

The Effect of Corners on the Complexity of Approximate Range Searching

Sunil Arya* Theocharis Malamatos† David M. Mount‡

January 31, 2007

Abstract

Given an n -element point set in \mathbb{R}^d , the range searching problem involves preprocessing these points so that the total weight, or for our purposes the semigroup sum, of the points lying within a given query range η can be determined quickly. In ε -approximate range searching we assume that η is bounded, and the sum is required to include all the points that lie within η and may additionally include any of the points lying within distance $\varepsilon \cdot \text{diam}(\eta)$ of η 's boundary.

In this paper we contrast the complexity of approximate range searching based on properties of the semigroup and range space. A semigroup $(S, +)$ is *idempotent* if $x + x = x$ for all $x \in S$, and it is *integral* if for all $k \geq 2$, the k -fold sum $x + \dots + x$ is not equal to x . Recent research has shown that the computational complexity of approximate spherical range searching is significantly lower for idempotent semigroups than it is for integral semigroups in terms of the dependencies on ε . In this paper we consider whether these results can be generalized to other sorts of ranges. We show that, as with integrality, allowing sharp corners on ranges has an adverse effect on the complexity of the problem. In particular, we establish lower bounds on the worst-case complexity of approximate range searching in the semigroup arithmetic model for ranges consisting of d -dimensional unit hypercubes under rigid motions. We show that for arbitrary (including idempotent) semigroups and linear space, the query time is at least $\Omega(1/\varepsilon^{d-2\sqrt{d}})$. In the case of integral semigroups we prove a tighter lower bound of $\Omega(1/\varepsilon^{d-2})$. These lower bounds nearly match existing upper bounds for arbitrary semigroups.

In contrast, we show that the improvements offered by idempotence do apply to smooth convex ranges. We say that a range is *smooth* if at every boundary point there is an incident Euclidean sphere that lies entirely within the range whose radius is proportional to the range's diameter. We show that for smooth ranges and idempotent semigroups, ε -approximate range queries can be answered in $O(\log n + (1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$ time using $O(n/\varepsilon)$ space. We show that this is nearly tight by presenting a lower bound of $\Omega(\log n + (1/\varepsilon)^{(d-1)/2})$. This bound is in the algebraic decision-tree model and holds irrespective of space.

*Department of Computer Science, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong. Supported by the Research Grants Council, Hong Kong, China (Project No. 610106). Email: arya@cs.ust.hk.

†Max-Planck-Institut für Informatik, Im Stadtwald, D-66123 Saarbrücken, Germany. Email: tmalamat@mpi-inf.mpg.de.

‡Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, Maryland 20742. Partially supported by the National Science Foundation under grant CCF-0635099. Email: mount@cs.umd.edu.

1 Introduction

Answering range queries is a problem of fundamental importance in spatial information retrieval and computational geometry. We assume that we are given a set of n points P in \mathbb{R}^d , where each point is associated with a weight from some commutative semigroup, and the objective is to preprocess these points so that, for any shape η from a given range space, it is possible to compute the semigroup sum of the points of $P \cap \eta$ efficiently. Range searching is among the most well studied problems in computational geometry. Excellent surveys have been written by Agarwal and Erickson [1] and Matoušek [20]. For many formulations of this problem, nearly matching asymptotic upper and lower bounds are known. For example, Matoušek [19] has shown that using m units of storage, where $n \leq m \leq n^d$, halfspace range counting queries can be answered in $O(n/m^{1/d})$ time. Nearly matching lower bounds for halfspace range searching have been proved by Brönnimann, Chazelle, and Pach [11].

Because of the high complexity of answering queries it is natural to consider approximation. In ε -approximate range searching we are given $\varepsilon > 0$ and a bounded range η of diameter $diam(\eta)$, and the sum is required to include all the points of P that lie within η and may additionally include any subset of P that lies within a distance of $\varepsilon \cdot diam(\eta)$ of η 's boundary. Arya and Mount [7] showed that for fixed d , approximate range queries over convex ranges can be answered in $O(\log n + 1/\varepsilon^{d-1})$ time with $O(n)$ space. The dependency on ε is significant, since as the dimension increases the $1/\varepsilon^{d-1}$ term dominates the query time in practice. Chazelle, Liu, and Magen [14] considered approximate halfspace and Euclidean ball range searching in the high-dimensional context. Ignoring polylogarithmic factors, they showed that it is possible to answer queries in $O(d/\varepsilon^2)$ time and $O(dn^{O(1/\varepsilon^2)})$ space. In fixed dimensions, in contrast, a natural goal is to achieve $O(\log n)$ query time and $O(n)$ space, while minimizing the ε -dependencies. Throughout, we treat both n and ε as asymptotic quantities and assume that $n \gg \varepsilon^{-1}$.

As mentioned above we are interested in computing some function of the weights of the points lying within a range. We assume that this function is the sum from a commutative faithful semigroup. (See Section 2.) Our broad goal is to understand how range and semigroup properties affect the computational complexity of approximate range searching. A semigroup is *idempotent* if $x + x = x$ for all semigroup elements x . It is *integral* if for all nonzero semigroup elements x and all natural numbers $k \geq 2$, the k -fold sum $x + \dots + x$ is not equal to x [15]. For example, (\mathbb{R}, \min) and $(\{0, 1\}, \vee)$ are both idempotent, and $(\mathbb{N}, +)$ is integral. Idempotence is relevant because of the way that most range searching algorithms work. At preprocessing time the algorithm implicitly computes the semigroup sum of a number of suitably chosen subsets of P , which we call *generators*. To answer a query η , the algorithm determines an (ideally small) subset of generators whose union is equal to $P \cap \eta$ and then returns their total sum. If the semigroup is idempotent these subsets may overlap, but for integral semigroups they must be disjoint.

Because of the constraint of disjointness, one would expect that range searching over integral semigroups would be harder than for idempotent semigroups. It is surprising, therefore, that this does not seem to be the case for exact range searching. For example, the lower bounds in the semigroup arithmetic model for exact halfspace range searching for idempotent semigroups [11] and integral semigroups [5] are both quite similar to the upper bound complexity [19]. Further, the difference in complexity decreases as dimension increases. This changes in the context of approximate range searching. In [5] we showed that the complexity of approximate range searching for Euclidean balls is much lower for idempotent semigroups than for integral semigroups. Assuming roughly linear space and ignoring polylogarithmic factors, our results there imply nearly matching

asymptotic upper and lower bounds of $O(1/\varepsilon^{d-O(1)})$ for range searching over integral semigroups and $O(1/\varepsilon^{d/2-O(1)})$ for idempotent semigroups. Thus, the exponent in the ε dependency is reduced by roughly half when the semigroup is idempotent.

This raises the question of whether other aspects of the problem formulation have similarly dramatic impacts on the computational complexity of ε -approximate range searching. The upper bounds of [5] made critical use of two properties of Euclidean balls: smoothness and rotational symmetry. In this paper we consider two alternative formulations that arise from relaxing these properties. The first involves ranges with sharp corners, and the second involves arbitrary smooth convex ranges. In both instances the aforementioned upper bounds of $O(\log n + 1/\varepsilon^{d-1})$ query time with $O(n)$ space apply [7]. We consider whether idempotence helps reduce query times for these two classes.

For the case of ranges with sharp corners, we consider the simple case of d -dimensional unit hypercube ranges under rigid motions, or *rotated unit hypercubes*. We show that the worst-case complexity of approximate range searching for these ranges is not significantly better, even in the idempotent case. Assuming linear space, we show that the worst-case query time in the semigroup arithmetic model is $\Omega(1/\varepsilon^{d-2\sqrt{d}})$. For integral semigroups, we show that the lower bound can be tightened to $\Omega(1/\varepsilon^{d-2})$. As in [5], our analysis of the integral case requires the assumption of *convex generators*, which states that the convex hull of each generator subset contains no other points of P .

In contrast, we show that the improvements offered by idempotence do apply to convex ranges that are sufficiently smooth. For $\kappa \geq 1$, say that a range η is κ -smooth if at any point x on the boundary of η , it is possible to place a Euclidean ball inside η that touches x and has radius at least $\text{diam}(\eta)/(2\kappa)$. Note that a Euclidean ball is 1-smooth, and a convex polytope is not κ -smooth for any finite κ . We show that for any fixed κ and any idempotent semigroup, ε -approximate range queries for κ -smooth ranges can be answered from a data structure of space $O(n/\varepsilon)$ in time $O(\log n + (1/\varepsilon)^{\frac{d-1}{2}})$. Further, we show that this query time is optimal by presenting a lower bound of $\Omega(\log n + (1/\varepsilon)^{\frac{d-1}{2}})$ on the complexity of range searching over the space of all κ -smooth convex ranges. This is proved in the algebraic decision-tree model, and so it holds irrespective of the amount of space used.

We consider space-time tradeoffs in some of our results. Rather than expressing our space and time tradeoffs in the conventional manner of query time as a function of space and data size, we adopt a notation that more clearly illustrates the incremental benefits of increased space [5]. Recall that n denotes the size of the point set, and let m denote the space of the data structure. Let $\rho = m/n$, which we call the *expansion ratio*, be the ratio of space to data size. Clearly $\rho \geq 1$, but the requirements of any particular data structure may imply other lower and upper bounds on ρ . (For example, one of our results assumes that space is at least n/ε , implying that $\rho \geq 1/\varepsilon$.) We express query time as a fraction, where the numerator gives the running time assuming linear space ($\rho = 1$) and the denominator gives the *tradeoff rate*, which can be interpreted as the rate with which query time decreases as a function of a multiplicative increase in space. For example, for exact halfspace range queries, the conventionally expressed asymptotic query time of $n/m^{1/d}$ would instead be expressed in our notation as $n^{1-1/d} / \rho^{1/d}$.

Here is a more detailed summary of our results:

- We present a lower bound for answering ε -approximate range queries for d -dimensional rotated unit hypercubes over arbitrary semigroups (and hence over idempotent semigroups). We show

that for any $0 \leq f \leq 1$, given $m = n(1/\varepsilon)^{f^2 d^2}$ units of storage (that is, $\rho = 1/\varepsilon^{f^2 d^2}$), the query time in the semigroup arithmetic model is at least

$$\Omega \left(\left(\frac{1}{\varepsilon} \right)^{d-2\sqrt{d}} / \left(\frac{1}{\varepsilon} \right)^{2fd} \right).$$

See Theorem 3.1(i) for details.

- We extend the above results to the case of integral semigroups. We show that under the assumption of convex generators the query time is at least

$$\Omega \left(\left(\frac{1}{\varepsilon} \right)^{d-2} / \left(\frac{1}{\varepsilon} \right)^{2fd} \right).$$

See Theorem 3.1(ii) for details.

- We present a lower bound for answering ε -approximate range queries over κ -smooth convex ranges and over any faithful semigroup. We show that, irrespective of space, the query time in the algebraic decision-tree model, is at least $\Omega \left(\log n + (1/\varepsilon)^{\frac{d-1}{2}} \right)$. See Theorem 4.1 for details.
- We show that an existing data structure for approximate spherical range searching [4] can be generalized to answer ε -approximate range queries for κ -smooth convex ranges over any idempotent semigroup. We show that, for any fixed κ , queries can be answered from a data structure of space $O(n/\varepsilon)$ in time $O \left(\log n + (1/\varepsilon)^{\frac{d-1}{2}} \right)$. See Theorem 4.2 for details. As in [7] we assume that ranges satisfy the *unit-cost test assumption*, which implies the primitive operations involving ranges can be computed in constant time. See Section 4.2 for a formal definition.

Table 1 summarizes both our results and recent results on the complexity of approximate range searching for Euclidean balls, κ -smooth, and rotated unit hypercube ranges in dimension d , assuming n points. To simplify and clarify the ε -dependencies with dimension, we have explicitly omitted factors involving $\log n$ and ε -factors that are independent of d . For the sake of comparison, we also provide known bounds for Euclidean balls and general convex ranges from [4], [5] and [7]. All these results assume either $O(n)$ or $O(n/\varepsilon)$ space, except the lower bound for κ -smooth ranges in the idempotent case, which holds independent of space.

The table shows that in the case of roughly linear space, approximate range searching can be performed most efficiently for Euclidean balls and smooth convex ranges over idempotent semigroups. It is not hard to see intuitively why this is to be expected. Approximate range searching in the semigroup arithmetic model is similar to a shape approximation problem, where the objective is to approximately cover a shape with a small number of precomputed canonical shapes (corresponding to the generator subsets). It is possible to approximately cover smooth ranges and Euclidean balls with $O(1/\varepsilon^{(d-1)/2})$ overlapping large Euclidean balls. This is not possible (within the given space bounds) when the range has sharp corners or when the semigroup is integral.

Converting this intuition into lower bound proofs involves considerable work. Our lower bound proofs for rotated unit hypercubes are based on the general framework developed by Chazelle [13] for exact simplex range searching and Brönnimann, Chazelle, and Pach [11] for exact halfspace range searching. Both papers are based on a number of geometric tools, such as the isoperimetric

		Idempotent	Integral
Euclidean Balls	Lower Bound	$1/\varepsilon^{\frac{d}{2}}$ [5]	$1/\varepsilon^d$ [5]
	Upper Bound	$1/\varepsilon^{\frac{d}{2}}$ [5]	$1/\varepsilon^d$ [4]
Rotated Unit Hypercubes	Lower Bound	$1/\varepsilon^{d-2\sqrt{d}}$ (new)	$1/\varepsilon^d$ (new)
	Upper Bound	$1/\varepsilon^d$ [7]	
κ -Smooth Convex Bodies (unit-cost test)	Lower Bound	$1/\varepsilon^{\frac{d}{2}}$ (new)	$1/\varepsilon^d$ [5]
	Upper Bound	$1/\varepsilon^{\frac{d}{2}}$ (new)	$1/\varepsilon^d$ [7]

Table 1: Query times (ignoring logarithmic and $1/\varepsilon^{O(1)}$ factors) for n points either linear or roughly linear space.

inequality, the slicing lemma, and Macbeath regions, which were tailored to these particular problems. In our case, these tools need to be adapted and generalized to our new setting. We believe that these generalizations, especially the generalizations of Macbeath regions may be of independent interest. Our lower bound for κ -smooth ranges is based on an entirely different decision-tree approach.

The remainder of the paper is organized as follows. In Section 2 we present preliminary definitions, which will be used throughout the paper. In Section 3 we give our main result, namely lower bounds on the complexity of approximate range searching for rotated unit hypercubes for both idempotent and integral semigroups. In Section 4 we present lower and upper bounds on the complexity of range searching for κ -smooth convex ranges over idempotent semigroups. Concluding remarks are given in Section 5.

2 Preliminaries

Before presenting our results we begin with some general definitions and assumptions. Throughout we assume that the dimension d is a fixed constant greater than 1, and treat n and ε as asymptotic quantities. Unless otherwise stated, we will use the term “constant” to refer to any fixed quantity, which may depend on d but not on n or ε . To avoid specifying the many real-valued constants that arise in our constructions and analyses, we will often hide them using asymptotic notation. For positive real x , we use the notation $O(x)$ (resp. $\Omega(x)$) to mean a quantity whose value is at most (resp. at least) cx for an appropriately chosen constant c .

Let $(S, +)$ be a commutative semigroup. Our lower bounds assume that the semigroup is *faithful*, meaning that any two identically equal linear forms have the same set of variables [13]. For example, $(\mathbb{N}, +)$, (\mathbb{R}, \min) , and $(\{0, 1\}, \vee)$ are faithful, but $(\{0, 1\}, + \bmod 2)$ is not. We will assume that each element of S can be stored in unit space, and that for any two elements $x, y \in S$, their semigroup sum $x + y$ can be computed in constant time. Let P be a set of n points in \mathbb{R}^d and let $w : P \rightarrow S$ be a function that assigns a semigroup value in S to each point in P . For any subset G of P , we define its weight $w(G) = \sum_{p \in G} w(p)$, where the summation is taken over the semigroup. Let \mathcal{Q} denote the set of query ranges in the range space. Recall that in the exact range searching problem, we are required to preprocess P so that for any query range $\eta \in \mathcal{Q}$, we can efficiently compute $w(P \cap \eta)$. Let η^+ denote the *expanded range*, consisting of all the points that lie within distance $\varepsilon \cdot \text{diam}(\eta)$ of η . A valid answer to an approximate range query is $w(P')$, where

P' is any subset of P satisfying $P \cap \eta \subseteq P' \subseteq P \cap \eta^+$.

Let us recall some of the basic elements of the *semigroup arithmetic model* for range searching [17, 21]. Given a set $\{x_1, \dots, x_n\}$ of n variables over S , a *generator* $G(x_1, \dots, x_n)$ is a linear form $\sum_{i=1}^n \alpha_i x_i$, where the values α_i are nonnegative integers, not all zero. A *storage scheme* is a set of generators $\{G_1, \dots, G_m\}$ satisfying the following property. For any query range $\eta \in \mathcal{Q}$, there exists a set $I_\eta \subseteq \{1, \dots, m\}$ and a set of labeled nonnegative integers $\{\beta_i : i \in I_\eta\}$ such that

$$w(P \cap \eta) = \sum_{i \in I_\eta} \beta_i G_i(w(p_1), \dots, w(p_n)) \quad (1)$$

holds for *any* weight function w . The *query time* for η is defined to be the size of the smallest such set I_η . In this model the *space* is the number of generators in the storage scheme. Intuitively, the generators correspond to partial sums that have been precomputed in the data structure. The query time in the semigroup arithmetic model counts the minimum number of semigroup operations on these generators needed to answer a query. The time for auxiliary operations such as determining which generators to use is ignored.

For our purposes the values α_i will be either 0 or 1, which means that we may identify each generator with a subset of the points P . In this context, the time to answer a query η is the smallest set of generators (subsets of P) whose union covers $P \cap \eta$. More formally, let \mathcal{G} be any storage scheme consisting of m generators. For any range $\eta \in \mathcal{Q}$, define $A_\eta \subseteq \mathcal{G}$ to be the smallest subset of generators of \mathcal{G} such that $\bigcup_{G \in A_\eta} G = P \cap \eta$. The worst-case query time for a storage scheme \mathcal{G} is the maximum of $|A_\eta|$ over all ranges $\eta \in \mathcal{Q}$. Given a space bound m , the worst-case query time complexity in the semigroup arithmetic model is the minimum worst-case query time over all storage schemes \mathcal{G} of size m [13]. The only modification necessary for approximate range searching is to define $A_\eta \subseteq \mathcal{G}$ to be the smallest set such that the union of the corresponding generators, $\bigcup_{G \in A_\eta} G$, contains all the points of P lying within η and none of the points lying outside of η^+ .

Let us define some other terms that will be used throughout the presentation. Let $\mathbb{U}^d = [0, 1]^d$ denote the unit hypercube in \mathbb{R}^d , and given a body K in \mathbb{R}^d , let $\mu(K)$ denote its Lebesgue measure. Let O denote the origin of the coordinate system. Given a point $p \in \mathbb{R}^d$, let $b(p, r)$ denote the closed Euclidean ball of radius r centered at p . Consider a compact (closed and bounded) convex body K in \mathbb{R}^d , and let u be a nonzero vector in \mathbb{R}^d . The *width* of K in the direction u is defined to be the orthogonal distance between the two supporting hyperplanes for K that are orthogonal to u . We denote this by $wid(K, u)$. (See Fig. 1(a).) The closed region bounded by two parallel $(d - 1)$ -dimensional hyperplanes in \mathbb{R}^d is called a *slab*, and its *width* is the orthogonal distance between these hyperplanes.




Fig. 1: Widths and caps.

Any hyperplane J that does not pass through the origin defines two closed halfspaces, one that contains the origin and one that does not. We refer to the first as the *inner halfspace*, denoted J^\leq and the second as the *outer halfspace*, denoted J^\geq . Given a halfspace H , let ∂H denote its bounding hyperplane. We say that two halfspaces are *parallel* if their bounding hyperplanes are parallel, and they both lie on the same sides of their respective bounding hyperplanes.

Consider a compact, convex body K and a halfspace H that has a nonempty intersection with K . Let u be a vector orthogonal to ∂H . (See Fig. 1(b).) The intersection $C = K \cap H$ is called the *cap* of K generated by H . The *width* of a cap, denoted $\text{wid}(C)$, is defined to be the width of C along its defining direction u . Given such a cap C of width w and a real $\lambda \geq 0$ we define the λ -*expansion* of C , denoted C^λ , to be the cap of K of width $\min(\lambda w, \text{wid}(K, u))$ generated by an appropriate translation of H . (See Fig. 1(c).) If λw exceeds the width of K in the direction u , then the expansion is simply equal to K itself.

The query ranges used in our lower bound analysis will consist of translates and rotations of a hypercube of side length $1/2$, which we will call *quads*. Our proof focuses on the portion of each quad that lies near one of its k -dimensional faces, for a suitably chosen k . To make this more precise, given an integer r , where $1 \leq r \leq d$, define an r -*corner* to be a sequence of r hyperplanes that are all mutually orthogonal and none of which passes through the origin. Given a d -corner $L = \langle J_1, J_2, \dots, J_d \rangle$, define the corresponding *quad*, denoted \square_L , to be the unique hypercube of side length $1/2$ that has a vertex at the intersection of the hyperplanes of L and which lies in the intersection of the corresponding inner halfspaces, $\bigcap_i J_i^\leq$. (See Fig. 2(a).) Note that such a hypercube is generally not axis-aligned. Given a quad \square_L , let \square_L^+ denote the ε -*expanded range* consisting of all the points of \square_L and all the points that lie within distance $\varepsilon \cdot \text{diam}(\square_L)$ of its boundary. (See Fig. 2(b).)




Fig. 2: The quad \square_L defined by $L = \langle J_1, J_2 \rangle$ (a) and the expanded quad (b).

Let \mathcal{Q} denote the set of all quads, which will form the ranges of our range space. Given n points and any storage scheme for \mathcal{Q} consisting of m generators, let $t(n, m)$ denote the worst-case query time in the arithmetic model, that is, the minimum number of generators needed to answer any query of \mathcal{Q} . The objective of the next section is to establish a lower bound on $t(n, m)$.

3 Lower Bound for Rotated Unit Hypercubes

In this section we prove the following theorem establishing a lower bound in the semigroup arithmetic model on the complexity of approximate range searching for unit hypercubes over any faithful semigroup.

Theorem 3.1 *Let $d \geq 2$ be a fixed dimension, and let f be a fixed parameter $0 \leq f \leq 1$. Consider ε -approximate range searching for n points over the range space of d -dimensional unit hypercubes for a weight function over any faithful semigroup, with $m = n(1/\varepsilon)^{f^2 d^2}$ units of storage (that is, $\rho = (1/\varepsilon)^{f^2 d^2}$). Then for all sufficiently small ε and sufficiently large n , we have the following.*

- (i) *For arbitrary (and hence idempotent) semigroups the worst-case query time in the semigroup arithmetic model is at least*

$$\Omega\left(\left(\frac{1}{\varepsilon}\right)^{d-2\sqrt{d}} \Big/ \left(\frac{1}{\varepsilon}\right)^{2fd}\right).$$

- (ii) *If the semigroup is integral, then under the assumption of convex generators the query time is at least*

$$\Omega\left(\left(\frac{1}{\varepsilon}\right)^{d-2} \Big/ \left(\frac{1}{\varepsilon}\right)^{2fd}\right).$$

Before giving the proof, let us consider some important special cases. When $f = 0$ and the semigroup is idempotent this implies that with $O(n)$ space the query time is at least $\Omega(1/\varepsilon^{d-2\sqrt{d}})$. For the integral case the corresponding bound is $\Omega(1/\varepsilon^{d-2})$. In order to bring the query time down to something approaching the $\Omega(1/\varepsilon^{(d/2)-O(1)})$ lower bound for Euclidean balls over idempotent semigroups, we consider $f = 1/4$. This leads to a query time of at least $\Omega(1/\varepsilon^{(d/2)-2\sqrt{d}})$. However, in this case the space increases dramatically to $O(n(1/\varepsilon)^{d^2/16})$, which is $O(n(1/\varepsilon)^{\Omega(d^2)})$. Thus, rotated unit hypercubes are significantly harder than Euclidean balls. This stands in stark contrast to exact range searching, where the two problems have very similar complexities as dimension increases.

Our lower bound is based on the general framework presented by Brönnimann, Chazelle, and Pach [11] (referred to henceforth as *BGP*) for establishing a lower bound on the complexity of halfspace range searching in the semigroup arithmetic model. Before presenting our methods, we begin with a high-level overview of this framework. The data point instance used in the proof consists of a set of uniformly distributed points in \mathbb{U}^d . This makes it possible to relate the volume of a convex subset of \mathbb{U}^d to the number of data points it contains.

The proof begins by defining a suitable set of query ranges and a suitable *region of interest* for each query range, which lies close to the range's boundary. The principal source of complexity in range searching is handling points that are close to the range's boundary. For halfspace range searching the region of interest is defined to be a thin slab close to the range's boundary. Recall that the goal in the arithmetic semigroup model is to relate the space (number of generators) to the query time (maximum number of generators needed to answer any query). In order to establish this relationship, a probability distribution is defined on the set of query ranges, and a lower bound is obtained by showing that if a generator covers a large number of points in the region of interest for some range, then it cannot be useful in this manner for many query ranges.

To understand the proof strategy better, we need to elaborate on this last point. Clearly, a worst-case query time of t implies that t generators suffice to cover all the data points in the region of interest. An easy implication is that at least half the points in the region of interest must be covered by generators contained within the ε -expanded query range that each cover at least a fraction of $1/2t$ of the points in the region of interest. We call these generators (absolutely) *fat* with respect to the query. This suggests the following definition of *usefulness*. Consider a random

variable that is zero if the generator is not fat with respect to the query, otherwise it is equal to the number of points in the region of interest covered by the generator. The usefulness of a generator is the expected value of this random variable over all ranges.

Techniques are presented in the BCP paper for establishing an upper bound on the usefulness of any generator. (We will describe these later.) Taking the product of this quantity with the number of generators yields an upper bound on the expected number of points covered by generators that are fat with respect to the query. By our earlier remarks, multiplying this by 2 gives an upper bound on the expected number of points in the region of interest. Using the fact that the data set is uniformly distributed, we can easily lower bound the expected number of points in the region of interest. Relating this back to the upper bound implies a lower bound on the number of generators necessary to achieve query time t .

We return to the BCP techniques for bounding the usefulness of a generator. These are designed specifically for halfspace queries. They involve an *isoperimetric inequality* and a *decomposition into convex bodies*. Their isoperimetric inequality bounds the probability that a fixed convex body is contained within a random query slab, as a function of the volume of the convex body and the thickness of the slab. Their decomposition involves breaking the convex hull of the generator into a small number of convex bodies, called *Macbeath regions*. The desired property is that if a generator is fat with respect to a query halfspace, then one of these bodies is contained completely within the halfspace. Furthermore, its volume is proportional to the generator's volume of intersection with the halfspace.

Our proof uses this framework as adapted to the context of approximate range searching [5]. However, we need to modify the techniques used to bound the usefulness of generators. Recall that our query ranges are hypercubes. We define our region of interest to be the intersection of r slabs associated with the facets of the hypercube, where r is a parameter between 1 and d , which is chosen to achieve the best bound. We modify BCP's isoperimetric inequality to bound the probability that a fixed convex body is contained within a randomly generated intersection of slabs. This is presented in Section 3.1. The decomposition of convex bodies needs to be generalized as well from halfspaces to the intersection of r mutually orthogonal halfspaces. The key elements are presented in Section 3.2, and mathematical technicalities upon which the method relies are formally established in Sections 3.2.1 and 3.2.2.

Once we have established the appropriate generalizations of the BCP techniques to our context, we then proceed with the main part of the analysis of the idempotent case, which we present in Section 3.3. One additional issue that arises only in the context of approximate range queries involves the thickness of the region of interest. In the BCP proof the thickness of the region of interest is adjusted to control the expected number of points that lie within the region. Thus, as n increases the thickness decreases. In approximate range searching, however, it is counterproductive to reduce the thickness of the region of interest to be smaller than the allowable approximation error. Following the approach of our earlier work [5], we handle this by first generating many separate instances of the same point distribution, called *replicants*, next we apply the lower bound argument separately to each replicant, and finally we combine the results. These three elements are presented in Sections 3.3.1, 3.3.2, and 3.3.3, respectively. Finally, in Section 3.4 we explain how to modify the analysis in the previous sections to obtain the lower bound for the integral case.

3.1 Generalized Isoperimetric Inequality

An important tool in our analysis is an appropriate generalization of Chazelle’s isoperimetric inequality [11, 13] to the context of r -corners. In this section we present this result.

Let O denote the origin of the coordinate frame in \mathbb{R}^d . For any hyperplane H not passing through O , let $q = (x_1, \dots, x_d)$ denote the point on H such that the segment \overline{Oq} is orthogonal to H . As in [13], define the measure of any set \mathcal{X} of hyperplanes as follows:

$$\int_{\mathcal{X}} dH = \int_{H \in \mathcal{X}} \frac{dx_1 \wedge \dots \wedge dx_d}{\|q\|^{d-1}}.$$

The choice of this measure is based on the fact that it is invariant under rigid motions of \mathbb{R}^d [13]. We will be applying this measure to the set of hyperplanes that pass through some bounded region of \mathbb{R}^d . It is well known that the total measure of such a set of hyperplanes is bounded [13], and therefore we can interpret the measure as a probability density by dividing it by the total measure of the region. We shall say that a hyperplane is *random* over a set \mathcal{X} of hyperplanes in \mathbb{R}^d of bounded measure if it is selected with the probability density associated with dH .

Throughout, unless otherwise specified we will use the term *hyperplane* to denote a $(d - 1)$ -dimensional hyperplane in \mathbb{R}^d . When dealing with hyperplanes of arbitrary dimension k , for $0 \leq k \leq d - 1$, we will often use the term *k-flat*. Let H' denote an arbitrary k -flat in \mathbb{R}^d . By using any rigid transformation ϕ that maps \mathbb{R}^k to H' , we can naturally transform the above measure on $(k - 1)$ -dimensional hyperplanes in \mathbb{R}^k to a measure on $(k - 1)$ -flats that lie on H' . Thus we can define the notion of a random $(k - 1)$ -flat on H' . Observe that the resulting measure will be invariant under rigid motions of \mathbb{R}^d that fix H' . Because the measure in \mathbb{R}^k is invariant under rigid motions, the choice of ϕ is unimportant. But for the sake of concreteness, let us assume that it maps the origin of \mathbb{R}^k to the point O' of H' that is closest to the origin of \mathbb{R}^d .

Recall that an r -corner is a sequence of r mutually orthogonal hyperplanes. Observe that an r -corner can be specified recursively by giving a sequence of flats $\langle f_1, f_2, \dots, f_r \rangle$, where for $1 \leq i \leq r$, the flat f_i is a $(d - i)$ -flat that lies on f_{i-1} , and f_0 is the full d -dimensional space \mathbb{R}^d . The associated r -corner $L_r = \langle J_1, J_2, \dots, J_r \rangle$ is defined by setting J_1 to f_1 , then for $i = 2, 3, \dots, r$, we define J_i to be the unique hyperplane that passes through f_i and is orthogonal to J_j , for $1 \leq j \leq i - 1$. (See Fig. 3.)




Fig. 3: The 3-corner determined by a sequence of flats.

In order to define a random r -corner consider the following random process, which generates a sequence of nested flats $\langle f_1, f_2, \dots, f_d \rangle$. Recall that $f_0 = \mathbb{R}^d$. For $i \geq 1$, assume inductively that

$\langle f_0, f_1, \dots, f_{i-1} \rangle$ have already been generated. Let O'_{i-1} be the closest point of f_{i-1} to the origin. Consider the set of $(d-i)$ -dimensional flats on f_{i-1} that intersect the ball of radius \sqrt{d} centered at O'_{i-1} . From this set select f_i randomly with the probability density associated with dH . We say that the r -corner associated with this sequence of flats $\langle f_1, f_2, \dots, f_r \rangle$ is *random*. Since we only consider hyperplanes that lie within some fixed distance of the origin, the resulting set is of bounded measure. Let dL_r denote the corresponding probability density, and let $dL = dL_d$.

Before presenting our generalization of the isoperimetric inequality, let us review Chazelle's original isoperimetric inequality. Let $\alpha > 0$ be a real parameter. Given any hyperplane J , let $S^\alpha(J)$ denote the slab consisting of points in \mathbb{R}^d whose distance from J is at most α . This slab has width 2α . Recall that \mathbb{U}^d denotes the unit hypercube in dimension d .

Lemma 3.1 (Chazelle [13]) *Given any compact convex body $K \subseteq \mathbb{U}^d$,*

$$\int_{S^\alpha(H) \supseteq K} dH = O\left(\frac{\alpha^{d+1}}{\mu(K)}\right).$$

When expressed in the language of probability theory this lemma implies that if H is selected randomly from a set of hyperplanes in \mathbb{R}^d whose measure is $\Theta(1)$, then the probability that the slab $S^\alpha(H)$ contains K is at most $O(\alpha^{d+1}/\mu(K))$.

Let s be an integer $1 \leq s \leq r$. Given any r -corner $L_r = \langle J_1, \dots, J_r \rangle$, we can associate it with a region, called an s -corner slab, which is the intersection of the α -distance slabs associated with the first s of its defining hyperplanes. More precisely, let $S_s^\alpha(L_r) = \bigcap_{i=1}^s S^\alpha(J_i)$ denote this s -corner slab. When $s = r$ we omit the subscript, that is $S^\alpha(L_r) = S_r^\alpha(L_r)$. We define the *width* of this s -corner slab to be 2α . We now establish our generalization of the isoperimetric inequality from hyperplane slabs to r -corner slabs.

Lemma 3.2 *Given any compact convex body $K \subseteq \mathbb{U}^d$,*

$$\int_{S^\alpha(L_r) \supseteq K} dL_r = O\left(\left(\frac{\alpha^{d+1}}{\mu(K)}\right)^r\right).$$

Proof: The integral on the left-hand side above represents the probability that the r -corner slab of width 2α associated with a random r -corner contains K . Let $L_r = \langle J_1, \dots, J_r \rangle$ be any r -corner. By our earlier discussion, L_r can be represented alternatively as a sequence of flats $\langle f_1, \dots, f_r \rangle$. By definition, the r -corner slab $S^\alpha(L_r)$ contains K if and only if $S^\alpha(J_i)$ contains K , for $1 \leq i \leq r$.

Since $J_1 (= f_1)$ is selected randomly from a set of $\Theta(1)$ measure, by Lemma 3.1 the probability that the slab $S^\alpha(J_1)$ contains K is at most $O(\alpha^{d+1}/\mu(K))$. Now consider any hyperplane f_1 such that $S^\alpha(f_1) \supseteq K$. Let K' be the orthogonal projection of K onto f_1 . Let $S_{f_1}^\alpha(f_2)$ denote the slab consisting of points in f_1 whose distance from f_2 is at most α . Observe that $S^\alpha(J_2) \supseteq K$ if and only if $S_{f_1}^\alpha(f_2) \supseteq K'$. Again, by observing that f_2 is selected randomly from a set of measure $\Theta(1)$, and applying Lemma 3.1, it follows that the probability that the slab $S_{f_1}^\alpha(f_2)$ contains K' is at most $O(\alpha^d/\mu(K'))$.

Since $S^\alpha(f_1) \supseteq K$ and this slab has width 2α , it follows that $\mu(K) \leq 2\alpha \cdot \mu(K')$. Therefore assuming that $S^\alpha(J_1)$ contains K , the conditional probability that $S_{f_1}^\alpha(f_2)$ contains K' (and hence of $S^\alpha(J_2)$ containing K) is at most $O(\alpha^{d+1}/\mu(K))$. Continuing in this manner, we see that the integral in the statement of the lemma is the product of r factors of $\alpha^{d+1}/\mu(K)$, ignoring constant factors in d . \square

This leads to the following lemma, which will be useful in our analysis.

Lemma 3.3 *Let r be any integer where $1 \leq r \leq d$. Given any compact convex body $K \subseteq \mathbb{U}^d$, we have*

$$\int_{S_r^\alpha(L) \supseteq K} dL = O\left(\left(\frac{\alpha^{d+1}}{\mu(K)}\right)^r\right).$$

Proof: The left-hand side is the probability that the r -corner slab associated with a random d -corner contains K . In view of the definition of dL and dL_r this quantity is the same as $\int_{S_r^\alpha(L) \supseteq K} dL_r$. The claim now follows from Lemma 3.2. \square

3.2 Macbeath Regions and Generalizations

The slicing lemma is another key in BCP's lower bound proof [11]. Intuitively it states that, given a convex body $K \subset \mathbb{R}^d$ of unit volume, and $0 < \beta < 1$, there exists a collection of $O((1/\beta)^{1-2/(d+1)})$ disjoint convex bodies, such that for any cap of K of volume at least β , one of these bodies lies entirely within the cap and its volume is proportional to the volume of the cap.

The concept of a Macbeath region was an integral element in the proof of the slicing lemma of [11]. Given a compact convex body K , a point $x \in K$, and a real number $\lambda > 0$, the set

$$M(x, \lambda) = x + \lambda((K - x) \cap (x - K))$$

is called a *Macbeath region*. Note that a Macbeath region is convex and centrally symmetric about the point x , and if $\lambda \leq 1$, then $M(x, \lambda) \subseteq K$. Macbeath regions have numerous applications in the theory of convex bodies [8–10].

Our ranges are not halfspaces, but rather involve the intersection of a number of halfspaces. In this section we establish Lemma 3.5 below, which is our generalization of the slicing lemma of [11] to the intersection of multiple halfspaces. Informally this lemma states that given a compact, convex body K and a volume parameter ρ , there exists a relatively small number of (overlapping) convex bodies contained within K satisfying the following property. Given any collection of halfspaces whose region of intersection with K has volume exceeding ρ , at least one of these convex bodies will fully contain this region of intersection, and it will not be significantly larger in the sense that it will be contained within a suitable expansion of each of the caps defined by these halfspaces. (Note that one difference between our slicing lemma and that of [11] is that the body is contained within the cap, whereas ours contain the cap. This distinction seems to be necessary for technical reasons arising from our approach.)

Throughout this section we define the real parameter $\beta_0 = 1/(2d)^{2d}$. Recall that given a cap C and positive real λ , the notation C^λ denotes the expansion of C by a width factor of λ . Before presenting our slicing lemma, we present the following a simpler version, which involves a single halfspace. (See Fig. 4(a) and (b).) It is a technical modification of the slicing lemma of [11].

Lemma 3.4 *There exist constants $c_1, c_2, c_3 > 1$ such that the following holds. Consider a compact, convex body $K \subset \mathbb{R}^d$ of volume v and a positive real $\rho \leq \beta_0 v$. Let $\beta = \rho/v$. There exists a collection of at most $c_1(1/\beta)^{1-2/(d+1)}$ convex bodies $K_1, K_2, \dots \subseteq K$ such that*

- (i) $\rho \leq \mu(K_i) \leq c_2\rho$, for all i .
- (ii) Let H be a halfspace, and let C denote the cap $K \cap H$. If $\mu(C) = \rho$ then there exists a K_j such that (a) $K_j \supseteq C$, and (b) $K_j \subseteq C^{c_3}$.




Fig. 4: Slicing lemmas

We provide a proof of the above lemma in Section 3.2.1. We will then apply this along with induction to prove the following slicing lemma. (See Fig. 4(c).) Its proof will be presented in Section 3.2.2.

Lemma 3.5 (Slicing Lemma) *There exist positive constants c_4 and λ such that the following holds. Consider a compact, convex body $K \subset \mathbb{R}^d$ of volume v , a positive real $\rho \leq v$, and a positive integer r . Let $\beta = \rho/v$. There exists a collection of at most*

$$c_4^r \left(\frac{1}{\beta} \right)^{1 - \frac{2}{d+1}} \log^{r-1} \left(\frac{2}{\beta} \right)$$

convex bodies $K_1, K_2, \dots \subseteq K$ such that given any r halfspaces H_1, H_2, \dots, H_r where $\mu(K \cap (H_1 \cap H_2 \cap \dots \cap H_r)) \geq \rho$, there exists a body K_j such that

- (i) $K_j \supseteq K \cap (H_1 \cap H_2 \cap \dots \cap H_r)$, and
- (ii) for $1 \leq i \leq r$, $K_j \subseteq C_i^\lambda$, where C_i^λ is the λ -expansion of the cap $C_i = K \cap H_i$.

3.2.1 Proof of Lemma 3.4

The proof of Lemma 3.4 follows from extensive but relatively straightforward modifications of the methods appearing in [8], [10], [11], and [16]. Before presenting the proof, we review a number of facts about Macbeath regions that will be useful for us. The first follows immediately from the definition of the Macbeath region. Throughout, K will denote a compact convex body, and H will denote a halfspace. Let x_H denote the center of mass of the section $K \cap \partial H$.

Lemma 3.6 *Let K be a compact, convex body, and let H be a halfspace. Let $C = K \cap H$ be a cap. Then $M(x_H, 1) \subseteq C^2$.*

The following lemma, proved by Ewald, Larman, and Rogers [16], is useful for establishing containment relationships between the Macbeath regions generated by suitably close points.

Lemma 3.7 *Let K be a compact, convex body. If for some $x, y \in K$, $M(x, \frac{1}{2}) \cap M(y, \frac{1}{2}) \neq \emptyset$, then $M(x, 1) \subset M(y, 5)$.*

The next result asserts that each suitably thin cap of K is contained within a Macbeath region centered on the cap's defining hyperplane. It was first proved by Ewald, Larman, and Rogers [16], but we use a variant due to Brönnimann, Chazelle, and Pach [11].

Lemma 3.8 *Let $K \subset \mathbb{R}^d$ be a compact convex body containing the ball of radius r centered at the origin. Let H be a halfspace not containing the origin such that the distance between ∂H and one of the supporting hyperplanes of K parallel to ∂H is at most $r/3$. Then $K \cap H \subseteq M(x_H, 3d)$.*

Given a compact convex body K and a real number $\beta > 0$, let K_β be the set of all points of K not contained in any halfspace H such that $\mu(K \cap H) = \beta$. K_β is called the *floating body*¹ of K for β . Bárány and Larman [10] established an upper bound on the volume of the region obtained by subtracting the floating body from the original convex body. This region is called the *wet part* of K .

Lemma 3.9 *Let $K \subset \mathbb{R}^d$ be a compact convex body of unit volume, and let $\beta > 0$ be any real number. Then $\mu(K \setminus K_\beta) = O\left(\beta^{\frac{2}{d+1}}\right)$.*

The following facts are easy to prove, and will be useful for the rest of our development. The first is a simple fact about centrally symmetric convex bodies. A proof appears in Bárány [8].

Lemma 3.10 *Let A and B be centrally symmetric convex bodies with center points a and b , respectively, where $B \subset A$. Then for any $\lambda > 1$, $b + \lambda(B - b) \subset a + \lambda(A - a)$.*

The following well known result, related to the John ellipsoid [18], states that any convex body can be transformed by a volume preserving transformation so that it is nested between two balls of similar radii. Recall that $b(O, r)$ denotes the Euclidean ball of radius r centered at the origin.

Lemma 3.11 *Given any compact convex body $K \subset \mathbb{R}^d$, there exists a volume-preserving affine transformation T such that $b(O, r) \subseteq T(K) \subseteq b(O, dr)$, for some $r > 0$.*

The next fact follows from straightforward geometric calculations. A proof can be found in Lemma 2.7 of [11]. (Recall the constant β_0 defined just prior to Lemma 3.5.)

Lemma 3.12 *Let $K \subset \mathbb{R}^d$ be a compact convex body of unit volume such that $b(O, r) \subseteq K \subseteq b(O, dr)$, for some $r > 0$. Let H be a halfspace such that $\mu(K \cap H) \leq \beta_0$. Then the width of the cap $K \cap H$ is at most $r/3$.*

We now present the proof of Lemma 3.4. As a convenience we assume that K has been scaled to unit volume. It follows that we need to establish properties (i) and (ii), with β written in place of ρ . Note that in light of Lemma 3.11, we may assume that K has been transformed so that $b(O, r) \subseteq K \subseteq b(O, dr)$, for some $r > 0$. Recall that for any halfspace H that intersects K , the center of mass of the section $K \cap \partial H$ is denoted by x_H .

Consider the following incremental process for constructing a set \mathcal{H} of halfspaces. Initially \mathcal{H} is the empty set. A general step of this process works as follows. If there exists a halfspace H such that $\mu(K \cap H) = \beta$ and $M(x_H, \frac{1}{2})$ does not intersect any of the Macbeath regions $\{M(x_{H'}, \frac{1}{2}) : H' \in \mathcal{H}\}$, then add H to \mathcal{H} . (Fig. 5(a) illustrates $M(x_H, \frac{1}{2})$.) If no such H exists the process terminates.

¹The term *floating body* originates from the physical analogy with a buoyant object that is floating on the flat surface of a body of water, so that the volume of the body lying beneath the water's surface is β .




Fig. 5: Entities used in the proof of Lemma 3.4.

We will show that the set of convex bodies

$$\mathcal{K} = \{M(x_{H'}, 15d) \cap K : H' \in \mathcal{H}\}$$

satisfies all the properties of this lemma.

Our first task is to bound the number of bodies in \mathcal{K} . We begin by showing that $|\mathcal{H}| = O((1/\beta)^{1-2/(d+1)})$. For each halfspace $H' \in \mathcal{H}$, associate a region $R(H') = M(x_{H'}, \frac{1}{2}) \cap H'$ with it. (See Fig. 5(b).) We claim that all the regions $R(H')$ for $H' \in \mathcal{H}$ are contained in the wet part, $K \setminus K_\beta$, and each has volume $\Omega(\beta)$. To prove the claim, note that since $\mu(K \cap H') = \beta$ (by definition of the floating body K_β) we have $K \cap H' \subseteq K \setminus K_\beta$. Also, by definition of a Macbeath region, $M(x_{H'}, 1) \subseteq K$. Therefore $R(H') = M(x_{H'}, \frac{1}{2}) \cap H' \subseteq K \cap H'$. It follows that $R(H') \subseteq K \setminus K_\beta$. Next we bound the volume of $R(H')$. Since $\mu(K \cap H') \leq \beta_0$, by Lemma 3.12 it follows that the width of the cap $K \cap H'$ is at most $r/3$. By Lemma 3.8, $M(x_{H'}, 3d) \supseteq K \cap H'$, and so $\mu(M(x_{H'}, 3d)) \geq \beta$. It follows that $\mu(M(x_{H'}, \frac{1}{2})) = \Omega(\beta)$. Since Macbeath regions are centrally symmetric and $\partial H'$ bisects $M(x_{H'}, \frac{1}{2})$ it follows that

$$\mu(R(H')) = \frac{1}{2} \mu\left(M\left(x_{H'}, \frac{1}{2}\right)\right) = \Omega(\beta).$$

Since the regions $R(H')$ for $H' \in \mathcal{H}$ are pairwise disjoint (by construction) and by Lemma 3.9, $\mu(K \setminus K_\beta) = O(\beta^{\frac{2}{d+1}})$, it follows by a standard packing argument that

$$|\mathcal{H}| \leq \frac{\mu(K \setminus K_\beta)}{\min_{H' \in \mathcal{H}} \mu(R(H'))} = O\left(\left(\frac{1}{\beta}\right)^{1-\frac{2}{d+1}}\right).$$

Because the elements of \mathcal{K} are in 1-1 correspondence with the elements of \mathcal{H} , the same bound applies to the number of bodies of \mathcal{K} .

To complete the proof, we show that \mathcal{K} satisfies conditions (i) and (ii) in the statement of Lemma 3.4. To establish (i) consider a convex body $K' = M(x_{H'}, 15d) \cap K$, where $H' \in \mathcal{H}$. We showed above that $M(x_{H'}, 3d) \supseteq K \cap H'$. Thus $K' \supseteq K \cap H'$. Since $\mu(K \cap H') = \beta$, it follows that $\mu(K') \geq \beta$.

Next, we show that $\mu(K')$ is $O(\beta)$. Recall that $M(x_{H'}, 1) \subseteq K$, which implies that $M(x_{H'}, 1) \cap H' \subseteq K \cap H'$. Therefore, $\mu(M(x_{H'}, 1) \cap H') \leq \mu(K \cap H') = \beta$. Recalling that Macbeath regions are centrally symmetric and that $\partial H'$ bisects $M(x_{H'}, 1)$ we have

$$\mu(M(x_{H'}, 1)) = 2\mu(M(x_{H'}, 1) \cap H') = O(\beta).$$

Thus, $\mu(K') \leq \mu(M(x_{H'}, 15d)) = O(\beta)$, as desired. This completes the proof of (i).

It remains to show that for any halfspace H with $\mu(K \cap H) = \beta$ there exists a convex body $K' \in \mathcal{K}$ that satisfies conditions (ii.a) and (ii.b). From the construction of \mathcal{H} , it is clear that there exists $H' \in \mathcal{H}$ such that $M(x_H, \frac{1}{2}) \cap M(x_{H'}, \frac{1}{2}) \neq \emptyset$. We claim that $K' = M(x_{H'}, 15d) \cap K$ is the desired body.

Let $C = K \cap H$. We first establish (ii.a), namely that $K' \supseteq C$. By the same reasoning that was used above to show that $M(x_{H'}, 3d) \supseteq K \cap H'$, we can show that $M(x_H, 3d) \supseteq K \cap H$. Since $M(x_H, \frac{1}{2}) \cap M(x_{H'}, \frac{1}{2}) \neq \emptyset$, by Lemma 3.7 we have $M(x_H, 1) \subset M(x_{H'}, 5)$. Since Macbeath regions are centrally symmetric, we can apply Lemma 3.10 with $A = M(x_{H'}, 5)$, $B = M(x_H, 1)$ and $\lambda = 3d$ to obtain $M(x_H, 3d) \subset M(x_{H'}, 15d)$. Since $M(x_H, 3d) \supseteq K \cap H$, it follows that $M(x_{H'}, 15d) \supseteq K \cap H$. Thus, $K' = M(x_{H'}, 15d) \cap K \supseteq K \cap H$, as desired.

To show (ii.b), we will use a similar argument as used for condition (ii.a). Recall the proof of condition (ii.a). By swapping the roles of H and H' , we have by Lemma 3.7 that $M(x_{H'}, 1) \subset M(x_H, 5)$. By Lemma 3.10 with $A = M(x_H, 5)$, $B = M(x_{H'}, 1)$ and $\lambda = 15d$, we obtain $M(x_{H'}, 15d) \subset M(x_H, 75d)$. From Lemma 3.6 we have $M(x_H, 1) \subseteq C^2$, where $C = K \cap H$. Therefore, $M(x_H, 75d) \cap K \subseteq C^{(1+75d)}$. Thus, $M(x_{H'}, 15d) \cap K \subseteq C^{(1+75d)}$, which establishes condition (ii.b), and completes the proof of Lemma 3.4.

3.2.2 Proof of Lemma 3.5

Now that we have established Lemma 3.4, we use it to prove Lemma 3.5. Before doing so we provide two preliminary lemmas. The first lemma shows that if a cap for some convex body has sufficiently high volume, then its expansion by a suitable constant contains the entire body.

Lemma 3.13 *There exists a constant $\lambda > 1$ such that the following holds. Let $K \subset \mathbb{R}^d$ be a compact, convex body of unit volume, and let H be a halfspace in \mathbb{R}^d . Let C denote the cap $K \cap H$. If $\mu(C) \geq \beta_0$, then $K \subseteq C^\lambda$.*

Proof: In view of Lemma 3.11, it suffices to prove this lemma for K such that $b(O, r) \subseteq K \subseteq b(O, dr)$, for some $r > 0$. Let H be a halfspace such that the cap $C = K \cap H$ has volume $\mu(C) \geq \beta_0$. We will show that $\text{wid}(C) \geq \beta_0 r / d^d$, which will imply that the λ -expansion of C contains K , for $\lambda = 2d^{d+1} / \beta_0$.

Let ω_d denote the volume of a d -dimensional unit ball. For the sake of contradiction, suppose that $\text{wid}(C) < \beta_0 r / d^d$. Since $K \subseteq b(O, dr)$, the $(d-1)$ -dimensional cross-sectional volume of the intersection of K with any hyperplane, is at most $\omega_{d-1}(dr)^{d-1}$. It follows that

$$\mu(C) < \omega_{d-1}(dr)^{d-1} \frac{\beta_0 r}{d^d} \leq \beta_0 \omega_d r^d,$$

where we have used the fact that $\omega_{d-1} \leq d\omega_d$. Since $b(O, r) \subseteq K$, we have $\omega_d r^d \leq \mu(K) = 1$. Thus $\mu(C) < \beta_0$, which contradicts our initial assumption. It follows that $\text{wid}(C) \geq \beta_0 r / d^d$, as desired. \square

The second lemma shows that two parallel caps (that is, bounded by parallel halfspaces) of similar volumes have similar widths. (Note that the converse is also true, but we will not need it.) The proof is similar to that of Theorem 7 in [10] by Bárány and Larman.

Lemma 3.14 *Let $K \subset \mathbb{R}^d$ be a compact, convex body, and let H_1 and H_2 be two parallel halfspaces, where $H_1 \subseteq H_2$. Consider the two caps $C_1 = K \cap H_1$ and $C_2 = K \cap H_2$, and suppose that $\mu(C_2) \leq 2\mu(C_1)$. Then $\text{wid}(C_2) \leq 2d \cdot \text{wid}(C_1)$.*

Proof: Note that it suffices to prove the lemma under the assumption that both ∂H_1 and ∂H_2 intersect K . Let u denote the unit vector orthogonal to the bounding hyperplanes of H_1 and H_2 . For $0 \leq z \leq 1$, let $J(z)$ denote the hyperplane $\{x : (x \cdot u) = z\}$. Through an appropriate translation, we may assume that $J(0)$ is a supporting hyperplane to K . Finally, by directing u properly, we may assume that $\partial H_1 = J(z_1)$ and $\partial H_2 = J(z_2)$, where $0 \leq z_1 \leq z_2$. (See Fig. 6.) Let V_1 be the maximum value of the $(d-1)$ -dimensional cross-sectional volume $\mu(J(z) \cap K)$ in the range $0 \leq z \leq z_1$. Clearly $\mu(C_1) \leq V_1 z_1$.




Fig. 6: Proof of Lemma 3.14.

Next we compute a lower bound on $\mu(C_2)$. Let y_1 and y_2 be any points in the cross-section $J(0) \cap K$ and $J(z_2) \cap K$, respectively. Let V_2 be the maximum value of $\mu(J(z) \cap K)$ in the range $0 \leq z \leq z_2$ and suppose that this maximum is realized at $z = z'$. Observe that C_2 contains the two cones with common base $J(z') \cap K$ and apexes at y_1 and y_2 , respectively. (See Fig. 6.) As the sum of the volume of these two cones is $V_2 z_2 / d$, it follows that $\mu(C_2) \geq V_2 z_2 / d$. Obviously $V_2 \geq V_1$ and therefore $\mu(C_2) \geq V_1 z_2 / d$. By our earlier remarks, $\mu(C_1) \leq V_1 z_1$, so we have $\mu(C_2) / \mu(C_1) \geq z_2 / (d z_1)$. By the statement of the lemma, $\mu(C_2) / \mu(C_1) \leq 2$. It follows that $z_2 \leq 2 d z_1$, as desired. \square

We now have all the pieces we need to prove Lemma 3.5. The idea is to apply Lemma 3.4 inductively for each halfspace. Consider applying this lemma directly to K where the value of ρ is proportional to $\mu(K \cap H_1)$. Given the first halfspace H_1 , it follows that there exists a convex body K_j that encloses the cap $K \cap H_1$. (See Fig. 7(a).) We then apply Lemma 3.4 recursively to each of the bodies of this first-level decomposition, which provides a second-level decomposition of K_j . By the lemma there exists a body in this second-level decomposition that contains the cap $K_j \cap H_2$. (See Fig. 7(b).) We continue in this manner for all r halfspaces.




Fig. 7: Proof of Lemma 3.5

One complication in this approach is that, although we have a bound on the final volume $K \cap (H_1 \cap H_2 \cap \dots \cap H_r)$, we do not have a bound on the volume of individual caps like $K \cap H_1$. For this reason, when we apply the construction of Lemma 3.4, we will do it for a geometrically increasing series of volume bounds, and then select the appropriate body at each step of the induction.

Let us define this multi-level decomposition more formally. Consider a compact, convex body K and let $v = \mu(K)$. Given a positive real $\rho \leq v$, a ρ -decomposition of K is a set of convex bodies, each contained within K , defined as follows.

Case I: If $\rho \geq \beta_0 v$, then it is $\{K\}$.

Case II: If $\rho < \beta_0 v$, then it is the union of $\{K\}$ together with the union of the sets of convex bodies described in Lemma 3.4 for

$$\rho' = \beta_0 v, \frac{\beta_0 v}{2}, \frac{\beta_0 v}{4}, \dots, \frac{\beta_0 v}{2^{s-1}},$$

where s is the smallest integer such that $\rho \geq \beta_0 v / 2^s$.

Note that the elements of the decomposition will generally overlap one another. Also, since $\beta_0 v / 2^{s-1} > \rho$, by property (i) of Lemma 3.4, it follows that the volume of each of these elements is at least ρ .

Next, we introduce the concept of an r -level decomposition of K , which involves recursively applying the decomposition to its own elements. As above, given K and ρ , and a positive integer r , we define an r -level ρ -decomposition of K as follows. We assume that $\rho \leq \mu(K)$. If $r = 1$, then it is just a ρ -decomposition of K . Otherwise, for $r > 1$, for each convex body K' in the ρ -decomposition of K recursively compute an $(r-1)$ -level ρ -decomposition of K' . The union of these decompositions over all K' is the final r -level decomposition.

We will show that the r -level ρ -decomposition of K satisfies the conditions of Lemma 3.5. Our first task is to bound the number of elements of an r -level decomposition. The proof is based on an induction on r and the bound on the single level decomposition of Lemma 3.4.

To simplify notation somewhat, throughout the proof let $f(d) = 1 - \frac{2}{d+1}$ denote the term in the exponent in the statement of the lemma. For the basis case, assume $r = 1$. If $\beta_0 v \leq \rho \leq v$, the ρ -decomposition consists of only one convex body, and the desired bound holds trivially. If $\rho < \beta_0 v$, the number of convex bodies in the ρ -decomposition is at most

$$1 + \sum_{i=0}^{s-1} c_1 \left(\frac{2^i}{\beta_0} \right)^{f(d)}.$$

(We have added 1 to account for K .) This is a geometric progression, and so is dominated by the largest term, which is $(2^{s-1}/\beta_0)^{f(d)}$. By definition of s , $\rho < \beta_0 v / 2^{s-1}$, and so up to a constant factor, the number of convex bodies in the ρ -decomposition is at most

$$\left(\frac{2^{s-1}}{\beta_0} \right)^{f(d)} \leq \left(\frac{v}{\rho} \right)^{f(d)},$$

which is within the stated bound for a suitable constant c_4 .

For the induction step, assume that $r > 1$. If $\beta_0 v \leq \rho \leq v$, the r -level ρ -decomposition of K consists of only one convex body. So the desired bound holds trivially. If $\rho < \beta_0 v$, consider the

set of convex bodies described in Lemma 3.4, where ρ' is set to $\beta_0 v / (2^i)$. The size of this set is at most $c_1 (2^i / \beta_0)^{f(d)}$. The volume of each convex body in this set is at most $\min\left(\frac{c_2 \beta_0}{2^i} v, v\right)$. By the induction hypothesis the size of the $(r-1)$ -level ρ -decomposition constructed for each such body is at most

$$c_4^{r-1} \left(\frac{c_2 \beta_0 v}{2^i \rho}\right)^{f(d)} \log^{r-2} \left(\frac{2v}{\rho}\right).$$

Therefore, the size of the r -level decomposition is at most

$$\left(\sum_{i=0}^{s-1} c_1 \left(\frac{2^i}{\beta_0}\right)^{f(d)} c_4^{r-1} \left(\frac{c_2 \beta_0 v}{2^i \rho}\right)^{f(d)} \log^{r-2} \left(\frac{2v}{\rho}\right)\right) + c_4^{r-1} \left(\frac{v}{\rho}\right)^{f(d)} \log^{r-2} \left(\frac{2v}{\rho}\right). \quad (2)$$

(The last term accounts for the $(r-1)$ -level decomposition of K .)

Simplifying Eq. (2), we find that the number of bodies in the r -level decomposition is at most

$$\begin{aligned} & \left(\sum_{i=0}^{s-1} c_1 c_4^{r-1} \left(\frac{c_2 v}{\rho}\right)^{f(d)} \log^{r-2} \left(\frac{2v}{\rho}\right)\right) + c_4^{r-1} \left(\frac{v}{\rho}\right)^{f(d)} \log^{r-2} \left(\frac{2v}{\rho}\right) \\ & \leq (s+1) c_1 c_4^{r-1} c_2^{f(d)} \left(\frac{v}{\rho}\right)^{f(d)} \log^{r-2} \left(\frac{2v}{\rho}\right). \end{aligned} \quad (3)$$

(The “+1” in the “ $(s+1)$ ” term is there to absorb the second term.)

Recall that $\rho < \beta_0 v / 2^{s-1}$, which yields $s+1 < \log(4\beta_0 v / \rho)$. Since $\beta_0 = 1/(2d)^{2d}$, we have $s+1 < \log(2v/\rho)$. Thus, by further simplifying Eq. (3) the number of bodies of the decomposition is at most

$$c_1 c_4^{r-1} c_2^{f(d)} \left(\frac{v}{\rho}\right)^{f(d)} \log^{r-1} \left(\frac{2v}{\rho}\right).$$

By making c_4 at least as large as $c_1 c_2^{f(d)}$, we find that the number of bodies of the decomposition is at most

$$c_4^r \left(\frac{v}{\rho}\right)^{f(d)} \log^{r-1} \left(\frac{2v}{\rho}\right),$$

which completes the induction proof.

Trivially, all the convex bodies in the r -level decomposition of K are contained within K . It remains to prove Properties (i) and (ii). Again, the proof is by induction on r .

For the basis case, assume $r = 1$. We need to show that for any halfspace H_1 with $\mu(K \cap H_1) \geq \rho$, there exists a convex body K_j in the 1-level ρ -decomposition of K such that (i) $K_j \supseteq K \cap H_1$ and (ii) $K_j \subseteq C_1^\lambda$, where C_1^λ is the λ -expansion of the cap $K \cap H_1$. Henceforth, let C_1 denote the cap $K \cap H_1$.

We consider two cases, based on $\mu(C_1)$. First, suppose that $\mu(C_1) \geq \beta_0 v$. Recall that the ρ -decomposition of K contains K . Clearly, $K \supseteq C_1$, so property (i) trivially holds. By Lemma 3.13, $K \subseteq C_1^\lambda$, for some constant λ , and so property (ii) also holds.

Next, suppose that $\rho \leq \mu(C_1) < \beta_0 v$. Recall that $\beta_0 v / (2^s) \leq \rho$, so $\beta_0 v / (2^s) \leq \mu(C_1) < \beta_0 v$. Therefore, for some i , $0 \leq i \leq s-1$, we must have

$$\frac{\beta_0 v}{2^{i+1}} \leq \mu(C_1) \leq \frac{\beta_0 v}{2^i}.$$

Consider expanding the cap C_1 by translating its defining hyperplane ∂H_1 until the volume of the cap is $\beta_0 v / (2^i)$. Let us denote this cap by C' . Clearly, $\mu(C') \leq 2\mu(C_1)$, and so it follows from Lemma 3.14 that the width of C' is at most $2d \cdot \text{wid}(C_1)$. Recall that the ρ -decomposition of K includes the set of convex bodies described in Lemma 3.4, where ρ' is set to $\beta_0 v / (2^i)$. As $\mu(C') = \beta_0 v / (2^i)$, it follows from Lemma 3.4 that there exists some K_j in this set such that $K_j \supseteq C'$ and $K_j \subseteq (C')^{c_3}$ for some constant c_3 . Since $C' \supseteq C_1$ we have $K_j \supseteq C_1$. Using the facts that the width of C' exceeds the width of C_1 by at most a constant factor, and $K_j \subseteq (C')^{c_3}$, we have $K_j \subseteq C_1^\lambda$ for a suitable constant λ .

For the induction step, assume that $r > 1$. Let H_1, H_2, \dots, H_r be any r halfspaces such that $\mu(K \cap (H_1 \cap H_2 \cap \dots \cap H_r)) \geq \rho$. Since $\mu(K \cap H_1) \geq \rho$, by the argument given for the basis case, there exists a $K' \subseteq K$ in the 1-level ρ -decomposition of K such that $K' \supseteq K \cap H_1$ and $K' \subseteq (K \cap H_1)^\lambda$. Since $K' \supseteq K \cap H_1$ we obtain

$$K \cap (H_1 \cap H_2 \cap \dots \cap H_r) \subseteq K' \cap (H_2 \cap \dots \cap H_r). \quad (4)$$

It follows that $\mu(K' \cap (H_2 \cap \dots \cap H_r)) \geq \rho$.

Applying the induction hypothesis, it now follows that the $(r-1)$ -level ρ -decomposition of K' must contain a convex body K_j such that $K_j \supseteq K' \cap (H_2 \cap \dots \cap H_r)$ and K_j is contained within each of the expanded caps $(K' \cap H_i)^\lambda$, for $2 \leq i \leq r$. Using Eq. (4), it follows that $K_j \supseteq K \cap (H_1 \cap H_2 \cap \dots \cap H_r)$, which proves property (i).

Next we prove property (ii). As just mentioned, we have $K_j \subseteq (K' \cap H_i)^\lambda$, for $2 \leq i \leq r$. Since $K' \subseteq K$, by definition, the width of cap $K' \cap H_i$ is no larger than the width of cap $K \cap H_i$. Therefore, for $2 \leq i \leq r$, $K_j \subseteq (K \cap H_i)^\lambda$. Also, by our earlier remarks, $K' \subseteq (K \cap H_1)^\lambda$ and $K_j \subseteq K'$, which implies that $K_j \subseteq (K \cap H_1)^\lambda$. Thus, for $1 \leq i \leq r$, we have $K_j \subseteq C_i^\lambda$, where C_i^λ is the λ -expansion of the cap $K \cap H_i$. This establishes property (ii), and completes the proof of Lemma 3.5.

3.3 Lower Bound for Rotated Hypercubes: The Idempotent Case

Having now established the necessary geometric tools, we are ready to tackle the proof of Theorem 3.1. This section is devoted to proving part (i) of this theorem, namely the lower bound in the idempotent case. As mentioned earlier, the general structure of the proof is similar to the proof we presented earlier for approximate spherical range searching [5], which in turn is an adaptation of the BCP proof [11] for exact halfspace range searching.

We begin with some of the basic concepts and properties that will be used throughout the proof. Recall that \mathbb{U}^d denotes the unit hypercube in \mathbb{R}^d . Let K be any convex body contained within \mathbb{U}^d , and let $k = |P \cap K|$. The BCP proof employs the notion of scattered point sets [11], which will be useful for us as well. A point set $P \subseteq \mathbb{U}^d$ is said to be *scattered* if the following holds for some constant $a' > 1$:

$$\frac{n}{a'} \mu(K) - \frac{\log n}{2} \leq k \leq a' n \mu(K) + \frac{\log n}{2}. \quad (5)$$

Intuitively scattered point sets behave like uniformly distributed point sets in the sense that the volume of a convex body is related to the number of points of the set it contains. This suggests the following lemma, which is a trivial variant of Lemma 2.1 in [11].

Lemma 3.15 (Brönnimann, Chazelle, Pach) *A random set of n points sampled uniformly and independently in \mathbb{U}^d is scattered with probability $1 - o(1)$.*

For our purposes it will be convenient to express scatteredness in terms of the following *scattered-point properties*, which (as we shall see below) follow from the above definition. For some constant $a > 1$:

Property 1: If $\mu(K) \geq \frac{a \log n}{n}$, then $k \geq \frac{n}{a} \mu(K)$.

Property 2: If $k \geq \log n$, then $k \leq a n \mu(K)$.

Property 3: If $k \geq \log n$, then $k \geq \frac{n}{a} \mu(K)$.

To see that these properties hold for any scattered set, first observe that the left inequality of Eq. (5) implies Property 1 (for $a \geq 3a'/2$), and the right inequality implies Property 2 (for $a \geq 2a'$). To see that Property 3 holds, we consider two cases. If the pre-condition of Property 1 holds, then the conclusion follows immediately. Otherwise, $\mu(K) < \frac{a \log n}{n}$, and since $k \geq \log n$ we have $\mu(K) < \frac{ak}{n}$. In either case the post-condition of Property 3 holds.

Let r denote an integer parameter, $1 \leq r \leq d$, whose value will be set later. Let ε denote the approximation error. We assume throughout that ε is a sufficiently small real number between 0 and 1. Recall from Section 2 that \mathcal{Q} denotes our range space, consisting of all quads, that is, hypercubes of side length $1/2$. Also recall that n denotes the number of points in the data set, $m \geq n$ denotes the number of generators, and $t = t(n, m)$ denotes the worst-case query time in the arithmetic model over all the ranges in \mathcal{Q} . Our bounds hold under the assumption that n is sufficiently large. As mentioned earlier, following the approach of our earlier work [5] our proof involves three parts: partitioning the set into a collection of disjoint sets called replicants, analyzing a single replicant, and then combining the results. These are presented in the next three subsections.

3.3.1 Decomposition into Replicants

We construct a set P of n data points for which we will argue that the query time must be sufficiently large for some range in \mathcal{Q} . As mentioned above in the overview of the proof, we will assume that P is composed of a collection of identical subsets, called *replicants*. Towards this end, let

$$n' = \frac{t}{\varepsilon^r} \log \frac{t}{\varepsilon^r}.$$

For simplicity we will assume that n is a multiple of n' . Consider any collection \mathcal{U} of n/n' unit hypercubes whose interiors are pairwise disjoint. Our set P consists of a scattered set of n' points placed in each of these hypercubes.

Let \mathcal{G} denote any set of m generators for P . For each hypercube in \mathcal{U} , consider the subset of generators that contains no point outside this hypercube. Let U' denote the hypercube that has the smallest such subset of generators. Let \mathcal{G}' denote this subset of generators. Clearly $|\mathcal{G}'| \leq m'$, where $m' = mn'/n$. Henceforth, we restrict attention to the subset of n' points $P' = P \cap U'$ in this replicant. Without loss of generality, we may take U' to be the unit hypercube \mathbb{U}^d .

The remainder of the proof consists of placing a lower bound on the number of generators needed to cover some quad of \mathcal{Q} that is contained within $U' (= \mathbb{U}^d)$, as a function of n' , m' , r , and ε . This lower bound will be presented in Section 3.3.2. To complete the proof, this bound will then be cast in terms of our original parameters n and m , and the value of r will be selected to produce the best lower bound, which is presented in Section 3.3.3.

Let \mathcal{L} denote the set of all d -corners L such that the corresponding ε -expanded quad, \square_L^+ , lies entirely within \mathbb{U}^d , and let \mathcal{Q}' denote the corresponding set of (unexpanded) quads, \square_L . Note that

two different d -corners may give rise to the same quad. (This might happen because they generate different vertices of the same hypercube, or because they generate the same vertex but through a different ordered sequence.) For the purposes of the proof, we consider these to be distinct elements of \mathcal{Q}' . Thus, we can associate a unique d -corner with each quad in \mathcal{Q}' , so that the hyperplanes of this d -corner all pass through a vertex of the quad.

For any $L \in \mathcal{L}$, let $A_L \subseteq \mathcal{G}$ denote the smallest set of generators that provides a valid answer to the query, that is,

$$P \cap \square_L \subseteq \bigcup_{G \in A_L} G \subseteq P \cap \square_L^+. \quad (6)$$

Clearly $t \geq |A_L|$. Since the hypercubes of \mathcal{U} have disjoint interiors and $\square_L^+ \subseteq \mathbb{U}^d$, it follows that the above inequality holds if P is replaced with P' and A_L is restricted to a subset of \mathcal{G}' .

3.3.2 Analysis of a Single Replicant

We have limited consideration to a single replicant, that is, we consider the subset of generators \mathcal{G}' of size at most m' that lie entirely within the unit hypercube \mathbb{U}^d , the subset of points P' of size n' that lie within \mathbb{U}^d , and the subset of quad ranges \mathcal{Q}' whose ε -expansions lie within \mathbb{U}^d .

Recall that we have a d -corner $L \in \mathcal{L}$ and its associated quad \square_L . We next consider how to define the region of interest. Let $L = \langle J_1, J_2, \dots, J_d \rangle$. Let b be a positive constant, whose value will be set later, and let $\alpha = b\varepsilon$. Recall the integer parameter r introduced earlier (whose value will be fixed later). For $1 \leq i \leq r$, let $J_i(\alpha)$ be the hyperplane that results by translating J_i by distance α towards the origin. Let S_i denote the slab bounded by the two parallel hyperplanes J_i and $J_i(\alpha)$. The *region of interest* for L , denoted R_L , is defined to be intersection of these slabs with \square_L , that is,

$$R_L = \square_L \cap \bigcap_{1 \leq i \leq r} S_i.$$

(See Fig. 8(a).) Recall that by definition, \square_L lies within the inner halfspaces of the hyperplanes J_i , which implies that R_L is nonempty.

Later in the proof, we will also make use of the following outer region, which will be convenient to define now. Define $J_i^{\geq}(\alpha)$ to be the halfspace that results by translating J_i^{\geq} by distance of α towards the origin. Let

$$R_L^{\geq} = \bigcap_{1 \leq i \leq r} J_i^{\geq}(\alpha). \quad (7)$$

(See Fig. 8(b).) Observe that since $S_i \subset J_i^{\geq}(\alpha)$ we have $R_L \subset R_L^{\geq}$.

As observed in [11], the complexity of (exact) halfspace range searching stems from the difficulty of covering points inside the range that lie close to its boundary. The same is true for approximate range searching. In order to make this precise, we introduce a quantity that corresponds roughly to the number of points lying within the region of interest for an average query. Let

$$\Phi = \int_{\mathcal{L}} |P' \cap R_L| dL, \quad (8)$$

where $dL = dL_d$ is the probability density defined earlier for d -corners. We will compute lower and upper bounds on Φ , which together will provide the desired lower bound on the worst-case query time t . Intuitively, if a generator covers a large number of points in R_L , then it cannot be useful in this manner for many queries.




Fig. 8: The regions of interest: (a) R_L and (b) R_L^\geq .

First we want to show that, through an appropriate choice of the constant b , we can apply the properties of scattered sets to lower bound the number of points in R_L and so provide a lower bound on Φ . For all sufficiently small ε , one can easily verify that for any $L \in \mathcal{L}$, the volume of the region of interest R_L satisfies

$$\mu(R_L) = \left(\frac{1}{2}\right)^{d-r} \alpha^r = \left(\frac{1}{2}\right)^{d-r} (b\varepsilon)^r \geq \frac{b^r \varepsilon^r}{2^d}.$$

Now, by setting $b = (8 \cdot 2^d a)^{1/r}$, where a is the constant used in the scattered set properties, we have

$$\mu(R_L) \geq 8a\varepsilon^r.$$

It is easy to verify that this exceeds $(a \log n')/n'$. Thus, by Property 1 of scattered point sets, it follows that $|P' \cap R_L| \geq (n'/a)(8a\varepsilon^r) = 8t \log(t/\varepsilon^r)$. Clearly the probability of generating the d -corners of \mathcal{L} is at least some constant, and so we have the following lower bound on Φ .

$$\Phi = \Omega\left(t \log \frac{t}{\varepsilon^r}\right). \quad (9)$$

Next, we compute an upper bound on Φ . To do this we will focus only on those generators that are most efficient in covering the region of interest R_L . Intuitively, in order to achieve a query time of t , an average generator that is useful in covering the points of R_L should cover or “consume” a fraction of roughly $\Omega(1/t)$ of these points. Based on this, we say that a generator $G \in \mathcal{G}'$ is *absolutely fat* with respect to a d -corner $L \in \mathcal{L}$ if $|G \cap R_L| > 4 \log(t/\varepsilon^r)$ and $G \subseteq \square_L^+$. Indeed, as shown in the next lemma, a constant fraction of points of R_L are covered by such generators.

Lemma 3.16 *For any d -corner $L \in \mathcal{L}$, a constant fraction of the points of P' lying within R_L are covered by generators in \mathcal{G}' that are absolutely fat with respect to L .*

Proof: Given $L \in \mathcal{L}$, recall the smallest set A_L of generators covering the query region \square_L defined just prior to the start of this section. Recall that $|A_L| \leq t$. Let \hat{A}_L denote the subset of A_L consisting of those generators that are absolutely fat with respect to L . Observe that the generators in the set $A_L \setminus \hat{A}_L$ cover at most $4t \log(t/\varepsilon^r)$ points of P' in R_L (because each such generator covers at most $4 \log(t/\varepsilon^r)$ points of P' in R_L and the number of these generators is bounded by $|A_L|$ which is at most t). As shown earlier, $|P' \cap R_L|$ is at least $8t \log(t/\varepsilon^r)$, which is at least twice the number

covered by the non-fat generators in A_L . Therefore, at least half of the points of $P' \cap R_L$ are covered by absolutely fat generators. \square

Recall the outer region R_L^\geq defined in Eq. (7). Since $R_L \subset R_L^\geq$, the above lemma implies that a constant fraction of the points of P' in R_L are covered by generators $G \in \mathcal{G}'$ that satisfy $|G \cap R_L^\geq| > 4 \log(t/\varepsilon^r)$ and $G \subseteq \square_L^+$ (where the last condition is needed to satisfy the approximation bound). Let us therefore define

$$\Delta_G = \left\{ L \in \mathcal{L} : |G \cap R_L^\geq| > 4 \log \frac{t}{\varepsilon^r} \text{ and } G \subseteq \square_L^+ \right\}.$$

If we take the definition of Φ in Eq. (8) but restrict consideration to only these generators and also change the order of integration and summation, then up to constant factors we have

$$\Phi \leq \sum_{G \in \mathcal{G}'} \int_{\Delta_G} |G \cap R_L| dL. \quad (10)$$

We will refer to the quantity $\int_{\Delta_G} |G \cap R_L| dL$ as the *usefulness* of generator G , denoted $u(G)$. Our main task now is to compute an upper bound on the usefulness of any generator $G \in \mathcal{G}'$, which will lead to a lower bound on t .

In order to apply our slicing lemma, it will be convenient to work with R_L^\geq , rather than R_L , since the former is expressed as the intersection of r halfspaces. Since $R_L \subset R_L^\geq$ we have

$$u(G) \leq \int_{\Delta_G} |G \cap R_L^\geq| dL. \quad (11)$$

Because each generator G contributing to Φ contributes a sufficiently large number of points within R_L^\geq , we may apply Property 2 of scattered points to bound the volume of the convex set $\text{conv}(G) \cap R_L^\geq$. Towards this end, consider any $G \in \mathcal{G}'$ and $L \in \Delta_G$. Clearly $G \subseteq P' \cap \text{conv}(G)$, and therefore

$$|G \cap R_L^\geq| \leq |P' \cap (\text{conv}(G) \cap R_L^\geq)|. \quad (12)$$

Since $|G \cap R_L^\geq| > 4 \log(t/\varepsilon^r) \geq \log n'$, we have $|P' \cap (\text{conv}(G) \cap R_L^\geq)| > \log n'$ (recall that we are dealing with a single replicant). Thus we may apply Property 2 (where the convex body is $\text{conv}(G) \cap R_L^\geq$) to obtain

$$|P' \cap (\text{conv}(G) \cap R_L^\geq)| \leq a n' \mu(\text{conv}(G) \cap R_L^\geq). \quad (13)$$

Thus, by combining Eqs. (11), (12), and (13) we obtain

$$u(G) \leq a n' \int_{\Delta_G} \mu(\text{conv}(G) \cap R_L^\geq) dL. \quad (14)$$

Our next task is to apply our version of the isoperimetric inequality (Lemma 3.3) to bound the above integral. In order to do this we will first apply our version of the slicing lemma (Lemma 3.5) to $\text{conv}(G)$ to produce suitably sized bodies to which the isoperimetric inequality can be applied. Observe that for all G in the integration domain, by Eq. (13) we have $\mu(\text{conv}(G) \cap R_L^\geq) \geq (1/an') |P' \cap (\text{conv}(G) \cap R_L^\geq)|$. Recalling that $|P' \cap (\text{conv}(G) \cap R_L^\geq)| > 4 \log(t/\varepsilon^r)$, and substituting the definition of n' for a single replicant, it follows that $\mu(\text{conv}(G) \cap R_L^\geq)$ is at least $4\varepsilon^r/(at)$. Recall that the corner L is defined by the hyperplanes $\langle J_1, J_2, \dots, J_d \rangle$. Now, by setting $\rho = 4\varepsilon^r/(at)$, we may apply

Lemma 3.5 to $\text{conv}(G)$, where the halfspaces are the defining halfspaces of R_L^\geq . For some constant λ , we obtain a collection of

$$O\left(\left(\frac{1}{\beta}\right)^{1-\frac{2}{d+1}} \log^{r-1} \frac{2}{\beta}\right)$$

convex bodies $K_1, K_2, \dots \subseteq \text{conv}(G)$ such that for some K_j :

- (i) $K_j \supseteq \text{conv}(G) \cap R_L^\geq$, and
- (ii) for $1 \leq i \leq r$, $K_j \subseteq C_i^\lambda$, where C_i^λ is the λ -expansion of the cap $\text{conv}(G) \cap J_i^\geq(\alpha)$. (Recall that $J_i(\alpha)$ is the i th defining hyperplane of R_L^\geq .)

Here $\beta = \rho/\mu(\text{conv}(G))$. Since $\text{conv}(G) \subseteq \mathbb{U}^d$, we obtain $\mu(\text{conv}(G)) \leq 1$. Thus $\beta \geq \rho$ and so $\beta \geq 4\varepsilon^r/(at)$.

Given a d -corner $L = \langle J_1, \dots, J_d \rangle$ and $\delta > 0$, recall that $S_r^\delta(L)$ denotes the intersection of slabs $\bigcap_{1 \leq i \leq r} S^\delta(J_i)$, where $S^\delta(J_i)$ denotes the slab of width 2δ centered about J_i . For some constant c , we claim that condition (ii) implies that K_j is contained within the r -corner slab $S_r^{c\varepsilon}(L)$. To see this, let J'_i be the hyperplane obtained by translating J_i by amount $\varepsilon \cdot \text{diam}(\square_L)$ away from the origin. (Note that this passes through a facet of the expanded range \square_L^+ .) Since $G \subseteq \square_L^+$, it follows that the cap $C_i = \text{conv}(G) \cap J_i^\geq(\alpha)$ lies between the hyperplanes $J_i(\alpha)$ and J'_i . Therefore, since $\alpha = \Theta(\varepsilon)$, by condition (ii), $K_j \subseteq S^{c\varepsilon}(J_i)$, for a suitable constant c . Since this holds for all i , where $1 \leq i \leq r$, the desired assertion follows. Thus,

$$\int_{\Delta_G} \mu(\text{conv}(G) \cap R_L^\geq) dL \leq \sum_j \mu(K_j) \int_{\Delta'_{K_j}} dL, \quad \text{where } \Delta'_{K_j} = \{L \in \mathcal{L} : S_r^{c\varepsilon}(L) \supseteq K_j\},$$

and where the sum is taken over the number of bodies described in Lemma 3.5.

We are finally at a point where we can apply our isoperimetric inequality. By Lemma 3.3 we obtain to within constant factors

$$\int_{\Delta_G} \mu(\text{conv}(G) \cap R_L^\geq) dL \leq \sum_j \mu(K_j) \left(\frac{\varepsilon^{d+1}}{\mu(K_j)}\right)^r = \sum_j \frac{\varepsilon^{(d+1)r}}{\mu(K_j)^{r-1}}.$$

Recall that $\mu(\text{conv}(G) \cap R_L^\geq) \geq 4\varepsilon^r/(at)$ and, by condition (i), $K_j \supseteq \text{conv}(G) \cap R_L^\geq$. It follows that $\mu(K_j) \geq 4\varepsilon^r/(at)$. Therefore, up to constant factors we have

$$\int_{\Delta_G} \mu(\text{conv}(G) \cap R_L^\geq) dL \leq \sum_j \frac{\varepsilon^{(d+1)r}}{(\varepsilon^r/t)^{r-1}} = \sum_j \varepsilon^{r(d-r+2)} t^{r-1}.$$

From our bound on the number of convex bodies K_j , we have up to constant factors

$$\begin{aligned} \int_{\Delta_G} \mu(\text{conv}(G) \cap R_L^\geq) dL &\leq \left(\frac{1}{\beta}\right)^{1-\frac{2}{d+1}} \left(\log^{r-1} \frac{1}{\beta}\right) \varepsilon^{r(d-r+2)} t^{r-1} \\ &\leq \left(\frac{t}{\varepsilon^r}\right)^{1-\frac{2}{d+1}} \left(\log^{r-1} \frac{t}{\varepsilon^r}\right) \varepsilon^{r(d-r+2)} t^{r-1} \\ &= t^{r-\frac{2}{d+1}} \cdot \varepsilon^{r(d-r+1+\frac{2}{d+1})} \log^{r-1} \frac{t}{\varepsilon^r}. \end{aligned}$$

Substituting this into Eq. (14) we achieve the desired upper bound on the usefulness of the generator G . Up to constant factors we have

$$u(G) \leq n' \cdot t^{r-\frac{2}{d+1}} \cdot \varepsilon^{r(d-r+1+\frac{2}{d+1})} \log^{r-1} \frac{t}{\varepsilon^r}.$$

Recalling that $|\mathcal{G}'| \leq m'$, we have our desired upper bound on Φ .

$$\Phi = O\left(n' m' t^{r-\frac{2}{d+1}} \varepsilon^{r(d-r+1+\frac{2}{d+1})} \log^{r-1} \frac{t}{\varepsilon^r}\right).$$

Now, by combining this with our lower bound on Φ given in Eq. (9) we obtain the following

$$n' m' t^{r-\frac{2}{d+1}} \varepsilon^{r(d-r+1+\frac{2}{d+1})} \log^{r-1} \frac{t}{\varepsilon^r} = \Omega\left(t \log \frac{t}{\varepsilon^r}\right). \quad (15)$$

3.3.3 Putting it All Together

We are now able to incorporate the above single-replicant results into the overall analysis. We will show that the lower bound on the query time given in Eq. (16) holds. We may assume that t is bounded from above by a polynomial in $1/\varepsilon$, since otherwise this lower bound holds trivially. It follows that $\log t = O(\log(1/\varepsilon))$. Recall that $m' = mn'/n$ and $n' = (t/\varepsilon^r) \log(t/\varepsilon^r)$. Substituting these values of m' and n' into Eq. (15) and then simplifying yields the following lower bound on the query time (up to constant factors)

$$t \geq \left(\frac{n}{m \varepsilon^{r(d-r-1+\frac{2}{d+1})} \log^r \frac{1}{\varepsilon}}\right)^{\frac{1}{r+1-\frac{2}{d+1}}} \geq \left(\frac{n}{m}\right)^{\frac{1}{r+1}} \left(\frac{1}{\varepsilon}\right)^{\frac{r(d-r-1)}{r+1}}, \quad (16)$$

which holds for all integer values $1 \leq r \leq d$.

Let us consider how to select r to achieve the strongest lower bound. To put the lower bound into a somewhat nicer form, let us express the space as $m = n(1/\varepsilon)^{f^2 d^2}$, where $f \geq 0$ is a real parameter. We obtain

$$t \geq \left(\frac{1}{\varepsilon}\right)^{\frac{r(d-r-1)-f^2 d^2}{r+1}} = \left(\frac{1}{\varepsilon}\right)^{\frac{rd-r^2-r-f^2 d^2}{r+1}}.$$

Let $E(r)$ denote the exponent of the $(1/\varepsilon)$ -term. Let $r_0 = \sqrt{f^2 d^2 + d} - 1$. Note that if $0 \leq f \leq 1$ then $0 \leq r_0 \leq d$. By examining the sign of the derivative of $E(r)$, it is easy to verify that $E(r)$ is an increasing function of r for $0 < r < r_0$, and is a decreasing function of r for $r > r_0$. Therefore $E(r)$ has a local maxima at $r = r_0$. Since r_0 may not be an integer and may not lie between 1 and d , we select r_I to obtain our best lower bound, where r_I is the smallest integer satisfying $r_I \geq r_0$ and $1 \leq r_I \leq d$. We have

$$\begin{aligned} E(r_I) &\geq E(r_0 + 1) = \frac{\left(\sqrt{f^2 d^2 + d}\right) \left(d - \sqrt{f^2 d^2 + d} - 1\right) - f^2 d^2}{\sqrt{f^2 d^2 + d} + 1} \\ &= d - 2\sqrt{f^2 d^2 + d} + \frac{\sqrt{f^2 d^2 + d}}{\sqrt{f^2 d^2 + d} + 1} \geq d - 2\sqrt{f^2 d^2 + d} \geq d - 2(fd + \sqrt{d}) \\ &= d(1 - 2f) - 2\sqrt{d}. \end{aligned}$$

This completes the proof of Theorem 3.1(i), the lower bound in the case of idempotent semigroups. In the next section we consider integral semigroups.

3.4 Lower Bound for Rotated Hypercubes: The Integral Case

In this section we prove part (ii) of Theorem 3.1. As mentioned earlier, throughout this section we assume that the set of generators \mathcal{G} satisfies the convex generator assumption, which states that for all $G \in \mathcal{G}$ we have $G = P \cap \text{conv}(G)$.

We recall that in the integral case, the generators used to answer a query are required to be disjoint. Intuitively, it seems that this requirement should reduce the usefulness of large generators. This idea was first formalized in [5], where it was used in the context of exact halfspace range searching and approximate spherical range searching. The proof proceeds along the same lines as in the idempotent case. The key idea is a different notion of generator *usefulness* that is more appropriate for the integral case. Toward this end, define a generator G to be *relatively fat* with respect to a d -corner $L \in \mathcal{L}$ if $|G \cap R_L| > \varepsilon^r |G|$ and $G \subseteq \square_L^+$. As in the idempotent case, it is possible to limit consideration to generators that are fat, but now both in the absolute and relative senses. Combining this with the generalized Macbeath region machinery results in the desired bound.

We now give the details, highlighting the differences in the analysis from the idempotent case. The proof is identical to that for the idempotent case up to Lemma 3.16, where it is argued that a constant fraction of the points in the region of interest are covered by generators that are absolutely fat. Recall that t denotes the worst-case query time in the arithmetic model over all the ranges. Also recall that at that point of the proof we have restricted attention to a subset P' of data points that satisfy the following conditions:

- (i) P' is a scattered set of n' points contained within \mathbb{U}^d , where $n' = (t/\varepsilon^r) \log(t/\varepsilon^r)$.
- (ii) Let \mathcal{G}' be the set of generators that contain only points from P' . Then $|\mathcal{G}'| \leq m'$ where $m' = mn'/n$.
- (iii) Let \mathcal{L} denote the set of all d -corners L such that the corresponding ε -expanded quad, \square_L^+ , lies entirely within \mathbb{U}^d , and let \mathcal{Q}' denote the corresponding set of (unexpanded) quads, \square_L . For any $L \in \mathcal{L}$, let $A_L \subseteq \mathcal{G}'$ denote the smallest set of generators that provide a valid answer for the query \square_L , that is,

$$P' \cap \square_L \subseteq \bigcup_{G \in A_L} G \subseteq P' \cap \square_L^+.$$

Then $t \geq |A_L|$.

Exploiting the fact that in the integral case the generators used to answer a query are disjoint allows us to strengthen Lemma 3.16 to include generators that are both absolutely and relatively fat.

Lemma 3.17 *For any d -corner $L \in \mathcal{L}$, a constant fraction of the points of P' lying within R_L are covered by generators in \mathcal{G}' that are both absolutely and relatively fat with respect to L .*

Proof: Given $L \in \mathcal{L}$, let $A_L \subseteq \mathcal{G}'$ denote the smallest set of generators that provide a valid answer for the query \square_L . That is,

$$P' \cap \square_L \subseteq \bigcup_{G \in A_L} G \subseteq P' \cap \square_L^+, \tag{17}$$

and the generators in A_L are all disjoint. Note that $|A_L| \leq t$. Let A_L^1 denote the subset of A_L consisting of those generators that are not absolutely fat with respect to L , and A_L^2 denote the subset of A_L consisting of those generators that are not relatively fat with respect to L . The proof of Lemma 3.16 shows that the generators in A_L^1 cover at most $4t \log(t/\varepsilon^r)$ points of P' in R_L . Since each generator $G \in A_L^2$ covers at most $\varepsilon^r |G|$ points of P' in R_L , the generators in A_L^2 together cover at most $\varepsilon^r \sum_{G \in A_L^2} |G|$ points of P' in R_L . Since the generators in A_L^2 are disjoint, we have $\sum_{G \in A_L^2} |G| \leq |P'|$. Recalling that $|P'| = (t/\varepsilon^r) \log(t/\varepsilon^r)$, it follows that the generators in A_L^2 cover at most $t \log(t/\varepsilon^r)$ points of P' in R_L . Recall that $P' \cap R_L = 8t \log(t/\varepsilon^r)$. Since the generators in A_L^1 cover at most $4t \log(t/\varepsilon^r)$ points of $P' \cap R_L$ and the generators in A_L^2 cover at most $t \log(t/\varepsilon^r)$ points of $P' \cap R_L$, it follows that a constant fraction of the points of $P' \cap R_L$ are covered by generators in the set $A_L \setminus (A_L^1 \cup A_L^2)$. As these generators are both absolutely and relatively fat, the lemma follows. \square

Arguing as in the idempotent case, but using generators that are both absolutely and relatively fat (i.e., Lemma 3.17 in place of Lemma 3.16), we can easily see that Eq. (10) holds with Δ_G redefined as follows:

$$\Delta_G = \left\{ L \in \mathcal{L} : |G \cap R_L^\geq| > \max \left(4 \log \frac{t}{\varepsilon^r}, \varepsilon^r |G| \right) \text{ and } G \subseteq \square_L^+ \right\}.$$

(Throughout our analysis, Δ_G will refer to this revised definition.) As before, we will refer to the quantity $\int_{\Delta_G} |G \cap R_L| dL$ as the *usefulness* of generator G , denoted $u(G)$. Our strategy again is to compute an upper bound on the usefulness of any generator $G \in \mathcal{G}'$, which will lead to a lower bound on t .

Using the scattered set properties, the fact that $R_L \subset R_L^\geq$, and the definition of Δ_G , we can establish Eqs. (11)–(14) as before. For convenience let us write Eq. (14) here again:

$$u(G) \leq a n' \int_{\Delta_G} \mu(\text{conv}(G) \cap R_L^\geq) dL. \quad (18)$$

Let $G \in \mathcal{G}'$ and $L \in \Delta_G$. Recall that in order to bound the integral in Eq. (18) via the generalized Macbeath region machinery (Lemma 3.5), we need to compute a lower bound on the ratio of $\mu(\text{conv}(G) \cap R_L^\geq)$ to $\mu(\text{conv}(G))$.

By Eqs. (12) and (13), we have $\mu(\text{conv}(G) \cap R_L^\geq) \geq (1/a n') |G \cap R_L^\geq|$. Using the definition of Δ_G , we have $|G \cap R_L^\geq| > \varepsilon^r |G|$. Thus $\mu(\text{conv}(G) \cap R_L^\geq) \geq (\varepsilon^r / a n') |G|$. Next we compute an upper bound on $\mu(\text{conv}(G))$. Obviously $|P' \cap \text{conv}(G)| \geq |G| > |G \cap R_L^\geq|$. By the definition of Δ_G , we have $|G \cap R_L^\geq| > 4 \log(t/\varepsilon^r) \geq \log n'$, and so $|P' \cap \text{conv}(G)| \geq \log n'$. We may, therefore, apply Property 3 of scattered point sets (with $K = \text{conv}(G)$) to obtain $\mu(\text{conv}(G)) \leq (a/n') |P' \cap \text{conv}(G)|$. By the assumption of convex generators, $G = P' \cap \text{conv}(G)$, and so $\mu(\text{conv}(G)) \leq (a/n') |G|$. Therefore

$$\beta = \frac{\mu(\text{conv}(G) \cap R_L^\geq)}{\mu(\text{conv}(G))} \geq \frac{(\varepsilon^r / a n') |G|}{(a/n') |G|} = \frac{\varepsilon^r}{a^2}.$$

Note that our lower bound on β is t times larger than in the idempotent case (ignoring constant factors). Thus, in the integral case, we can apply Lemma 3.5 with a larger value of β , which yields a smaller upper bound on the usefulness $u(G)$ of any generator G . The rest of the calculations are similar to that in the idempotent case and are omitted. We finally obtain (up to constant factors)

$$t \geq \left(\frac{n}{m \varepsilon^{r(d-r-1+\frac{2}{d+1})} \log^r \frac{1}{\varepsilon}} \right)^{\frac{1}{r}} \geq \left(\frac{n}{m} \right)^{\frac{1}{r}} \left(\frac{1}{\varepsilon} \right)^{d-r-1}, \quad (19)$$

which holds for all integer values $1 \leq r \leq d$. Expressing the space as $m = O(n \cdot (1/\varepsilon)^{f^2 d^2})$, where $f \geq 0$ is a real parameter, and optimizing for r , we can show that $t \geq (\frac{1}{\varepsilon})^{d(1-2f)^{-2}}$, which proves Theorem 3.1(ii).

4 Smooth Convex Ranges

In this section we present lower and upper bounds on the complexity of range searching for κ -smooth convex ranges over idempotent semigroups.

4.1 Lower Bound

We begin by presenting the lower bound for κ -smooth ranges. If κ is greater than some fixed constant, then for all sufficiently small ε and sufficiently large n the lower bound is $\Omega(\log n + (1/\varepsilon)^{(d-1)/2})$. This lower bound holds in the algebraic decision-tree model. Thus, it holds irrespective of space, and no space-time tradeoffs are possible. We present the bound in the following theorem, which includes the dependency on κ .

Theorem 4.1 *Let $d > 1$ be a fixed dimension, and consider any real $\kappa > 1$. Consider a range space consisting of all κ -smooth convex bodies and a weight function over any idempotent semigroup. Then for all sufficiently small ε and sufficiently large n , the worst-case query time for ε -approximate range searching among n points in the algebraic decision-tree model is at least*

$$\Omega\left(\log n + \left(\frac{\kappa - 1}{\varepsilon \kappa}\right)^{\frac{d-1}{2}}\right).$$

The proof is based on a construction that first builds a maximal set of $m = \Omega(((\kappa-1)/(\varepsilon\kappa))^{(d-1)/2})$ points on the unit sphere such that each pair is separated by a distance of $\Omega(\sqrt{\varepsilon})$. At each point x we create a ball of radius $1/\kappa$ that lies within the sphere and is tangent at x . For each of the 2^m subsets of points, we create a range consisting of the convex hull of the corresponding tangent balls. We show that no two of these ranges are equivalent to within an ε error, and so any decision tree must have at least 2^m leaves, which implies a depth of at least $\Omega(m)$, and hence query time of $\Omega(m)$.

Proof: The proof is based on a standard algebraic decision-tree argument. We show that any decision tree must distinguish between some number L of distinct possible outcomes, and hence the worst-case query time is $\Omega(\log L)$. Two ranges are ε -distinct if one contains a point that lies outside the other's ε expansion. By the faithfulness of the semigroup, distinct ranges must be handled differently by the decision tree.

First observe that the $\Omega(\log n)$ term is easy to prove. Consider n Euclidean ball ranges, each centered at a point, but so small that its ε expansion contains no other point. This bound holds irrespective of the point distribution.

To prove the other part of the bound, we generate a large collection of ε -distinct ranges. Given $\delta > 0$, we say that a set of points is δ -sparse if the distance between any pair of distinct points is at least δ . Let $\delta = \sqrt{6\varepsilon\kappa/(\kappa-1)}$. Let us assume that ε is small enough that $\delta < 1$. Consider any δ -sparse point set P on a unit sphere in \mathbb{R}^d centered at the origin O . It is well known that such a

set exists having at least n_δ points, where

$$n_\delta = \Omega\left(\frac{1}{\delta}\right)^{d-1} = \Omega\left(\frac{\kappa-1}{6\varepsilon\kappa}\right)^{\frac{d-1}{2}}.$$

Assume that $n \geq n_\delta$. The point set is constructed by taking n_δ points from P , and placing the remaining $n - n_\delta$ points arbitrarily (since they will play no role in the rest of the argument).

One approach to construct the ranges would be to consider the convex hulls of each of the $2^{n_\delta} - 1$ nonempty subsets of P . Although this would work, the resulting ranges would not be smooth. Our approach is to “round off” the sharp edges of these ranges. For each $p \in P$, let $b(p)$ be a ball of radius $1/\kappa$ that lies within the unit sphere and is tangent to the sphere at p . For each subset $P' \subseteq P$, let $\eta(P')$ be the convex hull of the union of $b(p)$ for $p \in P'$, that is,

$$\eta(P') = \text{conv}\left(\bigcup_{p \in P'} b(p)\right).$$

Let the range space \mathcal{Q} consist of the $2^{n_\delta} - 1$ ranges $\eta(P')$ corresponding to all the nonempty subsets P' of P . Note that the diameter of each range in \mathcal{Q} is at most 2, because they all lie within the unit sphere.

We assert that every pair of distinct ranges of \mathcal{Q} is ε -distinct. Consider two ranges arising from distinct subsets P' and P'' . There is some point y that is in one set and not the other. Assume without loss of generality that $y \in P'$ but $y \notin P''$. Clearly $y \in \eta(P')$, and so it suffices to show that y lies outside the ε expansion of $\eta(P'')$. To prove this fact, consider the line segment \overline{Oy} , and the hyperplane H that is orthogonal to \overline{Oy} and at distance 2ε from y . We assert that all the balls $b(x)$ for $x \in P''$ lie on the opposite side of H from y . It will follow that the convex hull of these balls also lies below H , implying that the minimum distance of y to $\eta(P'')$ is at least 2ε . Since $\text{diam}(\eta(P'')) \leq 2$ it follows that y lies outside the ε expansion of $\eta(P'')$, as desired.




Fig. 9: Lower bound for smooth convex ranges.

To prove this assertion, consider any x in P'' . The orthogonal projection of $b(x)$ onto line \overline{Oy} is a line segment. We will show that this segment’s closest endpoint to y is at least at distance 2ε , from which it follows that $b(x)$ lies entirely below H . Consider the intersection of $b(x)$ with

the plane containing x , y , and O . For the sake of illustration, imagine that the ray \overrightarrow{Oy} is directed upwards, let z be the highest point of $b(x)$, and let x' and z' be the orthogonal projections of x and z onto the line \overrightarrow{Oy} . (See Fig. 9.) Let $\theta = \angle xOy$. By basic trigonometry, the distance $\|Ox'\|$ is $\cos \theta$, and because $b(x)$ is of radius $1/\kappa$, we have $\|x'z'\| = (1/\kappa)(1 - \cos \theta)$. Thus, we find that the vertical distance from y to the upper endpoint of $b(x)$'s orthogonal projection is

$$\begin{aligned} \|z'y\| &= 1 - (\|Ox'\| + \|x'z'\|) \\ &= 1 - \cos \theta - \frac{1}{\kappa}(1 - \cos \theta) = \frac{\kappa - 1}{\kappa}(1 - \cos \theta). \end{aligned}$$

To bound this distance, we use the fact that if $\theta < 1$ then $1 - \cos \theta \geq \theta^2/3$. By the definition of θ we obtain

$$\|z'y\| \geq \frac{\kappa - 1}{\kappa} \cdot \frac{\theta^2}{3} \geq \frac{\kappa - 1}{\kappa} \cdot \frac{6\varepsilon\kappa/(\kappa - 1)}{3} = 2\varepsilon,$$

as desired.

Therefore, the decision tree must have at least $2^{n_\delta} - 1$ leaves, corresponding to these ε -distinct ranges. Its depth is at least $\log_2(2^{n_\delta} - 1) = \Omega(n_\delta)$, and this establishes the desired lower bound on the query time. \square

4.2 Upper Bound

In this section we present a data structure that can answer approximate nearest neighbor queries for idempotent semigroups for smooth convex ranges. Assuming any fixed degree of smoothness, the query time is quite similar to the lower bound proved in the previous section. This result demonstrates that the improvements offered by idempotence for spherical range searching also apply to ranges that are sufficiently smooth. Our results apply to the range space of all κ -smooth convex bodies, for any fixed $\kappa \geq 1$. (The hidden constants increase with κ .) As in [7] make the *unit-cost test assumption*, which states that given any ball or hypercube and range η , in constant time we can determine whether the ball or hypercube is contained within η^+ , is disjoint of η^+ , or neither. Here is our main result.

Theorem 4.2 *Consider a range space consisting of all κ -smooth convex bodies for any real constant $\kappa \geq 1$. Let P be a set of n points in \mathbb{R}^d . Let $0 < \varepsilon \leq 1/2$ be a real parameter. Then we can construct a data structure of $O(n/\varepsilon)$ space that allows us to answer ε -approximate range queries over any idempotent semigroup in time $O(\log n + (1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$. The time to construct the data structure is $O((n/\varepsilon^{(d+1)/2}) \log(n/\varepsilon) \log(1/\varepsilon))$.*

The data structure is similar in spirit to that described in [4], [5], and [6] for Euclidean balls. This structure is based on the concept of an *approximate Voronoi diagram*, or *AVD*, as described in [2] and [3]. This is a quadtree-like structure in which space is subdivided recursively until the leaf cells satisfy certain “separation properties” with respect to the surrounding points. Each node of the tree (internal and leaf) is responsible for handling query ranges that overlap the corresponding cell and whose diameter is proportional to the size of the cell. Each node is associated with a small collection of generators, each of which is the semigroup sum of points lying within a judiciously chosen Euclidean ball.

The critical fact that enables us to answer queries efficiently is that any κ -smooth range can be ε -approximated using $O(1/\varepsilon^{(d-1)/2})$ balls of radius $\Omega(\text{diam}(\eta)/\kappa)$, that are suitably placed just

touching the boundary of η^+ . The generator subsets are chosen so that it is possible to compute an appropriate subset of them that approximate each of these $O(1/\varepsilon^{(d-1)/2})$ balls. To avoid an excessive repetition of details, we will refer to results proved in our earlier works on approximate spherical range searching, which can be found in [4], [5], and [6]. The principal new elements arise from the lack of symmetry in arbitrary smooth ranges, but these obstacles are more technical than fundamental in nature.

Let P be a set of n points in \mathbb{R}^d and let $0 < \varepsilon \leq 1/2$ be a real parameter. We assume that the weights assigned to the points of P belong to an idempotent semigroup. Without loss of generality, we may assume that the set of points P has been scaled and translated to lie within a ball of radius $\varepsilon/2$ placed at the center of the unit hypercube \mathbb{U}^d . This allows us to easily dispense the case of a query η that is not contained within \mathbb{U}^d as follows. If η contains the center of \mathbb{U}^d , we output $w(P)$ (i.e., the semigroup sum of all the points in P), otherwise we output the special null symbol, which represents the empty set. Henceforth, we assume that $\eta \subseteq \mathbb{U}^d$, and show how to construct a data structure to handle such a query.

We will use the following terminology. A *quadtrees box* is defined recursively as \mathbb{U}^d or any d -cube that can be obtained by splitting a quadtrees box into 2^d identical subcubes by d axis-orthogonal hyperplanes passing through its center. We define a *cell* to be either a quadtrees box or the set-theoretic difference of two quadtrees boxes, an *outer box* and an *inner box*. The *size* of a cell is defined to be the side length of its outer box. Throughout, for a cell u , we define the following three items. We define s_u to be its size, we define r_u to be $s_u d/2$, and we define b_u to be the ball of radius r_u whose center coincides with the center of u 's outer box. (Note that $u \subseteq b_u$.) Finally, for any 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 . Because of the close association between each node and its cell, when it is clear from context, we will often use the same name to refer to both of them.

In Lemma 4.1 below we abstract the main features of the data structure used in [4–6] for spherical range queries. This data structure is based on a hierarchical subdivision of the unit hypercube into cells. We classify each cell u of this subdivision into three types. Let $\gamma \geq 16$ and $0 < f \leq 1$ be two real parameters.

Type-1: u enjoys no separation property in general.

Type-2: There exists a ball b'_u such that $|P \cap (\gamma b_u \setminus b'_u)| = O(1/f)$, and the ball $\gamma b'_u$ does not overlap u . (See Fig. 10).

Type-3: u is a quadtrees box (not the difference of two boxes), and there is an associated quadtrees box v such that $u \subseteq v$ and $|P \cap (\gamma b_v \setminus 8b_u)| = O(1/f)$. (See Fig. 10).

Note that these definitions are not mutually exclusive. The construction algorithm assigns a unique type to each cell.

To provide a somewhat better understanding of the usefulness of this classification, let us imagine for now that the $O(1/f)$ term is simply zero. In our application, we will set γ to a sufficiently large constant depending on the smoothness parameter κ . We first find a cell u such that the query range is roughly centered about u and satisfies certain properties, depending on u 's type. If u is of type 1, the range's diameter will be proportional to the diameter of u . Because of smoothness, such a range can be well approximated by a small collection of covering balls. If u is of type 2, the range's diameter can be arbitrarily small relative to u 's diameter (but is at most a constant factor larger). In this case, the problematic points are densely clustered within b'_u . Either b'_u lies entirely outside the range (in which case they do not contribute to the result) or else the




Fig. 10: Generic separation properties for cells. (Not drawn to scale.) The $O(1/f)$ points that do not satisfy the separation properties are shown as white points.

fact that $\gamma b'_u$ does not overlap u implies that they are at a certain distance from u and can be well approximated by a small collection of covering balls. Finally, if u is of type 3, the range will lie between a certain constant expansion of b_u and γb_v , where v is the quadtree box associated with u . In this case, as all the points in γb_v are very close to u , they can be treated as a lump and all inferred to lie within the range.

What about the $O(1/f)$ points that fail to satisfy the separation properties? These points arise as the leftovers of a sampling process (implicit in Lemma 4.1) that we employ to minimize the space requirements. They are called *pollutants*. The value f will be selected so that $1/f$ is proportional to the query time. This implies that we can simply inspect all of the pollutants one-by-one without increasing the overall asymptotic query time.

We now present the lemma, proved in [6], that establishes the existence of a data structure satisfying the desired separation properties for a given range. Although this lemma is provided without explicit reference to the type of range, the point q can be thought of as the center of the range and D as its diameter.

Lemma 4.1 (Generic Separation Properties) *Let P be a set of n points in \mathbb{R}^d . Let $\gamma \geq 16$ and $0 < f \leq 1$ be two real parameters. In $O(n\gamma^d \log(n\gamma) \log \gamma)$ time, it is possible to construct a data structure with $O(nf\gamma^d)$ cells of type-1, type-2, and type-3, respectively, such that the following holds. For any pair (q, D) , where q is a point in \mathbb{U}^d and $0 \leq D \leq \sqrt{d}$ is a real number, in $O(\log(n\gamma))$ time, we can find a cell u such that $q \in u$ and u satisfies one of the following properties.*

- (i) u is of type-1 and $\gamma r_u/4 < D \leq \gamma r_u/2$.
- (ii) u is of type-2 and $D \leq \gamma r_u/4$.
- (iii) u is of type-3 and $\gamma r_u/4 < D \leq \gamma r_v/4$, where v denotes the quadtree box associated with u .
(See properties of type-3 cells given above.)

In case (ii) the node u stores the ball b'_u and the set of points $P \cap (\gamma b_u \setminus b'_u)$. In case (iii) the node u stores the weight of points in the ball $8b_u$ (that is, $w(P \cap 8b_u)$) and the set of points $P \cap (\gamma b_v \setminus 8b_u)$.

We will now show how to apply the above lemma in our particular context of smooth convex ranges. Recall that the query range η being considered lies entirely within \mathbb{U}^d . Let $\text{diam}(\eta)$ denote its diameter. In order to apply the previous lemma, which is stated in terms of pair (q, D) , we need to introduce the notion of a center point q and diameter D for a convex range query. In particular

we shall assume that the query algorithm is given a pair (q_η, D_η) , where $q_\eta \in \eta$ is a point that is at distance at least $\text{diam}(\eta)/(2\kappa)$ from the boundary of η (by smoothness such a point exists), and D_η is the diameter of the range.

Throughout this section, let γ_κ denote the constant 72κ . (Any sufficiently large multiple of κ would work, and no attempt was made to optimize its value.) In the definition of type-1, type-2, and type-3 cells given above, let $\gamma = \gamma_\kappa = 72\kappa$, and let $f = \varepsilon^{(d-1)/2}$. That is, we replace the previous definitions of cell types with the following.

Type-1: u enjoys no separation property in general.

Type-2: There exists a ball b'_u such that $|P \cap (\gamma_\kappa b_u \setminus b'_u)| = O(1/\varepsilon^{(d-1)/2})$, and the ball $\gamma_\kappa b'_u$ does not overlap u .

Type-3: u is a quadtree box (not the difference of two boxes). There is an associated quadtree box v such that $u \subseteq v$ and $|P \cap (\gamma_\kappa b_v \setminus 8b_u)| = O(1/\varepsilon^{(d-1)/2})$.

We now tailor the generic separation properties of Lemma 4.1 to a form that can be conveniently applied to the space of κ -smooth convex ranges. The cases are illustrated in Fig. 11.

Lemma 4.2 (Separation Properties) *Let $\kappa \geq 1$ be a constant. Let \mathcal{Q} be a set of κ -smooth convex ranges. Let P be a set of n points in \mathbb{R}^d . In $O(n \log n)$ time, it is possible to construct a data structure with $O(n\varepsilon^{(d-1)/2})$ cells of type-1, type-2, and type-3, respectively, such that the following holds. For any range $\eta \in \mathcal{Q}$ that is contained within \mathbb{U}^d , in $O(\log n)$ time, we can find a cell u such that $q_\eta \in u$ and u satisfies one of the following properties.*

- (i) u is of type-1 and $\eta \subseteq \gamma_\kappa b_u$ and the diameter of η is at least $\gamma_\kappa r_u/4$.
- (ii) u is of type-2 and $\eta \subseteq \gamma_\kappa b_u$.
- (iii) u is of type-3, $8b_u \subseteq \eta \subseteq \gamma_\kappa b_v$, where v denotes the quadtree box associated with u . (See properties of type-3 cells given above.)

In case (ii) the node u stores the ball b'_u and the set of points $P \cap (\gamma_\kappa b_u \setminus b'_u)$. In case (iii) the node u stores the weight of points in the ball $8b_u$ (that is, $w(P \cap 8b_u)$) and the set of points $P \cap (\gamma_\kappa b_v \setminus 8b_u)$.




Fig. 11: Separation properties for cells for smooth ranges. (Not drawn to scale.) The $O(1/f)$ points that do not satisfy the separation properties are shown as white points.

Proof: We construct the data structure of Lemma 4.1 for $\gamma = \gamma_\kappa = 72\kappa$, and $f = \varepsilon^{(d-1)/2}$. Let $D = \text{diam}(\eta)$. By smoothness we have $b(q_\eta, D/(2\kappa)) \subseteq \eta \subseteq b(q_\eta, D)$. Also, since $\eta \subseteq \mathbb{U}^d$, it follows that $0 \leq D \leq \sqrt{d}$. By Lemma 4.1, for the query (q_η, D) , in $O(\log n)$ time, we can find a cell u such that $q_\eta \in u$ and u satisfies one of the three properties listed therein.

If u satisfies Lemma 4.1(i), then u is of type-1 and $\gamma_\kappa r_u/4 < D \leq \gamma_\kappa r_u/2$. By the triangle inequality, it follows that the distance of any point in η from the center of u is at most $D + r_u \leq \gamma_\kappa r_u/2 + r_u \leq \gamma_\kappa r_u$. Thus $\eta \subseteq \gamma_\kappa b_u$. Noting that $D > \gamma_\kappa r_u/4$, it follows that property (i) holds.

If u satisfies Lemma 4.1(ii), then u is of type-2 and $D \leq \gamma_\kappa r_u/4$. Arguing as for property (i), it follows that the distance of any point in η from the center of u is at most $D + r_u \leq \gamma_\kappa r_u/4 + r_u \leq \gamma_\kappa r_u$. Thus $\eta \subseteq \gamma_\kappa b_u$, and so property (ii) holds.

Finally, if u satisfies Lemma 4.1(iii), then u is of type-3 and $\gamma_\kappa r_u/4 < D \leq \gamma_\kappa r_v/4$, where v is the quadtree box associated with u . By the triangle inequality, it follows that the distance of any point on the boundary of η from the center of u is at least $D/(2\kappa) - r_u$. Since $D > \gamma_\kappa r_u/4$ and $\gamma_\kappa = 72\kappa$, this quantity is at least $8r_u$. Thus $8b_u \subseteq \eta$. The same calculation as for property (ii) implies that $\eta \subseteq \gamma_\kappa b_v$. This establishes (iii) and completes the proof. \square

During preprocessing, for each cell in the data structure of Lemma 4.2, we compute the weight of certain clusters and store them with the cell. To answer a query η , we first apply the above data structure to find the cell u that satisfies one of the three properties listed in Lemma 4.2. The cluster sums for the case of type-1 and type-2 cells are described in Lemma 4.5 below. Before presenting this lemma, we establish some useful geometric facts. The first is a useful trigonometric inequality due to Chan and Snoeyink [12]. The second is a technical result, which will be applied below in Lemma 4.5 for processing queries.

Lemma 4.3 (Chan and Snoeyink) *Let $\triangle xyz$ be a triangle with $\angle xzy = \theta$, $\angle yxz = \phi$, and $\angle xyz \geq \pi/2$. Then*

$$\|xy\| + \|yz\| \leq (1 + \sin \theta \sin \phi) \|xz\|.$$

Lemma 4.4 *Let $0 < \varepsilon \leq 1/2$ be a real parameter. Let b be a ball of radius r centered at o , x be a point on the boundary of b , and p be a point on \overline{ox} such that $\|px\| \geq r\varepsilon$. Let y be a point such that $\|oy\| \leq r/4$ and $\angle yox \leq \sqrt{\varepsilon}$. Let \hat{b} be a ball of radius at least $(r - \|oy\|)/(1 + \varepsilon/2)$ centered at y . Then $p \in \hat{b}$.*

Proof: Observe that $\hat{b} \subseteq b$. Let \hat{r} denote the radius of the ball \hat{b} . We are given that $\hat{r} \geq (r - \|oy\|)/(1 + \varepsilon/2)$. Since $\|oy\| \leq r/4$ and $\varepsilon \leq 1/2$, it follows that $\hat{r} \geq 3r/5$. Thus $o \in \hat{b}$. Let t denote the point of intersection of the boundary of \hat{b} with \overline{ox} (it is clear that such a point exists). By convexity, all the points on segment \overline{ot} are contained in \hat{b} . We will show that $\|ot\| \geq r(1 - \varepsilon)$. This will imply that the point p described in the statement of the lemma belongs to \hat{b} , which will complete the proof.

Let z denote the point on \overline{ox} such that $\|oy\| = \|oz\|$. Since $\|oy\| \leq r/4$ and $\angle yoz \leq \sqrt{\varepsilon}$, it follows that $\|yz\| \leq r\sqrt{\varepsilon}/4 \leq r/4$. Since $\|yt\| = \hat{r} \geq 3r/5$, we have $\|yt\| > \|yz\|$. By elementary trigonometry, it follows that $\angle ytz < \pi/2$. Applying the law of sines to the triangle $\triangle yzt$, we have

$$\frac{\sin \angle ytz}{\|yz\|} = \frac{\sin \angle yzt}{\|yt\|} \leq \frac{1}{\|yt\|}.$$

Since $\|yt\| = \hat{r} \geq 3r/5$, we obtain

$$\sin \angle ytz \leq \frac{\|yz\|}{\|yt\|} \leq \frac{r\sqrt{\varepsilon}/4}{3r/5} = \frac{5\sqrt{\varepsilon}}{12}.$$




Fig. 12: Proof of Lemma 4.4.

Consider the triangle $\triangle oyt$. Let θ denote $\angle yot$ and ϕ denote $\angle ytz$. Since $\theta \leq \sqrt{\varepsilon}$ and $\varepsilon \leq 1/2$, it follows that $\theta < \pi/4$. Also, we showed above that $\phi < \pi/2$ and $\sin \phi \leq 5\sqrt{\varepsilon}/12$. It follows that $\phi < \pi/4$. Thus $\angle oyt > \pi/2$. We can therefore apply Lemma 4.3 to $\triangle oyt$ to obtain

$$\begin{aligned} \|ot\| &\geq \frac{\|oy\| + \|yt\|}{1 + \sin \theta \sin \phi} \geq \frac{\|oy\| + (r - \|oy\|)/(1 + \varepsilon/2)}{1 + (\sqrt{\varepsilon})(5\sqrt{\varepsilon}/12)} = \frac{r + \varepsilon\|oy\|/2}{(1 + 5\varepsilon/12)(1 + \varepsilon/2)} \\ &\geq \frac{r}{(1 + 5\varepsilon/12)(1 + \varepsilon/2)}. \end{aligned}$$

Since $\varepsilon \leq 1/2$, we can easily verify that $\|ot\| \geq r(1 - \varepsilon)$, as desired. \square

We can now present the lemma that is used for processing cells of types 1 and 2. For these cases, the subset of data points of interest all lie inside a ball of radius r and the query ranges are known to be large relative to this ball. The lemma exploits the fact that such a range can be approximated as the union of a small number of balls, denoted \mathcal{B} . In particular, by precomputing the weight of points inside each of the balls of \mathcal{B} , we can answer such a query η by summing up the weights associated with the balls in \mathcal{B}_η .

Lemma 4.5 *Let $\beta > 0$ and $\kappa \geq 1$ be constants. Let \mathcal{Q} be a set of κ -smooth convex ranges. Let b be any ball of radius r . It is possible to find a set \mathcal{B} of $O(1/\varepsilon^{(d+1)/2})$ balls and store them in $O(1/\varepsilon^{(d+1)/2})$ space, such that for any range $\eta \in \mathcal{Q}$ whose diameter $\text{diam}(\eta)$ is at least βr , the following property holds: In $O((1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$ time, it is possible to find a subset $\mathcal{B}_\eta \subseteq \mathcal{B}$ of size $O(1/\varepsilon^{(d-1)/2})$ such that their union, $\bigcup_{\hat{b} \in \mathcal{B}_\eta} \hat{b}$, covers $\eta \cap b$ and is contained within η^+ . (Constants hidden in the O -notation depend only on β , κ , and d .)*

Proof: Let $r' = \beta r / (2\kappa)$. For each of the d coordinate axes, consider an infinite set of hyperplanes orthogonal to it such that the distance between successive hyperplanes is $r' / (16\sqrt{d})$. On each of these hyperplanes overlay a regular $(d-1)$ -dimensional grid of side length $r'\sqrt{\varepsilon} / (16\sqrt{d-1})$. Let \mathcal{G} denote the set of grid points on all these hyperplanes that are contained inside the ball $\tilde{b} = (1 + \beta/(2\kappa))b$. Note that $O(1)$ hyperplanes intersect \tilde{b} and each of these hyperplanes contains $O(1/\varepsilon^{(d-1)/2})$ grid points inside \tilde{b} . Thus $|\mathcal{G}| = O(1/\varepsilon^{(d-1)/2})$. For each point $y \in \mathcal{G}$, consider a set of concentric balls \mathcal{B}_y centered at y whose radii span the range from $3r'/5$ to r' , such that the radius of any two successive balls differs by a multiplicative factor of $1 + \varepsilon/2$. We store the balls in \mathcal{B}_y in

an array in order of increasing radius. Finally we define $\mathcal{B} = \cup_{y \in \mathcal{G}} \mathcal{B}_y$. Since $|\mathcal{G}| = O(1/\varepsilon^{(d-1)/2})$ and $|\mathcal{B}_y| = O(1/\varepsilon)$ for each $y \in \mathcal{G}$, it follows that $|\mathcal{B}| = O(1/\varepsilon^{(d+1)/2})$. Next we show that \mathcal{B} satisfies the property described in the lemma.

Let η be any range in \mathcal{Q} with diameter $D_\eta \geq \beta r$. For each $y \in \mathcal{G}$, first we check if the smallest ball in \mathcal{B}_y is contained within η^+ . If it is not, we ignore all the balls in \mathcal{B}_y . Otherwise, we use binary search to find the largest ball in \mathcal{B}_y that is contained within η^+ . Let B_η denote the set of balls found using binary search over all groups \mathcal{B}_y of concentric balls. Since it takes $O(\log(1/\varepsilon))$ time for each binary search, the total time to find B_η is $O((1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$. (By using the floor function, we can shave a $\log(1/\varepsilon)$ factor from this time, if we assume that we can compute the distance from any point inside η^+ to the closest point on $\partial\eta^+$ in constant time. The straightforward details of this relatively minor improvement are omitted.)

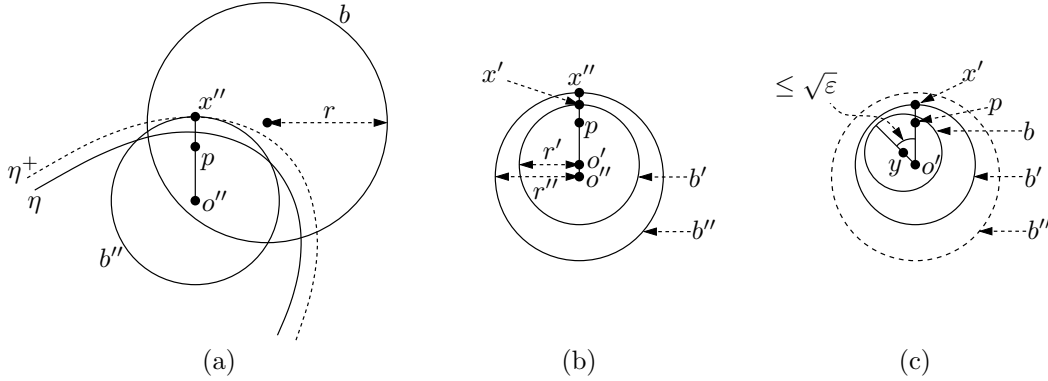


Fig. 13: Proof of Lemma 4.5.

It remains to show that $\eta \cap b \subseteq \cup_{\hat{b} \in \mathcal{B}_\eta} \hat{b} \subseteq \eta^+$. By construction, each ball in \mathcal{B}_η is contained inside η^+ , and so $\cup_{\hat{b} \in \mathcal{B}_\eta} \hat{b} \subseteq \eta^+$. To prove the first part, let p be a point in $\eta \cap b$. We will show that p is contained in some ball in \mathcal{B}_η . Let x'' denote the point on $\partial\eta^+$ that is closest to p . Let b'' denote the largest ball inside η^+ that is tangential to $\partial\eta^+$ at x'' . Let o'' denote the center of b'' and let r'' denote its radius. Clearly $p \in \overline{o''x''}$. As η is κ -smooth and η^+ is the ε -expansion of η , it follows that $r'' \geq D_\eta/(2\kappa) + \varepsilon D_\eta \geq D_\eta(1 + 2\varepsilon)/(2\kappa) = D_{\eta^+}/(2\kappa)$, where D_{η^+} denotes the diameter of η^+ . Since $D_{\eta^+} \geq \beta r$, we have $r'' \geq \beta r/(2\kappa)$. Also, since $p \in \eta$, it follows that $\|px''\| \geq \varepsilon D_\eta$. Since $D_{\eta^+} = D_\eta(1 + 2\varepsilon)$ and $\varepsilon \leq 1/2$, we have $\|px''\| \geq \varepsilon D_{\eta^+}/2$. Since $b'' \subseteq \eta^+$, we have $D_{\eta^+} \geq 2r''$. Thus $\|px''\| \geq r''\varepsilon$.

Let b' be the ball of radius $r' = \beta r/(2\kappa)$ obtained by scaling b'' about p as the origin. Let o' denote the center of b' and let x' be the point on $\partial b'$ that is closest to p . Clearly $b' \subseteq b''$, $p \in \overline{o'x'}$, and $\|px'\| \geq r'\varepsilon$.

We claim that there must be a point $y \in \mathcal{G}$, such that $\|o'y\| \leq r'/4$ and $\angle yo'x' \leq \sqrt{\varepsilon}$. To prove this, first observe that $p \in b$, and points o' and x' are both within distance r' from p . It follows that o' and x' are both contained in the ball $\tilde{b} = (1 + \beta/(2\kappa))b$. Consider the coordinate axis that make the smallest acute angle θ with segment $\overline{o'x'}$. Clearly $\cos \theta \geq 1/\sqrt{d}$. Recall that our construction places hyperplanes intersecting ball \tilde{b} that are orthogonal to this coordinate axis and separated by a distance of $r'/(16\sqrt{d})$. Since $\cos \theta \geq 1/\sqrt{d}$, these hyperplanes intersect $\overline{o'x'}$, separated by a distance of at most $r'/16$. Thus, there must be a hyperplane H in this set that intersects $\overline{o'x'}$ at a point whose distance from o' is between $r'/8$ and $3r'/16$. Let z denote this point of intersection.

Recall that we overlay a $(d-1)$ -dimensional grid on H of side length $r'\sqrt{\varepsilon}/(16\sqrt{d-1})$. Let y denote the grid point on H that is closest to z . It is easy to see that $\|zy\| \leq r'\sqrt{\varepsilon}/16$. By the triangle inequality, $\|o'y\| \leq \|o'z\| + \|zy\| \leq 3r'/16 + r'\sqrt{\varepsilon}/16 \leq r'/4$, since $\varepsilon \leq 1/2$. Also,

$$\sin \angle yo'x' \leq \frac{\|zy\|}{\|o'z\|} \leq \frac{r'\sqrt{\varepsilon}/16}{r'/8} = \frac{\sqrt{\varepsilon}}{2}.$$

It follows that $\angle yo'x' \leq (\pi/2) \sin \angle yo'x' \leq \sqrt{\varepsilon}$. This completes the proof of the claim.

Let \hat{b} denote the largest ball in \mathcal{B}_y that is contained within η^+ . Since $b' \subseteq \eta^+$, it is clear from our construction that the radius of \hat{b} is at least $(r' - \|o'y\|)/(1 + \varepsilon/2)$. Applying Lemma 4.4, it follows that $p \in \hat{b}$. Recalling that $\hat{b} \in \mathcal{B}_\eta$ completes the proof. \square

The previous lemma suggests a simple approach to answering approximate range queries for smooth ranges. First we construct the data structure of [4] for approximate spherical range queries, and then apply the data structure directly to the approximating balls described in the lemma. Unfortunately, this approach would result in an unacceptably high query time (larger by a factor of $O(1/\varepsilon^{(d-1)/2})$ assuming linear space). Instead, we will apply a more integrated solution. We now explain how to use the above lemma, to process the various types of cells.

Type-1 cells. By property (i) of Lemma 4.2, a type-1 cell u handles query ranges η such that $\eta \subseteq \gamma_\kappa b_u$ and $\text{diam}(\eta) \geq \gamma_\kappa r_u/4$. A cell of type-1 does not generally satisfy any separation property. But note that it only needs to handle ranges in its vicinity whose diameter is proportional to its own size (i.e., neither too small nor too big), which makes the task manageable.

We first discuss the preprocessing phase. Setting b to $\gamma_\kappa b_u$ and β to $1/4$, we apply Lemma 4.5 to obtain the set \mathcal{B} of $O(1/\varepsilon^{(d+1)/2})$ balls. For each ball $\hat{b} \in \mathcal{B}$, we compute $w(P \cap \hat{b})$ and associate it with \hat{b} . The space used for storing this information is on the order of the number of balls in \mathcal{B} , which is $O(1/\varepsilon^{(d+1)/2})$.

Using these cluster sums, we compute the answer for a query η by first finding the set $\mathcal{B}_\eta \subseteq \mathcal{B}$ of balls described in the statement of Lemma 4.5 and then outputting $\sum_{\hat{b} \in \mathcal{B}_\eta} w(P \cap \hat{b})$. To establish the correctness of this method, first observe that $\eta \subseteq b$ and the diameter of η is at least β times the radius of b . By Lemma 4.5, we obtain

$$\eta \cap b \subseteq \bigcup_{\hat{b} \in \mathcal{B}_\eta} \hat{b} \subseteq \eta^+.$$

Since $\eta = \eta \cap b$, it follows that the query is answered correctly.

Finally, we consider the query time. By Lemma 4.5, it takes $O((1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$ time to find the set \mathcal{B}_η . Recall from Lemma 4.2 that it takes $O(\log n)$ time to find the cell u . Thus, the total query time is $O(\log n + (1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$.

Type-2 cells. By property (ii) of Lemma 4.2, a type-2 cell u handles query ranges η such that $q_\eta \in u$ and $\eta \subseteq \gamma_\kappa b_u$. By the properties of type-2 cells mentioned earlier, the points $P \cap \gamma_\kappa b_u$ fall into two groups. Points in one group lie in b'_u and points in the other group (the *pollutants*) lie in $\gamma_\kappa b_u \setminus b'_u$. Recall that the number of pollutants is at most $O(1/\varepsilon^{(d-1)/2})$. Since the number of pollutants is few, preprocessing simply involves storing these points in a list. To answer query η we scan the list and identify the points that lie inside η .

Next we consider the set of points $P \cap b'_u$. Clearly these points are relevant to a query only if the range η overlaps the ball b'_u . By properties of type-2 cells, recall that the ball $\gamma_\kappa b'_u$ does not

overlap u . Since $q_\eta \in u$, it is clear that if η overlaps b'_u then $\text{diam}(\eta) \geq (\gamma_\kappa - 1)r'_u$, where r'_u is the radius of b'_u . Thus, η is large relative to b'_u and so we can use Lemma 4.5 to preprocess the ball b'_u and then answer such queries. To be precise, we set b to b'_u and β to $\gamma_\kappa - 1$, and apply Lemma 4.5 to obtain the set \mathcal{B} of $O(1/\varepsilon^{(d+1)/2})$ balls. For each ball $\hat{b} \in \mathcal{B}$, we compute $w(P \cap \hat{b})$ and associate it with \hat{b} . We answer a query η by finding the set $\mathcal{B}_\eta \subseteq \mathcal{B}$ of balls described in the statement of Lemma 4.5, and then outputting $\sum_{\hat{b} \in \mathcal{B}_\eta} w(P \cap \hat{b})$. We have $\eta \cap b \subseteq \cup_{\hat{b} \in \mathcal{B}_\eta} \hat{b} \subseteq \eta^+$, which implies that the query is answered correctly.

The space used is $O(1/\varepsilon^{(d-1)/2})$ for the pollutants and $O(1/\varepsilon^{(d+1)/2})$ for storing the structure of Lemma 4.5. The query time is $O(1/\varepsilon^{(d-1)/2})$ for scanning the list of pollutants. If η overlaps b'_u , it takes an additional $O((1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$ time for finding the clusters using Lemma 4.5. Thus the space used for cell u is $O(1/\varepsilon^{(d+1)/2})$ and the total query time, including the time for finding cell u , is $O(\log n + (1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$.

Type-3 cells. By property (iii) of Lemma 4.2, a type-3 cell u handles query ranges η such that $8b_u \subseteq \eta \subseteq \gamma_\kappa b_v$. Recall that a type-3 cell u is a quadtree box, and it is contained inside another quadtree box v such that $|P \cap (\gamma_\kappa b_v \setminus 8b_u)| = O(1/\varepsilon^{(d-1)/2})$. Thus, we can answer query η exactly by taking the sum of $w(P \cap 8b_u)$ and the weights of the pollutants in $P \cap (\gamma_\kappa b_v \setminus 8b_u)$ that lie inside η . Thus, by precomputing $w(P \cap 8b_u)$ and storing the list of the pollutants, we can answer queries in time $O(\log n + 1/\varepsilon^{(d-1)/2})$. The space used per cell is $O(1/\varepsilon^{(d-1)/2})$.

In summary, we have shown that for all three types of cells the query time is bounded by $O(\log n + (1/\varepsilon)^{(d-1)/2} \log(1/\varepsilon))$, and the space used per cell is bounded by $O(1/\varepsilon^{(d+1)/2})$. Since the total number of cells is $O(n\varepsilon^{(d-1)/2})$, it follows that the total space used by the data structure is $O(n/\varepsilon)$.

Finally, we make a few remarks regarding preprocessing time. The bottleneck is the computation of the cluster sums for type-1 and type-2 cells. Recall that each cluster is the intersection of the point set P with a suitable ball, and the total number of cluster sums is $O(n/\varepsilon)$. For the sake of simplicity, we assumed in our description that each cluster sum is computed exactly. However, this would take too long, and it can be easily replaced by an approximate computation (i.e., we allow errors $O(\varepsilon)$ -close to the boundary of the ball). For this purpose, we can use the standard BBD-tree data structure of Arya and Mount [7]. This structure takes $O(n \log n)$ time to construct and allows us to answer an ε -approximate convex range query in time $O(\log n + 1/\varepsilon^{d-1})$. Thus, the time to compute all the cluster sums approximately is $O((n \log n)/\varepsilon + n/\varepsilon^d)$.

We can improve the ε -dependency in the construction time significantly by using the more sophisticated data structure of [4], which is tailored to spherical range queries. This data structure uses a parameter γ , $2 \leq \gamma \leq 1/\varepsilon$, to control the space/time tradeoff. (The value γ here is not to be confused with the value $\gamma_\kappa = 72\kappa$ used in the earlier constructions.) The time to construct this data structure is $O(n\gamma^d \log(n/\varepsilon) \log(1/\varepsilon))$, and it allows us to answer an ε -approximate spherical range query in time $O(\log(n\gamma) + 1/(\varepsilon\gamma)^{d-1})$. To achieve our best bound, we construct this structure for $\gamma = 1/\varepsilon^{d/(2d-1)}$. The time to construct this structure and use it to compute all the cluster sums approximately is $O((n/\varepsilon^{(d+1)/2}) \log(n/\varepsilon) \log(1/\varepsilon))$. This completes the proof of Theorem 4.2.

5 Concluding Remarks

In this paper we have presented a number of results on the complexity of approximate range searching in spaces of constant dimension. In particular we have considered how semigroup properties, such as idempotence and integrality, interact with range shape properties, such as sharpness and smoothness, in determining the space and time complexities of the problem. We have shown that the advantages of idempotency do not apply to ranges with sharp corners. In particular, our results imply that for the range space of rotated unit hypercubes, arbitrary (including idempotent) semigroups, and linear space, the query time is $\Omega(1/\varepsilon^{d-2\sqrt{d}})$. For integral semigroups it is $\Omega(1/\varepsilon^{d-2})$. These lower bounds nearly match the upper bound of $O(\log n + (1/\varepsilon)^{d-1})$, which holds for arbitrary semigroups. In contrast, we showed that the improvements offered by idempotence do apply to smooth convex ranges. We presented a lower bound of $\Omega(\log n + (1/\varepsilon)^{(d-1)/2})$ in the algebraic decision-tree model of computation, and presented a nearly matching upper bound. Since the lower bound is in the algebraic decision-tree model, it holds irrespective of space.

There are a few problems that would be nice to resolve. Our lower bounds in the integral case (both here and in [5]) apply under the assumption of convex generators. Can this assumption be removed? The above lower bound on the complexity of range searching for rotated unit hypercubes contains a messy factor of \sqrt{d} in the exponent. An interesting question is whether this term can be eliminated, either in the case of rotated hypercubes or by considering some other natural range space involving ranges with sharp corners. Another problem is that of providing good upper bounds with space-time tradeoffs for rotated hypercube ranges.

This line of work provokes a number of general questions about the computational complexity of approximate range searching in spaces of constant dimension. In particular, what are the best ε -dependencies that can be achieved? How do semigroup properties and the nature of the range shapes affect the time and space complexity? What sorts of models, tools, and structures need to be developed to provide meaningful lower and upper bounds? In contrast to exact range searching, the landscape here is considerably richer and more complex.

6 Acknowledgments

We would like to thank the anonymous referees who read the conference version of this paper for their many valuable suggestions.

References

- [1] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, volume 223 of *Contemporary Mathematics*, pages 1–56. American Mathematical Society, Providence, RI, 1999.
- [2] S. Arya and T. Malamatos. Linear-size approximate Voronoi diagrams. In *Proc. 13th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 147–155, 2002.
- [3] S. Arya, T. Malamatos, and D. M. Mount. Space-efficient approximate Voronoi diagrams. In *Proc. 34th Annu. ACM Sympos. Theory Comput.*, pages 721–730, 2002.

- [4] S. Arya, T. Malamatos, and D. M. Mount. Space-time tradeoffs for approximate spherical range counting. In *Proc. 16th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 535–544, 2005.
- [5] S. Arya, T. Malamatos, and D. M. Mount. On the importance of idempotence. In *Proc. 38th Annu. ACM Sympos. Theory Comput.*, pages 564–573, 2006.
- [6] S. Arya, T. Malamatos, and D. M. Mount. Space-time tradeoffs for approximate spherical range counting. Technical Report CS-TR-4842, Dept. of Computer Science, University of Maryland, 2006.
- [7] S. Arya and D. M. Mount. Approximate range searching. *Comput. Geom. Theory Appl.*, 17:135–152, 2000.
- [8] I. Bárány. Intrinsic volumes and f -vectors of random polytopes. *Math. Annalen*, 285:671–699, 1989.
- [9] I. Bárány. The technique of M-regions and cap-coverings: A survey. In *Proc. III International Conference in Stochastic Geometry, Convex Bodies and Empirical Measures, Part II*, volume 65 of *Rendi. del Circ. Matemat. di Palermo*, pages 21–38, 2000.
- [10] I. Bárány and D. Larman. Convex bodies, economic cap coverings, random polytopes. *Mathematika*, 35:274–291, 1988.
- [11] H. Brönnimann, B. Chazelle, and J. Pach. How hard is halfspace range searching. *Discrete Comput. Geom.*, 10:143–155, 1993.
- [12] T. M. Chan and J. Snoeyink. Algorithms for approximate nearest-neighbor queries. Manuscript, 1995.
- [13] B. Chazelle. Lower bounds on the complexity of polytope range searching. *J. Amer. Math. Soc.*, 2:637–666, 1989.
- [14] B. Chazelle, D. Liu, and A. Magen. Approximate range searching in higher dimension. In *Proc. 16th Canad. Conf. Comput. Geom.*, pages 154–157, 2004.
- [15] J. Erickson. Space-time tradeoffs for emptiness queries. *SIAM J. Comput.*, 29:1968–1996, 2000.
- [16] G. Ewald, D. G. Larman, and C. A. Rogers. The directions of the line segments and of the r -dimensional balls on the boundary of a convex body in Euclidean space. *Mathematika*, 17:1–20, 1970.
- [17] M. L. Fredman. Lower bounds on the complexity of some optimal data structures. *SIAM J. Comput.*, 10:1–10, 1981.
- [18] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin Heidelberg, 1988. 2nd edition 1994.

- [19] J. Matoušek. Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.*, 10(2):157–182, 1993.
- [20] J. Matoušek. Geometric range searching. *ACM Comput. Surv.*, 26:421–461, 1994.
- [21] A. C. Yao. On the complexity of maintaining partial sums. *SIAM J. Comput.*, 14:277–288, 1985.