アルゴリズム 7－1
（1989. 5. 22）

n 点から分散, 直径最小の k 点を求めるアルゴリズム

今井浩　　加藤直樹
九州大学　　神戸商科大学

平面上の与えられた n 点から分散, 直径最小の k 点を求める問題に対して、それぞれ $O(k^2n + n\log n)$, $O(k^{2.5}n\log k + n\log n)$ 時間のアルゴリズムを与える。両問題で、高次の Voronoi 図を利用し、直径最小の k 点を求める問題では、幾何的制約を有する 2 部グラフ上での最大独立点集合問題を幾何情報を十分利用しながら用いる。その過程で、高次の Voronoi 図に関する新たな数え上げの結果も示す。

Algorithms for Finding k Points with Minimum Variance and Diameter

Hiroshi Imai　　Naoki Katoh
Kyushu University　　Kobe University of Commerce

The problems of finding a subset of k points with minimum variance and diameter among n given points in the plane are considered, and $O(k^2n + n\log n)$-time and $O(k^{2.5}n\log k + n\log n)$-time, respectively, algorithms are given. Both algorithms make use of higher order Voronoi diagrams, and, the algorithm for finding k points with minimum diameter utilizes an independent set algorithm on a bipartite graph obtained from points by geometric conditions. New extremal results on the higher order Voronoi diagram are also established, which are used in the analysis of the algorithms.

1 Introduction

We consider the problem of selecting a specified number of points, k, from a given set S, subject to some optimization criterion. Problems of this type often arise in statistical clustering and pattern recognition. From an algorithmic standpoint, these problems usually can be solved in time $O(n^{k+c})$, where n is the number of points in S and c a small constant. Observe that for arbitrary k this time complexity is exponential in the size of the input. Finding general methods to solve this problem for a wide variety of optimization criteria is a challenging and elusive goal and, except for a paper by Dobkin, Drysdale and Guibas [3], the study of this problem has been conducted mainly for fixed values of k. In this paper, we follow the lead of [3] and study the general case of the problem for several natural criteria of optimization. In particular, we give efficient algorithms for the following problems:

1. (*k-diameter*) Find a set of k points with a minimum diameter; the *diameter* of a set is the maximum distance between any two points of the set. Our algorithm takes $O(k^{2.5}n\log k + n\log n)$ time.

2. (*k-variance*) Find a set of k points with a minimum variance; the *variance* of a set is the sum of squares of the distances of all pairs of points in the set divided by the number of points in the set. Our algorithm takes $O(k^2n + n\log n)$ time.

The general idea behind our algorithms is the use of higher-order Voronoi diagrams. This approach is a fairly natural one, and it was also adopted by Dobkin, Drysdale and Guibas [3] for finding a smallest convex polygon containing at least k points. In deriving our algorithms, we add several new features to this basic approach, and also establish new combinatorial facts about the higher-order Voronoi diagrams, which may have independent interest.

2 The k-Variance Problem

Given a set S of n points in the plane, we wish to select k points from S with a minimum variance, where the variance of a set is defined as the sum of squares of the pairwise distances divided by the number of points in the set. We show that this problem can be solved efficiently by using the order k Voronoi diagram. We begin by deriving a useful inequality for the variance.

Let S_k^* be a set of k points with a minimum variance, and let $c^* \equiv (x^*, y^*)$ be the centroid of S_k^*, i.e.

$$x^* = \frac{1}{k}\sum_{p_i \in S_k^*} x_i \quad \text{and} \quad y^* = \frac{1}{k}\sum_{p_i \in S_k^*} y_i,$$

where x_i and y_i are the coordinates of p_i. Then we have the following relations:

$$\begin{aligned}
\mathrm{Var}(S_k^*) &= \frac{1}{k}\sum_{p_i,p_j \in S_k^*,\ i<j}((x_i - x_j)^2 + (y_i - y_j)^2) \\
&= \sum_{p_i \in S_k^*}(x_i^2 + y_i^2) - k(x^*)^2 - k(y^*)^2 \\
&= \sum_{p_i \in S_k^*}((x_i - x^*)^2 + (y_i - y^*)^2) \qquad (1) \\
&< \sum_{p_i \in S_k^*}((x_i - x)^2 + (y_i - y)^2), \qquad (2)
\end{aligned}$$

for any $(x, y) \neq (x^*, y^*)$. The last inequality follows from the fact that the function

$$f(x,y) = \sum_{p_i \in S_k^*}((x_i - x)^2 + (y_i - y)^2)$$

has a unique minimum, which occurs at (x^*, y^*). We can now prove the following lemma.

Lemma 2.1. Let $S_k^* \subseteq S$ be a set of k points with a minimum variance, and let $c^* \equiv (x^*, y^*)$ be the centroid of S_k^*. Then, the k nearest neighbors of c^* in S are precisely the points of S_k^*. Consequently, in the order k Voronoi diagram of S, there is a nonempty region associated with S_k^*.

Proof. Let p_j be the farthest point from c^* in S_k^*. Suppose, for a contradiction, that there exists a point $p_l \in S - S_k^*$ that is no farther from c^* than p_j. Define $S_k' = (S_k^* - \{p_j\}) \cup \{p_l\}$. Then, using the equation (1), we have

$$\begin{aligned}
\mathrm{Var}(S_k^*) &= \sum_{p_i \in S_k^*}((x_i - x^*)^2 + (y_i - y^*)^2) \\
&\geq \sum_{p_i \in S_k'}((x_i - x^*)^2 + (y_i - y^*)^2) \\
&> \mathrm{Var}(S_k'),
\end{aligned}$$

since $c^* \equiv (x^*, y^*)$ is not the centroid of S_k' (cf. inequality (2)). But this contradicts our assumption that S_k^* has the minimum variance. □

Theorem 2.1 clearly holds in arbitrary dimensions. We therefore have the following simple algorithm for solving the k-variance problem: construct the kth order Voronoi diagram of S; compute the variance of sets of k points associated with the regions of the Voronoi diagram; and select the set with a minimum variance. Given the kth order Voronoi diagram of S, we can compute the variance of various sets of points at constant time apiece, by using the relation (1). More precisely, let A and B be two neighboring regions in the Voronoi diagram. Then the sets of points associated with them, namely, $S \cap A$ and $S \cap B$, differ by exactly one point. Given the centroid of $S \cap A$, the centroid of $S \cap B$ can be easily computed in $O(1)$ time and, consequently, the relation (1) allows the variance of $S \cap B$ to be determined from the variance of $S \cap A$ in constant time. The running time of our algorithm for solving the k-variance problem is therefore dominated by the time for constructing the kth order Voronoi diagram. In particular, we have the following theorem.

Theorem 2.1. Given a set S of n points in a d-dimensional space, we can solve the k-variance problem for S in time $O(f_d(n,k))$, where $f_d(n,k)$ is the

time complexity of computing the kth order Voronoi diagram of n points in d dimensions. □

The time bounds for $f_d(n,k)$ are $f_2(n,k) = O(k^2 n + n \log n)$ [1] and $f_d(n,k) = O(n^{d+1})$ for $d \geq 3$.

3 The k-Diameter Problem

The k-diameter problem asks for a set of k points having a minimum diameter. In this section, we present an $O(k^{2.5} n \log k + n \log n)$ time algorithm to solve this problem. In Subsection 3.1, we show a connection between the k-diameter problem and the $(3k-3)$th order Voronoi diagram of S. In Subsection 3.2, we reduce our problem to the problem of finding a maximum independent set in a certain bipartite graph, and present an $O(k^{4.5} n \log k + n \log n)$ time algorithm for the problem. This bound is subsequently improved in Subsection 3.3 by exercising greater care in our use of the Voronoi diagram; in particular, this involves proving some new combinatorial bounds on the number of distinct pairs and triples whose points belong to the same region of a higher order Voronoi diagram.

3.1 The k-diameter Problem and Higher Order Voronoi Diagrams

Motivated by our success with the kth order Voronoi diagram in the k-variance problem, we are tempted to use the same idea for the present problem as well. Unfortunately, the following example shows that the kth order Voronoi diagram may not always help in solving the k-diameter problem. Given a regular hexagon $abcdef$, let $S(x)$ be a set of m points contained in a ball of radius ϵ centered at x, for $x \in \{a,c,e\}$, where $\epsilon > 0$ is a suitably small number. Similarly, let $S'(x)$ be a set of $m-1$ points lying inside the triangle $\triangle bdf$ at distance 2ϵ from x, for $x \in \{b,d,f\}$. It is easily seen that $S_k \equiv S(a) \cup S(c) \cup S(e)$ is the unique set of $k \equiv 3m$ points having a minimum diameter but there is no nonempty region associated with S_k in the order k Voronoi diagram of the given set of points; this follows because any point in the plane is closer to at least one point of $S'(b) \cup S'(d) \cup S'(f)$ than it is to all of S_k.

Consequently, if there is a relationship between the k-diameter problem and the Voronoi diagram, then it must exist only for orders higher than k. Fortunately, the following lemma shows that we need not look beyond diagrams of order $(3k-3)$. We introduce some notation.

The mth order Voronoi diagram of S is denoted by $V_m(S)$; recall that this diagram is a partition of the plane into convex regions, where each region is associated with a set of m points of S such that, for any point x in the region, the m nearest neighbors of x in S are precisely these m points. The set of points of S associated with a region of $V_m(S)$ is called a *voronoi set of* $V_m(S)$. Observe that a voronoi set of $V_m(S)$ has cardinality m and there are $O(mn)$ such sets (see e.g. Lee [6]). We now prove the key lemma of this section.

Lemma 3.1. Let $S_k \subseteq S$ be a set of k points with a minimum diameter, for $k \leq n/3$. Then S_k is contained in some voronoi set of $V_{3k-3}(S)$.

Proof. A set of points T is a voronoi set of $V_{|T|}(S)$ if and only if there exists a circle C containing T and no other point of S. To prove the lemma, we exhibit such a set $T \subseteq S$ of size $|T| \leq 3k-3$ such that $S_k \subseteq T$. Let C be the minimum enclosing circle of S_k. First consider the case where C is determined by two points of S_k, say, p and q. In this case C cannot contain any point of $S - S_k$ in its interior; if there were a point $s \in S - S_k$ lying in C, we could replace p, or q, by s to obtain another set of k points whose diameter is strictly smaller than the diameter of S_k. In this case, S_k is obviously contained in a voronoi set of $V_m(S)$ (in particular, the voronoi set associated with the Voronoi region containing the center of the circle C). So let us now assume that C is determined by three points of S_k. We divide the circle into three sectors by adding line segments from the center to the three points determining C. Observe that the diameter of each of the three sectors is less than the diameter of S_k. If C contains more than $3k-2$ points of S, then one of the sectors has at least k points, which gives a set of k points with a smaller diameter than S_k. Thus, C has at most $3k-3$ points, and the proof is completed. □

By Lemma 3.1, we can solve the k-diameter problem by examining all $\binom{3k-3}{k}$ subsets of each of the $O(kn)$ voronoi sets of $V_{3k-3}(S)$. This time complexity, although polynomial in n, is still exponential in k. In the following subsection, we develop a fully polynomial algorithm for the problem.

3.2 A Polynomial Time Algorithm

Given a nonnegative number d, let $G(d)$ be the graph defined as follows. The vertices of $G(d)$ are the points of S and there is an edge between two points, p_i and p_j, if $\|p_i - p_j\| \leq d$, where $\|\cdot\|$ denotes the Euclidean norm. We define the *lune* of p_i and p_j to be the intersection of two circles of radius $\|p_i - p_j\|$, one centered at p_i and the other at p_j. The graph $G(p_i,p_j)$ is then defined as the restriction of $G(\|p_i - p_j\|)$ to the lune of p_i and p_j: its vertices are the points of S that lie in the lune of p_i and p_j and two vertices are connected by an edge if their distance is at most $\|p_i - p_j\|$. Then the following observation is straightforward.

Observation 3.1. There is a set of k points in S whose diameter is determined by the points p_i and p_j if and only if there is a clique of size k in $G(p_i,p_j)$ that includes p_i and p_j.

Let $\overline{G}(p_i,p_j)$ be the complement graph of $G(p_i,p_j)$. Then a clique of size k in $G(p_i,p_j)$ corresponds to an independent set of the same size in $\overline{G}(p_i,p_j)$. In the following, we present an efficient algorithm for finding a maximum independent set of $\overline{G}(p_i,p_j)$. The following lemma is straightforward.

Lemma 3.2. The graph $\overline{G}(p_i,p_j)$ is bipartite.

Proof. Consider the two half-lunes obtained from the lune of p_i and p_j by drawing the line segment from p_i to p_j. Points in either half-lune form an independent

set of $\overline{G}(p_i, p_j)$ since the diameter of each half-lune is at most $\|p_i - p_j\|$. The points lying on opposite sides of the line from p_i to p_j form a bipartition of the vertices of $\overline{G}(p_i, p_j)$. □

The following lemma describes an efficient procedure for finding a maximum independent set of $\overline{G}(p_i, p_j)$.

Lemma 3.3. A maximum independent set of the graph $\overline{G}(p_i, p_j)$ can be found in $O(n^{1.5} \log n)$ time.

Proof. If G is a bipartite graph on n vertices, then a maximum independent set of G can be found in time $O(n^{1.5}Q(n)+n^{0.5}P(n))$ (see Imai and Asano [5]), provided that the following two operations, F (find) and D (delete), can be dynamically implemented for G in time $O(Q(n))$ after $O(P(n))$ time preprocessing:

F: given a vertex v, find an edge incident to v in the current graph;

D: given a vertex v, delete v and all its incident edges from the graph;

In our graph $\overline{G}(p_i, p_j)$, two points have an edge between them if and only if their distance is greater than $\|p_i - p_j\|$. Thus, we need a dynamic data structure for storing a set of point S such that the following operations are efficiently implemented: (F) given a point v, find a point of S whose distance from v is greater than $\|p_i - p_j\|$, and (D) delete a point from S. We use a data structure, called *circular hull*, that was recently proposed by Hershberger and Suri [4]. This data structure needs $O(n \log n)$ time for construction and $O(n)$ space for storage, and it supports the above operations at an amortized cost of $O(\log n)$ each. Consequently, a maximum independent set of $\overline{G}(p_i, p_j)$ can be computed in time $O(n^{1.5} \log n)$ time and $O(n)$ space. □

Observe, however, that we really do not need to compute a maximum independent set; we only need to check if the size of the independent set is at least k. The following lemma shows that if the number of vertices in $\overline{G}(p_i, p_j)$ is at least $2k$, then there always exists an independent set of size k.

Lemma 3.4. If $\overline{G}(p_i, p_j)$ has at least $2k$ vertices, then there exists an independent set of size at least k, and it can be found in linear time.

Proof. Consider the two half-lunes obtained from the lune of p_i and p_j by drawing the line segment from p_i to p_j. The set of points contained in the same half-lune forms an independent set. Since there are at least $2k$ points in the lune, one of the half-lunes must contain at least k points, which proves the lemma. □

Therefore, the problem of deciding whether $\overline{G}(p_i, p_j)$ contains an independent set of size k can be solved in time $O(n + k^{1.5} \log k)$. (In $O(n)$ time we count the number of points in the lune of p_i and p_j; if this number is at least $2k$, we have an independent set of size k, otherwise we use Lemma 3.3 to find a maximum independent set.) The k-diameter problem of S can be solved by repeating this process for each of the $O(n^2)$ pairs of S. This leads to the following theorem.

Theorem 3.1. The k-diameter problem for n points in the plane can be solved in $O(n^3 + k^{1.5}n^2 \log k)$ time and $O(n)$ space. □

By combining Theorem 3.1 with Lemma 3.1, we obtain the following theorem.

Theorem 3.2. The k-diameter problem for n points in the plane can be solved in $O(k^{4.5}n \log k + n \log n)$ time and $O(kn)$ space.

Proof. We compute the $(3k-3)$th order Voronoi diagram of S; this takes $O(k^2n + n \log n)$ time. By Lemma 3.1, we can solve the k-diameter problem of S by solving the problem for each of the voronoi sets of $V_{3k-3}(S)$. Since there are $O(kn)$ voronoi sets, each set has size $3k-3$, and we can solve the problem for each voronoi set in time $O(k^{1.5} \log k)$, the overall time complexity of our algorithm is

$$O(nk \cdot k^2 \cdot k^{1.5} \log k) \;=\; O(k^{4.5}n \log k),$$

plus the cost of computing the Voronoi diagram. This completes the proof. □

3.3 Improving the Time Complexity

Our final improvement to the time complexity comes from a more careful counting of the total number of pairs of points in all the voronoi sets of $V_{3k-3}(S)$. Theorem 3.1 derives its bound by checking $O(nk^3)$ pairs of points; $\binom{3k-3}{2}$ pairs in each of the $O(nk)$ voronoi sets. We show in this section that this bound is overly pessimistic and, in fact, we need to check only $O(nk)$ pairs. This fact will be proved in several stages. We begin with a definition.

Let $U_m(p_i)$ denote the union of the closed Voronoi regions of $V_m(S)$ whose voronoi sets contain $p_i \in S$. The following observation is quite straightforward.

Observation 3.2. $U_m(p_i)$ is simply connected and it is star-shaped with respect to p_i. □

We use b_{ij} to denote the perpendicular bisector of the line segment from p_i to p_j. Then we have the following key lemma.

Lemma 3.5. Let $m < n/2$; let $p_i, p_j \in S$ be two points such that $U_m(p_i) \cap U_m(p_j) \neq \emptyset$; and let U be a connected component of $U_m(p_i) \cap U_m(p_j)$. Then there is a Voronoi edge of $V_m(S)$ that is part of b_{ij} and that has a nonempty intersection with U.

Proof. We prove the lemma in two parts: (a) the boundary of U intersects with b_{ij} and (b) this intersection point is contained in a Voronoi edge of $V_m(S)$. To prove (a), consider a point $p \in U$ and assume without loss of generality that p and p_j are on opposite sides of b_{ij}. Let q be the point where the line segment from p to p_j meets b_{ij}. Since $U_m(p_j)$ is star-shaped with respect to p_j, the line segment from p to p_j lies in U_j. Therefore, if we move a point x from p to p_j, then the m nearest neighbors of x always include p_j. For any position of x between p and q, p_i must also be one of the m nearest neighbors of x since x is on the same side of b_{ij} as p_i. This implies that the segment from p to q is contained in $U_m(p_i)$, and consequently q is in the same component of $U_m(p_i) \cap U_m(p_j)$ as p,

namely, U. Since q lies on b_{ij}, this shows that U intersects b_{ij}. To complete the proof of (a), it remains to prove that b_{ij} cannot be completely contained inside U. If b_{ij} were to lie entirely within U, the m nearest neighbors of *any* point $x \in b_{ij}$ must include both p_i and p_j. We consider the two extreme points of b_{ij}, namely, $x = +\infty$ and $x = -\infty$, and denote by H^+ (resp. H^-) the half-plane determined by p_i and p_j and containing $+\infty$ (resp. $-\infty$). If p_i and p_j are among the m nearest neighbors of both $x = +\infty$ and $x = -\infty$, then neither H^+ nor H^- may contain more than m points. But that is impossible since $H^+ \cup H^-$ is the entire plane and $m < n/2$. This completes the proof of (a).

To prove (b), let u be the point where b_{ij} intersects the boundary of U, and assume without loss of generality that u belongs to the boundary of $U_m(p_i)$. Let e be an edge of the Voronoi diagram $V_m(S)$ that contains u, and let p_i and p_l be the two points whose bisector contains e. If $l = j$, then the claim easily follows: e is the desired edge. Otherwise, u is equidistant from p_i, p_j and p_l. We show that u is a Voronoi vertex of $V_m(S)$ determined by p_i, p_j, and p_l, as follows. Consider a point v close to the edge e but just outside the region $U_m(p_i)$. Since $e \subseteq b_{il}$, the set of m nearest neighbors of v is obtained by replacing p_i with p_l in the neighbor-set of u. That is, p_i is no longer one of the m nearest neighbors of v, but p_j is. This, however, cannot be the case for all positions of v along e: in particular, if v is chosen on the same side of b_{ij} as p_i, then v is closer to p_i that it is to p_j. Consequently, there must be a discontinuity at the point u where the bisectors b_{il} and b_{ij} meet. Therefore, the point u is a Voronoi vertex of $V_m(S)$ determined by p_i, p_j and p_l and there is a Voronoi edge incident to u that is part of b_{ij}. This completes the proof. □

Lemma 3.6. The total number of distinct pairs (p, q) such that both p and q are contained in a voronoi set of $V_m(S)$ is $O(mn)$.

Proof. If $m \geq n/2$, then the lemma is trivial; there are only $O(n^2)$ distinct pairs. Otherwise, observe that if p and q are in a voronoi set, then $U_m(p) \cap U_m(q) \neq \emptyset$. By Lemma 3.5, the number of such pairs is bounded above by the number of edges in the Voronoi diagram $V_m(S)$, which is $O(mn)$. □

Lemma 3.7. If $m < n/2$ and $p_i, p_j \in S$ are two distinct points, then neither $U_m(p_i)$ nor $U_m(p_j)$ is contained in the other.

Proof. If $U_m(p_i) \cap U_m(p_j) \neq \emptyset$, then Lemma 3.5 guarantees that b_{ij} contributes a Voronoi edge to $V_m(S)$, say, e_{ij}. The points in the neighborhood of e_{ij} lying on the same side as p_i are contained in $U_m(p_i)$ but not in $U_m(p_j)$ and vice versa, and therefore neither region can be contained in the other. □

Lemma 3.8. Let $m < n/2$; let $p_i, p_j, p_l \in S$ be three points such that $U_m(p_i) \cap U_m(p_j) \cap U_m(p_l) \neq \emptyset$; and let U be a connected component of $U_m(p_i) \cap U_m(p_j) \cap U_m(p_l)$. Then there is a Voronoi vertex v of $V_m(S)$ such that at least two of the three points determining v are in $\{p_i, p_j, p_l\}$, and v lies in U.

Proof. Let U_{ij} be the connected component of $U_m(p_i) \cap U_m(p_j)$ that contains U. By Lemma 3.5, there is a Voronoi vertex v in U_{ij} determined by p_i, p_j and some other point p_h. If v also belongs to U_l, then we are finished. Otherwise, since U_l cannot be completely contained in $U_m(p_i) \cap U_m(p_j)$ (cf. Lemma 3.7), the boundaries of U_l and U_{ij} intersect, say, at a point u. Since the boundaries of U_l and U_{ij} consist of Voronoi edges, u is necessarily a Voronoi vertex, and its determining points include p_l and p_i (resp. p_l and p_j) if u is common to the boundaries of $U_m(p_l)$ and $U_m(p_i)$ (resp. $U_m(p_l)$ and $U_m(p_j)$). This completes the proof. □

Lemma 3.9. The total number of distinct triples (p, q, r) such that p, q and r are contained in a voronoi set of $V_m(S)$ is $O(m^2 n)$.

Proof. If $m \geq n/2$, then the lemma is trivial: there are only $O(n^3)$ triples. Otherwise we proceed as follows. Notice that, for any triple (p, q, r) contained in a voronoi set, $U_m(p) \cap U_m(q) \cap U_m(r) \neq \emptyset$. By Lemma 3.8, we can associate each such triple with a Voronoi vertex of $V_m(S)$. We now need to show that a Voronoi vertex is not charged against too many distinct triples. First we observe that the number of different regions $U_m(x)$, for $x \in S$, containing a Voronoi vertex is $m + 2$; this follows because if R_1, R_2 and R_3 are the three Vororoi regions of $V_m(S)$ incident with a Voronoi vertex and if the set of m points associated with R_1 is $\{p_1, p_2, \ldots, p_m\}$, then only the regions $U_m(p_i)$, $i = 1, 2, \ldots, p_m$, may contain the Voronoi region R_1; furthermore, since the voronoi sets of two neighboring Voronoi regions differ by exactly one point, R_1, R_2 and R_3 together have only $m + 2$ different points. Now consider a Voronoi vertex v and suppose that p_a, p_b, p_c are the three points determining v. By Lemma 3.8, if a triple (p, q, r) is associated with v, then (a) $v \in U_m(p) \cap U_m(q) \cap U_m(r)$, and (b) at least two of p, q and r are in $\{p_a, p_b, p_c\}$. There are three ways to choose the two points from $\{p, q, r\}$ as possible determiners of v. For a fixed choice of two determiners, say, p and q, our accounting scheme may assign to v all the triples (p, q, r) that satisfy condition (a). However, as noted above, there are only $m + 2$ choices of x for which v lies in $U_m(x)$. Not counting p and q, we are left with m choices for r. Consequently, the vertex v gets associated with at most $3m$ triples. Since there are altogether $O(mn)$ Voronoi vertices in $V_m(S)$, the number of distinct triples is $O(m^2 n)$. □

Given two points p_i and p_j of S such that $U_m(p_i) \cap U_m(p_j) \neq \emptyset$, let S_{ij} be the union of all voronoi sets that contain both p_i and p_j.

Lemma 3.10. The sum of $|S_{ij}|$ over all pairs p_i and p_j with $U_m(p_i) \cap U_m(p_j) \neq \emptyset$ is $O(m^2 n)$.

Proof. If a point p_l belongs to S_{ij}, then the triple (p_i, p_j, p_l) is such that $U_m(p_i) \cap U_m(p_j) \cap U_m(p_l) \neq \emptyset$. In summing $|S_{ij}|$ over all pairs p_i and p_j, with $U_m(p_i) \cap U_m(p_j) \neq \emptyset$, we count each such triple three times. Thus, three times the number of distinct triples (p, q, r), with $U_m(p) \cap U_m(q) \cap U_m(r) \neq \emptyset$, is an upper bound on the sum of $|S_{ij}|$'s. This quantity is $O(m^2 n)$ as shown by Lemma 3.9. □

Lemma 3.11. The sets S_{ij}'s for all pairs (p_i, p_j) with $U_m(p_i) \cap U_m(p_j) \neq \emptyset$ can be computed in total time $O(m^2 n + n \log n)$.

Proof. We assume that the Voronoi diagram $V_m(S)$ is given in a standard form, where the edges of each face are listed in a cyclic order and the edges incident to each vertex are listed in an angularly sorted order. With each edge e, we store the two points p and q whose bisector contains e and, with each Voronoi vertex v, we store the three points p, q and r that determine v. Computing $V_m(S)$ with this representation takes $O(m^2n + n\log n)$ time [1].

Consider a set S_{ij} and a point p_l that belongs to it. Observe that $U_m(p_i) \cap U_m(p_j) \cap U_m(p_l) \neq \emptyset$. We define p_l to be a *type-one* point of S_{ij} if $U_m(p_i) \cap U_m(p_j) \subseteq U_m(p_l)$ and a *type-two* point otherwise. Observe that if p_l is a type-two point, then the boundaries of $U_m(p_l)$ and $U_m(p_i) \cap U_m(p_j)$ intersect since Lemma 3.7 states that $U_m(p_l) \not\subset (U_m(p_i) \cap U_m(p_j))$. Now, if p_l is a type-one point of S_{ij}, then p_l belongs to every voronoi set whose Voronoi region lies in $U_m(p_i) \cap U_m(p_j)$. Thus we can find all the type-one points of S_{ij} (there are only m of these) by choosing any voronoi set that contains both p_i and p_j. Next, if p_l is a type-two point, then the common intersection of the boundaries of $U_m(p_l)$ and $U_m(p_i) \cap U_m(p_j)$ is a Voronoi vertex of $V_m(S)$. Thus the type-two points of S_{ij} can be discovered by traversing the boundary of $U_m(p_i) \cap U_m(p_j)$; that is, these points are associated with the Voronoi vertices encountered during the traversal.

Since there is one-to-one correspondence between the pairs (p_i, p_j) for which we need to compute S_{ij} and the voronoi edges of $V_m(S)$ (cf. Lemma 3.5), we can explore all distinct pairs by processing all the edges of $V_m(S)$. If e is an edge that is part of the bisector b_{ij}, then we start the computation of S_{ij} by determining a Voronoi region that is incident to e and contains both p_i and p_j. (Notice that ordinarily an edge e is incident to only four regions.) We initialize S_{ij} to be the voronoi set associated with this region. (This takes care of all type-one points of S_{ij}.) We then traverse the boundary of $U_m(p_i) \cap U_m(p_j)$ by visiting adjacent Voronoi regions. Whenever we encounter a Voronoi vertex, we add a new point to S_{ij}; these are the type-two points of S_{ij}. (To check whether or not an edge is on the boundary of $U_m(p_i) \cap U_m(p_j)$, we only need to look at the two points that determine the edge: the edge is on the boundary if and only if p_i or p_j is one of its determiners.) That such a traversal correctly computes S_{ij} follows from the preceding discussion.

Finally, we need to establish the running time of our algorithm. The cost of the algorithm is dominated by the number of edges of $V_m(S)$ traversed during the algorithm. We claim that the algorithm traverses each edge of $V_m(S)$ at most $2m$ times. The proof of the claim is as follows. Given an edge e on the boundary of $U_m(p_i) \cap U_m(p_j)$, consider a point p_l such that e also belongs to the boundary of $U_m(p_i) \cap U_m(p_l)$. Clearly, the intersection of the two boundaries is a Voronoi vertex, and we proved in Lemma 3.9 that there are at most m choices of p_l. Similarly, there are at most m choices of p_l for which e also belongs to the boundary of $U_m(p_j) \cap U_m(p_l)$. Thus the edge e is traversed at most $2m$ times. Since there are $O(mn)$ Voronoi edges in $V_m(S)$, the bound claimed in the lemma follows. □

Theorem 3.3. The k-diameter problem for n points in the plane can be solved in $O(k^{2.5}n\log k + n\log n)$ time and $O(kn)$ space.

4 Concluding Remarks

Given a set of points S, we have presented efficient algorithms for selecting k points from S with a minimum diameter or variance. We also established new combinatorial bounds for the total number of pairs in all the voronoi sets. Specifically, a näive bound for the number of distinct pairs (p, q) such that p and q are contained in some voronoi set of $V_k(S)$ is $O(k^3n)$, where $|S| = n$; this follows because there are $O(nk)$ voronoi sets and $O(k^2)$ pairs per voronoi set. We show that the number of such pairs is in fact $O(kn)$. Generalizing this, we show that the number of distinct triples, with all three points contained in some voronoi set, is $O(k^2n)$. These improved bounds are used to speed up our algorithms, and we expect that they will find other applications.

The underlying theme of all our algorithms is the use of a higher-order Voronoi diagram. They appear to be a natural tool for this class of problems, although it would be interesting to see if fundamentally different techniques exist for these problems.

Acknowledgment

This paper is part of a joint work with Alok Aggarwal, Subhash Suri and the authors. A complete version will appear in [2].

References

[1] A. Aggarwal, L. J. Guibas, J. Saxe and P. Shor, "A Linear Time Algorithm for Computing the Voronoi Diagram of a Convex Polygon," *Proc. of the 19th Annual ACM Symposium on Theory of Computing*, 1987, pp. 39–45.

[2] A. Aggarwal, H. Imai, Naoki Kato and S. Suri, "Finding k Points with Minimum Diameter and Related Problems," to appear in *Proc. of the 5th Annual ACM Symposium on Computational Geometry*, 1989.

[3] D. P. Dobkin, R. L. Drysdale, III, and L. J. Guibas, "Finding Smallest Polygons," in *Advances in Computing Research*, Vol. 1, JAI Press, 1983, pp. 181–214.

[4] J. Hershberger and S. Suri, "Finding Tailored Partitions," *Proc. of the 5th Annual ACM Symposium on Computational Geometry*, 1989.

[5] H. Imai and T. Asano, "Efficient Algorithms for Geometric Graph Search Problems," *SIAM J. Computing*, Vol. 15, No.2, 1986, pp. 478–494.

[6] D. T. Lee, "On k-Nearest Neighbor Voronoi Diagrams in the Plane," *IEEE Transactions on Computers*, Vol. C-31, 1982, pp. 478–487.