

A SPECTRAL APPROACH TO LOWER BOUNDS WITH APPLICATIONS TO GEOMETRIC SEARCHING*

BERNARD CHAZELLE†

Abstract. We establish a nonlinear lower bound for halfplane range searching over a group. Specifically, we show that summing up the weights of n (weighted) points within n halfplanes requires $\Omega(n \log n)$ additions and subtractions. This is the first nontrivial lower bound for range searching over a group. By contrast, range searching over a semigroup (which forbids subtractions) is almost completely understood.

Our proof has two parts. First, we develop a general, entropy-based method for relating the linear circuit complexity of a linear map A to the spectrum of $A^T A$. In the second part of the proof, we design a “high-spectrum” geometric set system for halfplane range searching and, using techniques from discrepancy theory, we estimate the median eigenvalue of its associated map. Interestingly, the method also shows that using up to a linear number of help gates cannot help; these are gates that can compute *any* bivariate function.

Key words. lower bounds, eigenvalues, range searching, circuit complexity

AMS subject classifications. 68P05, 68Q20, 68R99, 51M99

PII. S0097539794275665

1. Introduction. Given n weighted points in the plane and n halfplanes, we consider the classical *halfplane range searching* problem, which is to compute the sum of the weights of the points within each of the given regions. If subtractions are not allowed (the semigroup model) the problem is almost completely solved [7, 11, 15]; see also [9, 16, 17, 19] for surveys of the vast literature on the subject. In the (commutative) group model, where subtractions are allowed, there is little evidence that any power should be gained beyond polylog speedups, but proving it has been elusive. In fact, in that model no superlinear lower bound has ever been established for any range searching problem of any kind. The problem is equivalent to asking for the nonmonotone circuit complexity of some fairly unwieldy linear transformation over the reals, so the lack of progress should not come as a big surprise.

This paper takes a first, modest step toward resolving this question. We establish a lower bound of $\Omega(n \log n)$ on the complexity of range searching with respect to n points and n halfplanes (given in advance). The model of computation is a straight-line program: each step performs a group operation of the form

$$z \leftarrow x \pm y,$$

where x and y are previously computed variables or input weights. The underlying group is assumed to be commutative. Note that it is easy to prove an $\Omega(n \log n)$ lower bound by reduction from sorting, but this says nothing about the number of times weights have to be added or subtracted. In the group model, memory accesses are not

*Received by the editors October 17, 1994; accepted for publication (in revised form) March 20, 1996. A preliminary version of this paper appeared as *A spectral approach to lower bounds*, in Proceedings of the 35th Annual IEEE Symposium on the Foundations of Computer Science (FOCS), IEEE Computer Society Press, Los Alamitos, CA, 1994, pp. 674–682.

<http://www.siam.org/journals/sicomp/27-2/27566.html>

†Department of Computer Science, Princeton University, Princeton, NJ 08544 (chazelle@cs.princeton.edu). This research was supported in part by NSF grant CCR-93-01254 and by The Geometry Center, University of Minnesota, a Science and Technology Center funded by the NSF, DOE, and Minnesota Technology, Inc.

charged; only group operations are. Note that this makes lower bounds even stronger. The program must work for all groups and all weight assignments.

This formulation of the group model is probably the most natural. But one might ask what happens if we extend the model by allowing extra computations free of charge. For example, can adding the same variable a large number of times, e.g., $z \leftarrow Mx$ (M integral), possibly help? How about encoding special functions in lookup tables? In general, we call a *free computation* any assignment of the form

$$z \leftarrow f(x, y),$$

where f is an arbitrary function. We show that allowing close to $n/4$ free computations cannot help. To put this result in perspective, one should note that over the reals $2n$ free computations suffice to make the problem trivial.

THEOREM 1.1. *Range searching with respect to n points and n halfplanes requires $\Omega(n \log n)$ group operations. This remains true even with the help of up to $n/4 - \varepsilon n$ free computations, for any fixed $\varepsilon > 0$. On the other hand, over the reals, the problem can be solved in linear time with only $2n - 1$ free computations.*

It is likely that the lower bound is far from optimal. The best known upper bound is slightly above $O(n^{4/3})$ [15], and in the semigroup model the best lower bound (for the on-line version) is also $\Omega(n^{4/3})$ [7]. On the brighter side, Theorem 1.1 provides the only lower bound known for the group model, so at least it is a step in the right direction.

Proof. The proof consists of two distinct parts. First, we establish a general *spectral lemma*, which asserts a lower bound of $\Omega((k - 2m) \log \lambda_k)$ on the linear circuit complexity of any linear transformation A (with integer coefficients) from \mathbf{R}^n to \mathbf{R}^n , where λ_k is the k th largest eigenvalue of $A^T A$, and m is the number of help gates. These are the circuit equivalent of free computations. (This shows that allowing up to roughly $k/2$ help gates does no good.) The lower bound holds for any value of k between $2m$ and n . This freedom is useful because often only a small range of the whole spectrum can be accurately estimated without too much effort.

In the second part of the proof, we design a *hard* instance of range searching by nonconstructive means. Then we use spectral methods from discrepancy theory to estimate the median eigenvalue of the quadratic form associated with the corresponding set system.

Remark 1. Our technique trivially implies an $\Omega(n \log n)$ lower bound for range searching over a finite projective plane. In general, the technique will yield a lower bound on any instance of range searching whose corresponding spectrum can be mapped out reasonably well. Powerful techniques in discrepancy theory [4], such as those used in section 3 of this paper, raise hope that more lower bounds can be derived by this approach.

Remark 2. A simple application of the spectral lemma is that computing Hx , where H is the $n \times n$ Sylvester Boolean matrix, takes $\Omega(n \log n)$ time even in the presence of about $n/2$ help gates. (All the eigenvalues λ_k of $H^T H$ are equal to n .) Note that the choice of ground field is crucial, since Hx can be computed in linear time over $\text{GF}(2)$ [2]. If we forbid help gates, the same bound can be obtained more simply by using Morgenstern's volume argument [18]. The spectral lemma works in a model that is ideally suited for range searching. If, instead of a group, the linear transformation operates over a ring or a field (like the discrete Fourier transform), then for the lemma to hold, the (nonhelp) gates must evaluate linear forms with bounded coefficients. This is the same limitation found in Morgenstern's result. Help gates can be thought of as a way of partly overcoming this limitation.

Remark 3. The proof of the spectral lemma is based on entropy considerations. By avoiding standard volume arguments, we are able to accommodate help gates. Indeed, a weakness of the volume argument of [18] is that it collapses even in the presence of a single help gate. Intuitively, the idea of that argument is to relate the work of the circuit to the volume of the ellipsoid into which the circuit “transforms” the unit sphere. Any such argument is vulnerable to even a single help gate, because any one of them has the ability to blow up the entire sphere. On the contrary, our entropy-based approach ensures that the “contribution” of a single help gate to the work of the circuit is always bounded, regardless of the gate’s power.

There has been a substantial amount of work in arithmetic circuit complexity—see surveys in [13, 20]—but, to our knowledge, nothing that allows us to tackle a geometric problem such as range searching. Most of the recent research in circuit complexity [5], including work involving help bits or oracle queries (variants of our help gates) [3, 6], has been mostly concerned with problems over finite fields and seems of little relevance to our problem.

2. Eigenvalues, entropy, and linear circuits. Let A be an $n \times n$ matrix with integer elements. A linear circuit for computing $y = Ax$, where $x \in \mathbf{R}^n$, is a directed acyclic graph with n input nodes $x = (x_1, \dots, x_n)$ and n output nodes $y = (y_1, \dots, y_n)$. The size of the circuit is its number of edges. A node is a gate that computes a real-valued function

$$f(z_1, z_2) = \alpha_1 z_1 + \alpha_2 z_2,$$

where $z_i \in \mathbf{R}$ and $\alpha_i \in \{-1, 0, 1\}$. In addition, we allow the presence of m help gates: these gates can evaluate *any* function $f(z_1, z_2)$ from \mathbf{R}^2 to \mathbf{R} . Recall that the matrix $M \stackrel{\text{def}}{=} A^\top A$ is diagonalizable and its eigenvalues are real:

$$\lambda_1 \geq \dots \geq \lambda_n \geq 0.$$

All logarithms are to the base 2.

SPECTRAL LEMMA. *Given any $1 \leq k \leq n$, any circuit for computing Ax has size at least $c(k - 2m) \log \lambda_k$, for some constant $c > 0$, where $m \leq k/2$ is the number of help gates and λ_k is the k th largest eigenvalue of $A^\top A$.*

Proof. Let K be the invariant subspace spanned by the eigenvectors for M associated with $\lambda_1, \dots, \lambda_k$. We ensure that K is of dimension exactly k by dropping some of the eigenvectors for λ_k , in case of multiplicity. Let $B_n(p, r)$ denote the Euclidean n -ball of radius r centered at p and let $V_n(r)$ be its volume. Consider the cubes of the form $\mathbf{Z}^n + [0, 1]^n$ that intersect $K \cap B_n(O, R)$, for some (large enough) parameter R . Let L be the set of centers of these cubes. Finally, let $H(x)$ denote the entropy of a random variable x with values distributed uniformly in L . We estimate the entropy of x (Lemma 2.1) and then we prove the key lemma, which says that if the spectrum of M is not too low we can suitably hash the image of x under A without losing much entropy (Lemma 2.2).

LEMMA 2.1.

$$H(x) \geq \log V_k(R) - \log V_k(\sqrt{n}).$$

Proof. Let $M_R^-(n, k)$ denote the minimum number, over all k -flats F containing the origin, of the n -cubes of the form $\mathbf{Z}^n + [0, 1]^n$ that intersect $F \cap B_n(O, R)$. Observe that the cubes to be counted cover the ball $B_k(O, R)$ embedded in F . Furthermore,

the intersection of F with each of these cubes fits into a k -ball of radius $\sqrt{n}/2$, so

$$M_R^-(n, k) \geq V_k(R)/V_k(\sqrt{n}).$$

The lemma follows from the fact that $|L| \geq M_R^-(n, k)$ and $H(x) = \log |L|$. \square

Another useful quantity, denoted by $M_r^+(n, m)$, is the maximum number over all m -flats F of the n -cubes of the form $\mathbf{Z}^n + [0, 1]^n$ that intersect $F \cap B_n(O, r)$. In the case $n - m = 0$, the cubes counted by $M_r^+(n, m)$ all lie within $B_n(O, r + \sqrt{n})$; therefore

$$(1) \quad M_r^+(n, n) \leq V_n(r + \sqrt{n}).$$

Assume now that $n - m > 0$. In the appendix we show that, for n and r large enough,

$$(2) \quad M_r^+(n, m) \leq 3^n V_m(r).$$

LEMMA 2.2. *If A is a matrix with real elements, such that $1 \leq \lambda_k \leq 2$, then¹*

$$H(\lfloor Ax \rfloor) \geq H(x) - \log V_n(5\sqrt{n}).$$

Proof. Let $x, x' \in L$ be such that $\lfloor Ax \rfloor = \lfloor Ax' \rfloor$. It follows that $\|A(x - x')\|_2 \leq \sqrt{n}$. Write x as the direct sum $x_0 + u$, where $x_0 \in K$, $u \in K^\perp$, and do the same with x' . Observe that

$$\|u - u'\|_2 \leq \|u\|_2 + \|u'\|_2 \leq \sqrt{n}.$$

By the variational characterization of eigenvalues, $\|A(x_0 - x'_0)\|_2 \geq \sqrt{\lambda_k} \|x_0 - x'_0\|_2$ and, because K^\perp contains $u - u'$ and is spanned by eigenvectors corresponding to $\lambda_j \leq \lambda_k$, we have $\|A(u - u')\|_2 \leq \sqrt{\lambda_k} \|u - u'\|_2$. It follows that (for $1 \leq \lambda_k \leq 2$)

$$\begin{aligned} \|x - x'\|_2 &\leq \|x_0 - x'_0\|_2 + \|u - u'\|_2 \\ &\leq \|A(x_0 - x'_0)\|_2 + \|u - u'\|_2 \\ &\leq \|A(x - x')\|_2 + \|A(u - u')\|_2 + \|u - u'\|_2 \\ &\leq \sqrt{n} + 3\|u - u'\|_2 \leq 4\sqrt{n}. \end{aligned}$$

Thus, the preimage of a fixed $z \in \mathbf{R}^n$ under $x \in L \mapsto \lfloor Ax \rfloor$ lies entirely in a ball $B_n(x, 4\sqrt{n})$, where $x \in L$, and therefore, the uniform distribution within that preimage has entropy at most $\log M_{4\sqrt{n}}^+(n, n)$. By (1) this does not exceed $\log V_n(5\sqrt{n})$. Standard identities on the entropy of joint distributions, namely,

$$H(x) = H(x, \lfloor Ax \rfloor) = H(\lfloor Ax \rfloor) + H(x | \lfloor Ax \rfloor),$$

complete the proof. \square

Let $z = (z_1, \dots, z_s)$ be the vector of \mathbf{R}^n whose coordinates are the intermediate variables computed by the gates. For convenience, we append the input variables at the beginning of the list ($z_j = x_j$, for $1 \leq j \leq n$) and the output variables at the end ($z_{s-n+j} = y_j$, for $1 \leq j \leq n$). We also assume that the list corresponds to a topological ordering of the DAG (directed acyclic graph), meaning that for any $j > n$,

$$z_j = \alpha_j z_{f(j)} + \beta_j z_{g(j)},$$

¹Given $z = (z_1, \dots, z_n) \in \mathbf{R}^n$, we use the shorthand $\lfloor z \rfloor$ for $(\lfloor z_1 \rfloor, \dots, \lfloor z_n \rfloor)$.

where $f(j) \leq g(j) < j$ and $|\alpha_j|, |\beta_j| \leq 1$. In the case of a help gate, z_j is an arbitrary real function of $z_{f(j)}$ and $z_{g(j)}$. Let

$$\mu_k = \lfloor \sqrt{\lambda_k} \rfloor.$$

The input x to the circuit is chosen so that $\tilde{x} = \mu_k x$ is a random variable uniformly distributed in L . We shall now assume that λ_k is large enough, which the spectral lemma obviously allows us to do. We now argue that $\lfloor z \rfloor$ has high entropy. Lemmas 2.4–2.6 will then show that only a large circuit can produce a “hashed” vector $\lfloor z \rfloor$ with that much entropy.

LEMMA 2.3.

$$H(\lfloor z \rfloor) \geq \log V_k(R) - \log V_k(\sqrt{n}) - \log V_n(5\sqrt{n}).$$

Proof. Because A is a linear map, the circuit outputs $B\tilde{x}$, where $B \stackrel{\text{def}}{=} A/\mu_k$, obviously satisfies the conditions of Lemma 2.2. Thus,

$$H(\lfloor z \rfloor) \geq H(\lfloor B\tilde{x} \rfloor) \geq H(\tilde{x}) - \log V_n(5\sqrt{n}).$$

The proof now follows from Lemma 2.1. \square

To be able to isolate the action of help gates, we devise the following artifice. Regard the outputs of the help gates, $z_{h(1)}, \dots, z_{h(m)}$, as new (help) variables, and express each z_j ($1 \leq j \leq s$) as a linear form over the set of variables $Z = \{x_1, \dots, x_n, z_{h(1)}, \dots, z_{h(m)}\}$. How this is done is best seen by induction. The input gates are linear forms over single variables. For any other gate z_j , if it is of the helping type, then the form is z_j itself. Otherwise, $z_j = \alpha_j z_{f(j)} + \beta_j z_{g(j)}$ and, by induction, z_j is a linear combination of two linear forms over Z and hence a linear form itself. (Note that even though the outputs of the help gates might be algebraically related, the $z_{h(j)}$ are added by adjunction and so are considered independent.) In the expression for z_j , let z_j^x (resp., z_j^y) denote the linear form obtained by taking only the nonhelp (resp., help) variables. A key step now is to look at z_j^x and z_j^y no longer as linear forms but as real functions $z_j^x(x_1, \dots, x_n)$ and $z_j^y(x_1, \dots, x_n)$. While the circuit computes $z_j = z_j^x + z_j^y$, we wish to monitor the information contents of the complex number,²

$$z_j^c \stackrel{\text{def}}{=} \lfloor z_j^x \rfloor + iz_j^y.$$

We denote by z^c the vector (z_1^c, \dots, z_s^c) . Our strategy is this: first, we establish that a small circuit can produce only a small entropy $H(z^c)$. We do this in three steps: Lemma 2.4 looks at the effect of the hashing on the entropy of the input variables. Lemmas 2.5 and 2.6 bound how much entropy the nonhelp and help gates, respectively, can inject into the vector of hashed variables. Finally, we show that $H(z^c)$ cannot be much smaller than $H(\lfloor z \rfloor)$, which by Lemma 2.3 is already known to be big.

LEMMA 2.4.

$$H(z_1^c, \dots, z_n^c) \leq 2n + \log V_k(R/\mu_k).$$

Proof. Let \mathcal{C} be the set of cubes of the form $(\mu_k \mathbf{Z})^n + [0, \mu_k]^n$. Any cube of \mathcal{C} that contains a point of L contains the entire unit cube centered at that point, and so

²We use complex numbers simply as a device for representing ordered pairs.

it intersects $K \cap B_n(O, R)$. Thus, the number of such cubes is at most $M_{R/\mu_k}^+(n, k)$, which, by (2), does not exceed $3^n V_k(R/\mu_k)$ so that (recall that $z_i^c = \lfloor x_i \rfloor$, for $i = 1, \dots, n$)

$$\begin{aligned} H(z_1^c, \dots, z_n^c) &= H(\lfloor \tilde{x}_1/\mu_k \rfloor, \dots, \lfloor \tilde{x}_n/\mu_k \rfloor) \\ &\leq 2n + \log V_k(R/\mu_k). \quad \square \end{aligned}$$

LEMMA 2.5. *For any nonhelp variable z_j , $n < j \leq s$, we have*

$$H(z_j^c \mid z_{f(j)}^c, z_{g(j)}^c) \leq 3.$$

Proof. Recall that $z_j^c = \lfloor z_j^x \rfloor + iz_j^y$. Obviously, since the imaginary part is completely determined by those of $z_{f(j)}^c$ and $z_{g(j)}^c$, we have

$$H(z_j^y \mid z_{f(j)}^y, z_{g(j)}^y) = 0.$$

To deal with the real part, we use the inequality $H(A \mid B) \leq H(A \mid C) + H(C \mid B)$ to derive

$$\begin{aligned} H(\lfloor z_j^x \rfloor \mid \lfloor z_{f(j)}^x \rfloor, \lfloor z_{g(j)}^x \rfloor) &\leq H(\lfloor z_j^x \rfloor \mid \lfloor \alpha_j z_{f(j)}^x \rfloor, \lfloor \beta_j z_{g(j)}^x \rfloor) \\ &\quad + H(\lfloor \alpha_j z_{f(j)}^x \rfloor \mid \lfloor z_{f(j)}^x \rfloor) \\ &\quad + H(\lfloor \beta_j z_{g(j)}^x \rfloor \mid \lfloor z_{g(j)}^x \rfloor). \end{aligned}$$

Given two random variables ξ, ξ' arbitrarily distributed in \mathbf{R} ,

$$H(\lfloor \xi + \xi' \rfloor \mid \lfloor \xi \rfloor, \lfloor \xi' \rfloor) \leq 1.$$

Intuitively, the only information missing is a one-bit carry. Similarly, given a fixed $\alpha \in \mathbf{Z}$, $|\alpha| \leq 1$, and a real random variable ξ , we have $H(\lfloor \alpha \xi \rfloor \mid \lfloor \xi \rfloor) \leq 1$. The lemma follows from the fact that $z_j = \alpha_j z_{f(j)} + \beta_j z_{g(j)}$. \square

LEMMA 2.6. *For any help variable z_j , $n < j \leq s$, we have $H(z_j^c \mid z_{f(j)}^c, z_{g(j)}^c) \leq 2(\log \mu_k + 1)$.*

Proof. We have $z_j^c = iz_j^y$, where z_j^y is an arbitrary function of $z_{f(j)}$ and $z_{g(j)}$. Regarding $z_{f(j)} = z_{f(j)}^x + iz_{f(j)}^y$, the only information we have at our disposal is $z_{f(j)}^c = \lfloor z_{f(j)}^x \rfloor + iz_{f(j)}^y$. There is no loss of information in the imaginary part. The same is not true of the real part, however. The key observation is that $z_{f(j)}^x$ is a linear form over x_1, \dots, x_n with integer coefficients. Thus, since $2\mu_k x_i$ is itself integral, so is $2\mu_k z_{f(j)}^x$. It follows that the fractional part of $z_{f(j)}^x$ can be one of only $2\mu_k$ possible values, and hence, $H(z_{f(j)}^x \mid \lfloor z_{f(j)}^x \rfloor) \leq \log \mu_k + 1$, from which the lemma follows. \square

We can use the last three lemmas to upper bound $H(\lfloor z \rfloor)$:

$$\begin{aligned} H(z^c) &= H(z_1^c, \dots, z_n^c) + \sum_{n+1 \leq j \leq s} H(z_j^c \mid z_1^c, \dots, z_{j-1}^c) \\ &\leq H(z_1^c, \dots, z_n^c) + \sum_{n+1 \leq j \leq s} H(z_j^c \mid z_{f(j)}^c, z_{g(j)}^c). \end{aligned}$$

By Lemmas 2.4, 2.5, and 2.6,

$$\begin{aligned} H(z^c) &\leq 2n + \log V_k(R/\mu_k) + 3(s - n - m) + 2m(\log \mu_k + 1) \\ &\leq 3s - n + \log V_k(R/\mu_k) + 2m \log \mu_k, \end{aligned}$$

and therefore,

$$\begin{aligned} H(\lfloor z \rfloor) &\leq H(z^c, \lfloor z \rfloor) = H(z^c) + H(\lfloor z \rfloor \mid z^c) \\ &\leq H(z^c) + \sum_{j=n+1}^s H(\lfloor z_j^x + z_j^y \rfloor \mid \lfloor z_j^x \rfloor + i z_j^y) \\ &\leq 4s + 2m \log \mu_k + \log V_k(R/\mu_k). \end{aligned}$$

Bringing the lower bound of Lemma 2.3 to bear, we derive

$$\begin{aligned} 4s &\geq -2m \log \mu_k - \log V_k(R/\mu_k) + \log V_k(R) \\ &\quad - \log V_k(\sqrt{n}) - \log V_n(5\sqrt{n}). \end{aligned}$$

Using the approximation [12],

$$V_d(r) = \frac{\pi^{d/2} r^d}{\Gamma(d/2 + 1)} \approx \frac{1}{\sqrt{\pi d}} \left(\frac{2e\pi}{d}\right)^{d/2} r^d,$$

and the fact that $\log V_d(rs) = \log V_d(r) + d \log s$, we find that

$$4s \geq -2m \log \mu_k + k \log \mu_k - \log V_k(\sqrt{n}) - \log V_n(5\sqrt{n}).$$

The last two terms add up to $O(n)$; therefore,

$$s \geq \frac{1}{8}(k - 2m) \log \lambda_k - O(n),$$

which establishes the spectral lemma. \square

3. Range searching over a group. Let P be the point set consisting of the vertices of a $(\sqrt{n} - 1) \times (\sqrt{n} - 1)$ square grid. Each point x_i of P is weighted by some real number, which by abuse of notation we also call x_i . Our goal is to exhibit n halfplanes h_1, \dots, h_n , and prove that computing the sum of the weights within each h_k , i.e., $\sum_{x_i \in h_k} x_i$, requires $\Omega(n \log n)$ time.

The model of computation is a straight-line program: each step performs a group operation of the form $z \leftarrow \alpha_1 x + \alpha_2 y$, where x and y are input weights or previously computed variables and $\alpha_i \in \{-1, 0, 1\}$. As a bonus, we also allow the use of close to $n/4$ instructions of the form $z \leftarrow f(x, y)$, where f is an arbitrary real function. The only requirement is that the same program should work for any assignment of real weights to the points. The analogy with the previous section is obvious. Let A be the $n \times n$ matrix whose k th row is the characteristic vector of $P \cap h_k$. By the spectral lemma, the lower bound of Theorem 1.1 follows directly from this lemma.

LEMMA 3.1. *There is a choice of n halfplanes, for which the matrix A of the corresponding set system is such that the k th largest eigenvalue of $A^T A$ is $n^{\Omega(1)}$, for some $k \geq n/2 - \epsilon n$, for any fixed $\epsilon > 0$.*

Proof. We use a nonconstructive configuration of halfplanes. Scale down the square grid so that it fits within $[1/\sqrt{n}, 1 - 1/\sqrt{n}]^2$. Let ω be the motion-invariant measure for lines: we normalize ω to provide a probability measure for the lines crossing $[0, 1]^2$. Given a halfplane h^+ bounded below by a nonvertical line h , consider the discrepancy function $f(h) \stackrel{\text{def}}{=} \sum_{x_i \in h^+} x_i$. A beautiful result of Alexander [1] (see [8, 10] for a simpler proof and various extensions) says that if $x_1 + \dots + x_n = 0$, then³

$$\int f^2(h) d\omega(h) \gg \frac{1}{\sqrt{n}} \|x\|_2^2.$$

³We use the notation \ll and \gg to denote inequality up to a constant factor.

Subdivide the space of lines crossing $[0, 1]^2$ into $N + O(n^2)$ regions within which the form $f(h)$ remains invariant. By choosing N large enough, say, $N = 2^n$, we can also ensure that the ω -area of N of these regions is exactly the same; call it σ , which is about $1/N$. The other $O(n^2)$ regions may have smaller areas. (Consider the arrangement in dual space to obtain this result.) Thus, the difference between integrating f^2 over the whole probability space and over the equal-area regions only is at most $O(n^2/N) \sup f^2$. Because $|f|$ cannot exceed

$$|x_1| + \dots + |x_n| \leq \sqrt{n} \|x\|_2,$$

the error is bounded by $O(n^3 \|x\|_2^2/N)$. This provides us with a good discrete approximation of the L^2 -norm of f . Indeed, let B be the $N \times n$ matrix whose rows are indexed by the N equal-area regions $\hat{\sigma}$ and are the characteristic vectors of the set of x_i 's appearing in (the unique form) $f(h)$, for $h \in \hat{\sigma}$. We have

$$\left| \|Bx\|_2^2 - \frac{1}{\sigma} \int f^2(h) d\omega(h) \right| = O(n^3) \frac{\|x\|_2^2}{N\sigma}.$$

But $\sigma = 1/N \pm O(n^2/N^2)$, so

$$\left| \|Bx\|_2^2 - N \int f^2(h) d\omega(h) \right| = O(n^3 \|x\|_2^2).$$

LEMMA 3.2.

$$\det B^T B = \Omega\left(N/\sqrt{n}\right)^{n-1}.$$

Proof. Let $\mu_1 \geq \dots \geq \mu_n \geq 0$ be the eigenvalues of $B^T B$ and let $\{v_i\}$ be an orthonormal eigenbasis, where v_i is associated with μ_i . We express $x = (\xi_1, \dots, \xi_n)$ in the basis $\{v_i\}$. The solution space of the system of equations, $x_1 + \dots + x_n = 0$ and $\xi_j = 0$ ($j < n - 1$), is of dimension at least 1. Since it lives in the (ξ_{n-1}, ξ_n) plane, it intersects the cylinder $\xi_{n-1}^2 + \xi_n^2 = 1$. For any point x of the intersection,

$$\|Bx\|_2^2 = \sum_{i=1}^n \mu_i \xi_i^2 = \mu_{n-1} \xi_{n-1}^2 + \mu_n \xi_n^2 \leq \mu_{n-1}.$$

This implies that for this unit vector x ,

$$\mu_{n-1} \geq N \int f^2(h) d\omega(h) - O(n^3 \|x\|_2^2) \gg \frac{N}{\sqrt{n}} - O(n^3),$$

and hence,

$$(3) \quad \mu_{n-1} \gg \frac{N}{\sqrt{n}}.$$

We need a lower bound on the smallest eigenvalue, but almost any one will do. With N being large enough, we can always assume that for each point x_i there exist two lines, each represented by a distinct row of B , that pass right above and below x_i . The contribution of these two rows to $\|Bx\|_2^2$ is of the form $\Phi^2 + (\Phi + x_i)^2$, which is always at least $x_i^2/2$. It follows that $\|Bx\|_2^2 \geq \frac{1}{2} \|x\|_2^2$, and hence, $\mu_n \geq 1/2$. The lemma follows from (3) and the fact that $\det B^T B$ is the product of the eigenvalues. \square

Of course, the set system B is much too big. Indeed, the map $x \mapsto Bx$ is actually trivial to compute. We use a nonconstructive argument to prove the existence of a hard $n \times n$ set system A . By the Binet-Cauchy formula,⁴

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

Therefore, there exists an $n \times n$ submatrix A of B such that

$$\begin{aligned} \det A^\top A &= \left| \det B \begin{pmatrix} j_1 & j_2 & \dots & j_n \\ 1 & 2 & \dots & n \end{pmatrix} \right|^2 \\ &\geq \binom{N}{n}^{-1} \det B^\top B = \Omega(1)^n \left(\frac{n}{eN}\right)^n \left(\frac{N}{\sqrt{n}}\right)^{n-1} \\ &\geq n^{n/2 - o(n)}, \end{aligned}$$

from which we find that

$$(4) \quad \log \det A^\top A \geq \left(\frac{n}{2} - o(n)\right) \log n.$$

By Morgenstern’s result [18], it follows easily that in the absence of any help gates, halfplane range searching requires $\Omega(n \log n)$ operations. To be able to deal with help gates we must collect more information about the spectrum of $A^\top A$. Let $\lambda_1 \geq \dots \geq \lambda_n \geq 0$ be the eigenvalues of $A^\top A$.

LEMMA 3.3. *For some constant $c > 0$,*

$$\lambda_k \leq \frac{cn^2 \log(k+1)}{k}.$$

Proof. Let L be a set of representative lines whose corresponding upper halfplanes define the sets encoded in the rows of A . Subdivide the unit square into a regular $r \times r$ grid of lines (r to be chosen later), and throw in a random sample of r lines chosen among L . Form the arrangement of these $3r$ lines and triangulate it. With high probability, no triangle is cut by more than $cn(\log r)/r$ lines of L , and none contains more than cn/r^2 points, for some constant c (assumed large enough for future purposes). For each triangle write the linear constraint expressing that the sum of the x_i ’s within the triangle is null. This gives us a set of $k_0 \leq cr^2$ linear constraints, called *canonical*. Assume that all are satisfied. Then, A can be rewritten in simpler form by means of a sparse matrix C . Specifically, by the zone theorem for line arrangements, we know that no line can cut more than cr triangles. Therefore, within the restriction to the constraint space, each row of A corresponds to a linear form with at most c^2n/r nonzero coefficients. Let C be the new matrix formed by those relevant entries. Note that no column (resp., row) of C contains more than $cn(\log r)/r$ (resp., c^2n/r) ones.

It is a standard result [14] that the spectral norm of a matrix Q satisfies

$$\|Q\|_s^2 \leq \left(\max_i \sum_j |q_{ij}|\right) \left(\max_j \sum_i |q_{ij}|\right),$$

⁴The notation refers to the matrix obtained by picking the rows indexed j_1, \dots, j_n in B .

and therefore the Rayleigh quotient $x^\top C^\top C x / x^\top x$ ($x \neq 0$) is at most $c^3 n^2 (\log r) / r^2$. As a result, for any x satisfying the canonical constraints and $\|x\|_2 = 1$, we have

$$(5) \quad \|Ax\|_2^2 \leq \frac{c^3 n^2 \log r}{r^2}.$$

Let $\{u_i\}$ be an orthonormal eigenbasis for $A^\top A$, where u_i is associated with λ_i . We express $x = (\xi_1, \dots, \xi_n)$ in the basis $\{u_i\}$. The solution space of the system of equations consisting of $\xi_j = 0$ ($j \geq k_0 + 2$) and the canonical constraints is of dimension at least 1. It is embedded in the $(k_0 + 1)$ -flat spanned by $(\xi_1, \dots, \xi_{k_0+1})$, so it intersects the cylinder $\xi_1^2 + \dots + \xi_{k_0+1}^2 = 1$. For any point x in the intersection, $\|x\|_2 = 1$ and

$$\|Ax\|_2^2 = \sum_{i=1}^n \lambda_i \xi_i^2 = \sum_{i=1}^{k_0+1} \lambda_i \xi_i^2 \geq \lambda_{k_0+1},$$

so by (5), $\lambda_{k_0+1} \leq c^3 n^2 (\log r) / r^2$. In the worst case, k_0 is proportional to r^2 , so the lemma is true for any k large enough. For small k , we can use the straightforward bound $\lambda_k \leq n^2$. \square

From the lemma we find that

$$\begin{aligned} \log \det A^\top A &= \sum_{i=1}^n \log \lambda_i \\ &\leq (n - k) \log \lambda_k + \sum_{j=1}^k \log \frac{cn^2 \log(j + 1)}{j} \\ &\leq (n - k) \log \lambda_k + k(2 \log n - \log k + \log \log k + c'). \end{aligned}$$

In view of (4) we find that we can set $k = n/2 - \varepsilon n$, for any fixed $\varepsilon > 0$, and still derive the lower bound $\log \lambda_k = \Omega(\log n)$, which proves Lemma 3.1. \square

We conclude that halfplane range searching requires $\Omega(n \log n)$ time, even in the presence of close to $n/4$ free computations. Notice that over the reals the problem can be solved in linear time with only $2n - 1$ free computations: the circuit is a tree of help gates whose leaves are the x_i 's and whose root "collects" the vector (x_1, \dots, x_n) and encodes it as a real. Then, with another n help gates, we can distribute the correct n outputs: the total number of help gates is $2n - 1$. This completes the proof of Theorem 1.1. \square

Appendix. We prove (2): $M_r^+(n, m) \leq 3^n V_m(r)$, for $n > m$. Let \mathcal{C} be the set of cubes counted by $M_r^+(n, m)$. Any cube $c \in \mathcal{C}$ has at least one $(n - m)$ -face intersecting $F \cap B_n(O, r)$ in exactly one point (call it q_c). This face is supported by an $(n - m)$ -flat which is specified by fixing exactly m integer coordinates. In other words, it is specified by an integral point p_c in the m -flat spanned by a set of m axes. Since by convexity such a point p_c corresponds to at most $2^{n-m} (n - m)$ -faces and at most 2^m cubes can share the same $(n - m)$ -face, counting the number of points p_c gives an upper bound on $M_r^+(n, m)$, up to a factor of 2^n . Of course, we can restrict the counting to the number of integral points that lie within any of the projections of $F \cap B_n(O, r)$ onto m -flats spanned by x_{i_1}, \dots, x_{i_m} . Furthermore, we can discount projections that map F to a flat of dimension less than m (because q_c is uniquely defined). Let E_{i_1, \dots, i_m} be the ellipsoid obtained by projecting $F \cap B_n(O, r)$ onto

the flat $(x_{i_1}, \dots, x_{i_m})$. We say that a point p_c is *peripheral* if it is the upper corner (upper with regard to all m dimensions) of a cube not fully contained in the projected ellipsoid, and we let N denote the total number of peripheral points. We have

$$(6) \quad M_r^+(n, m) \leq 2^n N + 2^n \sum_{1 \leq i_1 < \dots < i_m \leq n} \text{vol } E_{i_1, \dots, i_m}.$$

Let v_1, \dots, v_m be an orthonormal basis for F and let U be the $n \times m$ matrix whose columns are the v_i 's. (Note that $U^T U$ is the identity matrix.) The determinant of the $m \times m$ submatrix of U specified by the rows (i_1, \dots, i_m) is equal, in absolute value, to the volume of E_{i_1, \dots, i_m} divided by $V_m(r)$. (We need only sketch the proof of this simple fact: lift to F the principal vectors of the ellipsoid and scale them to unit length; this provides an orthonormal basis for F that satisfies the claim. Because the basis $\{v_i\}$ can be derived from it by a unitary transformation within F , the claim follows.) By the Binet-Cauchy formula, the determinant of $U^T U$ can be expressed as

$$\sum_{1 \leq i_1 < \dots < i_m \leq n} \left| \det U \begin{pmatrix} i_1 & i_2 & \dots & i_m \\ 1 & 2 & \dots & m \end{pmatrix} \right|^2,$$

and by Cauchy-Schwarz,

$$\sum_{i_1, \dots, i_m} \text{vol } E_{i_1, \dots, i_m} \leq V_m(r) \binom{n}{m}^{1/2} \sqrt{\det U^T U} = V_m(r) \binom{n}{m}^{1/2}.$$

On the other hand, in the flat $(x_{i_1}, \dots, x_{i_m})$, every cube of the unit lattice that has a vertex in E_{i_1, \dots, i_m} and that intersects $\partial E_{i_1, \dots, i_m}$ does so along at least one edge: at most four such edges can be collinear, and no more than 2^{m-1} cubes can charge the same edge. Projecting E_{i_1, \dots, i_m} by dropping one coordinate gives an $(m-1)$ -ellipsoid whose integral points are contained in an $(m-1)$ -ball of radius r . Within a factor of 2^{m+1} , the total number of such points is an upper bound on N . Each such point is the upper corner of a distinct unit cube in a ball of radius $r + \sqrt{m-1}$; therefore, $N \leq m 2^{m+1} \binom{n}{m} V_{m-1}(r + \sqrt{m-1})$, and by (6),

$$M_r^+(n, m) \leq m 4^n \binom{n}{m} V_{m-1}(r + \sqrt{m-1}) + 2^n \binom{n}{m}^{1/2} V_m(r).$$

As r goes to infinity the second term becomes dominant, and for n large enough, (2) follows.

Acknowledgments. I wish to thank Dick Lipton, Ran Raz, Avi Wigderson, and Andy Yao for helpful discussions. I also thank the referees for many useful comments and suggestions.

REFERENCES

[1] R. ALEXANDER, *Geometric methods in the study of irregularities of distribution*, *Combinatorica*, 10 (1990), pp. 115–136.
 [2] N. ALON, M. KARCHMER, AND A. WIGDERSON, *Linear circuits over GF(2)*, *SIAM J. Comput.*, 19 (1990), pp. 1064–1067.
 [3] A. AMIR, R. BEIGEL, AND W. GASARCH, *Some connections between bounded query classes and nonuniform complexity*, 5th Annual IEEE Structure in Complexity Theory Conference, IEEE Computer Society Press, Los Alamitos, CA, 1990, pp. 232–243.

- [4] J. BECK AND W. W. L. CHEN, *Irregularities of Distribution*, Cambridge Tracts in Mathematics 89, Cambridge University Press, Cambridge, 1987.
- [5] R. B. BOPPANA AND M. SIPSER, *The complexity of finite functions*, in Handbook of Theoretical Computer Science, Vol. A, Algorithms and Complexity, MIT Press/Elsevier, Cambridge/New York, 1990, pp. 757–804.
- [6] J. Y. CAI, *Lower bounds for constant depth circuits in the presence of help bits*, in Proc. 30th Annual Symp. Foundation Comput. Sci., (FOCS), IEEE Computer Society Press, Los Alamitos, CA, 1989, pp. 532–537.
- [7] B. CHAZELLE, *Lower bounds on the complexity of polytope range searching*, J. Amer. Math. Soc., 2 (1989), pp. 637–666.
- [8] B. CHAZELLE, *Geometric discrepancy revisited*, in Proc. 34th Annual IEEE Symp. Foundation Comput. Sci., (FOCS), IEEE Computer Society Press, Los Alamitos, CA, 1993, pp. 392–399.
- [9] B. CHAZELLE, *Computational geometry: A retrospective*, in Computing in Euclidean Geometry, 2nd ed., D.-Z. Du and F. Hwang, eds., World Scientific Press, River Edge, NJ, 1995, pp. 22–46.
- [10] B. CHAZELLE, J. MATOUŠEK, AND M. SHARIR, *An elementary approach to lower bounds in geometric discrepancy*, Discrete Comput. Geom., 13 (1995), pp. 363–381.
- [11] M. L. FREDMAN, *Lower bounds on the complexity of some optimal data structures*, SIAM J. Comput., 10 (1981), pp. 1–10.
- [12] M. GRÖTSCHEL, L. LOVÁSZ, AND A. SCHRIJVER, *Geometric Algorithms and Combinatorial Optimization*, Springer-Verlag, New York, 1988.
- [13] J. VON ZUR GATHEN, *Algebraic complexity theory*, Ann. Rev. Comput. Sci., 1988, pp. 317–347.
- [14] P. LANCASTER AND M. TISMENETSKY, *The Theory of Matrices*, 2nd ed., Academic Press, New York, 1985.
- [15] J. MATOUŠEK, *Range searching with efficient hierarchical cuttings*, Discrete Comput. Geom., 10 (1993), pp. 157–182.
- [16] J. MATOUŠEK, *Geometric range searching*, Tech. Report B-93-09, Free Univ. Berlin, 1993.
- [17] K. MEHLHORN, *Data Structures and Algorithms 3: Multidimensional Searching and Computational Geometry*, Springer-Verlag, Heidelberg, 1984.
- [18] J. MORGENSTERN, *Note on a lower bound of the linear complexity of the fast Fourier transform*, J. ACM, 20 (1973), pp. 305–306.
- [19] K. MULMULEY, *Computational Geometry: An Introduction Through Randomized Algorithms*, Prentice-Hall, Englewood Cliffs, NJ, 1994.
- [20] V. STRASSEN, *Algebraic complexity theory*, in Handbook of Theoretical Computer Science, Vol. A, Algorithms and Complexity, MIT Press/Elsevier, Cambridge/New York, 1990, pp. 633–672.