

Lower Bounds for Linear Degeneracy Testing

NIR AILON AND BERNARD CHAZELLE

Princeton University, Princeton, New Jersey

Abstract. In the late nineties, Erickson proved a remarkable lower bound on the decision tree complexity of one of the central problems of computational geometry: given n numbers, do any r of them add up to 0? His lower bound of $\Omega(n^{\lceil r/2 \rceil})$, for any fixed r , is optimal if the polynomials at the nodes are linear and at most r -variate. We generalize his bound to s -variate polynomials for $s > r$. Erickson's bound decays quickly as r grows and never reaches above pseudo-polynomial: we provide an exponential improvement. Our arguments are based on three ideas: (i) a geometrization of Erickson's proof technique; (ii) the use of error-correcting codes; and (iii) a tensor product construction for permutation matrices.

Categories and Subject Descriptors: F.2.0 [Analysis of Algorithms and Problem Complexity]: General

General Terms: Theory

Additional Key Words and Phrases: Computational geometry, linear decision trees, lower bounds

1. Introduction

Decision trees have often shown to be realistic and effective models for proving lower bounds on the complexity of fundamental geometric problems [Ben-Or 1983; Björner et al. 1992; Dobkin and Lipton 1979; Erickson 1999a, 1999b; Erickson and Seidel 1995; Grigoriev et al. 1996, 1997; Steele and Yao 1982; Yao 1997, 1995]. Testing degeneracy is one such example. The r -variate degeneracy testing problem is to decide whether, given a sequence of n reals x_1, \dots, x_n and a real linear polynomial over r variables $f \in \mathbf{R}[t_1, \dots, t_r]$, there exist distinct indices i_1, \dots, i_r such that $f(x_{i_1}, \dots, x_{i_r}) = 0$. Additional constraints might be imposed on the indices. The terminology owes to the problem formulation's suitability for checking the degeneracy of just about any geometric configuration. For example, general position of N points in \mathbf{R}^d can be reduced to $d(d+1)$ -variate degeneracy testing with respect to dn numbers; in this case, f is a $(d+1)$ -by- $(d+1)$ determinant with a row of ones, and constraints on the indices ensure that the entries of the

This work was supported in part by National Science Foundation grants CCR-998817, and CCR-0306283, ARO Grant DAAH04-96-1-0181, and NEC Research Institute.

Authors' Address: Princeton University, Department of Computer Science, 35 Olden Street, Princeton, NJ 08544, e-mail: {nailon, chazelle}@cs.princeton.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2005 ACM 0004-5411/05/0300-0157 \$5.00

determinant are, indeed, the coordinates of $d + 1$ distinct points. Similarly, we can formulate the degeneracy of Voronoi diagrams, power diagrams, algebraic varieties, real semi-algebraic sets, etc. Classical “bichromatic” problems also fall in that category: for example, checking incidence between points and hyperplanes (Hopcroft’s problem), rays and triangles, lines and spheres, etc. The list of problems studied in the literature that can be reduced to degeneracy testing is nearly endless.

Even the unconstrained version of r -variate linear degeneracy testing (r -LDT) is ubiquitous in the computational geometry landscape. This is looking among n numbers for a zero of $f(y) = a_0 + a_1y_1 + \dots + a_ry_r$ ($a_i \neq 0$ for $i > 0$). There is a vast collection of geometric problems known to be 3SUM-hard and 4SUM-hard, all of which are at least as hard as r -LDT (for $r = 3, 4$) via subquadratic reductions [Gajentaan and Overmars 1995]. Classical examples are separating line segments by a line, testing if a union of triangles is simply connected, checking for polygon containment under translation, minimizing the Hausdorff distance between segment sets, computing the Minkowski sum of two polygons, sorting the vertices of a line arrangement, etc. [Arkin et al. 1998; Barequet and Har-Peled 2001; Barrera 1996; Bose et al. 1993; de Berg et al. 1997; Matousek 1995]. Needless to say, the importance of elucidating the complexity of r -LDT can hardly be overstated.

While the problem, being a variant of SUBSET SUM, is clearly NP-complete, its parameterized complexity as a function of r is poorly understood (to put it charitably). The trivial $O(n^r)$ upper bound can be improved to $O((2n/r)^{\lceil r/2 \rceil})$ if r is odd and $O((2n/r)^{r/2} \log(n/r))$ if r is even [Erickson 1999a]. The idea for r even is to write $f = g - h$, where $g \in \mathbf{R}[t_1, \dots, t_{r/2}]$ and $h \in \mathbf{R}[t_{r/2+1}, \dots, t_r]$. We sort all possible values of h and store them in a table. By binary search we look up every possible value of g . Any successful search corresponding to a match with distinct indices is a certificate of degeneracy. In a nonuniform decision tree model, the extra log factor is not needed. In other words, for even n , there exists a linear decision tree for r -LDT of depth $O((2n/r)^{r/2})$. For r odd, the algorithm is slightly different: We write $f = g - h + \alpha_r t_r$, where $g \in \mathbf{R}[t_1, \dots, t_{\lfloor r/2 \rfloor}]$ and $f \in \mathbf{R}[t_{\lfloor r/2 \rfloor+1}, \dots, t_{r-1}]$. We sort all possible values of g and h , which are both $\lfloor r/2 \rfloor$ variate. Then for each x_i , we look for a value of g and a value of h satisfying $g - h + \alpha_r x_i = 0$ (being careful, again, that we use r distinct coordinates). This can be done by performing a linear scan of the two sorted lists. The total (uniform) running time is dominated by the n linear scans, resulting in $O(n(2n/r)^{\lfloor r/2 \rfloor})$. Both the even and odd versions can be implemented in a linear decision tree in which each internal node is associated with a linear polynomial over r variables. What makes this result particularly interesting is the existence of a matching lower bound.

The underlying model is the r -linear decision tree: each internal node v is assigned a linear n -variate polynomial q_v with at most r nonzero (real) coefficients; its outgoing edges are labeled $<$, $=$, or $>$. Leaves are labeled *yes* or *no*. To test the degeneracy of an input $x = (x_1, \dots, x_n) \in \mathbf{R}^n$, we evaluate $q_v(x)$ beginning at the root and follow outgoing edges in the obvious way until we reach a leaf, at which point we output its label: *yes* if the input is degenerate and *no* otherwise. Information theoretic Lower bounds of $\Omega(n \log n)$ on the depth of a tree deciding r -LDT in an unconstrained linear decision tree model (no restriction on the number of nonzero coefficients) are obtained by Dobkin and Lipton [1979], and under more general nonlinear models of computation by Steele and Yao [1982] and Ben-Or [1983]. Improving on previous work [Dietzfelbinger 1989; Fredman 1976], Erickson [1999a] proved that *any* r -linear decision tree for *any* r -LDT problem has depth $\Omega(n^{\lfloor r/2 \rfloor})$.

His proof is quite a tour de force. It is packed with ingenious, tightly coupled arguments, and its only downside is to offer little wiggle room to try out new ideas. In particular, extending the proof to s -linear trees for $s > r$ has long been elusive. Even the case $s = r + 1$, mentioned in Yao's list of major open problems in his 2000 DIMACS lecture [Yao 2000], has resisted all efforts. The contribution of this article, while far from closing the book on the problem, represents a significant advance on two fronts: (i) accommodating $s > r$ variables and (ii) allowing for large values of r .

—We prove a lower bound of $\Omega(nr^{-3})^{\lceil r/2 \rceil}$ on the depth of any r -linear decision tree for any r -LDT problem. This improves on Erickson's bound of $\Omega(nr^{-r})^{\lceil r/2 \rceil}$ from pseudopolynomial to exponential for large values of r . Indeed, if $r = r(n) > n^\varepsilon$, Erickson's bound can never exceed $n^{\log n / \log \log n}$, while ours is of the form $2^{n^{\Omega(1)}}$. The technical underpinning of this improvement is a new adversarial strategy based on error-correcting codes.

—By using a tensor product construction based on permutation matrices, we are able to generalize the lower bound to the s -linear decision tree model for $s > r$. We show that, for any instance of r -LDT, the tree depth is at least

$$\Omega(nr^{-3})^{\frac{2r-s}{2\lceil(s-r+1)/2\rceil}(1-\varepsilon_r)},$$

where $\varepsilon_r > 0$ tends to 0 as $r \rightarrow \infty$.

The exponential lower bound still holds for $s > r$. For any fixed $\varepsilon > 0$, the depth of an s -linear decision tree is $(nr^{-3})^{r^{\Omega(1)}}$, if $s \leq r + r^{1-\varepsilon}$. In the case $r > n^\varepsilon$, this gives a lower bound of $2^{n^{\Omega(1)}}$. Note that our bounds collapse if s is not $O(r)$. This is an obvious limitation of our method, but one must note that a dependency on s is inevitable. Indeed, our lower bound of $n^{\Omega(r)}$ for $s = r + O(1)$ cannot hold for arbitrary values of s . By a result of Meyer auf der Heide [1984], a decision tree of depth $O(n^4 \log n)$ exists for any instance of linear degeneracy testing over r variables, without a restriction on the number of nonzero coefficient of the tree polynomials ($s = n$). The existence of an unconstrained linear decision tree with depth $\text{poly}(n, r)$ deciding r -LDT is also implied by Meiser [1993].

Another contribution of this paper is methodological. To obtain our bounds requires a whole set of new algebraic arguments, but our starting point is essentially a geometrization of Erickson's method. The main benefit is to bypass the complicated machinery of infinitesimals found in Erickson [1999a], obviate the need for Tarski's transfer principle, and more generally do away with analytical arguments.

To make the proof more digestible, we begin our discussion with the geometric framework and then treat the case $s = r$. Next, we move on to the case $s = r + 1$, where we introduce the tensor product construction in its simplest form. Finally, we cover the general case.

2. A Geometric Framework for Lower Bounds

We consider the r -SUM problem: Given a point $x = (x_1, \dots, x_n) \in \mathbf{R}^n$, are there indices $i_1 < \dots < i_r$ such that $x_{i_1} + \dots + x_{i_r} = 0$? We wish to prove that any r -linear decision tree used to answer this question is of depth $\Omega(n^{\lceil r/2 \rceil})$. Choosing f to be the symmetric linear function on r variables simplifies the lower bound proof, but our results can be extended to any r -variate linear function. The generalization

is proven in Section 5. Each node v is associated with a polynomial q_v whose zeroes define a hyperplane, called a *query*. The set of query hyperplanes is denoted by \mathcal{Q} . It is not hard to see that if the decision tree is to be valid, \mathcal{Q} must include every one of the $\binom{n}{r}$ *canonical* hyperplanes $x_{i_1} + \cdots + x_{i_r} = 0$. Indeed, if such a hyperplane h^* is missing in \mathcal{Q} , then there exists a pair of points p_1 and p_2 such that p_1 lies on h^* (thus degenerate), p_2 is nondegenerate, and no hyperplane in the finite set \mathcal{Q} separates between p_1 and p_2 . Therefore, the decision tree cannot decide r -LDT.

The basic idea of the proof is to identify a “large” face C in the arrangement¹ $\mathcal{A}(\mathcal{Q})$ formed by \mathcal{Q} . The face C , called the *chamber*, may not necessarily be full-dimensional but

(C1) C must not be contained in any canonical hyperplane.

We also need a set \mathcal{H} of *critical* hyperplanes. These are canonical hyperplanes tangent to C such that

(C2) each $h^* \in \mathcal{H}$ has a designated point p_h on the boundary of C ;

(C3) no two points in the collection $\{p_h\}$ lie in the closure of the same face of the boundary of C .

LEMMA 2.1. *Any r -linear decision tree for the r -SUM problem is of depth at least $|\mathcal{H}|$.*

PROOF. The tree must lead to a *no* (respectively, *yes*) leaf for any input point $p_0 \in C$ (respectively, p_h , where $h^* \in \mathcal{H}$). For this reason, the path followed on input p_h must include a query hyperplane q_h that intersects, but does not contain, the segment $p_0 p_h$. Indeed, the same path would otherwise be followed for input p_0 . Since p_h is in the closure of a face of $\mathcal{A}(\mathcal{Q})$ that contains p_0 , the hyperplane q_h passes through p_h but does not intersect C . Now the crux is that by (C3) no query hyperplane can pass through more than one point p_h . \square

2.1. CRITICAL HYPERPLANES VIA ERROR-CORRECTING CODES. By padding the input if necessary, we can always assume that $n = rm$, for some integer m . This allows us to view a vector $h \in \mathbf{R}^n$ (and hence its hyperplane h^* through the origin) as an $r \times m$ real matrix M^h , whose rows are filled with the coordinates of h ; that is, $M_{ij}^h = h_{(i-1)m+j}$. A critical hyperplane being of the form $x_{i_1} + \cdots + x_{i_r} = 0$, its corresponding matrix has r ones and $n - r$ zeroes. The hyperplanes we will choose as critical hyperplanes will have a single one per row. Where to put the ones is dictated by an error-correcting recipe meant to ensure high “independence.” Throughout this section, we use the shorthand

$$r_0 = \left\lceil \frac{r}{2} \right\rceil. \quad (1)$$

Let q be the smallest prime greater than r , and let \mathcal{M} be a Reed–Solomon code [MacWilliams and Sloane 1977] of length $q - 1$ and distance $r - r_0 + 1$

¹*Terminology and Conventions:* Faces of polyhedra and arrangements are disjoint, relatively open sets. The intersection of the closures of any two faces is either empty or the closure of another face. Faces of dimension (codimension) 0 and 1 are called vertices and edges (cells, facets), respectively. Convex polyhedra are assumed to be closed. For $h \in \mathbf{R}^n$, the hyperplane h^* is defined as $\{x \mid \langle h, x \rangle = 0\}$. A hyperplane intersecting the boundary of a convex polyhedron but not its interior is *tangent* to the convex polyhedron.

over the finite field \mathbf{F}_q . This means that \mathcal{M} is a linear subspace of \mathbf{F}_q^{q-1} with the following combinatorial property: any nonzero vector in \mathcal{M} has at least $r - r_0 + 1$ nonzero coordinates. A constructive way to do this is to regard \mathbf{F}_q^{q-1} as the ring of polynomials $\mathbf{F}_q[X]$ modulo the polynomial $X^{q-1} - 1$. We then pick some primitive² $\beta \in \mathbf{F}_q$ and let \mathcal{M} be the ideal in this ring generated by the polynomial $(X - \beta)(X - \beta^2) \cdots (X - \beta^{r-r_0})$. This ideal has dimension $k = q - 1 - r + r_0$ with the desired distance property (see MacWilliams and Sloane [1977] for details). Now, define \mathcal{M}_r to be the linear subspace of \mathcal{M} defined by adding the constraints $\{x_i = 0 \mid r < i \leq q - 1\}$. (Note: This is *not* the same as chopping off the last $q - 1 - r$ coordinates.) In this way, we can think of \mathcal{M}_r as a linear code of length r , distance greater than $r - r_0$ and dimension at least $k - (q - 1 - r) = r_0$. Let v_1, \dots, v_{r_0} be an independent set of vectors in \mathcal{M}_r . Of course, by permuting coordinates and performing column operations, we can always assume that the set is in column echelon form, ie, the $r \times r_0$ matrix (v_1, \dots, v_{r_0}) consists of the $r_0 \times r_0$ identity matrix on top of some $(r - r_0) \times r_0$ matrix. Since \mathbf{F}_q is a prime field, we can naturally view the v_i 's as vectors in \mathbf{R}^r with coordinates in $\{0, \dots, q - 1\}$. We define \mathcal{L} as the set of vectors $n_1 v_1 + \dots + n_{r_0} v_{r_0}$ for all nonnegative integers $n_i \leq m/q r_0$. The upper bound is chosen so that all coordinates lie in $\{0, \dots, m - 1\}$. (Throughout this article, the notation $\text{span}(\mathcal{S})$ refers to the vector space spanned by \mathcal{S} over the reals.)

LEMMA 2.2. *Three facts: (i) the set \mathcal{L} consists of at least $(n/r^3)^{r_0}$ vectors in \mathbf{R}^r with coordinates in $\{0, \dots, m - 1\}$; (ii) the first r_0 coordinates of any vector in \mathcal{L} specify it uniquely; (iii) any nonzero vector in $\text{span}(\mathcal{L})$ has at least $r - r_0 + 1$ nonzero coordinates.*

PROOF. By Nagura's theorem [Nagura 1952], the interval $[x, 6x/5]$ contains a prime for any $x \geq 25$. This shows that $q r_0 \leq r^2$; therefore,

$$|\mathcal{L}| \geq (m/q r_0)^{r_0} \geq (n/r^3)^{r_0}.$$

Part (ii) comes from the echelon form of the matrix formed by (v_1, \dots, v_{r_0}) . To prove (iii), consider a nonzero element $\sum_{i=1}^{r_0} \alpha_i v_i$ of $\text{span}(\mathcal{L})$. The set of such vectors with at least r_0 zero coordinates can be expressed as a union of linear subspaces, each one defined by a set of homogeneous equations in the α_i 's with integer coefficients. Therefore, if the set is nonempty, it must contain a vector v with all its α_i 's integral and at least one of them not divisible by q . Reducing v modulo q gives us a nontrivial linear combination of the v_i 's. Since these vectors are independent over \mathbf{F}_q , it then follows that v is a nonzero vector of the code \mathcal{M}_r with at least r_0 zero coordinates. This contradicts the fact that \mathcal{M}_r has distance greater than $r - r_0$. \square

The set \mathcal{H} of critical hyperplanes is in bijection with \mathcal{L} . The hyperplane h^* corresponding to $\ell = (\ell_1, \dots, \ell_r) \in \mathcal{L}$ is defined by its matrix M^h : the coordinate ℓ_i indicates where to place the 1 in the i th row of the matrix, that is, $M_{ij}^h = 1$ (respectively, 0) if $j = \ell_i + 1$ (respectively, otherwise). By construction,

$$M^h(0, \dots, m - 1)^T \in \text{span}(\mathcal{L}). \tag{2}$$

²Meaning that the sequence $1, \beta, \beta^2, \dots, \beta^{q-2}$ has no repetitions.




FIG. 1. Output no if the input is p_0 , but yes if the input is p_h or $p_{h'}$. Both of these points lie on critical hyperplanes as well as in the flat $p_0 + W$.

The intersection $\cap \mathcal{H}$ of all the hyperplanes h^* in \mathcal{H} is a linear subspace of positive dimension. Indeed, it contains the vector

$$\underbrace{(1, \dots, 1)}_{n-n/r}, \underbrace{(1-r, \dots, 1-r)}_{n/r}.$$

Let \mathcal{K} denote the set of query hyperplanes that contain $\cap \mathcal{H}$. Note that $\mathcal{Q} \supseteq \mathcal{K} \supseteq \mathcal{H}$.

LEMMA 2.3. *Given any $q^* \in \mathcal{K}$, (i) $M^q (1, \dots, 1)^T = b(1, \dots, 1)^T$ for some real b , and (ii) $M^q (0, \dots, m-1)^T \in \text{span}(\mathcal{L})$.*

PROOF. Since $(\cap \mathcal{H})^\perp$ is the space spanned by the normals of hyperplanes in \mathcal{H} and $q \in (\cap \mathcal{H})^\perp$, $q = \sum_i \lambda_i h_i$, where $h_i^* \in \mathcal{H}$; therefore $M^q = \sum_i \lambda_i M^{h_i}$. But each M^{h_i} has a single one per row, and so $M^{h_i} (1, \dots, 1)^T = (1, \dots, 1)^T$; hence (i). Similarly, (ii) follows from (2). \square

2.2. THE CHAMBER. The query hyperplanes outside of \mathcal{K} intersect $\cap \mathcal{H}$ in lower-dimensional subspaces. Therefore, there exist $c_0 \in \cap \mathcal{H}$ and $\rho > 0$ such that the ball $B(c_0, \rho)$ centered at c_0 of radius ρ intersects none of the hyperplanes of $\mathcal{Q} \setminus \mathcal{K}$. By lying on every critical hyperplane the point c_0 is highly degenerate. Moving it by some vector ψ to be specified next changes all of that (Figure 1). We define the point

$$p_0 = c_0 + \psi \quad (3)$$

to be safely outside of the critical hyperplanes. To do that, we need a positive convex real function g , meaning one with positive second derivative; eg, $x \mapsto x^2 + 1$. For some fixed, small enough $\gamma > 0$, we define the vector $\psi \in \mathbf{R}^n$ by its matrix M^ψ :

$$M_{ij}^\psi = \begin{cases} \gamma g(j) & \text{if } i \leq r_0; \\ \gamma^2 g(j) & \text{otherwise.} \end{cases} \quad (4)$$

NOTE: γ is a scaling factor that is absolutely needed. The reason we use γ^2 , however, is in anticipation of the case $s > r$. We could use γ in this section instead.

LEMMA 2.4. *The point p_0 lies outside of any canonical hyperplane and any hyperplane of $\mathcal{Q} \setminus \mathcal{K}$.*

This implies that the decision tree must output *no* on input p_0 . Note, however, that p_0 might still lie on a query hyperplane.

PROOF. Recall that a canonical hyperplane h^* is one with an equation of the form $x_{i_1} + \dots + x_{i_r} = 0$. By choosing γ small enough, we can ensure that $\|\psi\|_2 \leq \rho/2$; therefore, the point p_0 lies inside $B(c_0, \rho)$, safely away from any hyperplane of $\mathcal{Q} \setminus \mathcal{K}$ (Figure 1). We have already observed that \mathcal{Q} must contain all of the canonical hyperplanes; therefore, the only danger is that p_0 lies on some canonical hyperplane h^* in \mathcal{K} . But this is impossible. Indeed, $c_0 \in \cap \mathcal{H} \subseteq h^*$, and so

$$\langle h, p_0 \rangle = \langle h, \psi \rangle = \sum_{j \in J} \gamma g(j) + \sum_{j \in J'} \gamma^2 g(j) > 0,$$

with $|J \cup J'| = r$. \square

The chamber C is the unique face of $\mathcal{A}(\mathcal{Q})$ that contains p_0 . To define the map $h^* \in \mathcal{H} \mapsto p_h \in \partial C$, we need to introduce the vector space W spanned by the $2r$ vectors $u_k, w_k \in \mathbf{R}^n$ ($k = 1, \dots, r$), defined (using the matrix notation) as follows:

$$M_{ij}^{u_k} = \begin{cases} 1 & i = k \\ 0 & \text{otherwise} \end{cases}, \quad M_{ij}^{w_k} = \begin{cases} j & i = k \\ 0 & \text{otherwise} \end{cases}.$$

In other words, W consists of vectors w such that $M_{ij}^w = \alpha_i^w + \beta_j^w$ for all i, j , for some real $\alpha_1^w, \dots, \alpha_r^w, \beta_1^w, \dots, \beta_r^w$. All the points p_h will lie on $C \cap (p_0 + W)$. The reason for this will be made clear in case (A) in the proof of Lemma 2.6. Given $h^* \in \mathcal{H}$, we define a vector

$$\varphi_h \in \psi + W \tag{5}$$

such that $M_{ij}^{\varphi_h} > 0$ (respectively, $M_{ij}^{\varphi_h} = 0$) if $M_{ij}^h = 0$ (respectively, $M_{ij}^h > 0$). Note that $\psi + W$ is not necessarily a vector space. One should think of M^{φ_h} as a mask: Its rows mark with zeroes the positions where M^h is 1 and fill the rest with positive entries. To see that such a vector φ_h actually exists, consider the i th row of the matrix M^{φ_h} . Let

$$\gamma_i = \begin{cases} \gamma & i \leq r_0 \\ \gamma^2 & i > r_0 \end{cases}.$$

It suffices to show that the row can satisfy constraints in t, u of the form $\gamma_i g(j) + t + u j = 0$ if j is equal to the one value j_0 where $M_{ij_0}^h = 1$, and $\gamma_i g(j) + t + u j > 0$ for any $j \neq j_0$. Feasibility is ensured by the fact that

$$\frac{g(j_0) - g(j)}{j_0 - j} < \frac{g(j') - g(j_0)}{j' - j_0}$$

for any $j < j_0 < j'$, which itself is a consequence of the mean-value theorem applied to the convex function g . It is immediate to check that

$$\lim_{\gamma \rightarrow 0^+} M_{ij}^{\varphi_h} = 0, \tag{6}$$

for all i, j . This implies that, by scaling down γ if necessary, we can ensure that $\|\varphi_h\|_2 < \rho/2$. We now define

$$p_h = c_0 + \varphi_h. \quad (7)$$

LEMMA 2.5. *The point p_h lies inside the critical hyperplane h^* and outside any hyperplane of $\mathcal{Q} \setminus \mathcal{K}$.*

PROOF. The second part follows directly from the fact that $\|\varphi_h\|_2 < \rho/2$. Recall that φ_h has zero coordinates precisely at the positions where h does not; therefore,

$$\langle h, p_h \rangle = \langle h, (c_0 + \varphi_h) \rangle = \langle h, \varphi_h \rangle = 0. \quad \square$$

Incidentally, note that the p_h 's are not strewn all across the boundary of \mathcal{C} : By (3)–(7), they all live in the low-dimensional flat $p_0 + W$.

LEMMA 2.6. *Given any $q^* \in \mathcal{Q}$ and $h^* \in \mathcal{H}$, if $\text{sgn}(\langle q, p_0 \rangle) \neq \text{sgn}(\langle q, p_h \rangle)$, then $\langle q, p_h \rangle = 0$ and $\text{sgn}\langle q, p_x \rangle = \text{sgn}\langle q, p_0 \rangle$ for any $x^* \in \mathcal{H}$ distinct from h^* . (In our notation, $\text{sgn}(y) = 1, 0$ or -1 if $y > 0, y = 0$ or $y < 0$, respectively.) In other words, collapsing p_0 to any p_h changes at most a single query from nonzero to zero, and leaves the sign of every other query unchanged.*

PROOF. Obviously we can assume that $q^* \in \mathcal{K}$, since otherwise q^* would miss $B(c_0, \rho)$ entirely and so $\langle q, p_0 \rangle$ and $\langle q, p_h \rangle$ would be nonzero and have the same sign. We distinguish between two cases.

(A) Each row of M^q has at least one nonzero element: Then, since $s = r$, each row has *exactly* one and, by Lemma 2.3(i), all the nonzero elements are equal to the same number, which without loss of generality we may assume to be 1. Therefore,

$$\begin{aligned} \langle q, p_0 \rangle &= \langle q, (c_0 + \psi) \rangle = \langle q, \psi \rangle \\ &= \sum_{i=1}^{r_0} \gamma g(j_i) + \sum_{i=r_0+1}^r \gamma^2 g(j_i) > 0. \end{aligned} \quad (8)$$

By the lemma's assumption, it follows that $\langle q, p_h \rangle \leq 0$. For any $x^* \in \mathcal{H}$, $\langle q, p_x \rangle = \langle q, (c_0 + \varphi_x) \rangle = \langle q, \varphi_x \rangle \geq 0$; therefore, $\langle q, p_h \rangle = 0$. Can any other $x^* \in \mathcal{H}$ also satisfy $\langle q, p_x \rangle = 0$? The answer is no. We will show that the equality $\langle q, p_x \rangle = 0$ characterizes exactly one hyperplane $x^* \in \mathcal{H}$, which therefore has to be h . To see why, recall that M^{φ_x} acts as a mask for the 1s in M^x . If M^x does not match M^q in each one of the first r_0 rows then, by (6), $\langle q, p_x \rangle = \langle q, \varphi_x \rangle = \text{tr } M^q (M^{\varphi_x})^T$, is of the form $C\gamma + O(\gamma^2)$, for some $C > 0$, and hence can never be 0 as long as we choose γ small enough. On the other hand, by Lemma 2.2(ii), if M^x matches M^q in each one of the first r_0 rows, then x is unique, as required. This completes the proof for case (A). Note that we never used the fact that each row $i > r_0$ of M^q has exactly one nonzero element. This will allow us to use the same proof verbatim in the next section, even though the only assurance we will then have on the matrix M^q is each of its first r_0 rows has a single 1.

(B) Some row of M^q is null: By Lemma 2.3(i),

$$M^q (1, \dots, 1)^T = 0,$$

and each row of M^q must then have at least two nonzero elements or none at all. It follows that the number of null rows is at least r_0 , and so, by Lemmas 2.2(iii)

and 2.3(ii), $M^q(0, \dots, m-1)^T = 0$. (This is where error correction kicks in.) As a result, $\langle q, w \rangle = 0$ for any $w \in W$. But, by (3)–(7), $p_h - p_0 = \varphi_h - \psi \in W$; therefore, $\langle q, p_h \rangle = \langle q, p_0 \rangle$, which contradicts the lemma's assumption. \square

It is immediate to verify that the chamber C and the points p_h satisfy the requirements C1–C3. If the chamber C lay within a canonical hyperplane, then so would p_0 , which would contradict Lemma 2.4; hence C1. Consider a query hyperplane $q^* \in \mathcal{Q}$ and a critical hyperplane $h^* \in \mathcal{H}$. By Lemma 2.6, there are only three possibilities: $p_0 p_h \subset q^*$, $p_0 p_h \cap q^* = \emptyset$, or $p_h \in q^*$ and $p_x \notin q^*$ for any distinct $x^* \in \mathcal{H}$. This proves that p_h lies in the closure of C . Combined with the fact that $p_h \in h^*$ (Lemma 2.5), this establishes condition C2. Finally, if two distinct points p_h and p_x lay in the same face of the closure of C , then some $q^* \in \mathcal{Q}$ would contradict the three possibilities above; hence C3. In view of Lemmas 2.1 and 2.2, we have proven.

THEOREM 2.7. *The depth of any r -linear decision tree for r -SUM is*

$$\Omega(nr^{-3})^{\lceil r/2 \rceil}.$$

3. The Case $s = r + 1$

What can go wrong with the previous proof if $s = r + 1$? The only place where the number s actually plays a role is in the proof of Lemma 2.6. Case (B) survives almost verbatim. The only problem is that the number of null rows is at least $r - \lfloor s/2 \rfloor$, which can be less than r_0 . We fix this by redefining r_0 so that it satisfies

$$1 \leq r_0 \leq r - \lfloor s/2 \rfloor. \quad (9)$$

In the present case, the setting $r_0 = \lfloor r/2 \rfloor$ will do.

Case (A) is far more difficult to fix. All the rows of M^q have exactly one nonzero element, except for one of them, i_0 , which has two nonzeros (the case of one nonzero in every row having already been handled). Again we can assume that all the nonzero elements are 1, except in row i_0 , where the elements are α and $1 - \alpha$, for some real $\alpha \notin \{0, 1\}$. Let

$$\gamma' = \begin{cases} \gamma & i_0 \leq r_0 \\ \gamma^2 & i_0 > r_0 \end{cases}.$$

Taking row i_0 into account, we can rewrite (8) as

$$\begin{aligned} \langle q, p_0 \rangle &= \langle q, \psi \rangle \\ &= \gamma' \alpha g(j_{i_0}) + \gamma' (1 - \alpha) g(j'_{i_0}) + \sum_{\substack{i=1 \\ i \neq i_0}}^{r_0} \gamma g(j_i) + \sum_{\substack{i=r_0+1 \\ i \neq i_0}}^r \gamma^2 g(j_i). \end{aligned}$$

If $i_0 > r_0$, then all is well. Indeed, by making γ small enough

$$\langle q, p_0 \rangle = \sum_{i=1}^{r_0} \gamma g(j_i) + O_q(\gamma^2) > 0. \quad (10)$$

(This is the reason we needed to use both γ and γ^2 in (4).) The remainder of the proof involves only the first r_0 rows of M^q , which happen to be as in case (A), and so it can be repeated verbatim.

The case $i_0 \leq r_0$ is a tougher nut to crack. In fact we have not found a way of tackling it directly. Consequently, our strategy is simply to modify \mathcal{H} so that this case cannot happen. Recall that, for the purpose of Lemma 2.6, we can assume that $q^* \in \mathcal{K}$. As we observed in the proof of Lemma 2.3, this implies that $q \in \text{span}(\mathcal{H}^*)$, where $\mathcal{H}^* = \{h \mid h^* \in \mathcal{H}\}$. Thus, our goal is to redefine³ a large set \mathcal{H} of critical hyperplanes so that, in addition to all the properties we expect of \mathcal{H} , the following should hold: If q is a vector of \mathbf{R}^n such that (i) with the exception of one row $i_0 \leq r_0$ each of the first r_0 rows of M^q consists of a single 1 with 0's everywhere else, and (ii) the exceptional row, i_0 , is null everywhere except for two entries summing up to 1, then q cannot be in the span of \mathcal{H}^* .

Recall from the construction of \mathcal{H} that the first r_0 rows of any M^h ($h^* \in \mathcal{H}$) completely determine the remaining rows. Furthermore, each one of the first r_0 rows can be chosen by placing a 1 arbitrarily between positions 1 and $m_0 = \lfloor m/qr_0 \rfloor$ and filling the rest of the row with 0's. So it suffices to concentrate on the first r_0 rows. Once we have the top r_0 rows, we use our Reed–Solomon code to fill in the bottom $r - r_0$ rows just as we did in the previous section.

An $r_0 \times a$ matrix is called *defective* if, with the exception of one row (called *anomalous*), each one consists of a single 1 with 0's everywhere else; furthermore the exceptional row is null everywhere except at two places. We postpone the proof of the next result.

LEMMA 3.1. *There exists a set \mathcal{P} of $r_0 \times m$ 0/1 matrices with exactly one 1 per row between positions 1 and m_0 such that no defective $r_0 \times m$ matrix belongs to $\text{span } \mathcal{P}$ and, for n large enough and any fixed $\varepsilon > 0$,*

$$|\mathcal{P}| \geq (nr^{-3})^{\lfloor r/2 \rfloor (1 - 1/\ln \lfloor r/2 \rfloor) (1 - \varepsilon)}.$$

In view of our previous discussion, this automatically implies a lower bound on the depth of $(r + 1)$ -linear trees. The theorem below does not indicate what happens for small values of r . A careful examination shows that we obtain nontrivial lower bounds for any $r \geq 6$.

THEOREM 3.2. *The depth of any $(r + 1)$ -linear decision tree for r -SUM is at least $(nr^{-3})^{\lfloor r/2 \rfloor - o(r)}$.*

3.1. THE TENSOR PRODUCT CONSTRUCTION. The problem fits into a general class of questions related to codes and combinatorial designs: How to build a large vector space that does not contain a family of forbidden vectors? In the case at hand, we start by building a “core” square matrix that satisfies the desired property and then show how to scale it up into an arbitrarily large rectangular matrix by using a suitable tensor product.

Let A (respectively, B) be an $r_0 \times a$ (respectively, $r_0 \times b$) real matrix. Following standard tensor notation, we write the element $A_{i,j}$ as A_j^i instead. We define a tensor product matrix operator

$$\otimes : \mathbf{R}^{r_0 \times a} \times \mathbf{R}^{r_0 \times b} \mapsto \mathbf{R}^{r_0 \times ab}$$

as follows: If $P = A \otimes B$, then $P_{j,k}^i = A_j^i B_k^i$. It is a mixed third-order tensor with two covariant indices and a single contravariant one.

³Although we are redefining (actually, only shrinking) \mathcal{H} , we are not choosing a different chamber.

This product extends to sets naturally. If \mathcal{A} (respectively, \mathcal{B}) is a set of $r_0 \times a$ (respectively, $r_0 \times b$) real matrices, then

$$\mathcal{A} \otimes \mathcal{B} = \{ A \otimes B \mid A \in \mathcal{A}, B \in \mathcal{B} \}.$$

Tensor exponentiation for sets is defined by

$$\mathcal{A}^{\otimes k} = \underbrace{\mathcal{A} \otimes \cdots \otimes \mathcal{A}}_{k \text{ times}}.$$

The elements of $\mathcal{A}^{\otimes k}$ belong to the vector space \mathcal{V}_k of mixed $(k + 1)$ st-order tensors with k covariant indices and 1 contravariant one. By fixing an ordering (say, lexicographic) of the covariant indices, we can interpret the tensors of \mathcal{V}_k as $r_0 \times a^k$ matrices, and vice-versa.

For our “core,” we choose permutation matrices. Let Π denote the set of $r_0 \times r_0$ 0/1 matrices with exactly one 1 per row and column. The lemma below gives our tensor product its *raison d’être*.

LEMMA 3.3. *No defective $r_0 \times r_0^k$ matrix can belong to the span of $\Pi^{\otimes k}$, for any $k \geq 1$.*

PROOF. In any matrix of $\text{span}(\Pi)$ each row and each column sum up to the same number, which can be assumed to be 1. Therefore, the anomalous row of a defective $r_0 \times r_0$ matrix consists of two entries, $\alpha, \alpha' \neq 0$ summing up to 1. Suppose that the column with the α also includes a set of ℓ ones for $\ell > 0$ (note that these can only be ones). Since the column sum is 1, we have $\alpha + \ell = 1$. This implies that α must be an integer and, since it is nonzero, $\alpha' = 1 - \alpha = \ell \geq 2$. But then the column with α' sums up to more than 1, which gives a contradiction. This implies that neither of the columns with α, α' has any other nonzero element. But then the $r_0 - 2$ other columns sum up to $r_0 - 1$, which exceeds the required count by one. This proves the lemma for $k = 1$.

For $k > 1$, we define the tensor homomorphism $h_l : \mathcal{V}_k \mapsto \mathcal{V}_{k-1}$, where

$$h_l(P)_{j_1, \dots, j_{l-1}, j_{l+1}, \dots, j_k}^i = \sum_{j=1}^{r_0} P_{j_1, \dots, j_{l-1}, j, j_{l+1}, \dots, j_k}^i.$$

Let M be a defective $r_0 \times r_0^k$ matrix. By definition, its anomalous row i_0 contains two nonzero elements: the two corresponding covariant k -tuple indices, being distinct, differ in at least one index l . Since $k > 1$, there exists at least one covariant index $l' \neq l$. We easily verify that $h_{l'}(M)$ is a defective $r_0 \times r_0^{k-1}$ matrix and

$$h_{l'}(\Pi^{\otimes k}) = \Pi^{\otimes(k-1)}.$$

The proof follows by simple linear algebra and induction. \square

To maximize its size, we choose the set $\mathcal{P} = \Pi^{\otimes k}$ for the largest k such that $r_0^k \leq m_0 = \lfloor m/qr_0 \rfloor$. Using Stirling’s approximation, we find that

$$|\mathcal{P}| = |\Pi|^k = (r_0!)^k \geq (nr^{-3})^{\lfloor r/2 \rfloor (1-1/\ln[r/2]) (1-\varepsilon)},$$

for any fixed $\varepsilon > 0$. Filling up each row with 0’s to get the proper of length m concludes the proof of Lemma 3.1. \square

4. The Case $s = r + 1$

We need a new idea to generalize the tensor product construction to higher values of s . We exploit the fact that the (hard part of the) lower bound involves only query hyperplanes whose normal vectors q are spanned by the normals h of the critical hyperplanes $h^* \in \mathcal{H}$. We use this to add combinatorial structure to the matrices M^q by redesigning the set $\cap \mathcal{H}$. We need to redefine r_0 so that the first r_0 rows of M^h can be grouped in equal-sized blocks. Of course, r_0 still needs to satisfy (9). The choice of $r_0 = \lambda \rho_0$ will do, where

$$\lambda = \left\lfloor \frac{s-r}{2} \right\rfloor + 1 \quad \text{and} \quad \rho_0 = \left\lfloor \frac{r - \lfloor s/2 \rfloor}{\lambda} \right\rfloor.$$

Note that this requires that s be not too large, say $s < \lfloor 3r/2 \rfloor$. Divide up the first r_0 rows of M^h into λ blocks of consecutive rows of ρ_0 rows each. To build up a matrix M^h of \mathcal{H} , we proceed as follows:

- Step 1.* Use the tensor construction of the previous section (the case $s = r + 1$) to produce the top ρ_0 rows of M^h . In carrying out the construction, of course, use permutation matrices of size $\rho_0 \times \rho_0$ instead of $r_0 \times r_0$. This gives us a set \mathcal{P} of matrices with the same properties as those of Lemma 3.1, except for the size of \mathcal{P} and the size of the matrices, now $\rho_0 \times m$.
- Step 2.* For each matrix of \mathcal{P} , make λ copies of it and stack them on top of one another to produce an $r_0 \times m$ matrix.
- Step 3.* Complete the bottom $r - r_0$ rows via Reed–Solomon as before.

LEMMA 4.1. *For any $q^* \in \mathcal{K}$, the top r_0 rows of M^q form an $r_0 \times m$ matrix made up of λ copies of the same $\rho_0 \times m$ matrix.*

PROOF. This is a simple consequence of the fact that $q \in \text{span}\{h \mid h^* \in \mathcal{H}\}$. \square

There is no need to revisit Lemma 2.6 in detail. Again, only case (A) is worth discussing: Each row of M^q has at least one nonzero element. By analogy with the case $s = r + 1$, if all the rows with more than one nonzero have indices greater than r_0 then inequality (10) holds and we are done.

Suppose now that at least one row $i_0 \leq r_0$ contains two or more nonzeros. By Lemma 4.1, the λ blocks that make up the top r_0 rows of M^q are identical. This shows that no block can have more than $\rho_0 + 1$ nonzeros. Indeed, any one of them did, then so would all of the others, and their combined contribution of nonzeros would be at least $(\rho_0 + 2)\lambda$. Added to the (at least) $r - r_0$ nonzeros of the bottom rows, this would give us a total of at least $(\rho_0 + 2)\lambda + r - r_0 > s$ nonzero coordinates in q , which is ruled out. So, the only possibility left is for each block to have exactly ρ_0 or $\rho_0 + 1$ nonzeros: the first case was handled in the proof of Lemma 2.6, while the second one was shown to be impossible in the last section because of the tensor product construction.

The new set \mathcal{P} is of size at least $(\rho_0!)^k$, where k is the largest integer such that $\rho_0^k \leq m/qr_0$. Elementary calculations show that

$$|\mathcal{P}| \geq (nr^{-3})^{\frac{2r-s}{2\lfloor (s-r)/2 \rfloor + 1} (1-\varepsilon_r)},$$

where $\varepsilon_r > 0$ tends to 0 as $r \rightarrow \infty$.

THEOREM 4.2. *The depth of any s -linear decision tree for r -SUM is at least*

$$(nr^{-3})^{\frac{2r-s}{2(s-r+1)/2} - o(r)}.$$

Note that for any $s \leq r + r^{1-\varepsilon}$, where $\varepsilon > 0$ is arbitrarily small constant, the depth is $(nr^{-3})^{r^{\Omega(1)}}$.

5. Generalization to Other Linear Degeneracy Tests

We define the following linear degeneracy problems:

r -SUM': Given a point $x \in \mathbf{R}^{r \times m}$, do there exist indices $i_1, \dots, i_r \leq m$ such that

$$\sum_{k=1}^r x_{ki_k} = 0 ?$$

Let f be a fixed r -variate linear polynomial of the form $f = \alpha_1 t_1 + \dots + \alpha_r t_r$, where $\alpha_k \neq 0$ for all $k = 1, \dots, r$.

f -SUM': Given a point $x \in \mathbf{R}^{r \times m}$, do there exist indices $i_1, \dots, i_r \leq m$ such that

$$\sum_{k=1}^r \alpha_k x_{ki_k} = 0 ?$$

Now fix some $c \in \mathbf{R}$.

(f, c) -SUM': Given a point $x \in \mathbf{R}^{r \times m}$, do there exist indices $i_1, \dots, i_r \leq m$ such that

$$\sum_{k=1}^r \alpha_k x_{ki_k} = c ?$$

Finally, the most general version of the linear degeneracy test we are considering, is

(f, c) -SUM: Given a point $x \in \mathbf{R}^n$, do there exist r distinct indices i_1, \dots, i_r such that

$$\sum_{k=1}^r \alpha_k x_{i_k} = c ?$$

OBSERVATION 5.1. *Theorems 2.7, 3.2 and 4.2 apply to r -SUM', because the critical hyperplanes have a single one in each row.*

CLAIM 5.2. *Let T be a linear decision tree deciding f -SUM' for some fixed n . If for all queries q of the tree we apply the following transformation, then T transforms into a tree deciding r -SUM': Assuming q is a query comparing $\sum_{i,j} \beta_{ij} x_{ij}$ to a constant z , replace β_{ij} with β_{ij}/α_i , for all i, j .*

CLAIM 5.3. *Let T be a linear decision tree deciding (f, c) -SUM'. If for all queries q of the tree we apply the following transformation, then T transforms into a tree deciding f -SUM': Assuming q is a query comparing $\sum_{i,j} \beta_{ij} x_{ij}$ to a constant z , replace z with $z - \frac{c}{r} \sum_{ij} (\beta_{ij}/\alpha_i)$.*

Observation 5.1 together with Claims 5.2 and 5.3 imply that our lower bounds apply for (f, c) -SUM', which is not harder than the most general problem we consider, (f, c) -SUM.

6. Open Problems

The techniques used in this paper break down when $r = n^{1/3}$ or when $s = \Omega(r)$. Unfortunately, even the cases $r < 6, s > r$ give trivial lower bounds, which are subsumed by the general $\Omega(n \log n)$. These cases are important, and we hope that the new techniques introduced here will return the focus to them. Also, it would be interesting to know if the lower bounds we get for $s > r$ (and especially for $s > r + 1$) are tight for big r .

ACKNOWLEDGMENTS. We wish to thank Jeff Erickson for enlightening discussions about his lower bound proof, and the referees for their helpful comments.

REFERENCES

- ARKIN, E., CHIANG, Y.-J., HELD, M., MITCHELL, J., SACRISTAN, V., SKIENA, S., AND YANG, T.-C. 1998. On minimum-area hulls. *Algorithmica* 21, 119–136.
- BAREQUET, G., AND HAR-PELED, S. 2001. Polygon containment and translational min-hausdorff-distance between segment sets are 3sum-hard. *Int. J. Comput. Geom. App.* 11, 465–474.
- BARRERA, A. 1996. Finding an $o(n^2 \log n)$ algorithm is sometimes hard. In *Proceedings of the 8th Canadian Conference on Computational Geometry* (Ottawa, Ont., Canada). 289–294.
- BEN-OR, M. 1983. Lower bounds for algebraic computation trees. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing*. ACM, New York, 80–86.
- BJÖRNER, A., LOVÁSZ, L., AND YAO, A. 1992. Linear decision trees: volume estimates and topological bounds. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*. ACM, New York, 170–177.
- BOSE, P., VAN KREVELD, M., AND TOUSSAINT, G. 1993. Filling polyhedral molds. In *Proceedings of the 3rd Workshop on Algorithms in Data Structures*. Lecture Notes in Computer Science, vol. 709. Springer-Verlag, New York, 210–221.
- DE BERG, M., DE GROOT, M., AND OVERMARS, M. 1997. Perfect binary space partitions. *Comput. Geom. Theory Appl.* 7, 81–91.
- DIETZFELBINGER, M. 1989. Lower bounds for sorting of sums. *Theoret. Comput. Sci.* 66, 137–155.
- DOBKIN, D., AND LIPTON, R. 1979. On the complexity of computations under varying set of primitives. *J. Comput. Syst. Sci.* 18, 86–91.
- ERICKSON, J. 1999a. Lower bounds for linear satisfiability problems. *Chi. J. Theoret. Comput. Sci.* 8.
- ERICKSON, J. 1999b. New lower bounds for convex hull problems in odd dimensions. *SIAM J. Comput.* 28, 1198–1214.
- ERICKSON, J., AND SEIDEL, R. 1995. Better lower bounds on detecting affine and spherical degeneracies. *Disc. Comput. Geom.* 13, 41–57.
- FREDMAN, M. 1976. How good is the information theory bound in sorting. *Theoret. Comput. Sci.* 1, 355–361.
- GAJENTAAN, A., AND OVERMARS, M. 1995. On a class of $o(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.* 5, 165–185.
- GRIGORIEV, D., KARPINSKI, M., MEYER AUF DER HEIDE, F., AND SMOLENSKY, R. 1996. A lower bound for randomized algebraic decision trees. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*. ACM, New York, 612–619.
- GRIGORIEV, D., KARPINSKI, M., AND VOROBOV, N. 1997. Lower bound on testing membership to a polyhedron by algebraic decision and computation trees. *Disc. Comput. Geom.* 17, 191–215.
- MACWILLIAMS, F., AND SLOANE, N. 1977. *The Theory of Error Correcting Codes*. North Holland.
- MATOUSEK, J. 1995. On geometric optimization with few violated constraints. *Disc. Comput. Geom.* 14, 365–384.
- MEISER, S. 1993. Point location in arrangements of hyperplanes. *Inf. Comput.* 106, 286–303.

- MEYER AUF DER HEIDE, F. 1984. A polynomial linear search algorithm for the n -dimensional knapsack problem. *J. ACM* 31, 668–676.
- NAGURA, J. 1952. On the interval containing at least one prime number. *Proc. Japan Acad.* 28, 177–181.
- STEELE, M., AND YAO, A. 1982. Lower bounds for algebraic decision trees. *J. Alg.* 3, 1–8.
- YAO, A. 1995. Algebraic decision trees and euler characteristics. *Theoret. Comput. Sci.* 141, 133–150.
- YAO, A. 1997. Decision tree complexity and betti numbers. *J. Comput. System Sci.* 55, 36–43.
- YAO, A. 2000. Why I'm an optimist. In *DIMACS Workshop on Intrinsic Complexity of Computation*. 10–13.

RECEIVED MARCH 2004; REVISED DECEMBER 2004; ACCEPTED DECEMBER 2004