Lower Bounds for Off-Line Range Searching

BERNARD CHAZELLE
Department of Computer Science
Princeton University
Princeton, NJ 08544, USA

1 Introduction

We establish three lower bounds for problems of the following kind: Given n weighted points in $\mathbf{R}^d$ and n axis-parallel boxes, compute the sum of the weights within each box. Problems of this sort have been extensively studied: see [6, 8, 13, 14, 16, 19] for surveys or general introductions to the subject of range searching. We prove that:

- If both additions and subtractions are allowed, then the problem requires $\Omega(n \log\log n)$ arithmetic operations. This is the first general result for the group model. Note, however, that it falls short of the best known upper bound of $O(n \log n)$. The proof uses the spectral method of [5]. This reduces the problem to that of finding a set system A such that the eigenvalues of $A^\mathsf{T}A$ are large. We do this nonconstructively by using a mixture of algebraic and probabilistic arguments. The key ingredient of the proof is a discrete version of Roth's method of orthogonal functions. This is a powerful technique from discrepancy theory, which we hope will find further use in complexity theory.

- If subtractions are disallowed (the semigroup model) then a much stronger lower bound can be established, i.e., $\Omega(n(\log n/\log\log n)^{d-1})$, which is nearly optimal. The semigroup model corresponds to the monotone arithmetic circuit complexity of the problem, so one should expect lower bounds to be easier to prove. Actually, the proof is surprisingly simple: it involves little more than the Chinese Remainder Theorem and basic properties of Halton-Hammersley sequences. The on-line version of the problem was treated in [4, 10] and required fairly complicated arguments. The off-line case was open.

- Our last result concerns the same problem as above, but with simplices replacing boxes. Again, a quasi-optimal lower bound can be established in the semigroup model, i.e., $\Omega(n^{2-2/(d+1)}(\log n)^{-5/2})$. A practical observation is that when d is large the bound is basically quadratic, which shows that the naive algorithm (checking which point belongs to which simplex, one pair at a time) is the method of choice. The proof makes use of recent results on Heilbronn's problem [3] and techniques from [7]. Again, the on-line version of the problem has been (almost completely) solved [3, 11, 18], while the off-line case was open. See also [9] for related results on Hopcroft's problem in two dimensions.

2 The Group Model

Given n weighted points in the plane, with weights chosen in an Abelian group $(G, +)$, and n axis-parallel boxes, we consider the problem of computing the sum of the weights of the points within each box. Obviously this is the same as computing Ax, where A is the incidence matrix of the associated set system and x is the vector of weights.

In the group model, a circuit (or straight-line program) encodes the map A and is required to compute Ax for any $x \in G^n$, where $(G, +)$ is an arbitrary Abelian group. There are two types of gates: A *regular gate* takes a pair (a, b) as input and it outputs $a + b$ or $a - b$. A *help gate* outputs $f(a, b)$, where f is *any* function from G^2 to G. The motivation behind the use of help gates is that often the group G can be embedded into a more complex structure, say, a ring or a field, and other operations might be possible. The use of help gates leaves open that possibility without restricting the generality of the model.

For the purpose of the proof we assume that G is the additive group of real numbers. Let $\lambda_1 \geq \cdots \geq \lambda_n$ be the eigenvalues of $A^\mathsf{T}A$. The *spectral lemma* of

This work was supported in part by NSF Grant CCR-93-01254 and The Geometry Center, University of Minnesota, an STC funded by NSF, DOE, and Minnesota Technology, Inc.

STOC '95, Las Vegas, Nevada, USA
© 1995 ACM 0-89791-718-9/95/0005..$3.50

[5] states that any circuit for computing Ax is of size $\Omega(\max_k (k-h)\log \lambda_k)$, where h is the number of help gates. Thus we construct a set of n points and a set of n axis-parallel boxes such that any eigenvalue of $A^\top A$ of rank slightly less than $n/4$ is at least $(\log n)^{\Omega(1)}$. This immediately implies that computing Ax requires a circuit of size $\Omega(n \log\log n)$, even in the presence of up to roughly $n/4$ help gates.

First, we build a large $N \times n$ set system B, from which we extract an $n \times n$ set system A that satisfies the lower bound. To construct B we use a set P of n points given by a (well-chosen) two-dimensional Halton-Hammersley subsequence; for boxes we take the southwest quadrants anchored at the N vertices of a very fine square grid. Let $\mu_1 \geq \cdots \geq \mu_n$ be the eigenvalues of $B^\top B$. (In the following we make no distinction between a set system and its incidence matrix.) We go through the following sequence, from which the lower bound follows by application of the spectral lemma:[1]

- **Step 1:** Show that $\mu_k \gg (n-k+1)N(\log n)/n^2$ (Lemma 2.3). By the Courant-Fischer characterization of eigenvalues, this entails estimating the minmax value of the Rayleigh quotient $\|Bx\|_2^2/\|x\|_2^2$ over all $(n-k+1)$-dimensional subspaces. By using Roth's method of orthogonal functions [2, 17] we derive a lower bound on the hybrid ratio $\|Bx\|_2/\|x\|_1$, for any $x \neq 0$ (Lemma 2.2). Unfortunately, we need the L^2 norm in the denominator. The standard inequalities relating the L^1 and L^2 norms are too crude for our purposes, and we need a probabilistic argument to produce the desired lower bound (Lemma 2.3).

- **Step 2:** Prove the existence of an $n \times n$ submatrix A of B such that $\det A^\top A = (\log n)^{n-o(n)}$ (Lemma 2.4). Step 1 yields a lower bound on $\det B^\top B = \prod_k \mu_k$. The Binet-Cauchy formula leads to the (nonconstructive) existence of A.

- **Step 3:** Show that the k-th largest eigenvalue λ_k of $A^\top A$ is at least $(\log n)^{\Omega(1)}$, for any k up to roughly $n/4$ (Lemma 2.5). Since $\det A^\top A = \prod_k \lambda_k$, such a bound follows from the previous step, provided that we can bound the low-ranked eigenvalues from above. This is done by exhibiting large enough invariant subspaces within which the spectral norm of the map $A^\top A$ is low.

Remark: Step 3 makes use (of all things!) of data structuring techniques for range searching. It is ironic that proving a lower bound on the complexity of range searching should require the use of data structures. But, of course, one consequence of this work is that low spectrum is a precondition for the existence of efficient data structures for linear maps, and hence for range searching. So, with hindsight it is not all that surprising.

[1] We shall use the notation $\gg$ or $\ll$ to denote inequality up to a constant multiplicative factor.

Theorem 2.1 *Range searching with respect to n points and n axis-parallel boxes requires $\Omega(n \log n \log n)$ group operations in the worst case. This remains true even in the presence of $n/4 - \varepsilon n$ help gates, for any fixed $\varepsilon > 0$.*

The theorem shows that up to $n/4$ help gates cannot help. On the other hand, it is easy to see that over the reals $2n-1$ help gates suffice to make the problem trivial. One should also note that without help gates the problem is easily solved in $O(n \log n)$ time on a RAM. The lower bound holds in any dimension higher than 1. The obvious open question is whether $\Theta(n \log\log n)$ is the right bound. The spectral method seems unlikely to provide an answer to this question. For example, one would expect the dimension of the ambient space to play a role, something which the method seems to rule out.

The Proof of Theorem 2.1

Let m be a large power of two and let $n = m/4$. The n-point set P is a subset of the classical bit-reversal m-point set:

$$Q = \Big\{ (1/(2m) + c(k), 1/(2m) + k/m) \,:\, 0 \leq k < m \Big\},$$

where $c(k) = \sum_{i \geq 0} b(i)/2^{i+1}$ and $\{b(i)\}$ is the binary expression for k, i.e., $k = \sum_{i \geq 0} b(i) 2^i$. For any $1 \leq k \leq \log n$, let X_k be the grid obtained by dividing $[0,1]^2$ into m axis-parallel rectangles of size $2^{-k} \times (2^k/m)$. Each cell σ of X_k is a rectangle of area $1/m$ that contains exactly one point q of Q. We say that q is *well-centered* for X_k if is lies near the center c_σ of σ; specifically, within the box $(\sigma + c_\sigma)/2$. A simple inductive proof shows that at least half the points of Q are well-centered for X_k. It follows that at least $m/4$ points of Q are each well-centered for at least $(\log n)/3$ grids X_k. We define P to consist of these $m/4$ points.

Consider the $(\sqrt{N}-1) \times (\sqrt{N}-1)$ square grid $\mathcal{G}$ covering $[0,1]^2$, where $N = (m^2+1)^2$. Each row of the $N \times n$ matrix B is the characteristic vector of the subset of P lying in the southwest quadrant anchored at a distinct grid point; in other words, for each grid point (x,y) there is a distinct row in B corresponding to the quadrant $(-\infty, x] \times (-\infty, y]$. Note that the N rows are not all distinct. Next, we show that the set system B has a high spectrum. We do this in two steps: we lower-bound successively the L^2 norm of Bx and the eigenvalues of $B^\top B$.

Lemma 2.2 *For any $x \in \mathbf{R}^n$,*

$$\|Bx\|_2 \gg \frac{1}{n}\sqrt{N \log n}\,\|x\|_1.$$

Proof: Fix $x = (x_1, \ldots, x_n) \in \mathbf{R}^n$: each x_i corresponds to a distinct point of P; for convenience we shall ignore this distinction. Without loss of generality, we can assume that

$$\|x\|_1 \leq 2 \sum_{x_i > 0} x_i. \tag{1}$$

Our approach is a variant of Roth's method of orthogonal functions [2, 17]. Given $1 \leq k \leq \log n$, we say that σ is *k-good* if it contains a well-centered point x_i and $x_i > 0$. We assign a weight to each grid point q of $\mathcal{G}$ as follows: Let σ be any cell of X_k that contains q;

- If σ is not uniquely defined (because q lies on its boundary) or if σ is not k-good, then assign q a weight of 0.
- Else, subdivide σ into four equal-size quadrants (similar to σ): Assign q a weight of 1 it it lies in the interior of the northeast or southwest quadrant; assign a weight of -1 it it lies in the interior of the northwest or southeast quadrant. It q lies elsewhere, assign it a weight of 0.

One might recognize in the weight assignment a modification of the standard two-dimensional Rademacher function. Let $g_k \in \mathbf{R}^N$ be the column vector of weights (with the coordinates in the same order as the corresponding rows of B). It is easily checked that the $\log n$ vectors g_k are orthogonal. Let G be the matrix whose columns are the g_k's and let u be the column vector of $\mathbf{R}^{\log n}$ whose coordinates are all ones. It follows that

$$\|Gu\|_2^2 = \sum_{k=1}^{\log n} \|g_k\|_2^2 \leq N \log n.$$

Summing separately over each k-good cell σ yields

$$g_k^\top Bx \gg \frac{N}{n} \sum_i \{\, x_i \in \text{ } k\text{-good cell of } X_k \,\}.$$

To see why, when summing up the weighted coordinates of Bx over each k-good cell, regroup each point with its three symmetric translates (one in each subquadrant) and apply the inclusion-exclusion formula. Finally, use the well-centeredness to argue for the presence of the factor N/n. We omit the details, which are straightforward. Since each $x_i > 0$ is well-centered for at least $(\log n)/3$ grids, then by (1)

$$(Gu)^\top Bx \gg \frac{N \log n}{n}\,\|x\|_1,$$

and by Cauchy-Schwarz,

$$\frac{N \log n}{n}\,\|x\|_1 \ll (Gu)^\top Bx \leq \|Gu\|_2 \cdot \|Bx\|_2$$

$$\leq \sqrt{N \log n}\,\|Bx\|_2.$$

□

We are now ready to complete the first step of our lower bound proof and estimate the eigenvalues of $B^\top B$ from below. Let $\mu_1 \geq \cdots \geq \mu_n \geq 0$ be the eigenvalues of $B^\top B$.

Lemma 2.3 *For any $1 \leq k \leq n$, the k-th largest eigenvalue of $B^\top B$ satisfies,*

$$\mu_k \gg \frac{(n-k+1)N \log n}{n^2}.$$

Proof: Let $\{v_i\}$ be an orthonormal eigenbasis for $B^\top B$, where v_i is associated with μ_i, and let F be the invariant subspace spanned by $v_k, \ldots, v_n$. By the Courant-Fischer theorem, we know that

$$\mu_k = \max_{0 \neq x \in F} \frac{\|Bx\|_2^2}{\|x\|_2^2}.$$

The difficulty is that in Lemma 2.2 the L^2 norm of Bx is bounded in terms of the L^1 norm of x. On the other hand, the inequality $\|x\|_1 \geq \|x\|_2$ is too weak for our purposes, so we argue (probabilistically) that the subspace F is "big" enough that it contains vectors whose L^1 and L^2 norms deviate from each other substantially. Specifically, we argue that the intersection of a large enough L^1-sphere with F must necessarily contain a short vector (in the L^2 sense). Of course, this will follow if the same is true of any one of the hyperplanes bounding the L^1-sphere in question. We show that a random bounding hyperplane satisfies this property.

Let ξ be the column vector obtained by expressing x in the basis $\{v_i\}$: we have $\xi = Qx$, where $Q = (q_{ij})$ is the orthogonal matrix whose rows are the eigenvectors v_i. Let $R = (r_{ij})$ be the matrix obtained by replacing each of the first $k-1$ rows of Q by a row of zeros. Now let $y = (y_1, \ldots, y_n)$ be a random vector chosen uniformly in $\{-1, 1\}^n$.

$$\begin{aligned}
\mathbf{E}\,\|Ry\|_2^2 &= \sum_{i=1}^n \mathbf{E}\left(\sum_{j=1}^n r_{ij} y_j\right)^2 \\
&= \sum_{i \geq k} \sum_{j=1}^n q_{ij}^2\, \mathbf{E}\, y_j^2 + \sum_{i \geq k} \sum_{j \neq j'} q_{ij} q_{ij'}\, \mathbf{E}\, y_j y_{j'} \\
&= \sum_{i \geq k} \sum_{j=1}^n q_{ij}^2 = n - k + 1.
\end{aligned}$$

This implies the existence of a vector $y \in \{-1,1\}^n$ such that $\|Ry\|_2^2 \geq n-k+1$, and therefore the $(n-k)$-flat defined by the equations,

$$\begin{cases} \xi_i = 0 & (1 \leq i < k) \\ (Qy)^\top \xi = \sqrt{n-k+1}, \end{cases}$$

cuts the unit-radius ball centered at the origin. Let x be a point of the intersection. Since $\xi = Qx$,

$$\|x\|_1 \geq y^\top x = (Qy)^\top \xi = \sqrt{n-k+1}.$$

Applying Lemma 2.2, we derive

$$\mu_k \geq \sum_{i=1}^{n} \mu_i \xi_i^2 = \|Bx\|_2^2 \gg \frac{N \log n}{n^2} \|x\|_1^2$$

$$\geq \frac{(n-k+1)N \log n}{n^2}.$$

□

The determinant of $B^\top B$ is the product of the eigenvalues, therefore

$$\det B^\top B \gg \Omega\Big(\frac{N \log n}{n^2}\Big)^n n! \,. \tag{2}$$

It is now easy to exhibit a hard set system A for orthogonal range searching and complete the second step of the lower bound proof.

Lemma 2.4 *There exists an $n \times n$ submatrix A of B such that*

$$\det A^\top A = (\log n)^{n-o(n)}.$$

Proof: By the Binet-Cauchy formula,

$$\det B^\top B = \sum_{1 \leq j_1 < \cdots < j_n \leq N} \left| \det B \begin{pmatrix} j_1 & j_2 & \cdots & j_n \\ 1 & 2 & \ldots & n \end{pmatrix} \right|^2.$$

Therefore by (2) there exists an $n \times n$ submatrix A of B such that

$$\begin{aligned} \det A^\top A &= \left| \det B \begin{pmatrix} j_1 & j_2 & \cdots & j_n \\ 1 & 2 & \ldots & n \end{pmatrix} \right|^2 \\ &\geq \binom{N}{n}^{-1} \det B^\top B \\ &= \Omega(1)^n \Big(\frac{n}{eN}\Big)^n \Big(\frac{n}{e}\Big)^n \Big(\frac{N \log n}{n^2}\Big)^n \\ &\geq (\log n)^{n-o(n)}. \end{aligned}$$

□

We can now move on to the last step of the lower bound proof. Let $\lambda_1 \geq \cdots \geq \lambda_n \geq 0$ be the eigenvalues of $A^\top A$.

Lemma 2.5 *For any fixed $\varepsilon > 0$ and any $k \leq n/4 - \varepsilon n$, we have $\lambda_k = (\log n)^{\Omega(1)}$.*

Proof: The intuition is this: write down some linear constraints which, if satisfied, allow us to express the map A by a matrix with only few ones. Then, by using standard matrix norm inequalities, argue that within the subspace satisfying the constraints, $\|Ax\|_2/\|x\|_2$ is always small. The variational characterization of eigenvalues allows us to conclude. We flesh out these ideas below.

Place the points of P in bijection with the leaves of a complete binary tree (from left to right): each node v of the tree is associated with the vertically sorted list N_v of the points stored at the leaves at or below v. This is a classical range tree construction [14]. Any set specified by a row of A can be partitioned into fewer than $\nu = \log n + 1$ subsets: each one is a prefix of a list N_v and all the relevant lists N_v are on different levels of the tree. Thus, we can create ν set systems, one for each level, such that their $n \times n$ incidence matrices $A_1, \ldots, A_\nu$ satisfy $A = \sum_i A_i$.

For every level i, perform the following operations. First, check whether some of the rows in A_i are identical: ρ identical rows all correspond to the same prefix in some N_v. Take the last element in the prefix (i.e., the point with highest y-coordinate) and duplicate it $\rho - 1$ times within the list N_v. Note that the total size of the augmented lists at level i is at most $2n$. Next, consider each list N_v (after the previous preprocessing) and subdivide it into contiguous lists of r (or fewer) points; r is a parameter to be specified later.

To summarize, at each level we have a collection of at most $2n/r$ lists of size exactly r along with a number of other lists of size less than r. We can now remove all the duplicates, as their presence was needed only to calibrate the size of the lists. (Note that in the process some lists might become empty.) Any subset specified by a row of A_i can be written as a union of full lists and a *remainder* set of size less than r. Thus, if H is the matrix whose rows are the characteristic vectors of each full list (over *all* A_i), then the restriction of A_i (viewed as a linear transformation) to $\operatorname{Ker} H$ can be expressed by an n-by-n matrix C_i, whose rows have each fewer than r ones. Because of the earlier duplication of points in the lists, note that similarly no column of C_i can have more than r ones. The rank of H is at most $2\nu n/r$, therefore

$$\operatorname{codim} \operatorname{Ker} H \leq \frac{2\nu n}{r}. \tag{3}$$

It is a standard result in matrix theory [12] that the spectral norm of a matrix M satisfies

$$\|M\|_s^2 \leq \Big(\max_i \sum_j |m_{ij}|\Big)\Big(\max_j \sum_i |m_{ij}|\Big),$$

and therefore, $\|C_i\|_s^2 \le r^2$. Using the fact that

$$\|C_i x\|_2 \le \|C_i\|_s \cdot \|x\|_2 \le r\|x\|_2,$$

we find that, for any $x \in \operatorname{Ker} H$, by the triangular inequality,

$$\|Ax\|_2 \le \sum_{i=1}^{\nu} \|A_i x\|_2 = \sum_{i=1}^{\nu} \|C_i x\|_2 \le r\nu\|x\|_2.$$

By the Courant-Fischer characterization of eigenvalues, we know that

$$\lambda_j = \min_F \max_{0 \ne x \in F} \frac{\|Ax\|_2^2}{\|x\|_2^2},$$

where the minimum is taken over all subspaces of dimension $n-j+1$. Fix j such that $\nu \le j \le n$; by (3), setting $r = \lceil 2\nu n/(j-1) \rceil$ makes the dimension of $\operatorname{Ker} H$ exceed $n-j$, therefore, for n large enough,

$$\lambda_j \le \max_{0 \ne x \in \operatorname{Ker} H} \frac{\|Ax\|_2^2}{\|x\|_2^2} \le r^2\nu^2 \le 5\Big(\frac{n \log^2 n}{j}\Big)^2.$$

Note that this inequality remains valid if $1 \le j < \nu$, because of the trivial upper bound, $\lambda_j \le n^2$. By Lemma 2.4,

$$\begin{aligned}(n-k+1)\log\lambda_k &\ge \log\det A^{\mathsf{T}}A - \sum_{j=1}^{k-1}\log\lambda_j \\ &\ge (n-o(n))\log\log n \\ &\quad - k(O(1)+2\log n - 2\log k \\ &\quad + 4\log\log n).\end{aligned}$$

Choosing $k = n/4 - \varepsilon n$, for any fixed $\varepsilon > 0$, gives $\lambda_k = (\log n)^{\Omega(1)}$. □

The lower bound for range searching follows directly from the spectral lemma, which proves Theorem 2.1. Note that over the reals the problem can be solved entirely with only $2n-1$ free computations: in $n-1$ steps encode the vector $(x_1, \ldots, x_n)$ into a real number. Then in another n steps, answer the n range queries. □

3 The Semigroup Model

In the semigroup version of range searching, subtractions are not allowed (or not defined). In other words, the model of computation is a straight-line program or circuit: each (charged) step is the form,

$$z \leftarrow x + y,$$

where x and y are previously computed variables or input weights. It is possible to weaken the model slightly (and hence, strengthen our results) by relaxing the assumption that the program should work for all commutative semigroups. It suffices to require that it should work for at least one faithful semigroup. This is a semigroup over which any two identically equal linear forms must have the same variables (though not necessarily the same coefficients) – see [3, 4, 18] for formal definitions.

Theorem 3.1 *Range searching with respect to n points and n axis-parallel boxes in $\mathbf{R}^d$ requires on the order of $n(\log n/\log\log n)^{d-1}$ semigroup operations.*

Theorem 3.2 *Range searching with respect to n points and n simplices in $\mathbf{R}^d$ requires on the order of $n^{2-2/(d+1)}(\log n)^{-5/2}$.*

Complicated arguments were used in [3, 4] to treat the on-line cases. As it turns out, the off-line case is considerably easier. The bound of Theorem 3.1 is fairly close to the known upper bound of $n(\log n)^{d-1}\alpha(n)$ [7]. The bound of Theorem 3.2 is within a polylogarithmic factor of the best current upper bound of $n^{2-2/(d+1)}(\log n)^{O(1)}$ [13].

The proofs of both theorems are based on a simple graph-theoretical lemma. Let $A = (a_{ij})$ denote the $n \times n$ incidence matrix of a range searching problem. Suppose that A has no $p \times q$ submatrix of ones. An equivalent formulation is to say that the corresponding bipartite graph has no (p, q) complete bipartite subgraph.

Lemma 3.3 *If A is an $n \times n$ incidence matrix with no $p \times q$ submatrix of ones, then the complexity of computing Ax over a semigroup is at least on the order of*

$$\frac{1}{pq}\Big(\sum_{i,j} a_{ij}\Big) - \frac{n}{p}.$$

Proof: Every gate of the circuit adds two linear forms together: we say that the gate is *heavy* if the linear form it outputs, $\sum_j \alpha_j x_j$ ($\alpha_j \in \mathbf{N}^+$), involves q variables x_j or more. Given a row i, let $A_i = \sum_j a_{ij}$. Because of faithfulness, the output gate g computing the form $\sum_j a_{ij}x_j$ is connected to A_i input variables x_{i_1}, x_{i_2}, etc. Consider a subtree T_i of the circuit graph with g as its root and $x_{i_1}, x_{i_2}, \ldots$ at its leaves. Note that the maximal subtrees of T_i with at most q leaves (i.e., those whose root's parent has more than q descending leaves) are disjoint. Each such subtree has one fewer 2-child nodes than leaves. There are at least A_i/q such subtrees, so they account for at most $A_i - A_i/q$ 2-child nodes. This shows that at least $A_i/q - 1$ internal nodes of T_i correspond to heavy gates. By faithfulness again, no heavy gate can provide a node for p trees T_i. Indeed, this would create a $p \times q$ submatrix of ones in A. The lower bound follows immediately. □

The Proof of Theorem 3.1

The set of input points is obtained from a *Halton-Hammersley* sequence [2]. Let $p_1 < p_2 < \cdots < p_{d-1}$ be consecutive primes. Any integer m has a unique decomposition in base p_k: $m = \sum_{i \geq 0} b_k(i) p_k^i$. We define the function,

$$x_k(m) = \sum_{i \geq 0} \frac{b_k(i)}{p_k^{i+1}}.$$

This allows us to construct the input point set:

$$P = \left\{ \left(x_1(m), \ldots, x_{d-1}(m), m/n \right) \; : \; 0 \leq m < n \right\}.$$

Any interval of the form $[M/p_k^j, (M+1)/p_k^j)$, where M is a nonnegative integer, is said to be of *type* (k, j). A box B is *special* if it is of the form $I_1 \times \cdots \times I_d$, where

- $B \subseteq [0,1]^d$;
- $I_1, \ldots, I_{d-1}$ are intervals of type $(1, j_1), \ldots, (d-1, j_{d-1})$, respectively, for some integers $j_1, \ldots, j_{d-1} \geq 0$;
- I_d is of the form $[Mp_1Q/n, (M+1)p_1Q/n)$, where $Q = p_1^{j_1} \cdots p_{d-1}^{j_{d-1}}$ and M is an integer.

We motivate this definition. First, observe that knowing which (k, j)-interval contains $x_k(m)$ amounts to the knowledge of the digits $b_k(0), \ldots, b_k(j-1)$. Thus, if we know that the point of P indexed by m lies in the box $I_1 \times \cdots \times I_{d-1} \times [0,1]$, then we know the residues classes of m modulo $p_1^{j_1}, \ldots, p_{d-1}^{j_{d-1}}$, respectively. By the Chinese Remainder Theorem, this implies that we know m modulo Q. It follows at once that each of the boxes $I_1 \times \cdots \times I_{d-1} \times [lQ/n, (l+1)Q/n)$ contains at most one point of P. Suppose that

$$p_1 \prod_{0<k<d} p_k^{j_k} \leq n. \tag{4}$$

Because $p_k^{j_k} \leq n$, for any $0 < k < d$, any special box B that satisfies (4) contains exactly p_1 points of P. The number N of such boxes is equal to

$$\begin{aligned} N &= \sum_{j_1, \ldots, j_{d-1} \geq 0} \prod_{0<k<d} p_k^{j_k} \left\lfloor \frac{n}{p_1 \prod_{0<k<d} p_k^{j_k}} \right\rfloor \\ &\geq \sum_{p_1^{j_1} \times \cdots \times p_{d-1}^{j_{d-1}} \leq n/2p_1} \left(\frac{n}{p_1} - \prod_{0<k<d} p_k^{j_k} \right) \\ &\geq \sum_{0 \leq m \leq \log n / \log p_{d-1} - 2} \; \sum_{j_1 + \cdots + j_{d-1} = m} \frac{n}{2p_1}. \end{aligned}$$

It follows that

$$N \gg \frac{n}{p_1} \left(\frac{\log n}{\log p_{d-1}} \right)^{d-1}.$$

By choosing p_1 to be around $(\log n)^{d-1}$, we can find n boxes that define a set system whose incidence matrix A has at least $n(\log n / \log p_{d-1})^{d-1}$ ones (note that we may have to pad with rows of zeroes). We will now show that A is square-free (i.e., has no 2×2 submatrix of ones).

Consider the intersection of the special box B with another special box B' with parameters $(j_1', \ldots, j_{d-1}')$. Without loss of generality, assume that $Q < Q'$. (Note that the case $Q = Q'$ corresponds to an empty intersection.) The intersection of two intervals of type (k, j) and (k, j'), $j' < j$, is either empty or an interval of type (k, j). This implies that $B \cap B'$ is a box $J_1 \times \cdots \times J_d$, where J_k $(k < d)$ is an interval of type $(k, \max\{j_k, j_k'\})$, and J_d has length at most $p_1 Q/n$. Assume that the box $B \cap B'$ contains a point of P and let m be its index. By the Chinese Remainder Theorem, m is completely specified modulo $\prod p_k^{\max\{j_k, j_k'\}}$, and hence modulo $p_i Q$, for some $0 < i < d$. We know that the point's d-th coordinate m/n lies in an interval J_d of length at most $p_1 Q/n$, therefore m is uniquely determined. Thus, two special boxes intersect in at most one point of P. It follows that A is square-free, and by Lemma 3.3, the proof of Theorem 3.1 is complete. □

The Proof of Theorem 3.2

We exhibit a set P of n points in $\mathbf{R}^d$ along with a collection $\{S_q\}$ of n slabs, such that: (i) each slab contains roughly $n^{1-2/(d+1)}$ points, and (ii) the intersection of any $k \geq \log n$ slabs contains at most a logarithmic number of points. The proof technique is similar to [7]; there are some differences, however, so we provide the details below.

Let H_q be the hyperplane of equation $\langle p, q \rangle - \|q\|_2^2 = 0$: this is the hyperplane normal to Oq passing through q. Fix a parameter w; we let S_q denote the slab of width w consisting of all the points at most $w/2$ away from H_q. To specify the collection of slabs $\{S_q\}$, it thus suffices to provide a set Q of n points q. First, we show that if any d of the points of Q are sufficiently spread apart (in a sense to be formalized below) then the corresponding slabs have a small intersection. Next, by appealing to the results on Heilbronn's problem of [3], we are able to exhibit a suitable set Q. Finally, throwing in a set of random points in the unit cube provides P and completes the construction of the set system. We now give the details of the construction of Q. (We use the notation $\mathrm{vol}_d(A)$ and $\mathrm{conv}(A)$ to refer to the d-dimensional volume and the convex hull of A, respectively.)

Lemma 3.4 *Let $q_1, \ldots, q_d$ be d points in $[0,1]^d$, and assume that the central projection q_i' of each q_i on the*

hyperplane $x_1 = 1$ also lies in $[0,1]^d$. Then

$$\mathrm{vol}_d \bigcap_{i=1}^{d} S_{q_i} \ll w^d \Big/ \mathrm{vol}_{d-1}\Big(\mathrm{conv}\,\{q'_1,\ldots,q'_d\}\Big).$$

Proof: Let $[u_1,\ldots,u_d]$ denote the matrix whose columns are the vectors u_i spanning the parallelepiped $\bigcap S_{q_i}$. Note that each u_i has the direction specified by the intersection of hyperplanes bounding the slabs S_{q_j}, for all $j \neq i$. We easily derive

$$\det\Big([u_1,\ldots,u_d]^\top [q_1,\ldots,q_d]\Big) = w^d \prod_i \|q_i\|_2,$$

and therefore

$$\begin{aligned}
\mathrm{vol}_d \bigcap_{i=1}^{d} S_{q_i} &= w^d\Big(\prod_i \|q_i\|_2\Big)/|\det[q_1,\ldots,q_d]| \\
&= w^d\Big(\prod_i \|q'_i\|_2\Big)/|\det[q'_1,\ldots,q'_d]| \\
&\ll w^d/|\det[q'_1,\ldots,q'_d]|,
\end{aligned}$$

from which the lemma easily follows. □

Choose an integer $m = \lfloor c_0 nw \rfloor$, for some constant $c_0 > 0$. By [3] (Theorem 4.10), we can place m points in $(1,0,\ldots,0)+[0,1]^{d-1}$ so that the convex hull of any $k \geq \log m$ points has $(d-1)$-dimensional volume at least $\Omega(k/m)$. For each such point q', place points on the segment Oq' at intervals of length w. This gives us $O(c_0 n)$ points q: if w is small enough then, for at least a constant fraction of them, the slab S_q intersects the cube $[0,1]^d$ in a polytope of volume $\Omega(w)$. By choosing c_0 large enough, we can find n such points, which thus form the set Q. To summarize, the set Q consists of n points such that for any $q \in Q$,

$$\mathrm{vol}_d\, S_q \cap [0,1]^d \gg w. \tag{5}$$

Also, for any $q_1,\ldots,q_k \in Q$, with $k \geq \log n$, either at least two q_i's have the same central projection q'_i, in which case $\mathrm{vol}_d \bigcap_{i=1}^{k} S_{q_i} = 0$, or else

$$\mathrm{vol}_{d-1}\,(\mathrm{conv}\,\{q'_1,\ldots,q'_k\}) \gg \frac{k}{m}.$$

By triangulating the convex hull of $q'_1,\ldots,q'_k$, using $O(k^{\lfloor (d-1)/2 \rfloor})$ simplices, we derive the existence of d points, say, $q'_1,\ldots,q'_d$, whose convex hull has volume at least $\Omega(k^{2-\lceil d/2 \rceil}/m)$. By Lemma 3.4, this shows that in all cases

$$\mathrm{vol}_d \bigcap_{i=1}^{k} S_{q_i} \leq \mathrm{vol}_d \bigcap_{i=1}^{d} S_{q_i} \ll nw^{d+1}k^{\lceil d/2 \rceil - 2}. \tag{6}$$

We define the point set P by choosing n points in $[0,1]^d$ at random uniformly and independently. Let $w = bn^{-2/(d+1)}(\log n)^c$, for some fixed $b > 0$ and $c = (3 - \lceil d/2 \rceil)/(d+1)$. Set $k = \lceil \log n \rceil$; we can ensure that k-wise intersections of slabs contain only $O(\log n)$ points. We use standard Chernoff bounds below; see e.g., [1]. By (6) there exists a constant $c_1 > 0$ such that the probability that $\bigcap_{i=1}^{k} S_{q_i}$ contains more than $c_1 n^2 w^{d+1} k^{\lceil d/2 \rceil - 2} = \Theta(\log n)$ points is less than $e^{-\Omega(b^{d+1}\log n)}$, which is less than n^{-d} for b large enough. Thus, with probability greater than $1 - \binom{n}{k} n^{-d}$, no k-wise intersection of slabs contains more than $O(\log n)$ points. Note that by (6) the factor of $\binom{n}{k}$ can be replaced by $\binom{n}{d}$, so the probability is actually greater than $1/2$.

Similarly, by (5) the probability that a given S_q contains fewer than $c_2 wn$ points is less than $e^{-\Omega(wn)} < e^{-n^{1/4}}$, for some constant $c_2 > 0$. So, with probability greater than $1/2$, all the slabs S_q contain $c_2 wn$ points or more. We derive the existence of a set P such that every slab S_q ($q \in Q$) contains $\Omega(wn)$ points and no subset of $\lceil \log n \rceil$ slabs has an intersection containing more than $O(\log n)$ points.

The set P and the slabs S_q form a set system A with no $p \times q$ submatrix of ones, where $p = \lceil \log n \rceil$ and $q = \alpha \log n$, for some constant α large enough. It follows from Lemma 3.3 that the semigroup complexity of the map $x \mapsto Ax$ is $\Omega(wn^2/\log^2 n)$. Since the constant c is at least $-1/2$, this bound is at least $\Omega(n^{2-2/(d+1)}(\log n)^{-5/2})$, which proves Theorem 3.2. Note that the polylogarithmic factor can be tightened a little by keeping the exact value of c. □

4 Concluding Remarks

One of the most intriging open problems is, of course, to improve the $\Omega(n \log\log n)$ lower bound, or to extend it to rings or fields. A more modest goal might be to increase the number of allowable help gates to, say, $(1-\varepsilon)n$. Also, the discrepancy theory literature is vast and rich in powerful mathematical techniques. To establish further links to complexity theory would be very interesting.

References

[1] Alon, N., Spencer, J. *The Probabilistic Method*, John Wiley & Sons, New York, 1992.

[2] Beck, J., Chen, W.W.L. *Irregularities of distribution*, Cambridge Tracts in Mathematics, 89, Cambridge Univ. Press, Cambridge, 1987.

[3] Chazelle, B. *Lower bounds on the complexity of polytope range searching*, J. Amer. Math. Soc., 2 (1989), 637–666.

[4] Chazelle, B. *Lower bounds for orthogonal range searching: II. the arithmetic model,* J. ACM, 37 (1990), 200–212 and 439–463.

[5] Chazelle, B. *A spectral approach to lower bounds,* Proc. 35th Ann. IEEE Symp. Foundat. Sci. (1994), 674–682.

[6] Chazelle, B. *Computational geometry: A retrospective,* Proc. 26th Ann. ACM Symp. Theory Comput. (1994), 75–94.

[7] Chazelle, B., Rosenberg, B. *Lower bounds on the complexity of simplex range reporting on a pointer machine,* Proc. 19th ICALP, LNCS 623, Springer-Verlag (1992), 439–449.

[8] Edelsbrunner, H. *Algorithms in Combinatorial Geometry,* Springer-Verlag, 1987.

[9] Erickson, J. *New lower bounds for Hopcroft's problem,* to appear in Proc. 11th Ann. ACM Symp. Comput. Geom., 1995.

[10] Fredman, M.L. *A lower bound on the complexity of orthogonal range queries,* J. ACM, 28 (1981), 696–705.

[11] Fredman, M.L. *Lower bounds on the complexity of some optimal data structures,* SIAM J. Comput., 10 (1981), 1–10.

[12] Lancaster, P., Tismenetsky, M. *The Theory of Matrices,* 2nd ed., Academic Press, 1985.

[13] Matoušek, J. *Geometric range searching,* Tech. Report B-93-09, Free Univ. Berlin, 1993.

[14] Mehlhorn, K. *Data Structures and Algorithms 3: Multidimensional Searching and Computational Geometry,* Springer-Verlag, Heidelberg, Germany, 1984.

[15] Morgenstern, J. *Note on a lower bound of the linear complexity of the fast Fourier transform,* J. ACM, 20 (1973), 305–306.

[16] Mulmuley, K. *Computational Geometry: An Introduction Through Randomized Algorithms,* Prentice-Hall, 1994.

[17] Roth, K.F. *On irregularities of distribution,* Mathematika, 1 (1954), 73–79.

[18] Yao, A.C. *On the complexity of maintaining partial sums,* SIAM J. Comput., 14 (1985), 277–288.

[19] Yao, F.F. *Computational Geometry,* in: Algorithms and Complexity, Handbook of Theoretical Computer Science, ed. J. van Leeuwen, Vol. A, Elsevier (1990), 343–389.