

Space-Efficient Approximate Voronoi Diagrams

Sunil Arya
Department of Computer
Science
The Hong Kong University of
Science and Technology
Clear Water Bay, Kowloon,
Hong Kong
arya@cs.ust.hk

Theocharis Malamatos*
Department of Computer
Science
The Hong Kong University of
Science and Technology
Clear Water Bay, Kowloon,
Hong Kong
tmalamat@cs.ust.hk

David M. Mount†
Department of Computer
Science and
Institute for Advanced
Computer Studies
University of Maryland
College Park, Maryland 20742
mount@cs.umd.edu

ABSTRACT

Given a set S of n points in \mathbb{R}^d , a (t, ϵ) -approximate Voronoi diagram (AVD) is a partition of space into constant complexity cells, where each cell c is associated with t representative points of S , such that for any point in c , one of the associated representatives approximates the nearest neighbor to within a factor of $(1 + \epsilon)$. Like the Voronoi diagram, this structure defines a spatial subdivision. It also has the desirable properties of being easy to construct and providing a simple and practical data structure for answering approximate nearest neighbor queries. The goal is to minimize the number and complexity of the cells in the AVD.

We assume that the dimension d is fixed. Given a real parameter γ , where $2 \leq \gamma \leq 1/\epsilon$, we show that it is possible to construct a (t, ϵ) -AVD consisting of

$$O(n\epsilon^{\frac{d-1}{2}} \gamma^{\frac{3(d-1)}{2}} \log \gamma)$$

cells for $t = O(1/(\epsilon\gamma)^{(d-1)/2})$. This yields a data structure of $O(n\gamma^{d-1} \log \gamma)$ space (including the space for representatives) that can answer ϵ -NN queries in time $O(\log(n\gamma) + 1/(\epsilon\gamma)^{(d-1)/2})$. (Hidden constants may depend exponentially on d , but do not depend on ϵ or γ).

In the case $\gamma = 1/\epsilon$, we show that the additional $\log \gamma$ factor in space can be avoided, and so we have a data structure that answers ϵ -approximate nearest neighbor queries in time $O(\log(n/\epsilon))$ with space $O(n/\epsilon^{d-1})$, improving upon the best known space bounds for this query time. In the case $\gamma = 2$, we have a data structure that can answer approximate nearest neighbor queries in $O(\log n + 1/\epsilon^{(d-1)/2})$ time using optimal $O(n)$ space. This dramatically improves the

*The work of the first two authors was supported by the Research Grants Council of Hong Kong, China under project number HKUST 6158/98E.

†This author's work was supported in part by the National Science Foundation under grant CCR-0098151.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'02, May 19-21, 2002, Montreal, Quebec, Canada.
Copyright 2002 ACM 1-58113-495-9/02/0005 ...\$5.00.

previous best space bound for this query time by a factor of $O(1/\epsilon^{(d-1)/2})$.

We also provide lower bounds on the worst-case number of cells assuming that cells are axis-aligned rectangles of bounded aspect ratio. In the important extreme cases $\gamma \in \{2, 1/\epsilon\}$, our lower bounds match our upper bounds asymptotically. For intermediate values of γ we show that our upper bounds are within a factor of $O((1/\epsilon)^{(d-1)/2} \log \gamma)$ of the lower bound.

1. INTRODUCTION

Given a set S of n points in \mathbb{R}^d , called *sites*, the Voronoi diagram of S is a partition of space into cells, such that each cell is the region of space consisting of all points that are closer to a particular site than to any other. Voronoi diagrams are among the most fundamental and well-studied objects in computational geometry.

Voronoi diagrams have numerous applications in areas such as pattern recognition and classification, machine learning, robotics, and graphics. Many of these applications are in high dimensions but, unfortunately, the complexity of Voronoi diagrams can be as high as $n^{\lceil d/2 \rceil}$ in dimension d . Constructing Voronoi diagrams can be complicated due either to numerical inaccuracies or degeneracies, which result when points are cocircular or nearly so. Also, although Voronoi diagrams implicitly encode the information of what site is closest to a given point, by themselves they are not suitable for use as a data structure for nearest neighbor searching. These shortcomings have led to consideration of simpler structures that can be used in place of the Voronoi diagram.

We begin with some definitions. For a real parameter $\epsilon > 0$, we say that a point $p \in S$ is an ϵ -nearest neighbor (ϵ -NN) of a point $q \in \mathbb{R}^d$, if the distance between q and p is at most $(1 + \epsilon)$ times the distance between q and its nearest neighbor in S . We assume that distances are measured in the Euclidean metric. An *approximate Voronoi diagram (AVD)* of S is defined to be a partition of space into cells, where each cell c is associated with a *representative* $r_c \in S$, such that r_c is an ϵ -NN for all the points in c [9]. More generally we may allow up to some given $t \geq 1$ representatives to be stored with each cell, and require that for any point in the cell, one of these t representatives is an ϵ -NN. We refer to such a decomposition as a (t, ϵ) -approximate Voronoi diagram.

Vleugels and Overmars [13] considered approximating the Voronoi diagram of a disjoint set of convex sites in \mathbb{R}^d for the

purpose of motion planning. Har-Peled [9] showed how to construct a $(1, \epsilon)$ -AVD of $O((n/\epsilon^d)(\log n) \log(n/\epsilon))$ size. A cell in this subdivision is the difference of two axis-aligned cubes (the inner cube is optional). Moreover, these cells are stored in a compressed quadtree data structure, which provides a simple and efficient method for answering ϵ -NN queries in $O(\log(n/\epsilon))$ time. This approach is very appealing because, like the Voronoi diagram, it defines a subdivision of space, but there are no problems with geometric degeneracies, and a point-location structure is provided. The construction and search structure are very simple and practical.

The principal shortcoming of Har-Peled's data structure is its size. Sabharwal et al. [12] reduced the size by a logarithmic factor. Arya and Malamatos [1] showed that it is possible to construct a $(1, \epsilon)$ -AVD of $O(n/\epsilon^d)$ size, and they provided a lower bound of $\Omega(n/\epsilon^{d-1})$ on the size of a $(1, \epsilon)$ -AVD. As with Har-Peled their cells are differences of two axis-aligned hyperrectangles. Further, they introduced the notion of allowing multiple representatives per cell and showed that given a real parameter $2 \leq \gamma \leq 1/\epsilon$, it is possible to construct a (t, ϵ) -AVD of $O(n\gamma^d)$ size, where $t = O(1/(\epsilon\gamma)^{(d-1)/2})$. By varying γ it is possible to achieve a tradeoff in space and query time.

To understand the significance of this line of work, it is useful to briefly review the evolution of results in the field of approximate nearest neighbor searching in Euclidean space. There have been two principal approaches. One direction has focused on eliminating exponential dependencies on dimension, as characterized by the work of Indyk and Motwani [10] and Kushilevitz et al. [11]. Although these methods provide space and query times without exponential dependence on dimension, their space requirements grow as $n^{O(1)}$, which may be too high for many applications.

The second direction, which has been developed in computational geometry, is to assume that d is a small constant and n is large. The resulting data structures have sizes that are nearly linear in n but may have factors in space and query time that grow as $(1/\epsilon)^{O(d)}$. Arya et al. [3] and, later, Duncan et al. [8] provided data structures that achieve $O((1/\epsilon)^d \log n)$ query time and use $O(n)$ space (independent of ϵ). These structures were optimal with respect to space, but the ϵ factors in the query times were far from optimal. The cells of the Voronoi diagram are convex polytopes, and results on approximating convex polytopes by Dudley [7] suggest that $(1/\epsilon)^{(d-1)/2}$ should be the proper bound. Indeed, Clarkson [6] showed that queries could be answered in $O((1/\epsilon)^{d/2} \log n)$ time. However Clarkson's space bounds were larger by a factor of at least $(1/\epsilon)^{d/2}$ and involved a parameter that depended on the geometric arrangement of the sites. Through an impressive combination of structures, Chan [5] showed that the query time could be reduced to $O((1/\epsilon)^{(d-1)/2} \log n)$ time using a data structure with space $O((1/\epsilon)^{(d-1)/2} n \log n)$. The recent results of Arya and Malamatos [1] further tightened these bounds by showing that the query time could be improved to $O(\log n + 1/\epsilon^{(d-1)/2})$ and removed the log factor from the space bound.

Finally, it seems that the constant factors are coming exactly in line with bounds suggested by Dudley's results. The one remaining deficiency is the significant factor of $O((1/\epsilon)^{(d-1)/2})$ in the space bound. In this paper we show that this factor can be eliminated altogether. In particular

we provide a data structure of size $O(n)$ that can answer approximate nearest neighbor queries in $O(\log n + 1/\epsilon^{(d-1)/2})$ time. Thus we simultaneously provide optimal space bounds while matching the query time suggested by Dudley's results. Our approach is based on a more space-efficient version of approximate Voronoi diagrams. Our approach differs from those of Har-Peled and Arya and Malamatos principally in how cells are generated and how representatives are chosen. Otherwise, it produces an AVD data structure with the same desirable characteristics.

More specifically, our main result is that, given a set of n sites S in \mathbb{R}^d , $0 < \epsilon \leq 1/2$, and a parameter $2 \leq \gamma \leq 1/\epsilon$, we show how to construct a (t, ϵ) -AVD, where $t = O(1/(\epsilon\gamma)^{(d-1)/2})$, consisting of

$$O(n\epsilon^{\frac{d-1}{2}} \gamma^{\frac{3(d-1)}{2}} \log \gamma)$$

cells. This yields a data structure of $O(n\gamma^{d-1} \log \gamma)$ space (including the space for representatives) that can answer ϵ -NN queries in time $O(\log(n\gamma) + 1/(\epsilon\gamma)^{(d-1)/2})$. The bounds described above on approximate nearest neighbor searching arise with $\gamma = 2$. In the case $\gamma = 1/\epsilon$, the additional $\log \gamma$ factor in space can be avoided, and we have a data structure that answers queries in time $O(\log(n/\epsilon))$ with space $O(n/\epsilon^{d-1})$. These results are presented in Theorem 1 and Corollary 1.1 in Section 3. Note that our bounds are superior to those of [9] and [1] for all values of γ .

We believe that the factor of $\log \gamma$ in the number of cells is an artifact of our proof technique. As evidence of this we provide a significantly different construction and a different analysis based on a charging argument in which this $\log \gamma$ factor in the number of cells is replaced by factor of $\log 2/(\epsilon\gamma)$. We feel that this alternative construction may be of independent interest. For $\gamma > \sqrt{1/\epsilon}$ this provides a superior space bound. These results are presented in Theorem 2 in Section 4.

We also present the first lower bounds on the worst-case number of cells in a (t, ϵ) -AVD for multiple representatives. We assume that cells are based on axis-aligned rectangles of bounded aspect ratio. These results show that our space bounds are tight in the two extremes, $\gamma = 2$ and $\gamma = 1/\epsilon$. For all intermediate values of γ our upper bounds on space are tight to within a factor of $O((1/\epsilon)^{(d-1)/2} \log \gamma)$. These results are presented in Section 5.

Our results come about from a few innovations. The principal reduction in space arises from a form of deterministic sampling based on the BBD-tree. By applying the construction of Arya and Malamatos on an appropriately sampled subset of points, we can decrease the number of cells significantly while increasing the number of representatives by only a constant factor. The main difficulty is avoiding an additional factor of $O(\log n)$ in the number of representatives, as would be expected from well known results on ϵ -nets. Our other innovation is to create cells more economically by considering a well-separated pair decomposition of the points, and concentrating the generation of cells along the bisector between well-separated pairs. Our alternative approach is based on a charging argument which shows that, after certain modifications to Arya and Malamatos's construction, the number of representatives for most of the cells is significantly smaller than the maximum possible. This enables us to reduce the number of cells drastically by combining cells with fewer representatives into larger cells.

2. PRELIMINARIES

Throughout we assume that the dimension d is a fixed constant, and the constants hidden in the asymptotic bounds may depend on d (but not on ϵ or γ). We assume that the set S of points has been scaled and translated to lie within a ball of radius $\epsilon/8$ placed at the center of the unit hypercube $[0, 1]^d$.

Let x and y denote any two points in \mathbb{R}^d . We use $|xy|$ to denote the Euclidean distance between x and y and \overline{xy} to denote the segment joining x and y .

We denote by $b(x, r)$ a ball of radius r centered at x , i.e., $b(x, r) = \{y : |xy| \leq r\}$. For a ball b and any positive real γ , we use γb to denote the ball with the same center as b and whose radius is γ times the radius of b , and \bar{b} to denote the set of points that are not in b . Given a set X of points and a point q , let $NN_q(X)$ be the distance to the nearest neighbor of q in X . If there are no points in X , then $NN_q(X)$ is defined to be infinity.

We briefly review the notions of well-separated pair decomposition and balanced box-decomposition trees, as they play an important role in our constructions.

2.1 The Well-Separated Pair Decomposition

Let S be a set of n points in \mathbb{R}^d . We say that two sets of points X and Y are *well-separated* if they can be enclosed within two disjoint d -dimensional balls of radius r , such that the distance between the centers of these balls is at least αr , where $\alpha \geq 2$ is a real parameter called the *separation factor*. If we consider joining the centers of these two balls by a line segment, the resulting shape resembles a *dumbbell*. The balls are the *heads* of the dumbbell. Define the *length* of a dumbbell to be the distance between the centers of the balls. A dumbbell *separates* two points x and y if x is contained in one head and y in the other.

A *well-separated pair decomposition (WSPD)* of S is a set $\mathcal{P}_{S,\alpha} = \{(X_1, Y_1), \dots, (X_m, Y_m)\}$ of pairs of subsets of S such that (i) for $1 \leq i \leq m$, X_i and Y_i are well-separated and (ii) for any distinct points $x, y \in S$, there exists a unique pair (X_i, Y_i) such that either $x \in X_i$ and $y \in Y_i$ or $x \in Y_i$ and $y \in X_i$. (We say that the pair (X_i, Y_i) *separates* x and y .) Callahan and Kosaraju [4] have shown that we can construct a WSPD containing $O(\alpha^d n)$ pairs in $O(n \log n + \alpha^d n)$ time. For each pair, their construction also provides the corresponding dumbbell. Throughout we will assume that the length of the dumbbell is *exactly* α times the radius of the dumbbell heads.

2.2 The BBD Tree

Let $U = [0, 1]^d$ denote a unit hypercube in \mathbb{R}^d . We define a *quadtree box* recursively as follows: (i) U is a quadtree box, and (ii) any hypercube obtained by splitting a quadtree box into 2^d equal parts is a quadtree box. The *size* of a quadtree box is its side length. A nice property of quadtree boxes is that any two quadtree boxes are either disjoint or one is contained inside the other.

The *balanced box-decomposition (BBD) tree* is a balanced 2^d -ary tree that compactly represents a hierarchical decomposition of space [3]. Each node of the tree is associated with a region of space called a *cell*, which is the difference of two quadtree boxes, an *outer box* and an (optional) *inner box*. The root of the tree is associated with U . The cell associated with any node is partitioned into disjoint cells, which are associated with the children of the node. (For details

see [3].) We define the *size* of a cell to be same as the size of its outer box.

The properties of the BBD tree, relevant to this paper, are given below.

- (i) A set S of n points can be stored in a BBD tree having $O(n)$ nodes and $O(\log n)$ depth.
- (ii) A collection \mathcal{C} of n quadtree boxes can be stored in a BBD tree having $O(n)$ nodes and $O(\log n)$ depth. The subdivision induced by its leaves is a refinement of the subdivision induced by the quadtree boxes in \mathcal{C} and, for any point q , we can determine the leaf containing q in $O(\log n)$ time.
- (iii) The number of cells of the BBD tree with pairwise disjoint interiors, each of size at least s , that intersect a ball of radius r is at most $O((1 + r/s)^d)$.
- (iv) Let $1 \leq t \leq n$. Consider the set of nodes of the BBD tree that have at most t points but whose parents have more than t points. The size of such a set is at most $O(n/t)$.
- (v) Suppose that we assign weights to a subdivision induced by an (unweighted) BBD tree. Let w be the maximum weight of a cell and let W denote the total weight of all the cells. Then it is possible to build a weighted BBD tree for exactly the same set of cells such that the following holds. Given $w \leq t \leq W$, consider a set of nodes of weight at most t but whose parents have weight more than t . The size of such a set is $O(W/t)$.

Properties (i) and (iii) are proved in [3]. Properties (ii) and (v) follow from a generalization of (i). Property (iv) follows from the balancing aspect of the BBD tree.

Throughout we use the following notation. Given a cell c , let s_c denote the size of c , and b_c be the ball of radius $s_c d/2$ whose center coincides with the center of c 's outer box (note that $c \subseteq b_c$).

3. SPACE REDUCTION BY SAMPLING

Let S be a set of n points in \mathbb{R}^d , and let $0 < \epsilon \leq 1/2$ and $2 \leq \gamma \leq 1/\epsilon$ be two real parameters. In this section we show how to construct a (t, ϵ) -AVD, consisting of $O(n\epsilon^{(d-1)/2} \gamma^{3(d-1)/2} \log \gamma)$ cells, where $t = O(1/(\epsilon\gamma)^{(d-1)/2})$.

The construction of AVDs described by Arya and Malamatos [1] is based on two fundamental constructions. The first is that given two concentric balls it is possible to construct a small set of nearest-neighbor representatives so that given a query point in the inner ball its approximate nearest neighbor outside the outer ball will be one of these representatives, and vice versa. A similar result holds for two disjoint balls. Lemmas 1 and 2 below indicate the sizes of these representative sets as a function of the separation of the two balls. The second construction is a subdivision of space into a collection of cells, such that for each cell, the points of S lying outside this cell satisfy certain separation properties. This is given in Lemma 3 below. In all three lemmas (adapted slightly for convenience), S is a set of n points in \mathbb{R}^d , and $0 < \epsilon \leq 1/2$ and $\gamma \geq 2$ are two real parameters.

LEMMA 1. Let b_1 and b_2 denote two concentric balls of radius r and γr , respectively. There exists a set $R \subseteq S$ consisting of

$$\left(1 + O\left(\frac{1}{\sqrt{\epsilon\gamma}}\right)\right)^{d-1}$$

points such that

(i) for any point $q \in b_1$, $NN_q(R) \leq (1 + \epsilon)NN_q(S \cap \overline{b_2})$, and

(ii) for any point $q \in \overline{b_2}$, $NN_q(R) \leq (1 + \epsilon)NN_q(S \cap b_1)$.

LEMMA 2. Let b_1 and b_2 be two disjoint balls of radius r_1 and r_2 , respectively, whose minimum distance of separation is at least ℓ' . Further, suppose that $\ell' \geq \max(r_1, r_2)/2$. Then there exists a set $R \subseteq S$ consisting of

$$\left(1 + O\left(\frac{\sqrt{r_1 r_2}}{\ell' \sqrt{\epsilon}}\right)\right)^{d-1}$$

points such that for any point $q \in b_1$, $NN_q(R) \leq (1 + \epsilon) \cdot NN_q(S \cap b_2)$.

LEMMA 3. It is possible to construct a subdivision consisting of $O(n\gamma^d)$ cells, where each cell c is the difference of two cubes and satisfies at least one of the following three properties:

(i) $|S \cap \gamma b_c| \leq 1$.

(ii) There exists a ball b'_c such that $S \cap \gamma b_c \subseteq b'_c$ and the ball $\gamma b'_c$ does not overlap c .

(iii) There exists a ball b'_c such that $S \cap \gamma b_c \subseteq b'_c$. Letting r_1 and r_2 denote the radius of balls b_c and b'_c , respectively, and ℓ' denote the minimum distance of separation between b_c and b'_c , then $\ell' \geq 2 \max(r_1, r_2)$ and $\ell' / \sqrt{r_1 r_2} \geq 5\sqrt{\gamma}$.

The three lemmas given above imply that, for a cell c in the subdivision of Lemma 3, we can choose a set of representatives of size $O((1/\epsilon\gamma)^{(d-1)/2})$ each from $\overline{\gamma b_c}$ and from b'_c . This forms the basis of the approach given by Arya and Malamatos. We can significantly reduce the number of cells by weakening the separation property in Lemma 3, allowing for $O((1/\epsilon\gamma)^{(d-1)/2})$ points in the region $\gamma b_c - b'_c$, and choosing all these points as additional representatives. A natural method for trying to achieve this goal is to construct the subdivision of Lemma 3 for a set of randomly sampled points of suitable size. Unfortunately, standard results in random sampling and ϵ -net theory imply that this would lead to an increase in the number of representatives by a factor that is logarithmic in n . We overcome this difficulty by devising a novel but simple sampling procedure based on the BBD-tree, which has the following property. Let nf denote the number of points of S that are sampled, where $0 < f \leq 1$ is a parameter. Then for any fat region (e.g., an Euclidean ball) that is free of sampled points, we can contract the region by a constant factor (say, 2), and the shrunken region is guaranteed to have at most $O(1/f)$ points of the original set. As we will see, this sampling procedure suffices for our application and leads to a large space improvement.

In Lemma 4, we present a simple version of this idea, intended to illustrate the basic space-reduction mechanism and, later, in Section 3.1, we apply this technique in full force.

LEMMA 4. Let S be a set of n points in \mathbb{R}^d , and let $\gamma \geq 2$ and $0 < f \leq 1$ be two real parameters. It is possible to construct a subdivision consisting of $O(nf\gamma^d)$ cells, where each cell c is the difference of two cubes and satisfies at least one of the following three properties:

(i) $|S \cap \gamma b_c| = O(1/f)$.

(ii) There exists a ball b'_c such that $|S \cap (\gamma b_c - b'_c)| = O(1/f)$ and the ball $\gamma b'_c$ does not overlap c .

(iii) There exists a ball b'_c such that $|S \cap (\gamma b_c - b'_c)| = O(1/f)$. Letting r_1 and r_2 denote the radius of balls b_c and b'_c , respectively, and ℓ' denote the minimum distance of separation between b_c and b'_c , then $\ell' \geq \max(r_1, r_2)/2$ and $\ell' / \sqrt{r_1 r_2} \geq \sqrt{\gamma}$.

PROOF. Let T denote the BBD tree for the set S of points. Let \mathcal{N} be the set of nodes of T that contain at most $1/f$ points (of S), but whose parents contain more than $1/f$ points. Let \mathcal{X} be the cells corresponding to these nodes. By BBD property (iv), $|\mathcal{X}|$ is $O(nf)$. Let $S' \subseteq S$ be the set of points obtained by sampling one point arbitrarily from each (non-empty) cell in \mathcal{X} . We construct the subdivision described in Lemma 3 for S' but using the value 2γ in place of γ in the lemma. Observe that since $|S'| = O(nf)$, by Lemma 3, the number of cells in this subdivision is $O(nf\gamma^d)$. We claim that any cell c in this subdivision satisfies the desired property.

If Case (i) of Lemma 3 holds, then $|S' \cap 2\gamma b_c| \leq 1$. Let $\mathcal{X}' \subseteq \mathcal{X}$ be the cells that overlap the ball γb_c , and contain at least one point of S . Since S' contains one point from each cell in \mathcal{X}' , all but at most one cell in \mathcal{X}' must intersect the boundary of $2\gamma b_c$. By BBD property (iii), the number of cells that overlap γb_c and intersect the boundary of $2\gamma b_c$ is bounded by a constant. Thus $|\mathcal{X}'| = O(1)$. Since a cell in \mathcal{X}' has at most $1/f$ points, the total number of points of S in γb_c is $O(1/f)$.

If Case (ii) of Lemma 3 holds, then there exists a ball b''_c such that $S' \cap 2\gamma b_c \subseteq b''_c$ and the ball $2\gamma b''_c$ does not overlap c . Define $b'_c = 2b''_c$. It follows that $\gamma b'_c$ does not overlap c . Next we show that $|S \cap (\gamma b_c - b'_c)| = O(1/f)$. Let $\mathcal{X}' \subseteq \mathcal{X}$ be the cells that overlap the region $\gamma b_c - b'_c$, and contain at least one point of S . Since S' contains one point from each cell in \mathcal{X}' , and $2\gamma b_c - b'_c/2$ contains no point of S' , it follows that each cell in \mathcal{X}' intersects the boundary of both $2\gamma b_c$ and γb_c , and intersects the boundary of both b'_c and $b'_c/2$. By BBD property (iii), the number of such cells is bounded by a constant. Since a cell in \mathcal{X}' has at most $1/f$ points, the total number of points of S in the region $\gamma b_c - b'_c$ is $O(1/f)$.

A similar argument shows that if Case (iii) of Lemma 3 holds, then the corresponding case holds here. The details are omitted. This completes the proof. \square

We construct the subdivision described in Lemma 4 for $f = (\epsilon\gamma)^{(d-1)/2}$. We assign representatives to the cells as follows. Let q be a point inside a cell c . Let b_c and b'_c be the balls defined in Lemma 4. Since c is contained within the ball b_c , applying Lemma 1(i), it follows that we can find a set R'_c consisting of $O(1/(\epsilon\gamma)^{(d-1)/2})$ points such that $NN_q(R'_c) \leq (1 + \epsilon)NN_q(S \cap \overline{\gamma b_c})$. We next consider the case when the nearest neighbor of q lies within γb_c . Note that one of the three cases given in the statement of Lemma 4 must hold. If Case (i) holds, then we define

$R_c'' = S \cap \gamma b_c$, and if Case (ii) or Case (iii) holds, then we define $R_c'' = S \cap (\gamma b_c - b'_c)$. Note that $|R_c''| = O(1/(\epsilon\gamma)^{(d-1)/2})$. It is clear that a point in $R_c' \cup R_c''$ is an ϵ -NN of q unless Case (ii) or (iii) holds and the nearest neighbor of q lies within b'_c . If Case (ii) (Case (iii), resp.) holds then, by Lemma 1(ii) (Lemma 2, resp.), it follows that we can find a set R_c''' consisting of $O(1/(\epsilon\gamma)^{(d-1)/2})$ points such that $NN_q(R_c''') \leq (1+\epsilon)NN_q(S \cap b'_c)$. Finally we assign the set of representatives R_c for c to be $R_c' \cup R_c''$ in Case (i) and $R_c' \cup R_c'' \cup R_c'''$ in Cases (ii) and (iii). Clearly, R_c has size $O(1/(\epsilon\gamma)^{(d-1)/2})$ and satisfies the desired property, namely, $NN_q(R_c) \leq (1+\epsilon)NN_q(S)$. Thus, we have thus shown how to construct a (t, ϵ) -AVD with $O(n\epsilon^{(d-1)/2}\gamma^{(3d-1)/2})$ cells for $t = O(1/(\epsilon\gamma)^{(d-1)/2})$.

3.1 Bisector-Sensitive Construction

In this section, we present a more sophisticated construction that reduces the space by nearly a factor of γ . The construction used in the previous section simply applied the existing construction of Arya and Malamatos to a suitable sample of the points. In this section we combine the sampling process with a new construction, which generates cells along the Voronoi bisectors of the well-separated pairs.

Let T denote the BBD tree for the set S of points. Let \mathcal{N} be the set of nodes of T that contain at most $1/f$ points (of S), but whose parents contain more than $1/f$ points. Here f is a parameter between 0 and 1, which will later be assigned a suitable value depending on ϵ and γ . Let \mathcal{X} be the cells corresponding to these nodes. By BBD property (iv), $|\mathcal{X}|$ is $O(nf)$. Let $S' \subseteq S$ be the set of points obtained by sampling one point arbitrarily from each (non-empty) cell in \mathcal{X} .

We construct a WSPD $\mathcal{P}_{S', \alpha}$ for S' , using separation factor $\alpha = 16$. Note that the number of pairs in $\mathcal{P}_{S', 16}$ is $O(nf)$. Let \mathcal{D}' denote the set of dumbbells corresponding to this WSPD. We expand both the heads of each dumbbell by a factor of two. Let \mathcal{D} denote the new set of dumbbells. Note that the dumbbells in \mathcal{D} have a separation factor of 8. For each dumbbell $P \in \mathcal{D}$, we compute a set of quadtree boxes \mathcal{C}_P as follows.

Let $X \subseteq S$ and $Y \subseteq S$ denote the set of points that are enclosed within the two heads of P . Let x and y denote the centers of the heads enclosing X and Y , respectively. Let $\ell = |xy|$, let z denote the center of the segment \overline{xy} , and let \mathcal{B}_P denote the set of balls of radius $2^i \ell$, for $4 \leq i \leq \lceil \log \gamma + 4 \rceil$, centered at z . Let V_P denote the set of points whose nearest neighbor in X and in Y are equidistant (i.e., the Voronoi bisector of X and Y). For a ball $b \in \mathcal{B}_P$, let \mathcal{C}_b be the set of quadtree boxes overlapping $b \cap V_P$ that have the largest size not exceeding $r_b^2 / (2048\gamma d \ell)$, where r_b denotes the radius of b . Let \mathcal{C}'_P be the set of quadtree boxes overlapping $b(z, 16\gamma\ell)$ that have the largest size not exceeding $\gamma\ell/d$. Let $\mathcal{C}_P = (\cup_{b \in \mathcal{B}_P} \mathcal{C}_b) \cup \mathcal{C}'_P$ and $\mathcal{C} = \cup_{P \in \mathcal{D}} \mathcal{C}_P$. Finally we store all the quadtree boxes in \mathcal{C} in a BBD tree T' . We will show that the subdivision induced by the leaves of T' , along with suitably chosen representatives, is the desired approximate Voronoi diagram.

We first bound the number of cells in the subdivision. The following lemma essentially shows that the Voronoi bisector of a well-separated pair (X, Y) behaves as a $(d-1)$ -dimensional hyperplane for the purpose of packing arguments. The proof is omitted, but is similar to one given in [2].

LEMMA 5. *Let X and Y be two sets of points enclosed within two disjoint d -balls of radius r , such that the distance between the centers of these balls is at least αr , where $\alpha \geq 4$. Let V denote the Voronoi bisector of X and Y , and let b be any d -ball of radius R . Then the number of quadtree boxes of size s that overlap $b \cap V$ is at most $O((1+R/s)^{d-1})$.*

By Lemma 5, it follows that for any ball $b \in \mathcal{B}_P$, $|\mathcal{C}_b|$ is $O((1+\gamma\ell/r_b)^{d-1})$. Also, by BBD property (iii), $|\mathcal{C}'_P| = O(1)$. It is now easy to see that $|\mathcal{C}_P|$ is $O(\gamma^{d-1})$. Since the number of dumbbells in \mathcal{D} is $O(nf)$, the total number of quadtree boxes, $|\mathcal{C}|$, is $O(nf\gamma^{d-1})$. Since the number of leaves in the BBD tree T' is $O(|\mathcal{C}|)$, this bound applies to the number of cells in the subdivision.

Lemma 7 describes an important property of the subdivision, which will enable us to choose representatives for the cells. We need the following technical lemma. The proof is omitted due to space limitations, but is proved by methods similar to those used in [1].

LEMMA 6. *Let c be a cell corresponding to a leaf of T' . Let x be a point in $S \cap 2\gamma b_c$. Let $S_c \subseteq S \cap 2\gamma b_c$ be the set of points p such that there is a dumbbell $P \in \mathcal{D}$ that separates x and p , and V_P intersects cell c . Then either $|S_c| = 0$ or there exists a ball b'_c such that $S_c \cup \{x\} \subseteq b'_c$ and which satisfies at least one of the following two properties:*

- (i) *The ball $8\gamma b'_c$ does not overlap c .*
- (ii) *$\ell' \geq 4 \max(r_1, r_2)$ and $\ell' / \sqrt{r_1 r_2} \geq 19\sqrt{\gamma}$, where ℓ' denotes the minimum distance of separation between b_c and b'_c , r_1 denotes the radius of b_c , and r_2 denotes the radius of ball b'_c .*

The proof is based on the fact that a dumbbell separating x from a point in S_c is either “far away” from c or it implies an upper bound on the size of c (because it generates c or a cell enclosing c). By considering the dumbbell that separates x from the point in S_c that is farthest from it, we can show that these two cases yield properties (i) and (ii), respectively.

LEMMA 7. *Let S be a set of n points in \mathbb{R}^d , and let $\gamma \geq 2$ and $0 < f \leq 1$ be two real parameters. It is possible to construct a subdivision consisting of $O(nf\gamma^{d-1})$ cells, where each cell c is the difference of two cubes and satisfies at least one of the following three properties. Let $S_c \subseteq S$ be the set of points that are the nearest neighbor of some point in c .*

- (i) *$|S_c \cap \gamma b_c| = O((1/f) \log \gamma)$.*
- (ii) *There exists a ball b'_c such that $|S_c \cap (\gamma b_c - b'_c)| = O((1/f) \log \gamma)$ and the ball $\gamma b'_c$ does not overlap c .*
- (iii) *There exists a ball b'_c such that $|S_c \cap (\gamma b_c - b'_c)| = O((1/f) \log \gamma)$. Letting ℓ' denote the minimum distance of separation between b_c and b'_c , r_1 denote the radius of b_c , and r_2 denote the radius of ball b'_c , then $\ell' \geq \max(r_1, r_2)$ and $\ell' / \sqrt{r_1 r_2} \geq 9\sqrt{\gamma}$.*

PROOF. If $|S_c \cap \gamma b_c| = 0$, then (i) trivially holds. So suppose that $|S_c \cap \gamma b_c| \geq 1$. First we show that if $|S' \cap 2\gamma b_c| = 0$ then (i) holds. Let $\mathcal{X}' \subseteq \mathcal{X}$ be the set of cells that overlap the ball γb_c , and contain at least one point of S . Since S' contains one point from each cell in \mathcal{X}' , all

the cells in \mathcal{X}' must intersect the boundary of $2\gamma b_c$. By BBD property (iii), the number of cells that overlap γb_c and intersect the boundary of $2\gamma b_c$ is bounded by a constant. Thus $|\mathcal{X}'| = O(1)$. Since a cell in \mathcal{X}' has at most $1/f$ points, the total number of points of S in γb_c is $O(1/f)$, and so clearly (i) holds.

In the remainder of the proof, we assume that $|S' \cap 2\gamma b_c| \geq 1$ and $|S_c \cap \gamma b_c| \geq 1$. Let (x, w) be the closest pair of points such that $x \in S' \cap 2\gamma b_c$ and $w \in S_c \cap 2\gamma b_c$. (In case of ties, choose any such pair.) Let $\beta = |xw|$. (Note that if $S' \cap S_c \cap 2\gamma b_c \neq \emptyset$, then $x = w$, and $\beta = 0$.)

Let $S'_c \subseteq S \cap 2\gamma b_c$ be the set of points p such that there is a dumbbell $P \in \mathcal{D}$ that separates x and p , and V_P intersects cell c . If $|S'_c| = 0$, then let b'_c be the ball of zero radius centered at x . Otherwise, by Lemma 6, there exists a ball b''_c such that $S'_c \cup \{x\} \subseteq b''_c$ and which satisfies either property (i) or (ii) listed therein. Let r''_2 be the radius of b''_c and let z denote its center. We distinguish two cases based on the closest distance L between z and cell c : (1) $L \geq s_c/16$ and (2) $L < s_c/16$.

Case 1: $L \geq s_c/16$.

Let $r_2 = \max(2r''_2, s_c/(16\gamma))$, and let $b'_c = b(z, r_2)$. Note that x is at a distance of at most $r_2/2$ from the center z of b'_c . We consider two subcases: (a) $\beta < r_2/64$ and (b) $\beta \geq r_2/64$.

Subcase (a): $\beta < r_2/64$.

It is easy to verify that b'_c satisfies either property (ii) or (iii). It remains to show that $|S_c \cap (\gamma b_c - b'_c)| = O((1/f) \log \gamma)$.

To this end, we first show that the cell $c' \in \mathcal{X}$ that contains a point $u \in S_c \cap (\gamma b_c - b'_c)$ has size at least $|zu|/(64d)$. For the sake of contradiction, suppose that the size of c' is less than $|zu|/(64d)$. Recall that S' contains one point from each cell in \mathcal{X} . Let u' be the point in c' that belongs to S' (note that u and u' may be the same point). By the triangle inequality, we get $|xu'| \geq |zu| - |uu'| - |xz|$. Since $|uu'| \leq |zu|/64$ and $|xz| \leq |zu|/2$ (because $|zu| \geq r_2$ and $|xz| \leq r_2/2$), it follows that $|xu'| \geq 31|zu|/64$.

Consider the dumbbell $P' \in \mathcal{D}'$ that separates x and u' . Let P be the dumbbell corresponding to P' in \mathcal{D} . Let A and B denote the heads of dumbbell P containing x and u' , respectively. Since the separation factor for the dumbbells in \mathcal{D} is 8, it follows that the radius r' of A and B is at least $|xu'|/10 \geq 31|zu|/640$. Recall that the heads of the dumbbells in \mathcal{D}' are enlarged by a factor of 2 to obtain the dumbbells in \mathcal{D} . Hence any point within distance $r'/2$ of x and u' will lie within A and B , respectively. Since $|wx| = \beta \leq r_2/64 \leq |zu|/64$, it follows that $w \in A$ and, since $|uu'| \leq |zu|/64$, it follows that $u \in B$. Since both w and u belong to S_c (i.e., they are the nearest neighbor of some point in cell c), it is not difficult to see that the Voronoi bisector V_P must intersect c . But then b''_c should have contained u , a contradiction.

Partition the points in $S_c \cap (\gamma b_c - b'_c)$ into groups, where the i th group has all the points whose distance from z lies in the interval $[2^i r_2, 2^{i+1} r_2]$. Since $r_2 \geq s_c/(16\gamma)$, and the maximum distance of a point in γb_c from z is some constant times γs_c , it is clear that the number of non-empty groups is $O(\log \gamma)$. Let $\mathcal{X}_i \subseteq \mathcal{X}$ be the set of cells that overlap some point in the i th group. By the above observation, the cells in \mathcal{X}_i have size at least $2^i r_2/(64d)$, and so by BBD property (iii), $|\mathcal{X}_i|$ is $O(1)$. Since each cell in \mathcal{X}_i has $1/f$ points, the number of points in the i th group is $O(1/f)$, which implies the desired bound on $|S_c \cap (\gamma b_c - b'_c)|$.

Subcase (b): $\beta \geq r_2/64$.

Consider a ball \tilde{b} of radius 32β centered at x . It is easy to show that for any dumbbell separating x and any point u outside this ball, the dumbbell head containing x also contains w . As in Subcase (a), we can now show that $|S_c \cap (\gamma b_c - \tilde{b})| = O((1/f) \log \gamma)$.

Next we claim that $|S_c \cap \gamma b_c \cap \tilde{b}| = O(1/f)$. To see this, let $\mathcal{X}' \subseteq \mathcal{X}$ be the set of cells that contain a point in $S_c \cap \gamma b_c \cap \tilde{b}$. Note that a cell in \mathcal{X}' must either overlap $2\gamma b_c$ or it must have size exceeding β/d (otherwise the distance between any pair of points in this cell that belong to S_c and S' , respectively, would be less than β , a contradiction). It follows from BBD property (iii) that $|\mathcal{X}'| = O(1)$, which implies the desired claim.

Thus we get $|S_c \cap \gamma b_c| = O((1/f) \log \gamma)$, which implies (i).

Case 2: $L < s_c/16$.

It is obvious that property (ii) given in Lemma 6 cannot apply to b''_c , so either b''_c satisfies property (i) given there or its radius must be zero. It follows that $r''_2 \leq L/(8\gamma)$. We consider two subcases: (a) $\beta < s_c/(256\gamma)$ and (b) $\beta \geq s_c/(256\gamma)$.

Subcase (a): $\beta < s_c/(256\gamma)$.

We set $b'_c = b(z, L/\gamma)$. Note that b'_c satisfies property (ii), or it is a ball of zero radius centered at a point inside cell c . We now show that $|S_c \cap (\gamma b_c - b'_c)| = O((1/f) \log \gamma)$.

Similar to Subcase (a) of Case 1, we can show that $|S_c \cap (\gamma b_c - b(z, s_c/(8\gamma)))| = O((1/f) \log \gamma)$. We omit the straightforward details.

We next claim that $|S \cap (b(z, s_c/(8\gamma)) - b'_c)| = O(1/f)$, which would imply (ii). Towards this end, we show that there are no points of S' in the region $b(z, s_c/(4\gamma)) - b'_c/2$. For contradiction, assume that there is a point $u \in S'$ in this region. Consider a dumbbell $P \in \mathcal{D}$ that separates x and u . Let ℓ be the length of P . By definition of well-separatedness, we get $4|xu|/5 \leq \ell \leq 4|xu|/3$. Applying the triangle inequality, it is easy to see that $3L/(8\gamma) \leq |xu| \leq s_c/(2\gamma)$, which implies that $3L/(10\gamma) \leq \ell \leq 2s_c/(3\gamma)$. It can be easily checked from the construction that all cells whose closest distance from x is less than $15\gamma\ell$ can have size at most $\gamma\ell/d$. Note that the closest distance from x to cell c is at most $L + L/(8\gamma) \leq 17L/16 \leq 4\gamma\ell$. Thus the size of cell c cannot exceed $2s_c/(3d)$, a contradiction.

Let $\mathcal{X}' \subseteq \mathcal{X}$ be the set of cells that overlap the region $b(z, s_c/(8\gamma)) - b'_c$, and contain at least one point of S . Since S' contains one point from each cell in \mathcal{X}' , and there are no points of S' in the region $b(z, s_c/(4\gamma)) - b'_c/2$, it follows that each cell in \mathcal{X}' intersects the boundary of both $b(z, s_c/(4\gamma))$ and $b(z, s_c/(8\gamma))$ or the boundary of both b'_c and $b'_c/2$. By BBD property (iii), the number of such cells is bounded by a constant. Noting that a cell in \mathcal{X}' has at most $1/f$ points, the desired result follows.

Subcase (b): $\beta \geq s_c/(256\gamma)$.

Using an argument similar to Subcase (b) of Case 1, we can show that (i) holds. We omit the straightforward details. This completes the proof. \square

In view of the similarity of Lemma 7 with Lemma 4, it should be clear that the same method for assigning representatives works for this subdivision too. We mention that the only difference is the value of f , which we set to $(\epsilon\gamma)^{(d-1)/2} \log \gamma$.

Given a query point q , we can determine the leaf of the BBD tree T' that contains q in $O(\log(n\gamma))$ time. By com-

putting the distance from q for each of the stored representatives, we can answer queries in $O(\log(n\gamma) + 1/(\epsilon\gamma)^{(d-1)/2})$ time. We summarize the main result of this section.

THEOREM 1. *Let S be a set of n points in \mathbb{R}^d , and let $0 < \epsilon \leq 1/2$ and $2 \leq \gamma \leq 1/\epsilon$ be two real parameters. We can construct an $(O(1/(\epsilon\gamma)^{(d-1)/2}), \epsilon)$ -approximate Voronoi diagram for S that consists of $O(n\epsilon^{(d-1)/2}\gamma^{3(d-1)/2}\log\gamma)$ cells, where each cell is the difference of two cubes. Moreover, for any query point, we can return its ϵ -NN in $O(\log(n\gamma) + 1/(\epsilon\gamma)^{(d-1)/2})$ time. Here the constants in the O -notation are independent of ϵ and γ .*

We obtain a family of data structures that can answer ϵ -NN queries in $O(\log(n\gamma) + 1/(\epsilon\gamma)^{(d-1)/2})$ time using space $O(n\gamma^{d-1}\log\gamma)$. Setting $\gamma = 2$ we obtain the most space-efficient solution in this family, which we present in the following corollary.

COROLLARY 1.1. *Given a set S of n points in \mathbb{R}^d , we can answer ϵ -NN queries in $O(\log n + 1/\epsilon^{(d-1)/2})$ time using a data structure of space $O(n)$.*

Remark: For the case of $\gamma = 1/\epsilon$, we can omit the sampling process altogether in Lemma 7, which allows us to save a $\log\gamma$ factor in space. By a straightforward extension of the approach given in this section, we can construct a $(1, \epsilon)$ -AVD of size $O(n/\epsilon^{d-1})$. Arya and Malamatos [1] have shown a lower bound of $\Omega(n/\epsilon^{d-1})$ on the size of a $(1, \epsilon)$ -AVD, assuming that the cells are differences of two axis-aligned hyperrectangles, which implies that our construction has optimal size.

4. LOWER SPACE BOUNDS BY BETTER CHARGING

Let S be a set of n points in \mathbb{R}^d , and let $0 < \epsilon \leq 1/2$ and $2 \leq \gamma \leq 1/\epsilon$ be two real parameters. In this section we show how to construct a (t, ϵ) -AVD, consisting of

$$O\left(n\epsilon^{\frac{d-1}{2}}\gamma^{\frac{3(d-1)}{2}}\log\frac{2}{\epsilon\gamma}\right)$$

cells, where $t = O(1/(\epsilon\gamma)^{(d-1)/2})$.

The construction proceeds as follows. First, we construct a (t, ϵ) -AVD consisting of $O(n\gamma^{d-1})$ cells, such that the total number of representatives summed over all cells is $O(n\gamma^{d-1}\log(2/(\epsilon\gamma)))$. Note that for this AVD, there is a huge gap between the number of representatives for a cell in the worst case, which is $O(1/(\epsilon\gamma)^{(d-1)/2})$, and their number averaged over all cells, which is only $O(\log(2/(\epsilon\gamma)))$. In the second step of the construction, we take advantage of this fact by combining cells that have fewer representatives into larger cells. This enables us to significantly reduce the number of cells and leads to the desired AVD.

Let \mathcal{D} be the set of dumbbells corresponding to the WSPD for S , using separation factor 8. For each dumbbell $P \in \mathcal{D}$, define the Voronoi bisector V_P , the set of balls \mathcal{B}_P , and the sets \mathcal{C}_P and \mathcal{C}_P of quadtree boxes, exactly as in Section 3.1. We store all the quadtree boxes in $\mathcal{C} = \cup_{P \in \mathcal{D}} \mathcal{C}_P$ in a BBD tree T . Using the same approach as in Section 3.1, we can show that the number of leaves in T is $O(n\gamma^{d-1})$.

We now describe a crucial property satisfied by the subdivision induced by the leaves of T , which enables us to select representatives for the cells economically. The proof is

omitted due to space limitations, but is proved by methods similar to those used in [1].

LEMMA 8. *Let S be a set of n points in \mathbb{R}^d , and let $\gamma \geq 2$ be a real parameter. It is possible to construct a subdivision consisting of $O(n\gamma^{d-1})$ cells, where each cell is the difference of two cubes, that satisfies the following properties. In the following, c is any cell in the subdivision and $S_c \subseteq S \cap \gamma b_c$ is the set of points that are the nearest neighbor of some point in c .*

- (a) *The number of cells, each of size at least s , that intersect a ball of radius r is at most $O((1+r/s)^d)$.*
- (b) *There is a constant $k > 1$ such that the ball $k\gamma b_c$ contains at least one point of S .*
- (c) *The cell c satisfies at least one of the following three properties.*
 - (i) $|S_c| \leq 1$.
 - (ii) *There exists a ball b'_c such that $S_c \subseteq b'_c$ and the ball $\gamma b'_c$ does not overlap c .*
 - (iii) *There exists a ball b'_c such that $S_c \subseteq b'_c$. Letting r_1 and r_2 denote the radius of balls b_c and b'_c , respectively, and ℓ' denote the minimum distance of separation between b_c and b'_c , then $\ell' \geq \max(r_1, r_2)$ and $\ell'/\sqrt{r_1 r_2} \geq 19\sqrt{\gamma}$.*

Let \mathcal{X} be any subdivision satisfying Lemma 8. Using Lemmas 1 and 2, we can easily show that $O(1/(\epsilon\gamma)^{(d-1)/2})$ representatives suffice for each cell. The following lemma gives a more complicated method for choosing cell representatives which, however, has the advantage that we can obtain a much better bound on the total number of representatives.

LEMMA 9. *Let S be a set of n points in \mathbb{R}^d , and let $0 < \epsilon \leq 1/2$ and $2 \leq \gamma \leq 1/\epsilon$ be two real parameters. Let $t = O(1/(\epsilon\gamma)^{(d-1)/2})$. Any subdivision \mathcal{X} satisfying Lemma 8 is a (t, ϵ) -AVD such that the total number of representatives summed over all cells of \mathcal{X} is $O(n\gamma^{d-1}\log(2/(\epsilon\gamma)))$.*

PROOF. (Sketch) Let c be any cell in \mathcal{X} and let $S'_c \subseteq S$ be the set of points that are the nearest neighbor of some point in c . Let q be any point in c . We will identify two sets of points $R'_c \subseteq S$ (outer representatives) and $R''_c \subseteq S$ (inner representatives) each consisting of $t = O(1/(\epsilon\gamma)^{(d-1)/2})$ points, such that $NN_q(R'_c) \leq (1+\epsilon)NN_q(S'_c \cap \gamma b_c)$ and $NN_q(R''_c) \leq (1+\epsilon)NN_q(S'_c \cap \gamma b_c)$. Finally, we will store the set of representatives $R_c = R'_c \cup R''_c$ with c . Clearly $NN_q(R_c) \leq (1+\epsilon)NN_q(S)$.

In this version of the paper, we only present our method for computing R'_c , the set of outer representatives, and show that $\sum_{c \in \mathcal{X}} |R'_c| = O(n\gamma^{d-1}\log(2/(\epsilon\gamma)))$. Using similar ideas, we can compute the inner representatives and show that this bound also applies to their total number. Details will appear in the full version.

Letting $k' = k + 1$, we have $S'_c \subseteq k'\gamma b_c$ because, by Lemma 8, the ball $k\gamma b_c$ contains at least one point of S and γ is at least 2. Let $S''_c = S'_c \cap (k'\gamma b_c - \gamma b_c)$. Let \mathcal{U}_c denote the set of quadtree boxes that have the largest size not exceeding $\epsilon\gamma^2 s_c/8$ and that contain at least one point of S''_c . Clearly $|\mathcal{U}_c| = O(1/(\epsilon\gamma)^d)$. Note that any

quadtrees box $u \in \mathcal{U}_c$ can be enclosed within a ball b_u of radius $r_u \leq \epsilon \gamma^2 s_c d / 16 \leq \gamma s_c d / 16$, which implies that the minimum distance of separation between the balls b_c and b_u , ℓ'_u , is at least $\gamma s_c d / 8$. It is now easy to verify that $\sqrt{r_1 r_u} / (\ell'_u \sqrt{\epsilon}) = O(1)$ and $\ell'_u \geq \max(r_1, r_u) / 2$, where $r_1 = s_c d / 2$ is the radius of ball b_c . Therefore, by Lemma 2, we can find a set $R_{c,u} \subseteq S'_c \cap u$ consisting of $O(1)$ points such that $NN_q(R_{c,u}) \leq (1 + \epsilon/2) NN_q(S'_c \cap u)$. It follows that the set $\tilde{R}'_c = \cup_{u \in \mathcal{U}_c} R_{c,u}$ has at most $O(1/(\epsilon\gamma)^d)$ points and satisfies $NN_q(\tilde{R}'_c) \leq (1 + \epsilon/2) NN_q(S'_c \cap \overline{\gamma b_c})$.

The main difficulty with using \tilde{R}'_c as the outer representatives for c is that their number can be too large. To remedy this, observe that $c \subseteq b_c$ and all the points in \tilde{R}'_c lie outside γb_c . Thus, we can apply Lemma 1 to obtain a set $R'_c \subseteq \tilde{R}'_c$ consisting of $O(1/(\epsilon\gamma)^{(d-1)/2})$ points such that $NN_q(R'_c) \leq (1 + \epsilon/4) NN_q(\tilde{R}'_c)$. Thus

$$\begin{aligned} NN_q(R'_c) &\leq (1 + \epsilon/4)(1 + \epsilon/2) NN_q(S'_c \cap \overline{\gamma b_c}) \\ &\leq (1 + \epsilon) NN_q(S'_c \cap \overline{\gamma b_c}), \end{aligned}$$

as desired.

We next present a charging argument to show that

$$\sum_{c \in \mathcal{X}} |R'_c| = O\left(n \gamma^{d-1} \log \frac{2}{\epsilon \gamma}\right).$$

Let \mathcal{D} be the set of dumbbells corresponding to the WSPD for S , using separation factor 8. Each dumbbell $P \in \mathcal{D}$ allocates a unit charge to each cell in a set $\tilde{\mathcal{C}}_P \subseteq \mathcal{X}$ defined as follows. Let $X \subseteq S$ and $Y \subseteq S$ denote the set of points that are enclosed within the two heads of P . Let x and y denote the centers of the heads enclosing X and Y , respectively. Let $\ell = |xy|$, let z denote the center of the segment \overline{xy} , and let $\tilde{\mathcal{B}}_P$ denote the set of balls of radius $2^i \ell$, for $0 \leq i \leq \lceil \log((40k'd)/(\epsilon\gamma)) \rceil$, centered at z . Let V_P denote the Voronoi bisector of X and Y . For a ball $b \in \tilde{\mathcal{B}}_P$, let $\tilde{\mathcal{C}}_b \subseteq \mathcal{X}$ be the set of cells overlapping $b \cap V_P$ that have size at least $r_b / (8k'\gamma d)$, where r_b denotes the radius of b . Finally, we let $\tilde{\mathcal{C}}_P = \cup_{b \in \tilde{\mathcal{B}}_P} \tilde{\mathcal{C}}_b$. We claim that: (i) the total charges allocated to all the cells in \mathcal{X} by the above procedure is $O(n \gamma^{d-1} \log(2/(\epsilon\gamma)))$ and, (ii) each cell $c \in \mathcal{X}$ receives enough charge to pay for the representatives stored with c .

To prove (i), note that Lemma 5 implies that for any ball $b \in \tilde{\mathcal{B}}_P$, $|\tilde{\mathcal{C}}_b|$ is $O(\gamma^{d-1})$. Since the number of balls in $\tilde{\mathcal{B}}_P$ is $O(\log(2/(\epsilon\gamma)))$, it follows that each dumbbell allocates a unit charge to at most $O(\gamma^{d-1} \log(2/(\epsilon\gamma)))$ cells in \mathcal{X} . Recalling that the number of dumbbells in \mathcal{D} is $O(n)$, (i) follows.

To prove (ii), we will show that cell c receives charge from at least $\Omega(|\tilde{R}'_c|)$ dumbbells. In light of (i) and the fact that $R'_c \subseteq \tilde{R}'_c$, this would clearly imply the desired bound on $\sum_{c \in \mathcal{X}} |R'_c|$. We first identify a set of points $\hat{R}'_c \subseteq \tilde{R}'_c$ such that $|\hat{R}'_c| = \Omega(|\tilde{R}'_c|)$ and the distance between any pair of points in \hat{R}'_c is at least $\epsilon \gamma^2 s_c / 16$. We can obtain the set \hat{R}'_c using the following procedure. We start with \hat{R}'_c being empty and consider the quadtree boxes in \mathcal{U}_c one by one. For each quadtree box u examined, we add any one point of $R_{c,u}$ to \hat{R}'_c , and eliminate the at most 3^d quadtree boxes in \mathcal{U}_c that share a common boundary with u from future consideration. We continue in this manner until all the quadtree boxes are eliminated. It is clear that this process yields a set \hat{R}'_c with the properties mentioned above. We next show that there are at least $|\hat{R}'_c| - 1$ distinct dumbbells that separate pairs

of points in $|\hat{R}'_c|$, and that each of these dumbbells allocates a unit charge to cell c . Clearly this would imply (ii).

Consider the following process for identifying a set of dumbbells. At each step, we find the dumbbell in \mathcal{D} that separates the closest pair of points among the remaining points in \hat{R}'_c . We then eliminate one of these two points and continue this process with the remaining points. We stop when only one point in $|\hat{R}'_c|$ remains. Obviously this process finds a set of $|\hat{R}'_c| - 1$ dumbbells. Noting that, at each step, each head of the dumbbell identified contains exactly one point, among the points in \hat{R}'_c that have not yet been eliminated, it is clear that all the dumbbells obtained are distinct.

It remains to show that each of these $|\hat{R}'_c| - 1$ dumbbells allocates a charge to cell c . In fact, we will show that any dumbbell P that separates two points $x, y \in \hat{R}'_c$ allocates a charge to c . Observe first that V_P must intersect c since x and y are both the nearest neighbor of some point in c . Let x' and y' denote the centers of the heads of P , let z denote the center of the line segment $\overline{x'y'}$, and let $\ell = |x'y'|$. Using the definition of well-separatedness and the triangle inequality, it follows that $|xx'| \leq \ell/8$, $|yy'| \leq \ell/8$, and $4|xy|/5 \leq \ell \leq 4|xy|/3$. Since x and y are both contained in $k'\gamma b_c$, it is easy to show that both ℓ , and the distance between z and any point in c , are less than $2k'\gamma s_c d$. Also, since the distance between x and y is at least $\epsilon \gamma^2 s_c / 16$, it follows that $\ell \geq \epsilon \gamma^2 s_c / 20$. This implies that the radius of the largest ball in $\tilde{\mathcal{B}}_P$ is at least $2k'\gamma s_c d$, and so there must be a ball in $\tilde{\mathcal{B}}_P$ that overlaps $c \cap V_P$. Let r_b denote the radius of the smallest ball $b \in \tilde{\mathcal{B}}_P$ that overlaps $c \cap V_P$. If b is the smallest ball in $\tilde{\mathcal{B}}_P$, then its radius r_b is ℓ , which is less than $2k'\gamma s_c d$. Otherwise, r_b cannot exceed twice the closest distance between z and $c \cap V_P$, which implies that r_b is less than $4k'\gamma s_c d$. Recalling that P allocates a charge to all cells overlapping $b \cap V_P$ that have size at least $r_b / (8k'\gamma d)$, it is clear that c receives a charge from P . This completes the proof. \square

To complete the construction, we weight each cell of this AVD with the associated number of representatives. We then build a weighted BBD tree for these cells. Letting $t = 1/(\epsilon\gamma)^{(d-1)/2}$, we apply property (v) of BBD trees to produce a truncated tree having

$$O\left(n \epsilon^{\frac{d-1}{2}} \gamma^{\frac{3(d-1)}{2}} \log \frac{2}{\epsilon \gamma}\right)$$

cells. For each cell in the resulting tree, the number of representatives is no greater than the sum of its associated weights. From this we have the following result.

THEOREM 2. *Let S be a set of n points in \mathbb{R}^d , and let $0 < \epsilon \leq 1/2$ and $2 \leq \gamma \leq 1/\epsilon$ be two real parameters. We can construct an $(O(1/(\epsilon\gamma)^{(d-1)/2}), \epsilon)$ -approximate Voronoi diagram for S that consists of*

$$O\left(n \epsilon^{\frac{d-1}{2}} \gamma^{\frac{3(d-1)}{2}} \log \frac{2}{\epsilon \gamma}\right)$$

cells, where each cell is the difference of two cubes. Moreover, for any query point, we can return its ϵ -NN in time $O(\log(n\gamma) + 1/(\epsilon\gamma)^{(d-1)/2})$. Here the constants in the O -notation are independent of ϵ and γ .

5. LOWER BOUNDS

Our lower bound construction is parameterized by the number of points n , the dimension d , the approximation factor ϵ and the desired number of representatives per cell t . Intuitively, to make the number of cells large we must create a situation in which many cells are small. To make cells as small as possible, we distribute $t + 1$ sites evenly on a $(k - 1)$ -dimensional sphere centered at the origin. The set of points in \mathbb{R}^d that are equidistant from these points lie on a linear subspace J of dimension $d - k$. Any cells that has a significant overlap with J has $t + 1$ contenders for the nearest neighbor for each point in the cell. If this cell is relatively close to the origin we can derive an upper bound on the size of such a cell. We will see that as k increases this size bound decreases. However, as k increases, the dimension of J decreases, and hence the number of overlapping cells decreases. Thus we face a tradeoff. Our construction takes the dimension k as a parameter, and then derives the value of k that produces the best lower bound. We make $\lfloor n/(t + 1) \rfloor$ widely distributed copies of this configuration to produce the final lower bound.

The main results of this section are provided below. The *aspect ratio* of a rectangle is defined to be the ratio between its longest and shortest side lengths.

THEOREM 3. *Consider AVDs in which each cell is an axis-aligned box of bounded aspect ratio or the difference of two such boxes.*

- (i) *Given t and ϵ , there exists a set of n sites in \mathbb{R}^d such that the number of cells in any (t, ϵ) -AVD for this set is*

$$\Omega \left(n \left(\frac{1}{\epsilon} \right)^{(d-1)+c-2\sqrt{2c(d-1)}} \right),$$

where $c = -\frac{\ln t}{\ln \epsilon}$ and $c \leq \frac{d-1}{2}$.

- (ii) *The ratio between the upper bound provided in Theorem 1 and this lower bound is*

$$O \left(\left(\frac{1}{\epsilon} \right)^{2\sqrt{2c(d-1)}-4c} \log \gamma \right) \leq O \left(\left(\frac{1}{\epsilon} \right)^{\frac{d-1}{2}} \log \gamma \right).$$

Remark: Recalling the observation that the $\log \gamma$ factor in the upper bound can be avoided if $\gamma = 1/\epsilon$, it follows that in the extreme case $t = 1$ ($c = 0$) the upper and lower bounds match asymptotically. Also, when γ is a constant, the factor $\log \gamma$ can be ignored, and hence it follows that in the other extreme case $t = 1/\epsilon^{(d-1)/2}$ ($c = (d-1)/2$) the upper and lower bounds match.

The remainder of this section is devoted to proving Theorem 3. Define H to be the linear subspace that is orthogonal to the diagonal vector whose coordinates are all 1, that is, $\vec{w} = (w_1, \dots, w_d) \in H$ if $\sum_i w_i = 0$. Let k , $1 \leq k \leq d - 1$, be an integer parameter to be fixed later, and let J be some $d - k$ dimensional linear subspace of H . Let K be the k dimensional orthogonal linear subspace of J . Let U be the intersection of a unit sphere with K . The principal properties of J , K and U are the following.

- (i) The angle between any coordinate vector and a nonzero vector of H is bounded below by a constant (depending on dimension). Because $J \subseteq H$, this applies to J as well.

- (ii) Given any set of points on U , every point on J is equidistant from all of these points.
- (iii) A set of points is δ -sparse if for each point in the set its nearest neighbor in the set is at distance at least δ . From standard results it follows that the unit sphere U has a θ -sparse set of cardinality $\Omega(1/\theta^{k-1})$.

Recalling that t is the number of representatives, define θ to be maximum such that there is a θ -sparse set on U of size at least $t + 1$. Let S_k denote this set. By the observation above θ is $\Omega(1/t^{1/(k-1)})$. Let $\alpha = 18\epsilon/\theta^2$. We show that no cell of an AVD of S can contain a large ball centered on J and close to the origin. Let B_J be the intersection of a ball of unit radius centered at the origin and J .

LEMMA 10. *Consider a (t, ϵ) -AVD for the set S , whose cells are the differences of axis-aligned rectangles of bounded aspect ratio. Consider a ball b of radius α centered at some point of B_J . Then no cell of the AVD can contain b .*

To sketch the proof we observe that if such a cell existed, then by property (ii) every point of S is a nearest neighbor to some point in the cell. Since S has $t + 1$ points, one point $p \in S$ is not a representative for this cell. Consider the point q on the boundary of the ball that is closest to p . Because S is θ -sparse, and the cell is relatively close to S , it can be shown that the distance of every other point in S to q exceeds the length $|qp|$ by ϵ , providing a contradiction.

Let \mathcal{C} denote the set of cells of the (t, ϵ) -AVD that intersect B_J . For each $c \in \mathcal{C}$, define c_α to be the set of points of the interior of c that are at distance at least α from c 's boundary. Clearly, if c_α intersects B_J , then c contains a ball of radius α that is centered at a point in B_J , and by the previous lemma, no such cell can be in the AVD. Otherwise, by the fact that AVD cells are fat and axis-aligned and by property (i) above, it follows that the diameter of the intersection of any cell and B_J is $O(\alpha)$. By applying a simple packing argument to B_J it follows that the number of cells of the AVD that intersect B_J is at least

$$\Omega \left(\frac{1}{\alpha} \right)^{d-k} = \Omega \left(\frac{\theta^2}{\epsilon} \right)^{d-k} = \Omega \left(\frac{1}{t^{2/(k-1)}\epsilon} \right)^{d-k},$$

for $t \leq (1/\epsilon)^{(k-1)/2}$. To complete the construction, we replicate this configuration of points $\lfloor n/(t + 1) \rfloor$ times at a sufficiently large distance from each other. Up to a constant factor depending on d , the total number of cells is at least $L = \frac{n}{t} (t^{2/(k-1)}\epsilon)^{k-d}$.

To simplify this expression, let us express t as $(1/\epsilon)^c$ for some constant c . Our constraint on t implies that $c \leq (k - 1)/2$. Thus we have the following lower bound.

$$L = n \left(\frac{1}{\epsilon} \right)^{(d-k)(1-\frac{2c}{k-1})-c} = n \left(\frac{1}{\epsilon} \right)^{(d-1)+c-(k-1)-2c\frac{d-1}{k-1}}.$$

Since k is under our control, to get the best lower bound we maximize the exponent, by setting $k - 1 = \sqrt{2c(d-1)}$. It is easy to verify that if $c \leq (d-1)/2$ then this satisfies our prior constraint on t . As a result we have the following lower bound, which establishes Theorem 3(i).

$$L = n \left(\frac{1}{\epsilon} \right)^{(d-1)+c-2\sqrt{2c(d-1)}}.$$

From Theorem 1 it follows that, ignoring constant factors and the $\log \gamma$ term, there exists a (t, ϵ) -AVD for our point set with the following number of cells. $U = n\epsilon^{(d-1)/2}\gamma^{3(d-1)/2}$. To achieve t representatives we set $\gamma = (1/\epsilon)(1/t)^{2/(d-1)}$ and set $t = (1/\epsilon)^c$ to get

$$U = n \left(\frac{1}{\epsilon} \right)^{(d-1)-3c}.$$

To compare the upper and lower bound, we compute their ratio

$$\rho(c) = \frac{U}{L} = \left(\frac{1}{\epsilon} \right)^{2\sqrt{2c(d-1)}-4c}.$$

To determine the worst-case ratio as a function of t , we maximize this expression over c yielding $c = (d-1)/8$ (and hence $t = (1/\epsilon)^{(d-1)/8}$). Note that this satisfies our earlier requirement that $c \leq (d-1)/2$. This yields the following upper bound on the ratio between the upper and lower bounds.

$$\rho(c) \leq \left(\frac{1}{\epsilon} \right)^{(d-1)/2}.$$

This establishes Theorem 3(ii).

6. REFERENCES

- [1] S. Arya and T. Malamatos. Linear-size approximate Voronoi diagrams. In *Proc. 13th ACM-SIAM Sympos. Discrete Algorithms*, pages 147–155, 2002.
- [2] S. Arya and D. M. Mount. Approximate range searching. *Computational Geometry: Theory and Applications*, 17:135–152, 2000.
- [3] S. Arya, D. M. Mount, N. Netanyahu, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *J. ACM*, 45:891–923, 1998.
- [4] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *J. ACM*, 42:67–90, 1995.
- [5] T. M. Chan. Approximate nearest neighbor queries revisited. *Discrete Comput. Geom.*, 20:359–373, 1998.
- [6] K. L. Clarkson. An algorithm for approximate closest-point queries. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 160–164, 1994.
- [7] R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *J. Approx. Theory*, 10:227–236, 1974.
- [8] C. A. Duncan, M. T. Goodrich, and S. G. Kobourov. Balanced aspect ratio trees: Combining the advantages of k -d trees and octrees. *J. Algorithms*, 33:303–333, 2001.
- [9] S. Har-Peled. A replacement for Voronoi diagrams of near linear size. In *Proc. 42 Annu. IEEE Sympos. Found. Comput. Sci.*, pages 94–103, 2001.
- [10] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, pages 604–613, 1998.
- [11] E. Kushilevitz, R. Ostrovsky, and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, pages 614–623, 1998.
- [12] Y. Sabharwal, S. Sen, and N. Sharma. Improved space bound for approximate Voronoi diagram. Manuscript, 2001.
- [13] J. Vleugels and M. Overmars. Approximating Voronoi diagrams of convex sites in any dimension. *Internat. J. Comput. Geom. Appl.*, 8:201–222, 1998.