

Tight Lower Bounds for Halfspace Range Searching

Sunil Arya, David M. Mount & Jian Xia

Discrete & Computational Geometry

ISSN 0179-5376


Volume 47

Number 4

Discrete Comput Geom (2012)

47:711-730

DOI 10.1007/s00454-012-9412-x



Your article is protected by copyright and all rights are held exclusively by Springer Science+Business Media, LLC. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your work, please use the accepted author's version for posting to your own website or your institution's repository. You may further deposit the accepted author's version on a funder's repository at a funder's request, provided it is not made publicly available until 12 months after publication.

Tight Lower Bounds for Halfspace Range Searching

Sunil Arya · David M. Mount · Jian Xia

Received: 24 August 2010 / Revised: 30 August 2011 / Accepted: 30 August 2011 /
Published online: 28 February 2012
© Springer Science+Business Media, LLC 2012

Abstract We establish two new lower bounds for the halfspace range searching problem: Given a set of n points in \mathbb{R}^d , where each point is associated with a weight from a commutative semigroup, compute the semigroup sum of the weights of the points lying within any query halfspace. Letting m denote the space requirements, we prove a lower bound for general semigroups of $\tilde{\Omega}(n^{1-1/(d+1)}/m^{1/(d+1)})$ and for integral semigroups of $\tilde{\Omega}(n/m^{1/d})$.

Our lower bounds are proved in the semigroup arithmetic model. Neglecting logarithmic factors, our result for integral semigroups matches the best known upper bound due to Matoušek. Our result for general semigroups improves upon the best known lower bound due to Brönnimann, Chazelle, and Pach. Moreover, Fonseca and Mount have shown that, given uniformly distributed points, halfspace range queries over idempotent semigroups can be answered in $O(n^{1-1/(d+1)}/m^{1/(d+1)})$ time in the semigroup arithmetic model. As our lower bounds are established for uniformly

A preliminary version of this paper appeared in the Proceedings of the 26th Annual Symposium on Computational Geometry, 2010, pp. 29–37.

S. Arya · J. Xia

Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong

S. Arya

e-mail: arya@cse.ust.hk

D.M. Mount (✉)

Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, USA

e-mail: mount@cs.umd.edu

Present address:

J. Xia

Yahoo! Labs Beijing, Beijing Yahoo! Global R&D Center, Beijing 100083, China

e-mail: johnx@yahoo-inc.com

distributed point sets, it follows that they also resolve the computational complexity of halfspace range searching over idempotent semigroups in this important special case.

Keywords Range searching · Lower bounds · Idempotence

1 Introduction

Given a set of n points P in \mathbb{R}^d , preprocess these points so that given any halfspace η , the number of points of P lying within η can be computed efficiently. This problem, called *halfspace range counting*, while simple to state, is among the most challenging problems in the field of computational geometry. It is a fundamental geometric search problem and a basic building block in a number of other search problems, such as simplex range searching and semialgebraic range searching [1, 18]. In this paper, we consider lower bounds on the computational complexity of this problem.

This problem is typically posed in a weighted setting. Each point is associated with a weight from a commutative semigroup (that is, a set over which is defined a commutative, associative binary operator, denoted “+”), and the output of the query is the semigroup sum of the points lying within the halfspace. Semigroup properties can be exploited by the data structure and may affect the computational complexity of the problem. 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 [13]. 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* (also known as *canonical subsets*). 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 are typically disjoint to avoid multiple counting. Because of the constraint of disjointness, one would expect that range searching over integral semigroups should be harder than for idempotent semigroups.

This problem has an extensive history within computational geometry. Early work focused on data structures of linear size. A number of methods based on partition-trees were proposed, including methods by Willard [21], Edelsbrunner and Welzl [12], and Yao [23]. A major breakthrough was achieved by Haussler and Welzl [16], who applied the concepts of VC-dimension and ε -nets to obtain a data structure with query time $O(n^\alpha)$, where $\alpha \approx 1 - \frac{1}{d(d-1)+1}$. In a sequence of papers, a number of new techniques were introduced, resulting in linear space (or near-linear space) algorithms by Welzl [20], Chazelle and Welzl [9], Matoušek and Welzl [19], and Chazelle et al. [8]. This culminated in a solution by Matoušek [17], which achieves linear space and the optimal query time of $O(n^{1-\frac{1}{d}})$. Chan [5] has shown that this same query time can be achieved with a conceptually simpler and

more flexible approach based on partition trees. For further information, see the excellent surveys by Agarwal and Erickson [1] and Matoušek [18]. All of these results are based on the use of disjoint generators and so apply to both integral and idempotent semigroups.

Since the query time is relatively high, it is natural to consider space-time trade-offs. By precomputing more generators, it is possible to cover ranges by a smaller number of generators, thus leading to lower query times. Letting m denote the space requirements, where $n \leq m \leq n^d$, the best upper bound for halfspace range searching is due to Matoušek [17], who has shown that, using m units of space, queries can be answered in $O(n/m^{1/d})$ time. This bound holds in both the standard real RAM model and the semigroup arithmetic model (defined below).

Lower bounds for this problem were established by Brönnimann, Chazelle, and Pach [4] in the semigroup arithmetic model. We refer to this paper throughout as BCP. They established a lower bound on the query time for general semigroups of $\tilde{\Omega}(n^{1-\frac{d-1}{d(d+1)}}/m^{1/d})$,¹ assuming m units of space. This bound holds for both idempotent and integral semigroups. The gap between the BCP lower bound and Matoušek's upper bound was investigated by Arya et al. [2], and, for the special case of integral semigroups, the lower bound was improved slightly to $\tilde{\Omega}(n/m^{\frac{d+1}{d^2+1}})$. Although the improvement was small, this result showed that semigroup properties are significant to the complexity of the problem.

In this paper, we take a substantial step in reducing the gap between the upper and lower bounds for halfspace range searching. In particular, we establish the following two new lower bounds for halfspace range searching in the semigroup arithmetic model:

- We establish an improved lower bound for general semigroups of $\tilde{\Omega}(n^{1-1/(d+1)}/m^{1/(d+1)})$.
- We establish an improved lower bound for integral semigroups of $\tilde{\Omega}(n/m^{1/d})$.

The latter result relies on the assumption that the convex hulls of the generators used to answer a query are pairwise disjoint. Although this assumption does not hold for Matoušek's incremental construction [17], it is satisfied by Chan's simpler partition-tree construction [5]. Subject to this assumption, our lower bound for integral semigroups is within a polylogarithmic factor of the best known upper bounds.

For general semigroups, it is useful to observe that, as with BCP, our lower bounds are established for uniformly distributed point sets. In contrast to many geometric problems, uniformly distributed point sets appear to represent a fairly hard case for halfspace range searching. Fonseca and Mount [10] recently showed that, given uniformly distributed points, halfspace range queries over idempotent semigroups can be answered in $O(n^{1-1/(d+1)}/m^{1/(d+1)})$ time in the semigroup arithmetic model. Thus, for the semigroup arithmetic model, our results nearly resolve the computational complexity of halfspace range searching over idempotent semigroups in the case of uniformly distributed point sets.

¹We use the notation $\tilde{\Omega}$ in place of Ω to indicate that we are ignoring factors of the form $\log^{O(1)} n$.

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. Unless otherwise stated, we will use the term “constant” to refer to any fixed quantity, which may depend on d but not on n . To avoid specifying many real-valued constants that arise in our analysis, 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 . It will also be convenient in our analysis to use the notations \leq (resp., \geq) to denote less than (resp., greater than) up to constant factors.

Let $(S, +)$ be a commutative semigroup. 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 $w(G)$ to be $\sum_{p \in G} w(p)$, where the summation is taken over the semigroup. Let \mathcal{Q} denote the set of query ranges in the range space. In the range searching problem, we are required to preprocess P so that, for any query range $\eta \in \mathcal{Q}$, we can efficiently compute $w(P \cap \eta)$.

Our lower bound proofs hold in the *semigroup arithmetic model* [14, 22]. We briefly recall the basic elements of this model for online range searching problems. (Details can be found in [1, 6, 7].) 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. Lower bounds in this model assume that the semigroup is *faithful*, meaning that any two identically equal linear forms have the same set of variables [6]. For example, $(\mathbb{N}, +)$, (\mathbb{R}, \min) , and $(\{0, 1\}, \vee)$ are faithful, but $(\{0, 1\}, + \text{ mod } 2)$ is not. 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 exist a set $I_\eta \subseteq \{1, \dots, m\}$ and a set of labeled nonnegative integers $\{\beta_i : i \in I_\eta\}$ such that


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

for any weight function w . The *query time* for η is defined to be the size of a 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.

Note that a storage scheme can exploit the properties of both the semigroup $(S, +)$ and the point set P , but it must work for *any* weight function.² For our purposes, it will be convenient to identify a generator with the subset of P corresponding to the variables in the linear form with positive integer coefficients. An easy consequence of the faithfulness and the definitions given above is that, for idempotent semigroups, the time to answer a query η in this model is the size of the smallest set of generators

²This requirement implies that the lower bounds in the semigroup arithmetic model do not apply to the range counting problem, which is the special case that arises when $w(p_i) = 1$, for all i .

Fig. 1 Widths, breadths, and caps



whose union is $P \cap \eta$ [6, 13]. 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 [6]. We can adapt the above definitions to integral semigroups by adding the restriction that, when taking unions in answering a query, the generators, viewed as subsets of P , are pairwise disjoint [13].

Let us define some other terms that will be used throughout the paper. Given a halfspace H , let ∂H denote its bounding hyperplane. Consider a compact (closed and bounded) convex body K , a halfspace H , and a vector u orthogonal to ∂H . The *width* of K in the direction u , denoted $\text{width}(K, u)$, is defined to be the orthogonal distance between the two supporting hyperplanes for K that are orthogonal to u (see Fig. 1(a)). The *breadth* of K orthogonal to u , denoted $\text{breadth}(K, u)$, is defined to be the maximum cross-sectional area of K orthogonal to u . Throughout, we use the term *area* to mean the $(d - 1)$ -dimensional Hausdorff measure. Note that for our purposes, width is always a one-dimensional measure, while breadth is a $(d - 1)$ -dimensional measure. As a convenience, given a halfspace H with boundary normal u , we define $\text{width}(K, H) = \text{width}(K, u)$ and $\text{breadth}(K, H) = \text{breadth}(K, u)$.

The closed region in \mathbb{R}^d bounded by two parallel $(d - 1)$ -dimensional hyperplanes is called a *slab*, and its *width* (or *thickness*) is the orthogonal distance between these hyperplanes. Let $\Delta > 0$ be a real parameter. Given any hyperplane J , let $S^\Delta(J)$ denote the slab consisting of points in \mathbb{R}^d whose distance from J is at most Δ . This slab has width 2Δ . For any halfspace H , let $S^\Delta(H) = S^\Delta(\partial H) \cap H$ denote the slab consisting of points in H that are within distance Δ of the bounding hyperplane of H .

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 . The intersection $C = K \cap H$ is called the *cap* of K generated by H . We refer to u , the vector orthogonal to ∂H , as the *defining direction* of cap C . The *width* (or *thickness*) of a cap, denoted $\text{width}(C)$, is defined to be the width of C along its defining direction u (see Fig. 1(b)).

Let $\mathbb{U} = [0, 1]^d$ denote the unit hypercube in \mathbb{R}^d , and for any body K in \mathbb{R}^d , let $\mu(K)$ denote its Lebesgue measure. The data point instance used in the proofs consists of a set of uniformly distributed points in \mathbb{U} . This will make it possible to relate the volume of any sufficiently large convex subset of \mathbb{U} to the number of data points it contains.




Fig. 2 Region of interest, fatness, and usefulness of generators (for $t = 5$)

3 Techniques

Our proofs are structurally similar to those in BCP. Their proof is based on a number of sophisticated geometric tools, such as the isoperimetric inequality, the slicing lemma, and Macbeath regions. In this section we provide a high level overview of the structure of their proof and provide a bit of intuition regarding the modifications leading to our improvements in the lower bounds. This will be formalized in Sect. 4.

The complexity of halfspace range searching arises from the difficulty of covering points close to the boundary of the query halfspace. To formalize this idea, in BCP, a suitable set of query halfspaces is considered, and for each halfspace H in this set, a *region of interest*, denoted R_H , is defined to be a thin slab close to the boundary of H (see Fig. 2(a)). Recall that the goal in the arithmetic 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 natural 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 query range, then it cannot be useful in this manner for many query ranges.

Let us consider how to formalize the intuitive notion of the “usefulness” of a generator. In order to achieve query time t , the storage scheme must have at most t generators that cover all the points of P , including those that lie in the region of interest (see Fig. 2(b)). In its simplest form, the *usefulness* of any generator G with respect to H can be defined to be $|G \cap R_H|$, the number of points it covers in the region of interest. In order to apply this notion effectively, however, it is necessary to add further restrictions on the generators that will be considered useful for a particular query. In BCP it was observed that, by a simple counting argument, in order to cover $P \cap R_H$ with t generators, a constant fraction of the points of $P \cap R_H$ belong to generators that cover at least a fraction $\Omega(1/t)$ of the points in the region of interest. We designate such generators as being (absolutely) *fat*, in the sense that they consume a sufficiently large number of points of the region of interest. (This is not to be confused with the conventional usage of the term “fat” in computational geometry [11].) For example, in Fig. 2(b), only generator G_3 fails to be fat, and hence it would not be considered among the useful generators. The BCP lower bound proof exploits the geometric properties of such fat generators.

There are, however, other ways of defining the notion of fatness. These lead to different geometric properties of the generators, which may be exploited to greater

success in the analysis. For example, suppose we have an integral semigroup (and thus generators cannot overlap), and consider the generator G_4 in Fig. 2(b). Such a generator satisfies the BCP notion of fatness, since it covers a fraction of $\Omega(1/t)$ of the points of the region of interest, but it covers a much higher fraction of points of the general set. Arya et al. [2] define a notion of a *relatively fat* generator to be one in which the ratio $|G \cap R_H|/|G|$ is not significantly smaller than the ratio between the volume of the region of interest and the total volume of the unit hypercube. If we were to attempt to cover too many points of the region of interest with absolutely fat generators that were not relatively fat, we would face a situation where points were being covered by multiple generators, which is not allowed in the integral case. In [2], these two notions of fatness are combined to establish a stronger lower bound for halfspace range searching for integral semigroups.


Even together, the notions of absolute and relative fatness are not strong enough to rule out configurations of generators that are clearly illegal. For example, in Fig. 2(c) all the generators intersecting R_H are both absolutely and relatively fat according to the definitions of BCP and [2]. But because they overlap, these generators would not be valid in an integral setting. What is needed is a stronger notion of fatness that does a better job of limiting the overlap between generators.

In this paper, we introduce a new characterization of relative fatness, which is based on cross-sectional area. Consider the generators G_1 through G_4 in Fig. 2(c). These generators satisfy both notions of fatness given earlier, but for each of these generators G_i , the cross-sectional area of its convex hull with respect to H ($\text{breadth}(\text{conv}(G_i), H)$) is very large relative to the fraction of points it covers within the region of interest. Intuitively, a large fraction of the points of the region of interest cannot be covered by such generators, since otherwise these generators must overlap each other within the unit hypercube. This intuition will be established formally in Lemma 4.11. It relies on our aforementioned assumption that the convex hulls of the generators used to answer a query are disjoint. As we shall see, the breadth constraint provides a much more effective method of characterizing the nonoverlapping nature of generators for integral semigroups.

Once the notion of usefulness has been specified, the analysis proceeds by establishing an upper bound on the expected usefulness of any generator. This, in turn, implies a lower bound on the total number of generators needed to achieve the given query time. A number of techniques are presented in BCP for bounding the usefulness of a generator. An important ingredient is something called the *slicing lemma*. Given a convex body K , the slicing lemma of BCP states that it is possible to generate a small collection of bodies, called *Macbeath regions*, all contained within K . These bodies have the property that, if K lies within a query halfspace H and K has a sufficiently large volume of intersection with R_H , then at least one of these bodies lies within $K \cap R_H$ and its volume is proportional to $\mu(K \cap R_H)$ (see Fig. 3).

A critical quantity in the BCP analysis is the number of convex bodies into which we need to decompose the convex hull K of a generator in order to satisfy the aforementioned property. We refine the slicing lemma of BCP, taking into consideration both the volume of intersection and thickness of the region of interest. This leads to a smaller number of Macbeath regions, and an improvement in the lower bound for range searching over general semigroups. This is given in Lemma 4.2. In the case

Fig. 3 Slicing lemma and Macbeath regions



of integral semigroups, we establish a much stronger version of the slicing lemma, which is tailored to work with this breadth-based notion of relative fatness. This is given in Lemma 4.5.

4 Lower Bound

The main results of this paper are summarized in the following theorem. Our first result for general semigroups improves the lower bound of Brönnimann, Chazelle, and Pach [4]. Moreover, since our bounds are established for uniformly distributed points, Fonseca and Mount’s result mentioned in the introduction implies that this result is almost optimal for this distribution and idempotent semigroups. Our second result is almost optimal in the worst-case for integral semigroups in light of Matoušek’s upper bound [17].

Theorem 4.1 *Let $d \geq 2$ be a fixed dimension. Consider a range space consisting of n points and all halfspaces in \mathbb{R}^d , where the weight function for the points is over any faithful semigroup. Then for all sufficiently large n , using $m \geq n$ units of storage,*

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

$$\Omega\left(\frac{n^{1-\frac{1}{d+1}}}{m^{\frac{1}{d+1}}} \cdot \frac{1}{\log^{1-\frac{1}{d+1}} n}\right) = \tilde{\Omega}\left(\frac{n^{1-\frac{1}{d+1}}}{m^{\frac{1}{d+1}}}\right).$$

- (ii) *If the semigroup is integral, then under the assumption that convex hulls of the generators used to answer a query are disjoint, the query time is at least*

$$\Omega\left(\frac{n}{m^{\frac{1}{d}}} \cdot \frac{1}{\log^{1+\frac{2}{d}} n}\right) = \tilde{\Omega}\left(\frac{n}{m^{\frac{1}{d}}}\right).$$

4.1 New Slicing Lemmas

In this section, we establish our new slicing lemmas tailored to the requirements of the idempotent and integral cases. To help us in that task, we will need the following lemma, which involves the useful concept of *Macbeath regions*. This lemma is an easy consequence of Lemma 2.7 in BCP. Recall that for any body K in \mathbb{R}^d , we let $\mu(K)$ denote its Lebesgue measure.

Lemma 4.1 Consider a compact convex body $K \subset \mathbb{R}^d$ and a parameter $0 < v \leq \mu(K)$. There exists a collection of pairwise disjoint convex bodies $K_1, K_2, \dots \subseteq K$ satisfying the following properties:

- (i) $\mu(K_i) = \Omega(v)$ for all i .
- (ii) Let H be a halfspace, and let C be the cap $K \cap H$. If $\mu(C) \geq v$, then there exists j such that $K_j \subseteq C$.

Using the above lemma, we now prove our slicing lemma that will be useful in the idempotent case. It differs from the one given in BCP in that it incorporates a condition on the cap width and, thereby produces a smaller number of convex bodies in cases of interest.

Lemma 4.2 Consider a compact convex body $K \subset \mathbb{R}^d$ of surface area $O(1)$, and two parameters $v > 0$ and $\Delta > 0$. There exists a collection of $O(\Delta/v)$ convex bodies $K_1, K_2, \dots \subseteq K$ satisfying the following property: Let H be a halfspace, and let C be the cap $K \cap H$. If $\text{width}(C) \leq \Delta$ and $\mu(C) \geq v$, then there exists j such that $K_j \subseteq C$ and $\mu(K_j) = \Omega(\mu(C))$.


Proof Let $K_\Delta \subseteq K$ denote the region consisting of points in K that are within distance Δ of K 's boundary. Observe that $\mu(K_\Delta) = O(\Delta)$, since the surface area of K is $O(1)$.

Consider the set \mathcal{C} consisting of caps C formed by intersecting K with any halfspace, such that C satisfies the properties $\text{width}(C) \leq \Delta$ and $\mu(C) \geq v$. We partition the set \mathcal{C} into groups based on the volume of the caps as follows. For $i > 0$, let \mathcal{C}_i be the subset of \mathcal{C} consisting of caps whose volume lies between $2^{i-1}v$ and $2^i v$. We claim that there exists a collection of $O(\Delta/(2^{i-1}v))$ convex bodies satisfying the properties mentioned in the lemma with respect to the caps in \mathcal{C}_i . The lemma then follows immediately from this claim by summing over all i .

To prove the claim, we apply Lemma 4.1, setting the parameter v in that lemma to $2^{i-1}v$. Clearly, this yields a collection \mathcal{K}_i of pairwise disjoint convex bodies, each contained within K and having volume $\Omega(2^{i-1}v)$. Further, for any cap $C \in \mathcal{C}_i$, one of these bodies is contained within C and has volume $\Omega(\mu(C))$ (because the volume of each body is $\Omega(2^{i-1}v)$ and $\mu(C) \leq 2^i v$). Thus, the collection \mathcal{K}_i satisfies all the required properties with respect to the caps in \mathcal{C}_i . However, the number of bodies in \mathcal{K}_i may be too large. We fix this by eliminating any convex body in \mathcal{K}_i that is not contained in K_Δ . Since the caps in \mathcal{C}_i have width at most Δ , such a body is in any case of no use in satisfying a cap of \mathcal{C}_i . Thus, the resulting collection continues to satisfy all the desired properties. Finally, since the bodies in \mathcal{K}_i are pairwise disjoint and are all contained in K_Δ , we have $|\mathcal{K}_i| = O(\mu(K_\Delta)/(2^{i-1}v)) = O(\Delta/(2^{i-1}v))$. This completes the proof of the claim and the lemma. \square

Before presenting our slicing lemma that will be useful for integral semigroups, we need to introduce a standard notion of fatness. A convex body $K \subset \mathbb{R}^d$ is said to be α -fat, for some $\alpha \geq 1$, if there exist two concentric balls B^+ and B^- such that $B^- \subseteq K \subseteq B^+$, and the ratio between the radii of B^+ and B^- is at most α . Further, we say that K is fat if it is α -fat, where α is a constant depending

Fig. 4 Proof of Lemma 4.4



only on d . The following well-known result states that any convex body can be transformed by a volume-preserving affine transformation into a fat convex body. (See, for example, [15].)

Lemma 4.3 *Given any compact convex body $K \subset \mathbb{R}^d$, there exists a volume-preserving affine transform T such that $T(K)$ is fat.*

The next lemma considers the effect of the above fattening transformation on the width of caps. We will find it useful in proving the slicing lemma.

Lemma 4.4 *Consider any compact convex body $K \subset \mathbb{R}^d$ of unit volume. Let H be a halfspace, and let C denote the cap $K \cap H$. Let $K' = T(K)$ and $H' = T(H)$, where T is the transformation of Lemma 4.3. Let $C' = K' \cap H'$. Then $\text{width}(C') = \Theta(\text{breadth}(K, H) \cdot \text{width}(C))$.*

Proof Let w_K denote the width of K along the defining direction of C and $w_{K'}$ denote the width of K' along the defining direction of C' (see Fig. 4). A straightforward geometric argument implies that $\mu(K) = \Theta(\text{breadth}(K, H) \cdot w_K)$. Since $\mu(K) = 1$, we have $w_K = \Theta(1/\text{breadth}(K, H))$. By Lemma 4.3, K' is fat and $\mu(K') = 1$. It follows that the width of K' in any direction is $\Theta(1)$. In particular, we have $w_{K'} = \Theta(1)$. Thus $w_{K'}/w_K = \Theta(\text{breadth}(K, H))$. By standard properties of affine transformations, we obtain $\text{width}(C')/\text{width}(C) = w_{K'}/w_K$. Thus $\text{width}(C')/\text{width}(C) = \Theta(\text{breadth}(K, H))$, as desired. \square

We are now ready to prove the slicing lemma that will be applied in the integral case.

Lemma 4.5 *Consider an α -fat compact convex body $K \subset \mathbb{R}^d$ and two parameters $\beta \geq 1$ and $\Delta > 0$. There exists a collection of $O(\beta \log \alpha)$ convex bodies $K_1, K_2, \dots \subseteq K$ satisfying the following property: Let H be a halfspace, and let C denote the cap $K \cap H$. If $\text{width}(C) \leq \Delta$ and $\mu(C) \geq \text{breadth}(K, H) \cdot \Delta/\beta$, then there exists j such that $K_j \subseteq C$ and $\mu(K_j) = \Omega(\mu(C))$.*

Proof Observe that all the relations stated in the lemma are invariant under uniform scaling, and hence, without loss of generality, we may assume that K is scaled to have unit volume. Consider the set \mathcal{C} consisting of caps C formed by intersecting K with some halfspace H , such that C satisfies the properties $\text{width}(C) \leq \Delta$ and

$\mu(C) \geq \text{breadth}(K, H) \cdot \Delta/\beta$. We partition the set \mathcal{C} into groups such that, for all caps in a group, the breadth of K orthogonal to the defining direction of the cap differs by a factor of at most 2. Since K is α -fat, it is clear that the ratio between the largest and smallest value of the breadth of K over all possible directions cannot exceed α^{d-1} , and thus we can obtain a partitioning with at most $O(\log \alpha)$ groups.

Let us focus on one such group of caps, for which the breadth of K orthogonal to the defining direction of any cap in it is between b and $2b$. We will show that there exists a collection of $O(\beta)$ convex bodies satisfying the properties mentioned in the lemma with respect to all the caps in this group. Since the total number of groups is $O(\log \alpha)$, the lemma will then directly follow from this.

Let $C = K \cap H$ denote any cap in this group. From the conditions in the lemma, we have $\mu(C) \geq \text{breadth}(K, H) \cdot \Delta/\beta \geq b\Delta/\beta$ and $\text{width}(C) \leq \Delta$. Let $K' = T(K)$ and $H' = T(H)$, where T is the transformation of Lemma 4.3. Let $C' = K' \cap H'$. Since T is volume-preserving, we have $\mu(C') = \mu(C) \geq b\Delta/\beta$. By Lemma 4.4, we have $\text{width}(C') = \Theta(\text{breadth}(K, H) \cdot \text{width}(C))$. Since $\text{width}(C) \leq \Delta$ and $\text{breadth}(K, H) \leq 2b$, we obtain $\text{width}(C') = O(b\Delta)$.

Note that the surface area of K' is $O(1)$. We invoke Lemma 4.2 on K' , setting the parameter Δ (upper bound on cap width) to a suitable constant times $b\Delta$, and parameter v (lower bound on cap volume) to $b\Delta/\beta$. We obtain a collection of $O((b\Delta)/(b\Delta/\beta)) = O(\beta)$ convex bodies in K' , such that one of these bodies is inside any cap C' of K' corresponding to a cap C of K in the group under consideration, and its volume is $\Omega(\mu(C'))$. It follows that the set of bodies in the original space corresponding to these bodies satisfies all the desired properties with respect to the caps in the group. This completes the proof of the claim and the lemma. \square

4.2 Lower Bound: The Idempotent Case

We begin with the proof of Theorem 4.1(i), namely the lower bound in the idempotent case. As mentioned earlier, our proof in this case is a simple adaptation of the one in BCP, differing only in its use of a more suitable slicing lemma. We present details for the sake of completeness.

Let P be a set of n points in the unit hypercube $\mathbb{U} = [0, 1]^d$. Let K be any convex body contained within \mathbb{U} , and let $k = |P \cap K|$. We say that P is *scattered* if the following two properties hold for some constant $a > 1$:

- (i) If $\mu(K) \geq a \log n/n$, then $k \geq (n/a)\mu(K)$.
- (ii) If $k \geq \log n$, then $k \leq an\mu(K)$.

Intuitively, scattered point sets behave like uniformly distributed points in the sense that the volume of a convex body is related to the number of points it contains. The following lemma asserts the existence of such point sets. It is proved in [3] and is a simple variant of Lemma 2.1 in BCP.

Lemma 4.6 *A random set of n points sampled uniformly and independently in \mathbb{U} is scattered with probability $1 - o(1)$.*

The data set P in the proof consists of a scattered set of n points in \mathbb{U} . We next describe the set of query ranges. Let O denote the origin of the coordinate system

in \mathbb{R}^d . Let q be any point in $\mathbb{R}^d \setminus \{O\}$. Let H_q denote the halfspace containing the origin, whose bounding hyperplane passes through q and is orthogonal to segment Oq . Let B be the ball of radius $1/4$ whose center coincides with the center of \mathbb{U} . We define the query set \mathcal{Q} to be the set of halfspaces $\{H_q : q \in B\}$.

Next, we define a probability measure over the set of queries. As in BCP, define the measure of any set \mathcal{H} of halfspaces as follows:

$$\int_{\mathcal{H}} dH_q = \int_{H_q \in \mathcal{H}} \frac{dx_1 \wedge \dots \wedge dx_d}{\|q\|^{d-1}},$$

where $q = (x_1, \dots, x_d)$. The choice of this measure is based on the fact that it is invariant under rigid motions [6]. It is easy to see that the measure of the query set \mathcal{Q} is $\Theta(1)$. Therefore, restricted to the set \mathcal{Q} , we can interpret this measure as a probability density by dividing it by the total measure of \mathcal{Q} . In our analysis, we will assume that a query halfspace is selected from \mathcal{Q} according to this probability density (we will say that the halfspace is *random* over \mathcal{Q}). Given any query halfspace $H \in \mathcal{Q}$, we see that for a suitable constant c_1 and any sufficiently small Δ (say $\Delta \leq 1/4$), $\mu(S^\Delta(H) \cap \mathbb{U}) \geq c_1 \Delta$.

The following isoperimetric inequality, which was proved by Chazelle [6], bounds the probability that a random slab encloses a given convex body. For convenience, we restate it in a slightly different form.

Lemma 4.7 *Let $K \subseteq \mathbb{U}$ be any compact convex body, let H be a random halfspace of \mathcal{Q} , and let $\Delta > 0$ be a real parameter. Then*

$$\mu(K) \cdot \Pr[K \subseteq S^\Delta(H)] = O(\Delta^{d+1}).$$

Let \mathcal{G} be any set of $m \geq n$ generators for P . Let $t = t(n, m)$ denote the worst-case query time in the semigroup arithmetic model over all the ranges of \mathcal{Q} . We will establish a lower bound on t , assuming that n is sufficiently large.

In BCP, it is observed that the complexity of halfspace range searching stems from the difficulty of covering points that lie close to the boundary of the halfspace. To formalize the idea, for each halfspace $H \in \mathcal{Q}$, define a *region of interest* $R_H \subseteq H$ to be $S^{\Delta_0}(H)$, where $\Delta_0 = c_0 t \log n/n$ for a suitable constant c_0 .

Let Φ denote the expected number of points of P in the region of interest R_H . That is,

$$\Phi = E[|P \cap R_H|].$$

We will compute lower and upper bounds on Φ , which together will provide the desired lower bound on the query time t . We may assume that $t = O(n/\log^2 n)$, since otherwise Theorem 4.1 holds trivially.

We claim that, for a suitable choice of c_0 , $|P \cap R_H| \geq 8t \log n$ for all $H \in \mathcal{Q}$. By the property of query halfspaces mentioned above, we have $\mu(R_H \cap \mathbb{U}) \geq c_1 \Delta_0 = c_1(c_0 t \log n/n)$, if $c_0 t \log n/n \leq 1/4$. Since $t = O(n/\log^2 n)$, this condition holds for sufficiently large n . Setting $c_0 = 8a/c_1$, where a is the constant used in the scattered set properties, we have $\mu(R_H \cap \mathbb{U}) \geq 8at \log n/n$. Since $t \geq 1$, we can apply property (i) of scattered point sets to obtain $|P \cap R_H| = |P \cap R_H \cap \mathbb{U}| \geq (n/a)\mu(R_H) \geq 8t \log n$, as desired. The following lower bound on Φ is now immediate.

Lemma 4.8 Φ is at least $\Omega(t \log n)$.

Next we compute an upper bound on Φ . Toward this end, it is helpful to focus on those generators that are most efficient in covering the region of interest R_H . Intuitively, in order to achieve a query time of t , a typical generator should cover a fraction of roughly $\Omega(1/t)$ of the points in R_H . This suggests the following definition. We say that a generator $G \in \mathcal{G}$ is *absolutely fat* with respect to a query $H \in \mathcal{Q}$ if $|G \cap R_H| > \log n$ and $G \subseteq H$. The following lemma, implicit in the BCP proof, shows that a constant fraction of the points in R_H are covered by such generators.

Lemma 4.9 For any query $H \in \mathcal{Q}$, a constant fraction of the points of P lying within R_H are covered by generators of \mathcal{G} that are absolutely fat with respect to H .

Proof Given $H \in \mathcal{Q}$, consider the set $A_H \subseteq \mathcal{G}$ of generators that allow us to answer query H within time t . Clearly, $|A_H| \leq t$. Let A_H^1 denote the subset of A_H consisting of generators that are not absolutely fat with respect to H . Observe that the generators in the set A_H^1 cover at most $t \log n$ points of $P \cap R_H$ (because there are at most t such generators and each covers at most $\log n$ points of $P \cap R_H$). Recalling that $|P \cap R_H|$ is at least $8t \log n$, it follows that at least $7/8$ of the points of $P \cap R_H$ are covered by the generators in $A_H \setminus A_H^1$. As all these generators are absolutely fat with respect to H , the lemma follows. \square

Consider any $G \in \mathcal{G}$ and $H \in \mathcal{Q}$. We define the *usefulness* of G with respect to H to be $|G \cap R_H|$ if G is absolutely fat with respect to H , and 0 otherwise. We denote this quantity by $u(G, H)$. Further, we define the *usefulness* of G , denoted $u(G)$, to be the expectation of $u(G, H)$, where the average is taken over the query distribution.

The above lemma implies that,

$$|P \cap R_H| \leq \sum_{G \in \mathcal{G}} u(G, H).$$

Therefore,

$$\Phi = \mathbb{E}[|P \cap R_H|] \leq \mathbb{E}\left[\sum_{G \in \mathcal{G}} u(G, H)\right] = \sum_{G \in \mathcal{G}} \mathbb{E}[u(G, H)] = \sum_{G \in \mathcal{G}} u(G).$$

In the remainder, we compute an upper bound on $u(G, H)$, and hence on $u(G)$, with the help of our version of the slicing lemma and Chazelle’s isoperimetric inequality. This will then allow us to upper bound Φ by using the above inequality.

Lemma 4.10 Consider any $G \in \mathcal{G}$ and $H \in \mathcal{Q}$ such that G is absolutely fat with respect to H . Let $K = \text{conv}(G)$. Then $|G \cap R_H| \leq n \mu(K \cap R_H)$ and $\mu(K \cap R_H) = \Omega(\log n/n)$.

Proof By the definition of absolute fatness, $|G \cap R_H| > \log n$. Clearly, $|P \cap K \cap R_H| \geq |G \cap R_H|$. Therefore $|P \cap K \cap R_H| > \log n$. Applying property (ii) of scattered point sets, we have $|P \cap K \cap R_H| \leq n \mu(K \cap R_H)$. Putting these inequalities together, the lemma follows. \square

Let G, H, K be as in the above lemma. Since G is absolutely fat with respect to H , by definition, $G \subseteq H$, and so $K \cap R_H$ is a cap of K . Clearly, this cap has width at most $\Delta_0 = O(t \log n/n)$, and its volume, by the above lemma, is $\Omega(\log n/n)$. Also, since $K \subseteq \mathbb{U}$, its surface area is $O(1)$. We apply our slicing lemma (Lemma 4.2) to K , setting v to a suitable constant times $\log n/n$ and Δ to Δ_0 . We obtain a collection of $O(\Delta/v) = O(t)$ convex bodies $K_1, K_2, \dots \subseteq K$ such that one of these bodies lies inside $K \cap R_H$ and has volume $\Omega(K \cap R_H)$. Thus,

$$\mu(K \cap R_H) \leq \sum_i \mu(K_i) \cdot \mathbb{I}[K_i \subseteq R_H],$$

where $\mathbb{I}[K_i \subseteq R_H]$ is an indicator random variable that takes the value 1 if $K_i \subseteq R_H$ and 0 otherwise. Applying Lemma 4.10, we obtain

$$u(G, H) = |G \cap R_H| \leq n \mu(K \cap R_H) \leq n \sum_i \mu(K_i) \cdot \mathbb{I}[K_i \subseteq R_H].$$

Since $u(G, H) = 0$, if G is not absolutely fat with respect to H , it follows that the inequality

$$u(G, H) \leq n \sum_i \mu(K_i) \cdot \mathbb{I}[K_i \subseteq R_H]$$

holds for all $G \in \mathcal{G}$ and $H \in \mathcal{Q}$. Taking expectation on both sides, we obtain

$$u(G) \leq n \sum_i \mu(K_i) \cdot \Pr[K_i \subseteq R_H] \leq n t \Delta_0^{d+1},$$

where we have used Lemma 4.7 (Chazelle's isoperimetric inequality) and the $O(t)$ bound on the number of convex bodies derived from our slicing lemma.

Recalling that $\Phi \leq \sum_{G \in \mathcal{G}} u(G)$ and the total number of generators is $|\mathcal{G}| = m$, we obtain

$$\Phi = O(m n t \Delta_0^{d+1}).$$

Combining this with the lower bound of $\Omega(t \log n)$ on Φ derived earlier and substituting the value $\Delta_0 = O(t \log n/n)$, we obtain

$$t \log n \leq m n t \left(\frac{t \log n}{n} \right)^{d+1},$$

which can be simplified to

$$t^{d+1} \geq \frac{n^d}{m} \cdot \frac{1}{\log^d n}.$$

Therefore,

$$t = \Omega \left(\frac{n^{\frac{d}{d+1}}}{m^{\frac{1}{d+1}}} \cdot \frac{1}{\log^{\frac{d}{d+1}} n} \right) = \tilde{\Omega} \left(\frac{n^{1-\frac{1}{d+1}}}{m^{\frac{1}{d+1}}} \right).$$

This completes the proof of Theorem 4.1(i).

4.3 Lower Bound: The Integral Case

In this section, we prove part (ii) of Theorem 4.1. As mentioned earlier, throughout this section, we assume that the convex hulls of the generators used to answer a query are disjoint. This assumption does not seem to be too restrictive. Although it does not hold for Matoušek’s data structure [17], it is satisfied by Chan’s partition-tree structure [5].

Our proof is structurally similar to that of Arya et al. [2], who first showed how to enhance the BCP framework to take advantage of disjointness. The key difference is that we use a different formulation for capturing disjointness, which enables us to obtain better bounds. We now present the details.

We define the data set P , query set \mathcal{Q} , worst-case query time t , slab thickness Δ_0 , region of interest R_H for each halfspace $H \in \mathcal{Q}$, set of generators \mathcal{G} , and Φ , exactly as in the idempotent case. To simplify the notation, for any generator G , we use $\text{breadth}(G, H)$ to denote the quantity $\text{breadth}(\text{conv}(G), H)$. As before, we can establish a lower bound of $\Omega(t \log n)$ on Φ . A key aspect of our strategy is to use a different notion of generator *usefulness*. To motivate our definition, recall from Sect. 3 that a generator is useful only if the fraction of points it covers within the region of interest is large relative to its breadth. Since the region of interest is of volume roughly Δ_0 , the fraction of points covered by generator G in R_H is about $|G \cap R_H|/(\Delta_0 \cdot n)$. This suggests that we are interested in generators for which $|G \cap R_H|$ is large relative to $\text{breadth}(G, H) \cdot \Delta_0 \cdot n$. For technical reasons, we will also need a normalization factor that is proportional to $\log(1/\Delta_0)$. Formally, we define a generator $G \in \mathcal{G}$ to be *relatively fat* with respect to a query $H \in \mathcal{Q}$ if

$$|G \cap R_H| \geq \frac{\text{breadth}(G, H) \cdot \Delta_0 \cdot n}{c \log(1/\Delta_0)},$$

for a suitable constant c to be specified later. Further, define generator G to be *fat* with respect to query H if G is absolutely fat with respect to H and either (i) $G \subset R_H$ or (ii) G is relatively fat with respect to H . We can now strengthen Lemma 4.9 for the integral case as follows.

Lemma 4.11 *Assume that the convex hulls of the generators used to answer a query are disjoint. For any query $H \in \mathcal{Q}$, a constant fraction of the points of P lying within R_H are covered by generators of \mathcal{G} that are fat with respect to H .*

The proof of Lemma 4.11 relies on the following important geometric observation.

Lemma 4.12 *Consider a halfspace $H \in \mathcal{Q}$ and any subset $\mathcal{G}_H \subseteq \mathcal{G}$ consisting of generators whose convex hulls are disjoint. Further, suppose that each generator $G \in \mathcal{G}_H$ satisfies the following properties:*

- (i) $G \subset H$,
- (ii) $G \cap S^{\Delta_0/2}(H) \neq \emptyset$, and
- (iii) $G \not\subset S^{\Delta_0}(H)$.

Then $\sum_{G \in \mathcal{G}_H} \text{breadth}(G, H) = O(\log(1/\Delta_0))$.




Fig. 5 Proof of Lemma 4.12

Proof To simplify the exposition, define *horizontal* to mean orthogonal to the d th coordinate axis. Through a suitable rotation of space, we may assume that ∂H is horizontal, and H is the halfspace consisting of points below ∂H (that is, having smaller d -coordinate values). Let J_0 denote the hyperplane parallel to ∂H and at distance $\Delta_0/2$ below it. Further, for $i \geq 1$, let J_i denote the hyperplane parallel to J_0 and at distance $2^{i-3}\Delta_0$ below it (see Fig. 5). Note that the number of hyperplanes in the set $\{J_i : i \geq 0\}$ that intersect $\mathbb{U} = [0, 1]^d$ is $O(\log(1/\Delta_0))$.

Let G be any generator of \mathcal{G}_H . Let $K_G = \text{conv}(G)$, $b_G = \text{breadth}(G, H)$, and let B_G be the cross-section that defines b_G (note area of B_G is b_G). Let \mathcal{G}'_H (resp., \mathcal{G}''_H) be the subset of \mathcal{G}_H consisting of generators G such that B_G lies below (resp., above) J_1 . It suffices to show that $\sum_{G \in \mathcal{G}'_H} b_G = O(\log(1/\Delta_0))$ and $\sum_{G \in \mathcal{G}''_H} b_G = O(1)$.

To establish the first bound, consider any generator $G \in \mathcal{G}'_H$. Observe that K_G intersects both J_0 and J_2 , since $G \cap S^{\Delta_0/2}(H) \neq \emptyset$ (Property (ii)) and $G \not\subset S^{\Delta_0}(H)$ (Property (iii)). Suppose that B_G lies between the hyperplanes J_i and J_{i+1} , for some $i \geq 1$. Consider the pyramid formed by joining any fixed point x in $K_G \cap J_0$ to all the points in the base B_G . Let B'_G be the surface patch formed by intersecting this pyramid with the plane J_i . It follows from convexity that $B'_G \subseteq K_G$. Also, since the distance of x from J_i is at least half the distance of x from J_{i+1} , it follows that the area of B'_G is $\Omega(b_G)$.

Because the convex hulls of the generators of \mathcal{G}'_H are disjoint, the associated set of surface patches $\{B'_G : G \in \mathcal{G}'_H\}$ are also disjoint. Since these patches lie on the portions of the hyperplanes $J_i, i \geq 1$, that lie inside \mathbb{U} , and the number of these hyperplanes that intersect \mathbb{U} is $O(\log(1/\Delta_0))$, it follows that $\sum_{G \in \mathcal{G}'_H} b_G = O(\log(1/\Delta_0))$. This completes the proof of the first bound.

The second bound is established by a similar argument. In this case, the points x are chosen to lie on J_2 , and the surface patches are taken only from the hyperplanes J_1 and J_0 . Since there are only a constant number of hyperplanes involved, the bound obtained for this case is $O(1)$. This completes the proof. \square

We can now prove Lemma 4.11 by using the above observation. The proof relies crucially on showing that, among the set of generators used to answer query H , the ones that are neither contained in R_H nor relatively fat do not cover more than a constant fraction of the points of $P \cap S^{\Delta_0/2}(H)$. It can be shown that the convex hulls of such generators intersect both bounding hyperplanes of the slab $S^{\Delta_0}(H) \setminus S^{\Delta_0/2}(H)$. We show that the convex hull disjointness assumption implies that the sum of the breadths associated with these generators is at most $O(\log(1/\Delta_0))$. Recall that generators that are not relatively fat cover “few” points in R_H relative to their associated breadths. Using these two facts then allows us to obtain the desired bound on the number of points covered by these generators.

Proof of Lemma 4.11 Given $H \in \mathcal{Q}$, consider the set $A_H \subseteq \mathcal{G}$ of generators that allow us to answer the query H within time t . Clearly, $|A_H| \leq t$. Consider the slab $S^{\Delta_0/2}(H)$. Note that $S^{\Delta_0/2}(H) \subset S^{\Delta_0}(H) = R_H$. Just as we showed $|P \cap R_H| \geq 8t \log n$ (after Lemma 4.7), we can show that $|P \cap S^{\Delta_0/2}(H)| \geq 4t \log n$. Applying property (ii) of scattered point sets, we can easily show that $|P \cap R_H| = O(t \log n)$. We claim that the generators of A_H that are not fat cover at most $2t \log n$ points of $P \cap S^{\Delta_0/2}(H)$. (Here and throughout this proof, fatness, absolute fatness, and relative fatness of generators are understood to be with respect to H .) Clearly this implies that the fat generators of A_H cover at least $2t \log n$ points of $P \cap S^{\Delta_0/2}(H)$, and hence a constant fraction of the points of $P \cap R_H$.

To prove the claim, consider the subset $A_H^0 \subseteq A_H$ consisting of generators that are not fat and cover some point of $P \cap S^{\Delta_0/2}(H)$. Let A_H^1 denote the subset of A_H^0 consisting of generators that are not absolutely fat. By definition, each such generator covers at most $\log n$ points of $P \cap R_H$. Therefore, the generators of A_H^1 cover at most $t \log n$ points of $P \cap R_H$, and hence of $P \cap S^{\Delta_0/2}(H)$.

Let $A_H^2 = A_H^0 \setminus A_H^1$. Since any generator $G \in A_H^2$ is absolutely fat but not fat, and covers some point of $P \cap S^{\Delta_0/2}(H)$, it is easy to see that it satisfies the following:

- (i) $G \subset H$,
- (ii) $G \cap S^{\Delta_0/2}(H) \neq \emptyset$,
- (iii) $G \not\subset R_H$, and
- (iv) G is not relatively fat (i.e., $|G \cap R_H| < (\text{breadth}(G, H) \cdot \Delta_0 \cdot n) / (c \log(1/\Delta_0))$).

Also, by the assumption in the lemma, the convex hulls of the generators in A_H^2 are disjoint. Clearly, the generators of A_H^2 satisfy all the preconditions of Lemma 4.12, and therefore $\sum_{G \in A_H^2} \text{breadth}(G, H) = O(\log(1/\Delta_0))$. Since the generators of A_H^2 are not relatively fat, it follows that the number of points of $P \cap S^{\Delta_0/2}(H)$ covered by these generators is at most

$$\sum_{G \in A_H^2} |G \cap R_H| \leq \sum_{G \in A_H^2} \frac{\text{breadth}(G, H) \cdot \Delta_0 \cdot n}{c \log(1/\Delta_0)} = \frac{\Delta_0 n}{c \log(1/\Delta_0)} \sum_{G \in A_H^2} \text{breadth}(G, H).$$

We know that $\sum_{G \in A_H^2} \text{breadth}(G, H) = O(\log(1/\Delta_0))$ and $\Delta_0 = O(t \log n/n)$, and so by choosing the constant c (in the definition of relatively fat generators) to be sufficiently large, we can ensure that the generators of A_H^2 cover at most $t \log n$ points of $P \cap S^{\Delta_0/2}(H)$. Combining the contributions from the generators of A_H^1 and A_H^2 ,

it follows that the generators of A_H^0 together cover at most $2t \log n$ points of $P \cap S^{\Delta_0/2}(H)$. This completes the proof of the claim and the lemma. \square

Next we define a notion of generator usefulness based on the above lemma, which will help us derive the lower bounds. Consider any $G \in \mathcal{G}$ and $H \in \mathcal{Q}$. We define the *usefulness* of G with respect to H , denoted $u(G, H)$, as in the idempotent case, except that we use “fatness” in place of “absolute fatness”. As before, let $u(G) = E[u(G, H)]$. Arguing as before, but using Lemma 4.11 in place of Lemma 4.9, we obtain

$$\Phi \leq \sum_{G \in \mathcal{G}} u(G).$$

Consider any $G \in \mathcal{G}$ and $H \in \mathcal{Q}$ such that G is fat with respect to H . Let $K = \text{conv}(G)$. If $G \subset R_H$, then by Lemma 4.10,

$$u(G, H) = |G \cap R_H| \leq n \mu(K \cap R_H) = n \mu(K).$$

Otherwise, G must be relatively fat with respect to H . In this case, applying Lemma 4.10, we have

$$\mu(K \cap R_H) \geq \frac{|G \cap R_H|}{n} \geq \frac{\text{breadth}(G, H) \cdot \Delta_0}{c \log(1/\Delta_0)}.$$

Thus, $K \cap R_H$ is a cap of K of width at most Δ_0 and volume $\Omega((\text{breadth}(G, H) \cdot \Delta_0) / \log(1/\Delta_0))$. Also, since $K \subseteq \mathbb{U}$ and, by Lemma 4.10, $\mu(K) \geq \mu(K \cap R_H) = \Omega(\log n/n)$, it is easy to show that K is α -fat for $\alpha = O(n/\log n)$. We apply our slicing lemma (Lemma 4.5) to K , setting β to a suitable constant times $\log(1/\Delta_0)$ and Δ to Δ_0 . We obtain a collection of $O(\beta \log \alpha) = O(\log(1/\Delta_0) \cdot \log n) = O(\log^2 n)$ convex bodies $K_1, K_2, \dots \subseteq K$ such that one of these bodies lies inside $K \cap R_H$ and has volume $\Omega(\mu(K \cap R_H))$. Arguing as before, we obtain

$$u(G, H) \leq n \sum_i \mu(K_i) \cdot \mathbb{I}[K_i \subseteq R_H].$$

By definition, $u(G, H) = 0$ if G is not fat with respect to H . Combining all the cases, we have

$$u(G, H) \leq n \left(\mu(K) \cdot \mathbb{I}[K \subseteq R_H] + \sum_i \mu(K_i) \cdot \mathbb{I}[K_i \subseteq R_H] \right)$$

holds for all $G \in \mathcal{G}$ and $H \in \mathcal{Q}$.

Taking expectation on both sides and arguing as before, we obtain

$$u(G) \leq n \log^2(n) \Delta_0^{d+1},$$

which yields

$$\Phi = O(m n \log^2(n) \Delta_0^{d+1}).$$

Substituting $\Delta_0 = O(t \log n/n)$ and combining this with the $\Omega(t \log n)$ lower bound on Φ , we obtain

$$t \log n \leq m n \log^2(n) \left(\frac{t \log n}{n} \right)^{d+1},$$

which can be simplified to

$$t^d \geq \frac{n^d}{m} \cdot \frac{1}{\log^{d+2} n}.$$

Therefore,

$$t = \Omega\left(\frac{n}{m^{\frac{1}{d}}} \cdot \frac{1}{\log^{\frac{d+2}{d}} n}\right) = \tilde{\Omega}\left(\frac{n}{m^{\frac{1}{d}}}\right).$$

This completes the proof of Theorem 4.1(ii).

5 Conclusions

In this paper we have established new lower bounds for halfspace range searching in the semigroup arithmetic model, for both general (including idempotent) and integral semigroups. Neglecting logarithmic factors, our result for integral semigroups matches the best known upper bound due to Matoušek [17] and Chan [5]. However, it relies on the assumption that the convex hulls of the generators used to answer a query are pairwise disjoint. This assumption is satisfied by methods, like Chan’s, that are based on a straightforward application of partition trees. Our result for general semigroups improves upon the best known lower bound due to Brönnimann, Chazelle, and Pach [4], and matches an upper bound by Fonseca and Mount [10] for halfspace range searching in the semigroup arithmetic model over uniformly distributed points.

The near tightness of the lower bound in the integral case is somewhat surprising, given the weakness of the semigroup arithmetic model (where the cost of computing the generators used to answer a query is not counted). It is natural to wonder whether our assumption about the disjointness of the generator hulls can be overcome. It is certainly central to our approach, but perhaps there are other ways of constraining overlap between generators without resorting to this assumption.

The other obvious open problem is to determine the computational complexity of halfspace range searching over idempotent semigroups. Arya et al. [2] have shown that idempotence can be exploited to dramatically improve query times for approximate range searching, but their methods cannot be easily extended to the exact case. If one believes that uniformly distributed point sets are reflective of hard instances of halfspace range searching, the results of Fonseca and Mount [10] offer hope that halfspace range searching over idempotent semigroups may be strictly easier than over integral semigroups. Unfortunately, their construction only shows the existence of the desired number of generators; it does not provide a concrete method of computing them efficiently.

Acknowledgements We would like to thank the referees for their valuable suggestions.

S.A. was supported by the Research Grants Council, Hong Kong, China under project number 610110.

D.M.M. was partially supported by the National Science Foundation under grants CCR-0635099 and CCF-1117259 and the Office of Naval Research under grant N00014-08-1-1015.

J.X. was supported by the Research Grants Council, Hong Kong, China under project number 610106.

References

1. Agarwal, P.K., Erickson, J.: Geometric range searching and its relatives. In: Chazelle, B., Goodman, J.E., Pollack, R. (eds.) *Advances in Discrete and Computational Geometry*. Contemporary Mathematics, vol. 223, pp. 1–56. American Mathematical Society, Providence (1999)
2. Arya, S., Malamatos, T., Mount, D.M.: On the importance of idempotence. In: Proc. 38th Annu. ACM Sympos. Theory Comput., pp. 564–573 (2006)
3. Arya, S., Malamatos, T., Mount, D.M.: The effect of corners on the complexity of approximate range searching. *Discrete Comput. Geom.* **41**(3), 398–443 (2009)
4. Brönnimann, H., Chazelle, B., Pach, J.: How hard is halfspace range searching? *Discrete Comput. Geom.* **10**, 143–155 (1993)
5. Chan, T.M.: Optimal partition trees. In: Proc. 26th Annu. ACM Sympos. Comput. Geom., pp. 1–10 (2010)
6. Chazelle, B.: Lower bounds on the complexity of polytope range searching. *J. Am. Math. Soc.* **2**, 637–666 (1989)
7. Chazelle, B.: Lower bounds for orthogonal range searching, II: The arithmetic model. *J. Assoc. Comput. Mach.* **37**, 439–463 (1990)
8. Chazelle, B., Sharir, M., Welzl, E.: Quasi-optimal upper bounds for simplex range searching and new zone theorems. *Algorithmica* **8**, 407–429 (1992)
9. Chazelle, B., Welzl, E.: Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete Comput. Geom.* **4**(5), 467–489 (1989)
10. da Fonseca, G.D., Mount, D.M.: Approximate range searching: The absolute model. *Comput. Geom.* **43**(4), 434–444 (2009)
11. de Berg, M., Katz, M.J., van der Stappen, A.F., Vleugels, J.: Realistic input models for geometric algorithms. *Algorithmica* **34**, 81–97 (2002)
12. Edelsbrunner, H., Welzl, E.: Halfplanar range search in linear space and $O(n^{0.695})$ query time. *Inf. Process. Lett.* **23**(6), 289–293 (1986)
13. Erickson, J.: Space-time tradeoffs for emptiness queries. *SIAM J. Comput.* **29**, 1968–1996 (2000)
14. Fredman, M.L.: Lower bounds on the complexity of some optimal data structures. *SIAM J. Comput.* **10**, 1–10 (1981)
15. Grötschel, M., Lovász, L., Schrijver, A.: *Geometric Algorithms and Combinatorial Optimization*. Algorithms and Combinatorics, vol. 2, 2nd edn. Springer, Berlin (1993)
16. Haussler, D., Welzl, E.: Epsilon-nets and simplex range queries. *Discrete Comput. Geom.* **2**, 127–151 (1987)
17. Matoušek, J.: Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.* **10**(2), 157–182 (1993)
18. Matoušek, J.: Geometric range searching. *ACM Comput. Surv.* **26**, 421–461 (1994)
19. Matoušek, J., Welzl, E.: Good splitters for counting points in triangles. In: Proc. 5th Annu. ACM Sympos. Comput. Geom., pp. 124–130 (1989)
20. Welzl, E.: Partition trees for triangle counting and other range searching problems. In: Proc. 4th Annu. ACM Sympos. Comput. Geom., pp. 23–33 (1988)
21. Willard, D.E.: Polygon retrieval. *SIAM J. Comput.* **11**, 149–165 (1982)
22. Yao, A.C.: On the complexity of maintaining partial sums. *SIAM J. Comput.* **14**, 277–288 (1985)
23. Yao, F.F.: A 3-space partition and its applications. In: Proc. 15th Annu. ACM Sympos. Theory Comput., pp. 258–263 (1983)