EXPECTED TIME ANALYSIS OF ALGORITHMS IN COMPUTATIONAL GEOMETRY

Luc Devroye

McGill University
School of Computer Science
805 Sherbrooke Street West
Montreal, Canada H3A 2K6

We give a brief inexhaustive survey of recent results that can be helpful in the expected time analysis of algorithms in computational geometry. Most fast average time algorithms use one of three principles: bucketing, divide-and-conquer (merging), or quick elimination (throw-away). To illustrate the different points, the convex hull problem is taken as our prototype problem. We also discuss searching, sorting, finding the Voronoi diagram and the minimal spanning tree, identifying the set of maximal vectors, and determining the diameter of a set and the minimum covering sphere.

1. INTRODUCTION.

There has been an increasing interest in the study and analysis of algorithms in computational geometry; a recent survey paper by Toussaint (1980) has 168 references. Most of the emphasis has been placed on the study of the worst-case complexity of various algorithms under several computational models. It is well-known that many algorithms perform considerably better on the average than predicted by the worst-case analyses. In this note, we would like to point out a few recent developments in the analysis of the expected complexity of some algorithms. To keep the general discussion simple and yet insightful we make a few convenient assumptions.

The Assumptions.

The input data $X_1, \ldots, X_n$ can be considered as a sequence of independent identically distributed R^d-valued random vectors with common density f. (Here the unrealistic assumption is that real numbers can be stored in a computer.)

An algorithm takes time $T = T(X_1, \ldots, X_n)$, a Borel measurable function of the input data, and it halts with probability one, i.e. $T < \infty$ almost surely.

The <u>common operations</u> (+,-,/,*,mod,compare,move) take time uniformly bounded over all values of the operands. For example, a*b or a mod b take time bounded by a constant not depending upon a or b. (Once again, this is unrealistic, because the multiplication or comparison of two real numbers takes infinite time.)

We note here that the floor of a number x is x- x mod 1, so that it too is a constant time operation. The assumptions made here are not entirely unreasonable: real numbers can be approximated very accurately, and the common operations take constant time for all operands that can be stored in a word (i.e., whose value does not exceed a large threshold value). In many situations, the benefit of being able to use the rich theory for random variables with a density, far outweighs the disadvantages of our assumptions.

Expected Time.

We are interested in the expected time E(T) taken by certain algorithms. Obviously, E(T) depends upon f and n because the averaging is done over all random samples of size n drawn from the density f.

Fast Expected Time Algorithms.

In many applications, fast expected time algorithms can be obtained by the <u>bucketing principle</u>: find the smallest rectangle C covering $X_1,...,X_n$; we recall here that a rectangle in R^d is a d-fold product of intervals of the type [a,b]. Now, divide C into equal-sized rectangles (buckets), and solve the problem by traveling from bucket to bucket while performing some local operations.

We will also discuss the dramatic savings in expected time that can be obtained by the proper application of the <u>divide-and-conquer principle</u>. Finally, it is sometimes possible to take a superficial look at the data and to eliminate without much work all useless points. The more involved work is then performed on the reduced data sequence: this will be called the <u>quick elimination</u> (or: <u>throw-away</u>) <u>principle</u>.

Throughout the paper, we will use n for the sample size. For behaviour as $n \to \infty$, we use the Landau symbols O, o, ∿, and the newer notation Ω, Θ. If a_n and b_n are two positive number sequences, then $a_n = O(b_n)$ $(\Omega(b_n))$ when there exist positive numbers n_0 and K_0 such that for all $n \geq n_0$, $a_n \leq K_0 b_n$ $(a_n \geq K_0 b_n)$. Next, $a_n = \Theta(b_n)$ if $a_n = O(b_n)$ and $a_n = \Omega(b_n)$. Finally, $a_n = o(b_n)$ if $\lim_{n \to \infty} a_n/b_n = 0$, and $a_n \sim b_n$ when $\lim_{n \to \infty} a_n/b_n = 1$.

The paper is not meant as a survey of algorithms, or as a state-of-the-art article describing the fastest known expected time algorithm. Its sole purpose is to give selected analytical results that can be used to obtain upper bounds for E(T), or to make inferences about the asymptotic behavior of E(T). But the reader should be aware that even when E(T) is small, it is possible that the distribution of T is such that the algorithm is undesirable. While it is true, by Markov's inequality, that

$$P(T \geq cE(T)) \leq \frac{1}{c} , \text{ all } c > 0,$$

this bound is often too crude to be practical. What one usually wants is a property such as: $P(T-E(T) > cE(T)) \to 0$ as $n \to \infty$, all $c > 0$. This requires usually a deeper analysis. For example, if Var(T) is the variance of T, we have

$$P(T-E(T) > cE(T)) \leq \frac{Var(T)}{c^2 E^2(T)} , c > 0,$$

and the right-hand-side tends to 0 when $Var(T)/E^2(T)$ has a zero limit.

We will not discuss <u>all</u> the fast expected time algorithms. For example, to find the convex hull of $X_1,...,X_n$ in R^2 under certain computational models, at least c n log n time is needed (Avis, 1979; Yao 1981). Jarvis' algorithm (Jarvis, 1973) takes expected time O(nE(N)) where N is the number of convex hull points. In the design of this algorithm, no special care is taken to obtain fast expected time behavior. Nevertheless, for certain distributions E(N)=O(1) (see Carnal (1970); this is true for multivariate t-distributions, etc.), so that Jarvis' algorithm runs in linear expected time for a fairly large class of distributions.

2. THE BUCKETING PRINCIPLE.

Let C be the smallest closed rectangle covering $X_1,...,X_n$ and let C be divided into m^d equal-sized rectangles (buckets) where m=int($^{1/d}$) the floor of $n^{1/d}$. Bucket memberships can thus be obtained for all data points in time O(n). Often one keeps track of these memberships by using m^d linked lists, one per bucket, so that O(n) space is used.

The obvious application in R^1 involves sorting $X_1,\ldots,X_n$. Here one empties the buckets from left to right and performs a subsequent sort within each bucket, if necessary. If this subsequent sort is a comparison-based sort (e.g., heapsort, bubble sort, shell sort, merge sort or quicksort) with average time $g(n)$ (this number is independent of f, since we have a comparison-based sort), then the overall average time for sorting is

$$E(T) = 0(n) + E(\sum_{i=1}^{n} g(N_i))$$

where $N_1,\ldots,N_n$ are the cardinalities of the n buckets. Devroye and Klincsek (1981) addressed the question of when $E(T)=0(n)$. They showed that when $g(u)/u \uparrow \infty$ as $u \to \infty$, $g(u)/u^2$ is nonincreasing, and g is convex, then $E(T)=0(n)$ if and only if f has compact support (i.e., f=0 outside some finite interval) and

$$\int g(f(x))\, dx < \infty .\qquad(1)$$

Notice that they put no continuity or boundedness assumptions on f. In experiments, Akl and Meijer (1982) found that for sufficiently smooth densities, bucket sort (with slight ad hoc improvements) compares favorably with even the best version of quicksort.

These results do not contradict the fact that if $X_1,\ldots,X_n$ have to be sorted by using comparisons only, that the average number of comparisons is at least $\Omega(n \log n)$, because our model allows the floor function to be counted as a constant time operation.

Consider now the following generalization of the previous result: travel from bucket to bucket, performing within the j-th bucket operations taking expected time bounded between $ag(N_j)$ and $bg(N_j)$ when N_j is given. Here $0 < a \le b < \infty$. Once again, the expected time of the entire algorithm is

$$E(T) = \Theta(n+E(\sum_{i=1}^{m^d} g(N_i))).$$

Assume that $g(u)/u \uparrow \infty$ and $g(u)/u^K \downarrow 0$ for some positive K as $u \to \infty$. Also, assume that g is convex. Then $E(T)=0(n)$ if and only if f has compact support and (1) holds (Devroye, 1981a). These basic results have several applications. We cite just two examples.

Examples.

1. Searching in constant expected time.

    Assume that $X_1,\ldots,X_n$ are stored in the bucket data structure given above, and that we are presented with X_Z (where Z is uniformly distributed over $1,\ldots,n$). We have to determine the index i such that $X_i=X_Z$. This is the classical problem of successful search. If the X_i's are stored in the buckets in order of arrival, then the expected search time is

$$\Theta\left(E(\frac{1}{n}\sum_{i=1}^{m^d} N_i^2 +1)\right).$$

By a simple extension of the previous result, we see that $E(T)=0(1)$ if and only if f has compact support and $\int f^2(x)\, dx < \infty$. If within each bucket the data are organized into a binary search tree rather than a linked list, by considering one of the coordinates of the X_1's as the key for sorting, then $E(T)=0(1)$ if and only if f has compact support and

$$\int f(x)\, \log_+ f(x)\, dx < \infty, \text{ where } \log_+ u = \max(0, \log(u)).$$

2. Convex hull algorithms that are based upon sorting.

    The convex hull of $X_1,\ldots,X_n$ is a subsequence $X_{i_1},\ldots,X_{i_k}$ of $X_1,\ldots,X_n$ such that for all X_{i_j} there exists a hyperplane through X_{i_j} and all X_q's, $q \neq i_j$, belong to the same closed halfspace determined by this hyperplane. In R^2, it can be obtained from $X_1,\ldots,X_n$ as follows: (i) Find a point x that belongs to the interior of the convex hull of $X_1,\ldots,X_n$. Sort all the X_i's according to the polar angles of X_i-x (by using the bucket sort described above). This yields a polygon P. (ii) Visit all vertices of P in turn by pushing them on a stack. Pop the stack when non-convex-hull points are encountered. In essence, this is Graham's algorithm (1972) with a modification in the sorting method that is used. Step (ii) takes time $0(n)$. The average time taken by (i) is $0(n)$ when the density

of the polar angle of $X_i - x$ is square integrable. Since x itself is
a random vector, one must be careful before making any inference
about f. Nevertheless, it is sufficient that f is bounded and has
compact support.

The previous applications have one feature in common: the times taken
by the algorithms on individual buckets just depend upon the number and/
or position of the data points within these buckets (and not on, say, the
number of data points in neighboring buckets). In more involved problems,
we cannot avoid looking at neighboring buckets. For example, consider
the class of "closest point problems" in R^d (Shamos and Hoey, 1975) such
as: find all nearest neighbor pairs, construct the Voronoi graph, find
the minimal spanning tree, etc. (see Bentley and Friedman (1979) for other
applications). Shamos (1978) and Weide (1978) discuss many applications
of the bucketing principle, and Bentley, Weide and Yao (1980) give a
fairly comprehensive treatment of the time analysis of bucketing algo-
rithms for closest point problems. We take the liberty to cite a couple
of examples from their study:

Examples.

1. The all-nearest-neighbor problem.

All nearest neighbor pairs can be found in O(n log n) time
(worst-case) (Shamos and Hoey, 1975). Weide (1978) proposed a
bucketing algorithm in which for a given X_i, a "spiral search" is
started in the bucket of X_i, and continues in neighboring cells,
in a spiraling fashion, until no data point outside the buckets
already checked can be closer to X_i than the closest data point
already found. Bentley et al. (1980) showed that Weide's
algorithm halts in average time O(n) when there exists a bounded
open convex region B such that the density f of X_1 is 0 outside B
and satisfies $0 < \inf_{B} f(x) \leq \sup_{B} f(x) < \infty$.

2. The Voronoi diagram.

The Voronoi diagram in R^2 can be found in time O(n log n)
(worst-case) (Shamos (1978), Horspool (1979), Brown (1979)).
Bentley et al. (1980) have a bucketing algorithm that uses spiral
search and has some additional features. The Voronoi diagram can
be found in average time O(n) when d=2 and the density f of X_1
satisfies the condition of Example 1. From the Voronoi diagram,
the convex hull can be obtained in linear time (Shamos, 1978).

3. The minimal spanning tree.

For a graph (V,E), Yao (1975) and Cheriton and Tarjan (1976) give
algorithms for finding the minimal spanning tree (MST) in worst-case
time $O(|E| \log \log |V|)$. The Euclidean minimal spanning tree (EMST)
of n points in R^d can therefore be obtained in O(n log log n) time
if we can find a supergraph of the EMST with O(n) edges in
O(n log log n) time. Yao (1982) suggested to find the nearest
neighbor of each point in a critical number of directions; the
resulting graph has O(n) edges and contains the MST. This nearest
neighbor search can be done by a slight modification of the algo-
rithm in Example 1. Hence, the EMST can be found in expected time
O(n log log n) for any d and for all distributions given in Example 1.
The situation is a bit better in R^2. We can find a planar super-
graph of the EMST in expected time O(n) (such as the Delaunay
triangulation (the dual of the Voronoi diagram), the Gabriel graph,
etc.) and then apply Cheriton and Tarjan's (1976) O(n) algorithm
for finding the MST of a planar graph. Thus, in R^2 and for the
class of distributions given in Example 1, we can find the EMST
in linear expected time.

Finally, we should mention a third group of bucketing algorithms, where
special buckets are selected based upon a global evaluation of the con-
tents of the bucket. For example, assume that not more than a_n buckets
are selected according to some criterion (from the approximately n
original buckets) in time O(n), and that only the data points within the
selected buckets are considered for further processing. If N is the
number of selected points, then we assume that "further processing" takes
time O(g(N)) for a given function g. Because the global evaluation
procedure is not specified, we should assume the worst-case, and this
leads to the study of the order statistics of the cardinalities of the
buckets. The following results can be found in Devroye (1981b). When
M is the maximum of n i.i.d. Poisson (1) random variables, then
$E(M) \sim \log n/ \log \log n$. The same is true if M is the maximum of
$N_1, \ldots, N_n$, where N_i is the cardinality of $[\frac{i-1}{n}, \frac{i}{n}]$ and the data is
$U_1, \ldots, U_n$, a sequence of i.i.d. uniform (0,1) random variables. Using
tight bounds on the upper and lower tails of M, one can show that

$$E(g(N)) = O(g(a_n \frac{\log n}{\log \log n}))$$

where $a_n \geq 1$, g is nondecreasing, $g(x)=0(1+x^{\beta})$ (some $\beta > 0$),
$\sup\limits_{x>0} \frac{g(cx)}{g(x)} < \infty$ (all $c > 1$), and the X_i's have a bounded f with compact
support.

Example.

The convex hull in R^2.

Shamos (1979) suggested to construct the convex hull in R^2 in
the following fashion: mark all the nonempty extremal buckets in
each row and column (the extremes are taken in the north and south
directions for a column, and east and west directions for a row);
mark all the adjacent buckets in the same rows and columns; apply
Graham's $0(n \log n)$ convex hull algorithm to all the points in the
marked buckets. It is clear that $a_n=0(\sqrt{n})$ and that the expected
time of the algorithm is $0(n)+0(E(g(N)))$ = where $g(n)=n \log n$. This
is

$$0(n) + 0(\sqrt{n} \frac{(\log n)^2}{\log \log n}) = 0(n).$$

Furthermore, the expected time spent on determining the bucket
membership divided by the total expected time tends to 1.

3.   THE DIVIDE-AND-CONQUER PRINCIPLE.

A problem of size n can often be split into two similar subproblems
of size approximately equal to n/2, and so forth, until subproblems are
obtained of constant size for which the solutions are trivially known.
For example, quicksort (Sedgewick (1977, 1978)) is based on this principle.
The average time here is $\Theta(n \log n)$, but, unfortunately enough, since the
sizes of the subproblems in quicksort can take values $0,1,2,\ldots$ with equal
probabilities, the worst-case complexity is $\Theta(n^2)$. One can start in the
other direction with about n equal-sized small problems, and marry sub-
solutions in a pairwise manner as in mergesort. Because of the controlled
subproblem size, the worst-case complexity becomes $\Theta(n \log n)$ (Knuth, 1975).
Both principles will be referred to as divide-and-conquer principles.
They have numerous applications in computational geometry with often
considerable savings in expected time. The first general discussion of
their value in the design of fast expected time algorithms can be found
in Bentley and Shamos (1978).

Let us analyze the divide-and-conquer algorithms more formally.
Assume that $X_1,\ldots,X_n$ are R^d-valued independent random vectors with common
distribution, and that we are asked to find $A_n=A(X_1,\ldots,X_n)$, a subset of
$X_1,\ldots,X_n$, where A(.) satisfies:

1) $A(x_1,\ldots,x_n)=A(x_{\sigma(1)},\ldots,x_{\sigma(n)})$, for all $x_1,\ldots,x_n \in R^d$,
     and all permutations $\sigma(1),\ldots,\sigma(n)$ of $1,\ldots,n$.

2) $x_1 \in A(x_1,\ldots,x_n) \Rightarrow x_1 \in A(x_1,\ldots,x_i)$ for all $x_1,\ldots,x_n \in R^d$,
     and all $i \leq n$.

The convex hull satisfies these requirements. If $Q_1(x),\ldots,Q_{2d}(x)$ are
the open quadrants centered at $x \in R^d$, then we say that X_1 is a maximal
vector of $X_1,\ldots,X_n$ if some quadrant centered at X_i is empty (i.e., con-
tains no X_j, $j \neq i$, $j \leq n$). The set of maximal vectors also satisfies the
given requirements. Let N=cardinality(A_n). For $p \geq 1$, we known by
Jensen's inequality that

$$E(N^p) \geq (E(N))^p.$$

In the present context, we would like an inequality in the opposite
direction. For random sets A_n satisfying 1) and 2), and under very weak
conditions on the behavior of E(N), we have

$$E(N^p) = 0((E(N))^p)$$

(Devroye, 1983). For example, it suffices that E(N) is nondecreasing,
or that E(N) is regularly varying at infinity. Also, if $E(N) \leq a_n \uparrow$,
then $E(N^p)=0(a_n^p)$. In essence, the results of Devroye (1983) imply that
under weak conditions on the distribution of X_1, $E(N^p)$ and $(E(N))^p$ are
comparable. The same is true for other nonlinear functions of N. For
example, if $E(N) \leq a_n \uparrow$, then $E(N \log(N+e))=0(a_n \log a_n)$. Thus, the
knowledge of E(N) allows us to make statements about other moments of N.
Here are some known results about E(N).

Examples.

1. A_n is the convex hull. X_1 has a density f.

(i)   E(N)=o(n) (Devroye, 1981b).

(ii) If f is normal, then $E(N)=0((\log n)^{(d-1)/2})$ (Raynaud, 1970).
      For d=2, $E(N) \sim 2\sqrt{2\pi \log n}$ (Renyi and Sulanke, 1963, 1964).

(iii) If f is the uniform density in the unit hypersphere of R^d, then $E(N)=0(n^{(d-1/d+1)})$ (Raynaud, 1970).

(iv) If f is the uniform density on a polygon of R^2 with k vertices, then $E(N) \sim \frac{2k}{3} \log n$ (Renyi and Sulanke, 1963, 1968).

(v) If f is a radial density, see Carnal (1970). For example, if f is radial, and $P(||X_1||>u)=L(u)/u^r$ where $r \geq 0$ and L is slowly varying (i.e., $L(cx)/L(x) \to 1$ as $x \to \infty$, all $c > 0$), then $E(N) \to c(r) > 0$. If $P(||X_1||>u) \sim c(1-u)^r$ for some $c,r >$ as $u \uparrow 1$, and $P(||X_1||>1) = 0$, then $E(N) \sim c(r) n^{1/(2r+1)}$ for some $c(r) > 0$.

2. $\underline{A_n \text{ is the set of maximal vectors}}$.

When X_1 has a density and the components of X_1 are independent, then $E(N)$ is nondecreasing (Devroye, 1980) and $E(N) \sim 2^d(\log n)^{d-1}/(d-1)!$ (Barndorff-Nielsen, 1966; Devroye, 1980).

A_n can be found by the following merging method. Assume for the sake of simplicity that $n=2^k$ for some integer $k > 1$.

1. Let $A_{1i}=A(X_i)$, $1 \leq i \leq n$. Set $j \leftarrow 1$.

2. Merge consecutive A_{ji}'s in a pairwise manner (A_{j1} and A_{j2}; A_{j3} and A_{j4}; etc.)

3. Set $j \leftarrow j+1$. If $j > k$, terminate the algorithm ($A_n=A_{k1}$). Otherwise, go to 2.

We assume that merging and editing of A_{ji} and $A_{j\ i+1}$ with cardinalities k_1 and k_2 can be done in time bounded from above by $g(k_1)+g(k_2)$ for some nondecreasing positive-valued function g, and assume that $E(g(N)) \leq b_n \uparrow$ where, as before, N=cardinality(A_n). Then the given algorithm finds A_n in expected time

$$0(n \sum_{j=1}^{2n} b_j/j^2).$$

If the merging and editing take time bounded from below by $a(g(k_1)+(k_2'))$ and $E(g(N)) \geq sb_n$ where g and b_n are defined above, and $a,s > 0$ are constants, then we take at least

$$\gamma n \sum_{j=1}^{n} b_j/j^2$$

expected time for some $\gamma > 0$ and all n large enough (Devroye, 1983).

Thus, the divide-and-conquer method finds A_n in linear expected time if and only if

$$\sum_{j=1}^{\infty} b_j/j^2 < \infty .$$

Examples.

1. The set of maximal vectors.

Merging of two sets of maximal vectors can be achieved in quadratic time by pairwise comparisons, for any dimension d. We can thus take $g(u)= u^2$ in the previous analysis if we merge in this way. If $E(N) \sim a_n \uparrow \infty$, then we can check that the divide-and-conquer algorithm runs in linear expected time if and only if

$$\sum_{j=1}^{\infty} a_j^2/j^2 < \infty .$$

2. Convex hulls in R^2.

Two convex hulls with ordered vertices can be merged in linear time into a convex hull with ordered vertices (Shamos, 1978). Thus if $E(N)=0(a_n)$ and $a_n \uparrow \infty$, then

$$\sum_{j=1}^{\infty} a_j/j^2 < \infty \qquad (2)$$

is sufficient for the linear expected time behavior of the divide-and-conquer algorithm given here. When $\lim \inf E(N)/a_n > 0$, then (2) is also necessary for linear expected time behavior. Notice here that (2) is satisfied when, say, $a_n=n/\log^{1+\delta} n$ or $a_n=n/(\log n.\log^{1+\delta} \log n)$ for some $\delta > 0$. This improves the sufficient condition $a_n=n^{1-\delta}$, $\delta > 0$, given in Bentley and Shamos (1978).

3. Convex hulls in R^d.

Merging can trivially be achieved in polynomial (n^{d+1}) time for two convex hulls with total number of vertices equal to n. When $E(N) \leq a_n \uparrow$ and

$$\sum_{j=1}^{\infty} a_j^{d+1}/j^2 < \infty,$$

we can achieve linear expected time. This condition is fulfilled for the normal density in R^d and the uniform density on any hypercube of R^d.

4. THE QUICK ELIMINATION (THROW-AWAY) PRINCIPLE.

In extremal problems (e.g., find the convex hull, find the minimal covering ellipse, etc.) many of the data points can be eliminated from further considerations without much work. The remaining data points then enter the more involved portion of the algorithm. Often the worst-case time of these elimination algorithms is equal to the worst-case time of the second part of the algorithm used on all n data points. The expected time is sometimes considerably smaller than the worst-case time. We illustrate this once again on our prototype problem of finding the convex hull.

Examples.

   1.   The convex hull.

Assume that we seek the extrema $e_1,\ldots,e_m$ in m carefully chosen directions of R^d, form the polyhedron P formed by these extrema, and eliminate all X_i's that belong to the interior of the extremal poly-hedron P. The remaining X_i's are then processed by a simple worst-case $O(g(n))$ convex hull algorithm. What can we say about the expected time of these algorithms? An expected time of $o(g(n))$ would indicate that the elimination procedure is worthwhile on large data sequences. We could also say that the elimination procedure achieves 100% asymptotical efficiency. In Devroye (1981c) it is shown that this happens when (i) the open half-spaces defined by the hyperplanes through the origin perpendicular to the e_i's cover R^d except possibly the origin; and (ii) X_1 has a radial density f, where

$$\alpha(u) = \inf\{t:P(||X_1||>t)=u\}$$

is slowly varying at 0 ($\lim_{t\downarrow 0} \alpha(tu)/\alpha(t) = 1$, all u>0), and $\alpha(u) \to \infty$ as u↓0. Condition (i) holds when the e_i's are determined by the d+1 vertices of the regular (d+1)-vertex simplex in R^d centered at the origin. One could also take 2d directions defined by $(0,0,\ldots,0,\pm1)$, etc. Condition (ii) is satisfied by the normal density and a class of radial exponential densities (Johnson and Kotz, pp. 298).

The previous result can be sharpened in specific instances. For example, if f is normal, g(n)=n log n, and (i) holds, then the expected time is $0(n)$. Furthermore, the expected time spent by the algorithm excluding the elimination is o(n) (Devroye, 1981c).

When d=2 and f is bounded away from 0 and infinity on a non-degenerate rectangle of R^2 (f=0 elsewhere), and $e_1,\ldots,e_8$ are equi-spaced directions, then the expected time of the elimination algorithm is $0(n)$ even when $g(n)=n^2$ (Devroye and Toussaint, 1981).

Eddy (1977) has given a slightly different elimination algorithm in which the number of directions and the directions themselves depend upon the data. Akl and Toussaint (1978) report that in R^2, for certain distributions, almost all the elimination algorithms achieve extremely fast average times provided that $e_1,\ldots,e_m$ are easily computed (e.g., they are axial or diagonal directions).

2.   Finding a simple superset.

Assume that we wish to find $A_n=A(X_1,\ldots,X_n)$ in the following manner: (i) Find a set $B_n=B(X_1,\ldots,X_n)$ where B_n is guaranteed to contain A_n, in expected time T_n; (ii) Given that cardinality (B_n) is equal to N, find A_n from B_n in worst-case time bounded by g(N). Note that the expected time of the entire elimination algorithm is bounded by

$$T_n + E(g(N)).$$

2.1. A_n is the convex hull, B_n is the set of maximal vectors.

We discussed some distributions for which $T_n = 0(n)$. In R^2, step (ii) can be executed with $g(n)=n^2$ (Jarvis' algorithm, 1973) or g(n)=n log n (Graham's algorithm, 1972). Thus, the entire algorithm takes expected time $0(n)$ when $E(N \log_+ N)=0(n)$ or $E(N^2)=0(n)$, according to the algorithm selected in step (ii). When the components of X_1 are independent and X_1 has a density, then these conditions are satisfied by the results of Devroye (1980, 1983) given in section 3. The linearity is not lost in this case in R^d even when $g(n)=n^{d+1}$ in step (ii).

2.2. A_n is the diameter of $X_1,\ldots,X_n$; B_n is the convex hull.

$A_n = \{X_i, X_j\}$ is called a diameter of $X_1,\ldots,X_n$ when $||X_k - X_m||$ $\leq$ $||X_i - X_j||$ for all $1 \leq k \leq m \leq n$, $(k,m) \neq (i,j)$. Given B_n, some A_n can be found by comparing all $\binom{N}{2}$ distances between points in B_n, (see Bhattacharya (1980) for an in-depth treatment of the diameter problem, and a survey of earlier results). But B_n can be found in linear expected time for many distributions. In such cases, our trivial diameter algorithm runs in linear expected time provided that $E(N) = O(\sqrt{n})$. Assume for example that f is the uniform density in the unit hypersphere of R^d, then the trivial diameter algorithm runs in linear expected time if and only if (i) the convex hull can be found in linear expected time, and (ii) $d \leq 3$ (section 3, example 1 (iii)).

2.3 A_n is the minimum covering circle, B_n is the convex hull.

The minimum area circle in R^2 covering $X_1,\ldots,X_n$ has either three convex hull points on its perimeter, or has a diameter determined by two convex hull points. Again, it can be found (trivially) from the convex hull in worst-case time $O(n^4)$ (see Elzinga and Hearn (1972, 1974), Francis (1974) and Shamos (1978) for $O(n^2)$ algorithms and subsequent discussions). Thus, A_n can be identified in linear expected time if the convex hull B_n can be found in linear expected time, and if $E(N) = O(\sqrt{n})$ (or $O(n^{1/4})$, if the trivial algorithm is used). In fact, by using the $O(n \log n)$ algorithms of Preparata (1977) or Shamos (1978), the condition on $E(N)$ can be relaxed even further.

5. REFERENCES.

[1]   S.G. Akl and H. Meijer, Recent advances in hybrid sorting algorithms, Utilitas Mathematica, 21 (1982) 325-343.

[2]   S.G. Akl and G.T. Toussaint, A fast convex hull algorithm, Information Processing Letters, 7 (1978) 219-222.

[3]   Avis, D., On the complexity of finding the convex hull of a set of points, Technical Report SOCS 79.2, School of Computer Science, McGill University (1979).

[4]   O. Barndorff-Nielsen and M. Sobel, On the distribution of the number of admissible points in a vector random sample, Theory of Probability and its Applications, 11 (1966) 249-269.

[5]   J.L. Bentley and J.H. Friedman, Data Structures for range searching, Computing Surveys, 11 (1979) 398-409.

[6]   J.L. Bentley and M.I. Shamos, Divide and conquer for linear expected time, Information Processing Letters, 7 (1978) 87-91.

[7]   J.L. Bentley, B.W. Weide and A.C. Yao, Optimal expected-time algorithms for closest point problems, ACM Transactions of Mathematical Software, 6 (1980) 563-580.

[8]   Bhattacharya, B., Applications of computational geometry to pattern recognition problems, Ph.D. Dissertation, School of Computer Science, McGill University (1980).

[9]   K.Q. Brown, Voronoi diagrams from convex hulls, Information Processing Letters, 9 (1979) 227-228.

[10] H. Carnal, Die konvexe Hülle von n rotationssymmetrische verteilten Punkten, Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete 15 (1970) 168-176.

[11] D. Cheriton and R.E. Tarjan, Finding minimum spanning trees, SIAM Journal of Computing, 5 (1976) 724-742.

[12] L. Devroye, A note on finding convex hulls via maximal vectors, Information Processing Letters, 11 (1980) 53-56.

[13] Devroye, L., Some results on the average time for sorting and searching in R^d, Manuscript, School of Computer Science, McGill University (1981a).

[14] L. Devroye, On the average complexity of some bucketing algorithms, Computers and Mathematics with Applications, 7 (1981b) 407-412.

[15] L. Devroye, Moment inequalities for random variables in computational geometry, Computing, 30 (1983) 111-119.

[16] L. Devroye, How to reduce the average complexity of convex hull finding algorithms, Computers and Mathematics with Applications, 7 (1981c) 299-308.

[17] L. Devroye and T. Klincsek, Average time behavior of distributive sorting algorithms, Computing, 26 (1981) 1-7.

[18] L. Devroye and G.T. Toussaint, A note on linear expected time algorithms for finding convex hulls, Computing, 26 (1981) 361-366.

[19] W.F. Eddy, A new convex hull algorithm for planar sets, ACM Transactions of Mathematical Software, 3 (1977) 398-403.

[20] J. Elzinga and D. Hearn, The minimum covering sphere problem, Management Science, 19 (1972) 96-104.

[21] J. Elzinga and D. Hearn, The minimum sphere covering a convex polyhedron, Naval Research Logistics Quarterly, 21 (1974) 715-718.

[22] Francis, R.L. and White, J.A., Facility Layout and Location: An Analytical Approach (Prentice-Hall, 1974).

[23] R. Graham, An efficient algorithm for determining the convex hull of a finite planar set, Information Processing Letters, 1 (1972) 132-133.

[24] Horspool, R.N., Constructing the Voronoi diagram in the plane, Technical Report SOCS 79.12, School of Computer Science, McGill University (1979).

[25] R.A. Jarvis, On the identification of the convex hull of a finite set of points in the plane, Information Processing Letters, 2 (1973) 18-21.

[26] Johnson, N.L. and Kotz, S., Distributions in Statistics: Continuous Multivariate Distributions (John Wiley, New York, 1972).

[27] Knuth, D., The Art of Computer Programming, vol. 3: Sorting and Searching (Addison-Wesley, Reading, Mass., 2nd Ed., 1975).

[28] Lipton, R.J. and Tarjan, R.E., Applications of a planar separator theorem, 18th Annual IEEE Symposium on the Foundation of Computer Science (1977) 162-170.

[29] Preparata, F., Steps into computational geometry, Report R-760, Coordinated Sciences Laboratory, University of Illinois (1977).

[30] H. Raynaud, Sur le comportement asymptotique de l'enveloppe convexe d'un nuage de points tirés au hasard dans R^n, Comptes Rendus de l'Académie des Science de Paris, 261 (1965) 627-629.

[31] A. Renyi and R. Sulanke, Uber die konvexe Hülle von n zufällig gewählten Punkten I, Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete, 2 (1963) 75-84.

[32] A. Renyi and R. Sulanke, Uber die konvexe Hülle von n zufällig gewählten Punkten II, Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete 3 (1964) 138-147.

[33] A. Renyi and R. Sulanke, Zufällige konvexe Polygone in einem Ringgebiet, Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete 9 (1968) 146-157.

[34] R. Sedgewick, The analysis of quicksort programs, Acta Informatica, 7 (1977) 327-355.

[35] R. Sedgewick, Implementing quicksort programs, Communications of the ACM, 21 (1978) 847-857.

[36] Shamos, M.I., Computational geometry, Ph.D. Dissertation, Dept. of Computer Science, Yale University (1978).

[37] Shamos, M.I., Seminar given at McGill University (1979).

[38] Shamos, M.I. and Hoey, D., Closest-point problems, Proceedings of the 16th IEEE Symposium on the Foundations of Computer Science (1975) 151-162.

[39] Toussaint, G.T., Pattern recognition and geometrical complexity, Proceedings of the 5th International Conference on Patter Recognition and Image Processing (Miami, Florida, 1980).

[40] Weide, B.W., Statistical methods in algorithm design and analysis, Ph.D. Dissertation, Department of Computer Science, Carnegie-Mellon University (1978).

[41] A.C. Yao, An $O(|E|\log\log|V|)$ algorithm for finding minimum spanning trees, Information Processing Letters, 4 (1975) 21-23.

[42] A.C. Yao, On constructing minimum spanning trees in k-dimensional space and related problems, SIAM Journal on Computing, 11 (1982) 721-736.

[43] A.C. Yao, A lower bound to finding convex hulls, Journal of the ACM, 28 (1981) 780-787.