New Lower Bounds for Convex Hull Problems in Odd Dimensions*

Jeff Erickson

Computer Science Division
University of California
Berkeley, CA 94720-1776
jeffe@cs.berkeley.edu

Abstract

We show that in the worst case, $\Omega(n^{\lceil d/2\rceil -1} + n\log n)$ sidedness queries are required to determine whether the convex hull of n points in $\mathbb{R}^d$ is simplicial, or to determine the number of convex hull facets. This lower bound matches known upper bounds in any odd dimension. Our result follows from a straightforward and completely constructive adversary argument. A key step in the proof is the construction of a quasi-simplicial n-vertex polytope with $\Omega(n^{\lceil d/2\rceil -1})$ degenerate facets. Using similar techniques, we also derive a constructive proof of Erickson and Seidel's lower bounds for detecting affine degeneracies in arbitrary dimensions. As a related result, we show that detecting simplicial convex hulls in $\mathbb{R}^d$ is $\lceil d/2\rceil$SUM-hard, in the sense of Gajentaan and Overmars. While it has been known for several years that d-dimensional convex hulls can have $\Omega(n^{\lfloor d/2\rfloor})$ facets, the previously best lower bound for either of the problems we consider is only $\Omega(n\log n)$.

1 Introduction

The construction of convex hulls is one of the most basic and well-studied problems in computational geometry [4, 5, 6, 7, 9, 10, 15, 18, 20, 21, 22, 29, 24, 26, 27]. Over twenty years ago, Graham described an algorithm that constructs the convex hull of n points in the plane in $O(n\log n)$ time [15]. The same running time was first achieved in three dimensions by Preparata and Hong [21]. Yao [30] proved a lower bound of $\Omega(n\log n)$ on the complexity of identifying the convex hull vertices, in the quadratic decision tree model. This lower bound was later generalized to the algebraic decision tree and algebraic computation tree models by Ben-Or [3]. It follows that both Graham's scan and Preparata and Hong's algorithm are optimal in the worst case. If the output size f is also taken into account, the lower bound drops to $\Omega(n\log f)$ [18], and a number of algorithms match this bound both in the plane [18, 5, 4] and in three dimensions [10, 8].

Chazelle [7] describes an algorithm for constructing convex hulls in $\mathbb{R}^d$ in time $O(n^{\lfloor d/2\rfloor} + n\log n)$. Since an n-vertex polytope in $\mathbb{R}^d$ can have $\Omega(n^{\lfloor d/2\rfloor})$ facets [31], Chazelle's algorithm is optimal in the worst case. Several faster algorithms are known when the output size is also considered; see for example [4, 5, 9, 10, 18, 26].

In this paper, we consider convex hull problems for which the output size is a single integer, or even a single bit, although the convex hull itself may be large. We show that in the worst case, $\Omega(n^{\lceil d/2\rceil -1} + n\log n)$ sidedness queries are required to decide whether the convex hull of n points in $\mathbb{R}^d$ is simplicial, or to determine the number of convex hull facets. This matches known upper bounds when d is odd [7]. The only lower bound previously known for either of these problems is $\Omega(n\log n)$, following from the techniques of Yao [30] and Ben-Or [3].

Our lower bounds follow from a straightforward adversary argument. We start by constructing a set whose convex hull contains a large number of independent degenerate facets. To obtain the adversary configuration, we perturb this set to eliminate the degeneracies, but in a way that the degeneracies are still "almost there". An adversary can reintroduce any one of the degenerate facets, by moving its vertices back to their original position.

Our argument is similar to earlier arguments of Erickson and Seidel [12] and the author [11]. In contrast to these results, however, our proofs are completely constructive. In fact, our proof technique yields a much simpler constructive proof of Erickson and Seidel's lower

*This research was partially supported by a GAANN Fellowship.

bounds for affine degeneracy detection in arbitrary dimensions.

The paper is organized as follows. Section 2 contains definitions and some preliminary results. In Section 3, we derive relative complexity results. Section 4 contains the proof of our main theorem. We discuss extensions of our model of computation in Section 5. We discuss the relevance of our model of computation in Section 6. Section 7 contains some related results. Finally, in Section 8, we summarize and suggest directions for further research.

2 Geometric Preliminaries

2.1 Definitions

We begin by reviewing basic terminology. For a more detailed introduction to the theory of convex polytopes, we refer the reader to Ziegler [31] or Grünbaum [16].

The *convex hull* of a set of points is the smallest convex set that contains it. A *polytope* is the convex hull of a finite set of points. A hyperplane h *supports* a polytope if the polytope intersects h and lies in a closed halfspace of h. The intersection of a polytope and a supporting hyperplane is called a *face* of the polytope. The *dimension* of a face is the dimension of the smallest affine space that contains it; a face of dimension k is called a *k-face*. The faces of a polytope are also polytopes. Given a d-dimensional polytope, its $(d-1)$-faces are called *facets*, its $(d-2)$-facets are called *ridges*, its 1-faces are called *edges*, and its 0-faces are called *vertices*.

A polytope is *simplicial* if all its facets, and thus all its faces, are simplices. A polytope is *quasi-simplicial* if all of its ridges are simplices, or equivalently, if its facets are simplicial polytopes. A *degenerate facet* of a quasi-simplicial polytope is any facet that is not a simplex.

The basic computational primitive that we consider is the *sidedness query*: Given $d+1$ points $p_0, p_1, \ldots, p_d \in \mathbb{R}^d$, does the point p_0 lie "above", on, or "below" the oriented hyperplane determined by the other d points? Algebraically, the result of a sidedness query is given by the sign of the following $(d+1) \times (d+1)$ determinant, where p_{ij} denoted the jth coordinate of p_i.

$$\begin{vmatrix} 1 & p_{01} & p_{02} & \cdots & p_{0d} \\ 1 & p_{11} & p_{12} & \cdots & p_{1d} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & p_{d1} & p_{d2} & \cdots & p_{dd} \end{vmatrix}$$

The value of this determinant is $d!$ times the signed volume of the simplex spanned by the points. The *orientation* of a simplex $(p_0, p_1, \ldots, p_d)$ is the result of a sidedness query on its vertices (in the order presented). If the orientation is zero, we say that the simplex is *degenerate*. A set of points is *affinely degenerate* if any $d+1$ of its elements lie on a single hyperplane, or equivalently, if the set contains the vertices of a degenerate simplex. The convex hull of an affinely nondegenerate set of points is simplicial. (The converse is not true in general—consider the regular octahedron in $\mathbb{R}^3$.)

2.2 The Weird Moment Curve

The *weird moment curve* in $\mathbb{R}^d$, denoted $\omega_d(t)$, is the set of points

$$\omega_d(t) = (t, t^2, \ldots, t^{d-1}, t^{d+1}),$$

where the parameter t ranges over the reals. The weird moment curve is similar to the standard moment curve $(t, t^2, \ldots, t^{d-1}, t^d)$, except that the degree of the last coordinate is increased by one.

If we project the weird moment curve down a dimension by dropping the last coordinate, we get a standard moment curve. Since every set of points on the standard moment curve is in convex position, every set of points on the d-dimensional weird moment curve in convex position if $d \geq 3$. Similarly, since every set of points on the standard moment curve is affinely nondegenerate, no d points on the d-dimensional weird moment curve lie on a single $(d-2)$-flat. It follows that the convex hull of any set of points on the weird moment curve is quasi-simplicial.

Lemma 2.1. *Let $x_0 < x_1 < \cdots < x_n$ be real numbers. The orientation of the simplex $(\omega_d(x_0), \omega_d(x_1), \ldots, \omega_d(x_d))$ is given by the sign of $\sum_{i=0}^d x_i$. In particular, the simplex is degenerate if and only if $\sum_{i=0}^d x_i = 0$.*

Proof: The orientation of the simplex $(\omega_d(x_0), \omega_d(x_1), \ldots, \omega_d(x_d))$ is given by the sign of the determinant of the following matrix.

$$M = \begin{bmatrix} 1 & x_0 & x_0^2 & \cdots & x_0^{d-1} & x_0^{d+1} \\ 1 & x_1 & x_1^2 & \cdots & x_1^{d-1} & x_1^{d+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & x_d & x_d^2 & \cdots & x_d^{d-1} & x_d^{d+1} \end{bmatrix}$$

The determinant of M is an antisymmetric polynomial of degree $\binom{d+1}{2}+1$ in the variables x_i, and it is divisible by $(x_i - x_j)$ for all $i < j$. It follows that

$$\frac{\det M}{\prod_{i<j}(x_j - x_i)}$$

is a symmetric polynomial of degree one, and we easily observe that its leading coefficient is 1. (This polynomial is well-defined, since the x_i are distinct.) The only such polynomial is $\sum_{i=0}^{d} x_i$. $\square$

This result, or at least its proof, is hardly new. If we replace the weird moment curve by any polynomial curve, the orientation of a simplex is given by the sign of a *Schur polynomial* [23]. A determinantal formula for Schur polynomials was discovered by Jacobi in the mid-1800's [17].

The next lemma characterizes degenerate convex hull facets on the weird moment curve. The result is similar to Gale's evenness condition [14], which describes which vertices of a cyclic polytope form its facets.

Lemma 2.2. *Let X be a set of real numbers, and let $x_0, x_1, \ldots, x_d \in X$, such that $\sum_{i=0}^{d} x_i = 0$. The points $\omega_d(x_0), \omega_d(x_1), \ldots, \omega_d(x_d)$ are the vertices of a degenerate facet of $\operatorname{conv}(\omega_d(X))$ if and only if for any two elements $y, z \in X \setminus \{x_0, x_1, \ldots, x_d\}$, the number of elements of $\{x_0, x_1, \ldots, x_d\}$ between y and z is even.*

Proof: Let h be the hyperplane passing through the points $\omega_d(x_0), \omega_d(x_1), \ldots, \omega_d(x_d)$. Expanding the appropriate sidedness determinant, we find that the point $\omega_d(x)$ lies above, on, or below h according to the sign of the polynomial

$$f(x) = \left(x + \sum_{i=1}^{d} x_i\right) \prod_{i=1}^{d}(x - x_i) = \prod_{i=0}^{d}(x - x_i).$$

The hyperplane h supports $\operatorname{conv}(\omega_d(X))$ if and only if $f(x)$ has the same sign for all $x \in X \setminus \{x_0, x_1, \ldots, x_d\}$.

The polynomial $f(x)$ has degree $d+1$, and vanishes at each x_i. Thus, the sign of $f(x)$ changes at each x_i. In more geometric terms, the weird moment curve crosses the hyperplane h at each of the points $\omega_d(x_i)$. It follows that $f(y)$ and $f(z)$ both have the same sign if and only if an even number of x_i's lie between y and z. $\square$

3 $\lceil d/2 \rceil$SUM-Hardness

Gajentaan and Overmars [13] define the class of 3SUM-hard problems, all of which are harder than the following base problem.

> 3SUM: Given a set of real numbers, do any three sum to zero?

Formally, a problem is 3SUM-hard if there is a sub-quadratic reduction from 3SUM to the problem in question. Thus, a sub-quadratic algorithm for any 3SUM-hard problem would imply a sub-quadratic algorithm for 3SUM, and a sufficiently powerful quadratic lower bound for 3SUM would imply similar lower bounds for every 3SUM-hard problem. Examples of 3SUM-hard problems include several degeneracy detection, separation, hidden surface removal, and motion planning problems in two and three dimensions.

More generally, we will say that a problem is *rSUM-hard* if there is a $o(n^{\lceil r/2 \rceil})$-time reduction to it from the following problem.

> rSUM: Given a set of real numbers, do any r sum to zero?

Many higher-dimensional versions of 3SUM-hard problems are rSUM-hard for larger values of r. For example, detecting affine degeneracies in $\mathbb{R}^d$ is dSUM-hard. The problem rSUM can be solved in time $O(n^{(r+1)/2})$ when r is odd, or in time $O(n^{r/2} \log n)$ when r is even. We conjecture that these algorithms are optimal. However, the best lower bound in any general model of computation, for any fixed r, is only $\Omega(n \log n)$ [3].

Theorem 3.1. *Deciding whether the convex hull of n points in $\mathbb{R}^d$ is simplicial is $\lceil d/2 \rceil$SUM-hard.*

Proof: We describe the proof explicitly only for the case $d = 5$; generalizing the proof to higher dimensions is straightforward.

Given a set of integers $X = \{x_1, x_2, \ldots, x_n\}$, we first replace them with the larger set $X' = \{x_1^\flat, x_1^\sharp, x_2^\flat, x_2^\sharp, \ldots, x_n^\flat, x_n^\sharp\}$, where $x_i^\flat = x_i - 2^{-i}$ and $x_i^\sharp = x_i + 2^{-i}$ for all i. We then consider the points $\omega_5(X')$ obtained by lifting X' onto the weird moment curve in $\mathbb{R}^5$. To prove the theorem, it suffices to show that the points $\omega_5(X')$ have a simplicial convex hull if and only if no three elements of X sum to zero.

The sum of any six elements of X' can only be an integer if the various "fudge factors" $\pm 2^{-i}$ cancel out. This is only possible if the six elements consist of three matched pairs $a^\flat, a^\sharp, b^\flat, b^\sharp, c^\flat, c^\sharp$ for some $a, b, c \in X$. Thus, if no three elements of X sum to zero, then no six elements of X' sum to zero. In this case, Lemma 2.1 implies that no six points in $\omega_5(X')$ lie on a single hyperplane, so the convex hull of $\omega_5(X')$ must be simplicial.

Conversely, suppose that $a+b+c = 0$ for some $a, b, c \in X$. This immediately implies $a^\flat + a^\sharp + b^\flat + b^\sharp + c^\flat + c^\sharp = 0$, and thus, by Lemma 2.1, the corresponding points in $\omega_5(X')$ all lie on a single hyperplane. Moreover, by Lemma 2.2, this hyperplane supports a facet of the convex hull of $\omega_5(X')$, since no other elements of X' lie in the intervals $(a^\flat, a^\sharp)$, $(b^\flat, b^\sharp)$, or $(c^\flat, c^\sharp)$. Thus, the convex hull of $\omega_5(X')$ is not simplicial. $\square$

The best lower bound we can ever hope to derive using this reduction is $\Omega(n^{\lceil d/4 \rceil} + n \log n)$, which is significantly smaller than the best known upper bound

$O(n^{\lfloor d/2 \rfloor} + n \log n)$, except for the single case $d = 5$. In particular, Theorem 3.1 tells us absolutely nothing about the four-dimensional case, since we already have a lower bound of $\Omega(n \log n)$ in all dimensions.

In an earlier paper [11], we derived an $\Omega(n^{\lceil r/2 \rceil})$ lower bound for rSUM in the r-linear decision tree model, matching the non-uniform upper bounds. In this model, decisions are based on the signs of arbitrary affine combinations of r or fewer input variables. Unfortunately, since the reduction described by the previous theorem does not follow this model, we do not automatically get similar lower bounds for convex hull simpliciality. In the next section of the paper, we derive such lower bounds directly.

4 The Main Lower Bound

Our main result is based on the following combinatorial construction.

Lemma 4.1. *For all n and d, there is a quasi-simplicial polytope in $\mathbb{R}^d$ with $O(n)$ vertices and $\Omega(n^{\lceil d/2 \rceil - 1})$ degenerate facets.*

Proof: First consider the case when d is odd, and let $r = (d-1)/2$. Without loss of generality, we assume that n is a multiple of r. Let X denote the following set of $n + 2n/r = O(n)$ integers.

$$X = \{-rn, -rn+r, \ldots, -r; r, r+1, 2r, 2r+1, \ldots, n, n+1\}$$

We can specify a degenerate facet of $\omega_5(X)$ as follows. Arbitrarily choose r elements $a_1, a_2, \ldots, a_r \in X$, all positive multiples of r. Let $a_0 = -\sum_{i=1}^r a_i$, let $b_0 = a_0 + r$, and for all $i > 0$, let $b_i = a_i + 1$. Each a_i and b_i is a unique element of X, and no element of X lies between a_i and b_i for any i. The points $\omega_d(a_i)$ and $\omega_d(b_i)$ all lie on a single hyperplane by Lemma 2.1, since

$$\sum_{i=0}^r a_i + b_i = 2 \sum_{i=0}^r a_i = 0.$$

Moreover, since any pair of elements of $X \setminus \{a_i, b_i\}$ has an even number of elements of $\{a_i, b_i\}$ between them, Lemma 2.2 implies that these points are the vertices of a single facet of $\operatorname{conv}(\omega_d(X))$. There are at least $\binom{n/r}{r} = \Omega(n^r)$ ways of choosing such a degenerate facet.

In the case where d is even, let $r = d/2 - 1$, and assume without loss of generality that n is a multiple of r. Let X be the following set of $n + 2n/r + 1 = O(n)$ integers.

$$X = \{-n - rn, -n - rn + r, \ldots, -n - r; \\ r, r+1, 2r, 2r+1, \ldots, n, n+1; 2n\}.$$

Using similar arguments as above, we easily observe that the polytope $\operatorname{conv}(\omega_d(X))$ has $\Omega(n^r)$ degenerate facets, each of which has $\omega_d(2n)$ as a vertex. □

This result is the best possible when d is odd, since an odd dimensional n-vertex polytope has at most $O(n^{(d-1)/2})$ facets [31]. In the case where d is even, the best known upper bound is $O(n^{d/2})$, which is a factor of n bigger than the result we prove here. We conjecture that this upper bound is tight. However, if we consider only sets of points on the weird moment curve, the bound given in the lemma is tight. That is, the convex hull of any set of n points on ω_d has at most $O(n^{\lceil d/2 \rceil - 1})$ degenerate facets.

We now prove the main result of the paper.

Theorem 4.2. *Any decision tree that decides whether the convex hull of a set of n points in $\mathbb{R}^d$ is simplicial, using only sidedness queries, must have depth $\Omega(n^{\lceil d/2 \rceil - 1} + n \log n)$.*

Proof: Let X be the set of numbers described in the proof of Lemma 4.1, and let $X' = X + 1/(2d+2) = \{x + 1/(2d+2) \mid x \in X\}$. Initially, the adversary presents the set of points $\omega_d(X')$. Since $\sum_{i=0}^d x'_i$ is always a half-integer, this point set is affinely nondegenerate, so its convex hull is simplicial.

It suffices to consider the case where d is odd. Let $r = (d-1)/2$. Choose $a'_0, b'_0, a'_1, b'_1, \ldots, a'_r, b'_r \in X'$ so that $\sum_{i=0}^r (a'_i + b'_i) = 1/2$ and no other elements of X' lie between a'_i and b'_i for any i. The corresponding points $\omega_d(a'_i), \omega_d(b'_i)$ form a *collapsible simplex*. To collapse it, the adversary simply moves the points back to their original positions in $\omega_d(X)$. Lemmas 2.1 and 2.2 imply that the collapsed simplex forms a degenerate facet of the new convex hull. Since the sum of any other $(d+1)$-tuple changes by at most $1/2 - 1/(2d+2)$, no other simplex changes orientation. In other words, the only way for an algorithm to distinguish between the original configuration and the collapsed configuration is to perform a sidedness query on the collapsible simplex.

Thus, if an algorithm does not perform a separate sidedness query on every collapsible simplex, then the adversary can introduce a degenerate facet that the algorithm cannot detect. There are $\Omega(n^{\lceil d/2 \rceil - 1})$ collapsible simplices, one for each of the degenerate facets in $\operatorname{conv}(\omega_d(X))$.

Finally, the $n \log n$ term follows from the algebraic decision tree lower bound of Ben-Or [3]. □

A three-dimensional version of our construction is illustrated in Figure 1.

Our lower bound matches known upper bounds when d is odd [7]. We emphasize that if the points are known

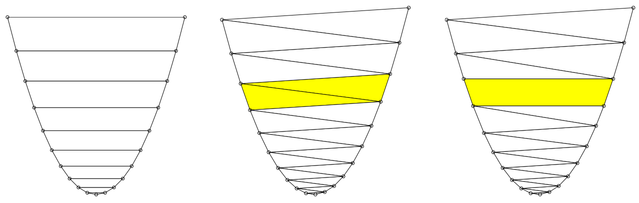

Figure 1. Our adversary construction in three dimensions. Bottom views of (a) a quasi-simplicial polytope with $\Omega(n)$ degenerate facets, (b) the simplicial adversary polytope with one collapsible simplex highlighted, and (c) the corresponding collapsed polytope.

in advance to lie on the weird moment curve, this problem can be solved in $O(n^{\lceil d/4 \rceil})$ time if $\lceil d/2 \rceil$ is odd, and in $O(n^{\lceil d/4 \rceil} \log n)$ time if $\lceil d/2 \rceil$ is even, by an algorithm that uses more complicated queries.

The convex hull of the adversary configuration $\omega_d(X')$ has $\lceil d/2 \rceil - 1$ more facets than the convex hull of any collapsed configuration. Thus, we immediately have the following lower bound.

Theorem 4.3. *Any decision tree that computes the number of convex hull facets of a set of n points in $\mathbb{R}^d$, using only sidedness queries, must have depth $\Omega(n^{\lceil d/2 \rceil - 1} + n \log n)$.*

A simple modification of our argument implies the following "output-sensitive" version of our lower bound.

Theorem 4.4. *Any decision tree that decides whether the convex hull of a set of n points in $\mathbb{R}^d$ is simplicial or computes the number of convex hull facets, using only sidedness queries, must have depth $\Omega(f)$ when d is odd, and $\Omega(f^{1-2/d})$ when d is even, where f is the total complexity of the convex hull.*

Proof: We construct a modified adversary polytope as follows. We start by constructing an adversary polytope with f faces, exactly as described in the proof of Theorem 4.2. When d is odd, this polytope is the convex hull of $\Theta(f^{2/(d-1)})$ points on the wired moment curve, and has $\Omega(f)$ collapsible simplices. When d is even, the polytope is the convex hull of $\Theta(f^{2/d})$ points and has $\Omega(f^{1-2/d})$ collapsible simplices.

By introducing a new vertex extremely close to the relative interior of any facet of a simplicial polytope, we can split that facet into d smaller facets. Each such split increases the number of polytope faces by $2^d - 2$. To bring the number of vertices of our adversary polytope up to n, we choose some facet and repeatedly split it in this fashion. The augmented polytope has at most $f + (2^d - 2)n = O(f)$ faces. To avoid moving the new vertices, we no longer consider a simplex collapsible if it shares a vertex with the split facet, but this eliminates less than a constant fraction of the collapsible simplices. □

5 Other Computational Primitives

In this section, we identify a general class of computational primitives which, if added to our model of computation, do not affect our lower bounds. In fact, even if we allow any finite number of these primitives to be performed at no cost, the number of required sidedness queries is the same. These primitives include comparisons between coordinates of input points in any number of directions, comparisons between coordinates of hyperplanes defined by d-tuples of points, and in-sphere queries.

The primitives we consider are all *algebraic queries*. The result of an algebraic query is given by the sign of a multivariate *query polynomial*, evaluated at the coordinates of the input. If the sign is zero (resp. nonzero), we say that the input is *degenerate* (resp. *nondegenerate*) with respect to that query. For example, a set of points is affinely degenerate if and only if it is degenerate with respect to some sidedness query.

A *projective transformation* of $\mathbb{R}^d$ (or more properly, of the projective space $\mathbb{RP}^d$) is any map that takes hyperplanes to hyperplanes. If we represent the points of $\mathbb{R}^d$ in homogeneous coordinates, a projective transformation is equivalent to a linear transformation of $\mathbb{R}^{d+1}$. In Stolfi's two-sided projective model [28], projective maps preserve (or reverse) the orientation of every simplex in $\mathbb{R}^d$, and thus preserve the combinatorial structure of convex hulls. See Chapter 14 of [28].

Let X be the set of numbers described in the proof of Lemma 4.1. We call an algebraic query *allowable* if for some projective transformation ϕ, the configuration $\phi(\omega_d(X))$ is nondegenerate with respect to that query. Our choice of terminology is justified by the following theorem.

Theorem 5.1. *Any decision tree that decides whether the convex hull of n points in $\mathbb{R}^d$ is simplicial, using only sidedness queries and a finite number of allowable queries, requires $\Omega(n^{\lceil d/2 \rceil - 1})$ sidedness queries in the worst case.*

Proof: If *some* projective transformation makes $\omega_d(X)$ nondegenerate with respect to an algebraic query, then *almost every* projective transformation (*i.e.*, all but a measure zero subset) makes $\omega_d(X)$ nondegenerate. Thus, for any finite set of allowable queries, almost every projective transformation makes $\omega_d(X)$ nondegenerate with respect to all of them. Let ϕ be such a transformation.

If $\phi(\omega_d(X))$ is nondegenerate with respect to some finite set of allowable queries, then for all X' in an open neighborhood of X in $\mathbb{R}^n$, the configuration $\phi(\omega(X'))$ is also nondegenerate with respect to that set of queries.

The theorem now follows from a slight modification of the proof of Theorem 4.2. Let $\varepsilon > 0$ be some sufficiently small real number. The set $\phi(\omega_d(X+\varepsilon))$ has a simplicial convex hull, but has $\Omega(n^{\lceil d/2 \rceil + 1})$ collapsible simplices, each corresponding to a degenerate facet in $\phi(\omega_d(X))$. No allowable query can distinguish between $\phi(\omega_d(X+\varepsilon))$ and any collapsed configuration, or even between $\phi(\omega_d(X+\varepsilon))$ and $\phi(\omega_d(X))$. $\square$

We characterize allowable queries algebraically as follows. Consider the degenerate configuration $\omega_d(X)$ as a single point in the configuration space $\mathbb{R}^{dn}$. Any algebraic query induces a surface in configuration space, consisting of all configurations that are degenerate with respect to that query. Since any projective map ϕ can be represented by a $(d+1) \times (d+1)$ matrix with determinant ± 1, the set of projectively transformed configurations $\phi(\omega_d(X))$ forms a $(d^2 + 2d)$-dimensional algebraic variety in configuration space. Any query whose surface does not completely contain this variety is allowable.

We give below a (nonexhaustive!) list of allowable queries. We leave the proofs that these queries are in fact allowable as easy exercises.

- Comparisons between points in any fixed direction are allowable. In fact, we can allow the input points to be pre-sorted in any finite number of fixed directions. Seidel describes a similar result in three dimensions [25, Theorem 5]. We emphasize that the directions in which these comparisons are made must be fixed in advance. No matter how we transform the adversary configuration, there is always *some* direction in which a point comparison can distinguish it from a collapsed configuration.

- More generally, deciding which of two points is hit first by a hyperplane rotating around a fixed $(d-2)$-flat is allowable. We can even pre-sort the points by their cyclic orders around any finite number of fixed $(d-2)$-flats. If the $(d-2)$-flat is "at infinity", then "rotation" is just translation, and we have the previous notion of point comparison. We can interpret this type of query in dual space as a comparison between the intersections of two hyperplanes with a fixed line. Again, we emphasize that the $(d-2)$-flats must be fixed in advance.

- Sidedness queries in any fixed lower-dimensional projection are allowable. This is a natural generalization of point comparisons, which can be considered sidedness queries in a one-dimensional projection. We can even specify in advance the complete order types of the projections onto any finite number of fixed affine subspaces. As a technical point, we would not actually include this information as part of the input, since this would drastically increase the input size. Instead, knowledge of the projected order types would be "hard-wired" into the algorithm.

- "Second-order" comparisons between vertices of the dual hyperplane arrangement, in any fixed direction, are also allowable. Such a query can be interpreted in the primal space as a comparison between the intersections of two hyperplanes, each defined by a d-tuple of input points, with a fixed line. To prove that such a query is allowable, it suffices to observe that a projective transformation of the primal space induces a projective transformation of the dual space, and vice versa. Note that a second-order comparison is algebraically equivalent to a sidedness query if the two d-tuples share $d-1$ points.

- In-sphere queries are allowable. Given $d+2$ points, an in-sphere query asks whether the first point lies "inside", on, or "outside" the oriented sphere determined by the other $d+1$ points. Similarly, in-sphere queries in any fixed lower-dimensional projection are allowable.

- Distance comparisons between pairs of points or pairs of projected points are allowable. More generally, comparing the measures of pairs of simplices of dimension less than d — for example, comparing the areas of two triangles — defined either by the original points or by any fixed projection, are allowable.

On the other hand, comparing the volumes of arbitrary simplices of *full* dimension is not allowable. In any projective transformation of $\omega_d(X)$, all of the degenerate simplices have the same (zero) volume. It is not possible to collapse a simplex in any adversary configuration while maintaining the order of the volumes of the other collapsible simplices.

6 Our Models vs. Real Convex Hull Algorithms

A large number of convex hull algorithms rely (or can be made to rely) exclusively on sidedness queries. These include the "gift-wrapping" algorithms of Chand and

Kapur [6] and Swart [29], the "beneath-beyond" method of Seidel [24], Clarkson and Shor's [10] and Seidel's [27] randomized incremental algorithms, Chazelle's worst-case optimal algorithm [7], and the recursive partial-order algorithm of Clarkson [9].

Seidel's "shelling" algorithm [26] and the space-efficient gift-wrapping algorithms of Avis and Fukuda[1] [2] and Rote [22] require only sidedness queries and second-order comparisons. Matoušek [20] and Chan [4] improve the running times of these algorithms (in an output-sensitive sense), by finding the extreme points more quickly. Clarkson [9] describes a similar improvement to a randomized incremental algorithm. Since every point in our adversary configuration is extreme, our lower bound still holds even if the extremity of a point can be decided for free. We are not suggesting that the computational primitives used by these algorithms cannot be used to break our lower bounds; only that the ways in which these primitives are currently applied are inherently limited.

Chan [4] describes an improvement to the gift-wrapping algorithm, using ray shooting data structures of Agarwal and Matoušek [1] and Matoušek and Schwarzkopf [19] to speed up the pivoting step. In each pivoting step, the gift-wrapping algorithm finds a new facet containing a given ridge of the convex hull. In the dual, this is equivalent to shooting a ray from a vertex of the dual polytope along one of its outgoing edges. The dual vertex that the ray hits corresponds in the primal to the new facet. A single pivoting step tells us the orientation of $n-d$ simplices, all of which share the d vertices of the new facet. However, at most one of these simplices can be collapsible, since two collapsible simplices share at most $d/2$ vertices. Thus, even if we allow a pivoting step to be performed in constant time, our lower bound still holds.

There are a few convex hull algorithms which seem to fall outside our framework, most notably the divide and conquer algorithm of Chan, Snoeyink, and Yap [5]. The two-dimensional version of their algorithm uses sidedness queries, along with first-, second-, and even *third*-order comparisons; higher-dimensional versions use even more complex primitives.

7 Related Results

7.1 Affine Degeneracies

Using similar methods, we can derive a very simple proof of the following theorem of Erickson and Seidel [12]. In the original proof of this theorem, the existence of the adversary set was proven nonconstructively. Here, we derive it explicitly.

[1] at least if Bland's pivoting rule is used

Theorem 7.1. *Any decision tree that decides whether a set of n points in $\mathbb{R}^d$ is affinely nondegenerate, using only sidedness queries, must have depth $\Omega(n^d)$. If $d \geq 3$, this lower bound holds even when the points are known in advance to be in convex position.*

Proof: Let X denote the set of integers from $-dn$ to n, and let $X' = X + 1/(2d+2)$. The adversary initially presents the point set $\omega_d(X')$. This point set is affinely nondegenerate, since the expression $\sum_i x_i'$ is always a half-integer.

Choose arbitrary distinct positive elements $x_1, x_2, \ldots, x_d \in X$, and let $x_0 = -\sum_i x_i$; this is also an element of X. Then the points $\omega_d(x_i')$ form a collapsible simplex. To collapse it, the adversary just shifts the points back down to $\omega_d(x_i)$; the collapsed simplex is obviously degenerate. Since the expression $\sum_{i=0}^{d} x_i'$ changes by at most $1/2 - 1/(2d+2)$ for any other simplex, no other simplex changes orientation.

Thus, if an algorithm does not perform a sidedness query on every collapsible simplex, the adversary can introduce an affine degeneracy that the algorithm cannot detect. There are at least $\binom{n}{d} = \Omega(n^d)$ such simplices. If $d \geq 3$, the original point set and each collapsed point set is in convex position. $\square$

The previous theorem easily generalizes to allow additional queries, as described in Section 5.

Theorem 7.2. *Any decision tree that decides whether a set of n points in $\mathbb{R}^d$ is affinely nondegenerate, using only sidedness queries and a finite number of allowable queries, requires $\Omega(n^d)$ sidedness queries in the worst case. If $d \geq 3$, this lower bound holds even when the points are known in advance to be in convex position.*

We note in passing that adding allowable queries to our model allows us to prove nontrivial quadratic lower bounds for several 3SUM-hard problems. Previous techniques [12, 11] resulted in models that were too weak to solve many of these problems.

7.2 Circular Degeneracies

We also easily prove the following related theorem first proven in [12]. A set of points in the plane is *circularly degenerate* if any four points lie on a circle. The basic computational primitive used to detect circular degeneracies is the *in-circle query*: Given four points, is the first points inside, on, or outside the oriented circle defined by the other three points? In-circle queries are typically answered by lifting the points to the unit paraboloid $z = x^2 + y^2$, or stereographically projecting them onto a sphere, and performing a three-dimensional sidedness query.

Theorem 7.3. *Any decision tree that decides whether n points in $\mathbb{R}^2$ is circularly degenerate, using only incircle queries, must have depth $\Omega(n^3)$.*

Proof: Four points $(a, a^2), (b, b^2), (c, c^2), (d, d^2)$ on the unit parabola are cocircular if and only if $a+b+c+d = 0$. Let X be the set of integers from $-n$ to n. There are $\Theta(n^3)$ 4-tuples in X whose sums are zero. The adversary presents a set of points on the unit parabola with x-coordinates taken from the set $X + 1/8$. This set is non-degenerate and has $\Omega(n^3)$ collapsible 4-tuples. □

We can extend the model of computation in a similar fashion as before, but with a different set of new queries. A *linear fractional transformation* of the plane (or more formally, of the Riemann sphere $\mathbb{CP}^1$) is any transformation that maps circles to circles. If we represent the points of $\mathbb{R}^2$ in complex homogeneous coordinates — representing $(x, y) \in \mathbb{R}^2$ by any complex multiple of $(1 + 0i, x + yi) \in \mathbb{C}^2$ — then a linear fractional transformation is equivalent to a linear transformation of $\mathbb{C}^2$.

We say that a query is *circularly allowable* if some linear fractional transformation of the set (X, X^2) is non-degenerate with respect to that query, where X is the set of numbers described in the proof of Theorem 7.3. Circularly allowable queries include first- and second-order comparisons and sidedness queries, but do not include comparisons between arbitrary incircle determinants.

Arguments similar to those in Section 5 give us the following theorem.

Theorem 7.4. *Any decision tree that decides whether n points in $\mathbb{R}^2$ is circularly degenerate, using only incircle queries and a finite number of circularly allowable queries, requires $\Omega(n^3)$ incircle queries in the worst case.*

Erickson and Seidel give a nonconstructive proof that $\Omega(n^{d+1})$ in-sphere queries are require to detect spherical degeneracies in $\mathbb{R}^d$ [12]. We have been unable to generalize our constructive proof of the two-dimensional case to higher dimensions. A direct proof for the general case would follow from the construction of a set of numbers having $\Omega(n^{d+1})$ $(d+2)$-tuples in the zeroset of a certain symmetric polynomial. For example, in three dimensions, we need $\Omega(n^4)$ 5-tuples in the zeroset of the polynomial

$$1 + \sum_{1 \le i \le 5} x_i^2 + \sum_{1 \le i < j \le 5} x_i x_j.$$

8 Conclusions and Open Problems

We have presented new lower bounds on the worst-case complexity of detecting simplicial convex hulls or counting convex hull facets, in a fairly natural model of computation. Our lower bounds follow from a simple adversary argument, based on the construction of a convex polytope with a large number of degenerate features. In order to be correct, any algorithm must individually check that each of those degenerate features is not present in the input. Similar arguments give us new simple proofs of lower bounds for degeneracy detection problems.

A number of open problems remain to be answered. While our lower bounds match existing upper bounds in odd dimensions, where is still a gap when the dimension is even. A first step in improving our lower bounds would be to improve the combinatorial bounds in Lemma 4.1. Is there a quasi-simplicial four-dimensional polytope with n vertices and $\Omega(n^2)$ degenerate facets? Simple variations on the weird moment curve will not suffice, since an "evenness condition" like Lemma 2.2 always forces the number of degenerate facets to be linear.

A common application of convex hull algorithms is the construction of Delaunay triangulations and Voronoi diagrams. Are $\Omega(n^{\lceil d/2 \rceil})$ in-sphere queries required to decide if the Delaunay triangulation is simplicial (*i.e.*, really a triangulation)? Again, a first step is to construct a Delaunay triangulation with $\Omega(n^{\lceil d/2 \rceil})$ independent degenerate features.

Another major open problem is to strengthen the lower bounds in which our lower bounds hold. Quadratic lower bounds for either the five-dimensional convex hull problem or the two-dimensional affine degeneracy problem in stronger models of computation would imply similar lower bounds for a number of other 3SUM-hard problems. While the lower bounds we prove here and in earlier papers [12, 11] are in fairly natural models, there are still 3SUM-hard problems that cannot even be solved in these models. For example, one of the simplest problems for which our techniques fail is finding the minimum area triangle determined by a set of points in the plane. In order to prove a useful lower bound for this problem, we must consider a model that allows comparison of signed triangle areas. It seems impossible to apply our "collapsible simplex" adversary argument in such a model; a radically new idea is called for.

Ultimately, of course, we would like a lower bound bigger than $\Omega(n \log n)$ that holds in some general model of computation, such as algebraic decision trees or algebraic computation trees.

References

[1] P. K. Agarwal and J. Matoušek. Ray shooting and parametric search. *SIAM J. Comput.*, 22(4):794–806, 1993.

[2] D. Avis and K. Fukuda. A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. *Discrete Comput. Geom.*, 8:295–313, 1992.

[3] M. Ben-Or. Lower bounds for algebraic computation trees. In *Proc. 15th Annu. ACM Sympos. Theory Comput.*, pages 80–86, 1983.

[4] T. M. Y. Chan. Output-sensitive results on convex hulls, extreme points, and related problems. In *Proc. 11th Annu. ACM Sympos. Comput. Geom.*, pages 10–19, 1995.

[5] T. M. Y. Chan, J. Snoeyink, and C.-K. Yap. Output-sensitive construction of polytopes in four dimensions and clipped Voronoi digrams in three. In *Proc. 6th ACM-SIAM Sympos. Discrete Algorithms*, pages 282–291, 1995.

[6] D. R. Chand and S. S. Kapur. An algorithm for convex polytopes. *J. ACM*, 17:78–86, 1970.

[7] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993.

[8] B. Chazelle and J. Matoušek. Derandomizing an output-sensitive convex hull algorithm in three dimensions. Technical report, Dept. Comput. Sci., Princeton Univ., 1992.

[9] K. L. Clarkson. More output-sensitive geometric algorithms. In *Proc. 35th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 695–702, 1994.

[10] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[11] J. Erickson. Lower bounds for linear satisfiability problems. In *Proc. 6th ACM-SIAM Sympos. Discrete Algorithms*, pages 388–395, 1995.

[12] J. Erickson and R. Seidel. Better lower bounds on detecting affine and spherical degeneracies. *Discrete Comput. Geom.*, 13(1):41–57, 1995.

[13] A. Gajentaan and M. H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom. Theory Appl.*, 5:165–185, 1995.

[14] D. Gale. Neighborly and cyclic polytopes. In V. Klee, editor, *Convexity*, volume VII of *Proc. Symposia in Pure Mathematics*, pages 225–232. Amer. Math. Soc., 1963.

[15] R. L. Graham. An efficient algorithm for determining the convex hull of a finite planar set. *Inform. Process. Lett.*, 1:132–133, 1972.

[16] B. Grünbaum. *Convex Polytopes.* Wiley, New York, NY, 1967. Revised edition, V. Klee and P. Kleinschmidt, editors, *Graduate Texts in Mathematics*, Springer-Verlag, in preparation.

[17] C. G. J. Jacobi. De functionibus alternantibus earumque divisione per productum e differentiis elementorum conflatum. *J. Reine Angew. Mathematik*, 22:360–371, 1841. Reprinted in *Gesammelten Werke* III, G. Reimer, Berlin, 1884.

[18] D. G. Kirkpatrick and R. Seidel. The ultimate planar convex hull algorithm? *SIAM J. Comput.*, 15:287–299, 1986.

[19] J. Matoušek and O. Schwarzkopf. On ray shooting in convex polytopes. *Discrete Comput. Geom.*, 10(2):215–232, 1993.

[20] J. Matoušek. Linear optimization queries. *J. Algorithms*, 14:432–448, 1993. The results combined with results of O. Schwarzkopf also appear in *Proc. 8th ACM Sympos. Comput. Geom.*, 1992, pages 16–25.

[21] F. P. Preparata and S. J. Hong. Convex hulls of finite sets of points in two and three dimensions. *Commun. ACM*, 20:87–93, 1977.

[22] G. Rote. Degenerate convex hulls in high dimensions without extra storage. In *Proc. 8th Annu. ACM Sympos. Comput. Geom.*, pages 26–32, 1992.

[23] I. Schur. *Über eine Klasse von Matrizen, die sich einer gegebenen Matrix zuordnen lassen.* Thesis, Berlin, 1901. Reprinted in *Gesammelte Abhandlungen*, Springer, 1973.

[24] R. Seidel. A convex hull algorithm optimal for point sets in even dimensions. M.Sc. thesis, Dept. Comput. Sci., Univ. British Columbia, Vancouver, BC, 1981. Report 81/14.

[25] R. Seidel. A method for proving lower bounds for certain geometric problems. In G. T. Toussaint, editor, *Computational Geometry*, pages 319–334. North-Holland, Amsterdam, Netherlands, 1985.

[26] R. Seidel. Constructing higher-dimensional convex hulls at logarithmic cost per face. In *Proc. 18th Annu. ACM Sympos. Theory Comput.*, pages 404–413, 1986.

[27] R. Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete Comput. Geom.*, 6:423–434, 1991.

[28] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations.* Academic Press, New York, NY, 1991.

[29] G. F. Swart. Finding the convex hull facet by facet. *J. Algorithms*, 6:17–48, 1985.

[30] A. C. Yao. A lower bound to finding convex hulls. *J. ACM*, 28:780–787, 1981.

[31] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, 1994.