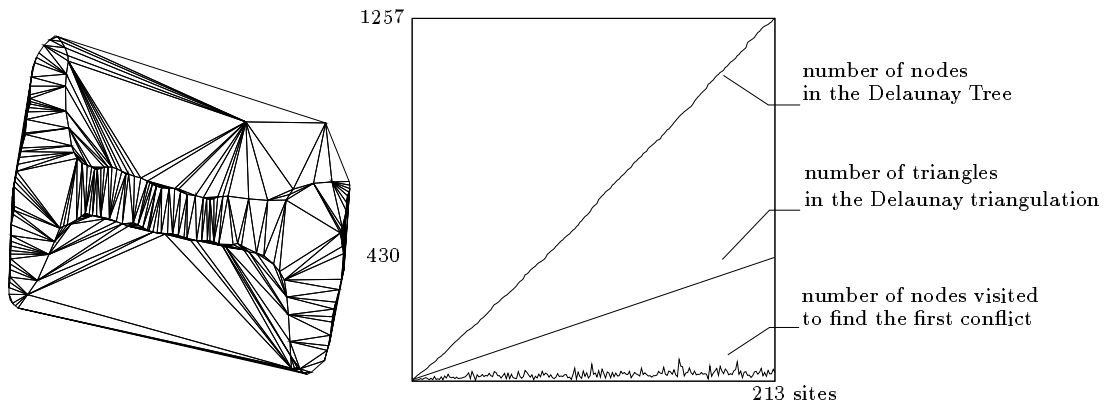


Figure 8: A non-convex curve

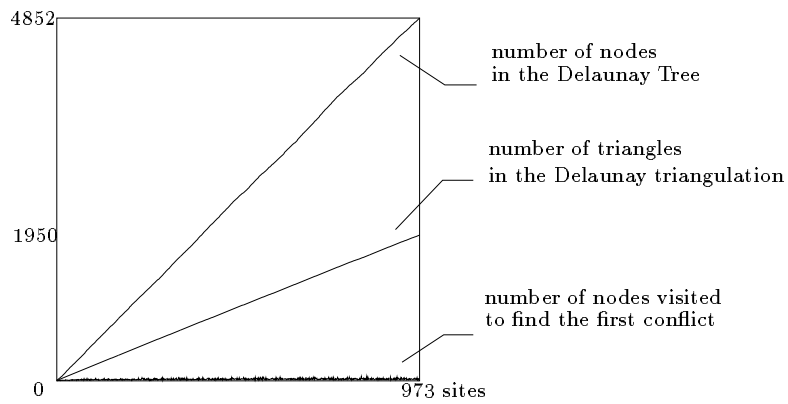


Figure 9: An ellipsis

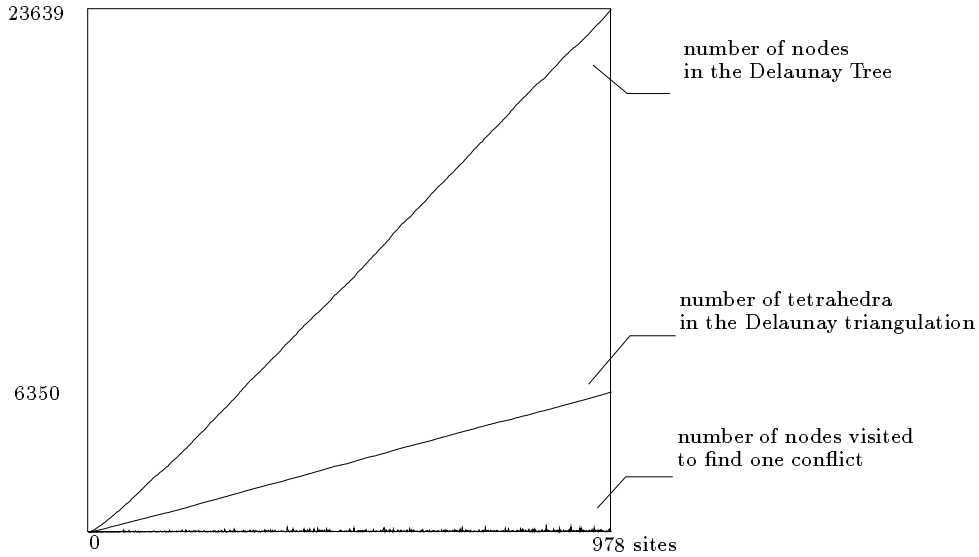


Figure 10: Random sites in 3-space

In the two dimensional case, the ratio of the number of nodes in the Delaunay Tree to the size of the output is less than 3, in any example. In the space, this ratio becomes 4 in the linear case, and only 2 in the quadratic case.

Inserting a site in the Delaunay Tree

The chosen parameter to evaluate the cost of inserting a site is the number of nodes visited by Procedure `location` to find the first conflict. According to Remark 2.2, the remainder cost of this procedure consists of an output sensitive search in the Delaunay triangulation, and the cost of Procedure `creation` also depends on the number of modifications in the Delaunay triangulation.

The cost of locating a site in the Delaunay Tree appears to be very small compared with the size of the Delaunay triangulation.

Some empirical investigations have been done to evaluate the constants. In the plane, the variations of the number of nodes visited to find the first conflict can be roughly (after smoothing) assimilated to the variations of the function $x \mapsto 3 \log_2 x$, in all examples. In the space, in the linear case, the function is $x \mapsto 7 \log_2 x$, which has been explained in Remark 3.3.

We have computed in section 3 the expected number of nodes visited by Procedure `location` to find all the simplices in conflict with a new site M , but no theoretical analysis has been done for the number of nodes visited to find the first conflict. We can see that, even in the quadratic case, this number remains logarithmic. Moreover, the constant appears to be better than in the linear case, since it is about 4.





Figure 11: A closed surface (a heart)




Figure 12: A quadratic example

5 Conclusion

We have shown that the Delaunay Tree of n points can be constructed in randomized expected time $O(n \log n)$ in the plane and $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ in d -space for $d > 2$. Its randomized expected size is $O\left(n^{\lceil \frac{d}{2} \rceil}\right)$ in dimension $d \geq 2$. The Delaunay Tree allows to compute the Delaunay triangulation of n points within the same bounds, which is optimal for fixed d . A byproduct of this result is that the Delaunay Tree can be used to locate a point inside a Delaunay triangulation in randomized expected time $O(\log n)$ in the plane and $O\left(n^{\lceil \frac{d}{2} \rceil - 1}\right)$ in d -space.

An important point is that these results hold whatever the point distribution may be.

The algorithm is extremely simple and, moreover, the numerical computations involved are also very simple.

Experimental results in 2 and 3-dimensional spaces, for uniform as well as degenerate distributions of points, have provided strong evidence that this algorithm is very effective in practice.

In [DMT], the Delaunay Tree has been made fully dynamic in the planar case, with a $O(\log \log n)$ expected complexity for removing a site.

Further investigations have been done to generalize the Delaunay Tree to solve other applications efficiently. The k -Delaunay Tree can be used to compute higher

order Voronoi diagrams [BDT], and the IDAG is a more general data structure that can be used to design a large class of on-line randomized algorithms [BDS⁺92].

Acknowledgements

The authors would like to thank Olivier Devillers for a careful reading of the paper, and Jean-Pierre Merlet for supplying us with his interactive drawing preparation system J^Pdraw.

We are also very grateful to the anonymous referee who suggested many improvements in the earlier version of the paper.

References

- [AEH85] T. Asano, M. Edahiro, H. Imai, and M. Iri. Practical use of bucketing techniques in computational geometry. In G.T. Toussaint, editor, *Computational Geometry*, pages 153–196. North Holland, 1985.
- [Aur91] F. Aurenhammer. Voronoi diagrams — a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3), September 1991.
- [BDS⁺92] J-D. Boissonnat, O. Devillers, R. Schott, M. Teillaud, and M. Yvinec. Applications of random sampling to on-line algorithms in computational geometry. *Discrete and Computational Geometry*, 8:51–71, 1992.
- [BDT] J-D. Boissonnat, O. Devillers, and M. Teillaud. A semi-dynamic construction of higher order Voronoi diagrams and its randomized analysis. *Algorithmica*. To appear. Available as Technical Report INRIA 1207. Abstract published in Second Canadian Conference on Computational Geometry 1990 in Ottawa.
- [Bow81] A. Bowyer. Computing Dirichlet tessellations. *The Computer Journal*, 24(2), 1981.
- [BT86] J-D. Boissonnat and M. Teillaud. A hierarchical representation of objects: The Delaunay Tree. In *Second ACM Symposium on Computational Geometry in Yorktown Heights*, June 1986.
- [CS89] K.L. Clarkson and P.W. Shor. Applications of random sampling in computational geometry, II. *Discrete and Computational Geometry*, 4(5), 1989.
- [DMT] O. Devillers, S. Meiser, and M. Teillaud. Fully dynamic Delaunay triangulation in logarithmic expected time per operation. *Computational Geometry Theory and Applications*. To appear. Available as Technical Report INRIA 1349. Abstract published in LNCS 519 (WADS'91, august 1991).
- [Dwy91] R.A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Discrete and Computational Geometry*, 6:343–367, 1991.
- [Ede87] H. Edelsbrunner. *Algorithms on Combinatorial Geometry*. Springer-Verlag, 1987.
- [For87] S. Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2:153–174, 1987.
- [GKS92] L.J. Guibas, D.E. Knuth, and M. Sharir. Randomized incremental construction of Delaunay and Voronoi diagrams. *Algorithmica*, 7(4):381–413, 1992.

- [GS78] P.J. Green and R. Sibson. Computing Dirichlet tessellations in the plane. *The Computer Journal*, 21, 1978.
- [GS85] L.J. Guibas and J. Stolfi. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *ACM Transactions on Graphics*, 4(2):74–123, April 1985.
- [Kle80] V. Klee. On the complexity of d-dimensional Voronoi diagrams. *Archiv der Mathematik*, 34:75–80, 1980.
- [Lee82] D.T. Lee. On k -nearest neighbor Voronoi diagrams in the plane. *IEEE Transactions on Computers*, C-31:478–487, 1982.
- [LS80] D.T. Lee and B.J. Schachter. Two algorithms for constructing a Delaunay triangulation. *International Journal of Computer and Information Sciences*, 9(3), 1980.
- [MMO91] K. Mehlhorn, S. Meiser, and C. Ó’Dúnlaing. On the construction of abstract Voronoi diagrams. *Discrete and Computational Geometry*, 6:211–224, 1991.
- [Mul89] K. Mulmuley. On obstruction in relation to a fixed viewpoint. In *IEEE Symposium on Foundations of Computer Science*, pages 592–597, 1989.
- [Pas82] I. Paschinger. *Konvexe Polytope und Dirichletsche Zellenkomplexe*. PhD thesis, Institut für Mathematik, Universität Salzburg, Austria, 1982.
- [PH77] F.P. Preparata and S.J. Hong. Convex hulls of finite sets of points in two and three dimensions. *Communications of the ACM*, 2(20):87–93, 1977.
- [PS85] F.P. Preparata and M.I. Shamos. *Computational Geometry : an Introduction*. Springer-Verlag, 1985.
- [Rog64] C.A. Rogers. *Packing and Covering*. Cambridge University Press, 1964.
- [Sei81] R. Seidel. A convex hull algorithm optimal for point sites in even dimensions. Technical Report 14, Departement of Computer Science, University British Columbia, Vancouver, BC, 1981.
- [Sei82] R. Seidel. The complexity of Voronoi diagrams in higher dimensions. In *20th Annual Allerton Conference Communication, Control, Computing*, pages 94–95, 1982.
- [Sha78] M.I. Shamos. *Computational Geometry*. PhD thesis, Department of Computer Science, Yale University, (USA), 1978.