A Bound on Local Minima of Arrangements that implies the Upper Bound Theorem

Kenneth L. Clarkson
AT&T Bell Laboratories
Murray Hill, New Jersey 07974
e-mail: clarkson@research.att.com

Abstract

This paper shows that the i-level of an arrangement of hyperplanes in E^d has at most $\binom{i+d-1}{d-1}$ local minima. This bound follows from ideas previously used to prove bounds on $(\leq k)$-sets. Using linear programming duality, the Upper Bound Theorem is obtained as a corollary, giving yet another proof of this celebrated bound on the number of vertices of a simple polytope in E^d with n facets.

1 Introduction

We will need some terminology for arrangements, similar to that in Edelsbrunner's text[3]. Let $\mathcal{A}(H)$ be a simple arrangement of a set H of n hyperplanes in E^d. For $h \in H$, let the upper halfspace h^+ be the open halfspace bounded by h that contains $(\infty, 0, \ldots, 0)$, and let the lower halfspace h^- be the other open halfspace bounded by h. Say that $x \in E^d$ is above $h \in H$ if $x \in h^+$, and below h if $x \in h^-$. The i-level of $\mathcal{A}(H)$ is the boundary of the set of points that are below no more than i hyperplanes of H. Thus for example the 0-level of $\mathcal{A}(H)$ is the boundary of the convex polytope $\mathcal{P}(H) = \bigcap_{h \in H}(h^+ \cup h)$. The maximum number of vertices of $\mathcal{A}(H)$ on its i-level is a combinatorial problem of long standing. While some results have long been known for $d = 2$ [4], and recently sharpened slightly[8], only relatively recently have nontrivial bounds been known for the general problem in higher dimensions. These results are stated in a dual form, concerning *k-sets* of sets of points. One related result is that the maximum total number of vertices on all i-levels, for $i \leq k$, is $\Theta(n^{\lfloor d/2 \rfloor} k^{\lceil d/2 \rceil})$, a $(\leq k)$-*set* bound[2].

Using similar techniques, Mulmuley then showed that the number of local minima on levels $i \leq k$ is $O(k^d)$, where a local minimum is point of the i-level such that all points on the i-level in a neighborhood of the point have a larger

x_1 coordinate. Call a local minimum on the i-level an *i-minimum*, or an $(\leq k)$-minimum if $i \leq k$. An i-minimum is a vertex, and so Mulmuley's result is a bound on a class of vertices of the i-level. Note that the 0-minimum of H is the solution $x^*(H)$ of the linear programming problem $\min\{x_1 \mid x \in \mathcal{P}(H)\}$. In addition to bounding the number of $(\leq k)$-minima, Mulmuley showed some bounds on related quantities, and conjectured that the number of i-minima is $O(i^{d-1})$ for every i [7]. This conjecture is confirmed here by the bound $\binom{i+d-1}{d-1}$, proven in the next section using the same technique as for bounds on $(\leq k)$-sets and $(\leq k)$-minima.

This i-minima bound is of course not new for $i = 0$ and $i = 1$, and it isn't even new for $i = n/2$: using (projective or polar) duality, it is equivalent to the preliminary observation for $d = 2$ that forms the basis of a bound on the number of vertices on the $n/2$-level in E^2 [4]. Thus the contribution here is mostly one of observed connections and new proofs, and not new theorems.

Section 3 uses ideas of linear programming duality to show that the bound on i-minima readily implies the celebrated Upper Bound Theorem for convex polytopes[6, 1]. Here we mean only the upper bound of that theorem, and do not characterize the polytopes for which the bound is tight.

2 The bound for i-minima

Some preliminary notation: for a set S, let $\binom{S}{k}$ denote the collection of subsets of S of size k, so $|\binom{S}{k}| = \binom{|S|}{k}$. We will sometimes use the coordinate-wise partial order on E^d where $x, y \in E^d$ have $x > y$ if $x_i > y_i$ for $i = 1 \ldots d$.

The bound for i-minima follows from the following well-known properties of solutions of linear programming problems.

Lemma 2.1 *Any arrangement $\mathcal{A}(H)$ has at most one 0-minimum $x^*(H)$, and if it exists, there is $B \subset H$ of size d with $x^*(B) = x^*(H)$.*

Proof. Omitted; the second statement follows from Helly's theorem, as applied to the upper halfspaces of H and the halfspaces $\{x_1 \leq q\}$, for all q smaller than the first coordinate of $x^*(H)$. □

Call the set B promised by the lemma a *basis* $b(H)$ of H. We can extend the notations $\mathcal{P}(H)$, $x^*(H)$, and $b(H)$ to subsets of H in the obvious way; however, for many $G \subseteq H$, the linear programming problem $\mathcal{LP}(G)$, of finding $\min\{x_1 \mid x \in \mathcal{P}(G)\}$, may be unbounded, or have many solutions, and even if $x^*(G)$ is unique, there may not be a unique basis $b(G)$. To apply the lemma and bound i-minima, the definitions of $x^*(G)$ and $b(G)$ are extended below to all $G \subseteq H$, using lexicographic orders, such that every $G \subseteq H$ has a unique basis.

A point $x = (x_1, \ldots, x_d)$ is lexicographically (lex) smaller than point $y = (y_1, \ldots, y_d)$, written $x \prec y$, if $x_i < y_i$ for the smallest i at which their coordinates differ. For sufficiently small $\epsilon > 0$ we have $x \prec y$ if and only if $x \cdot b_\epsilon < y \cdot b_\epsilon$, where $b_\epsilon = (1, \epsilon, \epsilon^2, \ldots, \epsilon^{d-1})$.

We broaden the definition of local minimum to include vertices that have lexicographically minimal (lexmin) coordinates in a neighborhood of the i-level. Thus for $G \subset H$, if the associated linear programming problem $\mathcal{LP}(G)$ has a bounded solution, then $x^*(G)$ exists and is unique. Note that a basis $b(G)$ yielding $x^*(G)$ also exists.

Extend the definition of $x^*(G)$ to the unbounded case as follows: choose a sufficiently small value K so that all vertices v of $\mathcal{A}(H)$ have all coordinates larger than K. Define $\underline{x}^*(G)$ as the lexmin point in $\mathcal{P}(G)$ with all coordinates no smaller than K.

With these definitions, all $G \subseteq H$ have a 0-minimum $\underline{x}^*(G)$, which is the same as the initial definition when $\mathcal{LP}(G)$ has a unique vertex with minimum x_1 coordinate. It remains to appropriately extend the notion of basis $b(G)$. Here again lexicography is useful.

Given a set S of integers $\{i \mid 1 \leq i \leq n\}$, the lexicographic order on $\binom{S}{k}$ is as follows: for $A, B \in \binom{S}{k}$, order A and B so that $A = \{a_1 \ldots, a_k\}$ and $a_1 < a_2 \cdots < a_k$ and similarly order $B = \{b_1 \ldots, b_k\}$. Now $A \prec B$ if and only if $a_i < b_i$ at the smallest index i at which they differ.

We impose a lexicographic order on $\binom{H}{d}$ by numbering the hyperplanes of H arbitrarily from 1 to n and then saying $A, B \in \binom{H}{d}$ have $A \prec B$ if and only if the associated sets of numbers A' and B' have $A' \prec B'$.

To define the basis $b(G)$ for $G \subset H$, let $b(G)$ denote the lexmin $B \in \binom{G}{d}$ so that $\underline{x}^*(B) = \underline{x}^*(G)$. Note that some of the hyperplanes determining $\underline{x}^*(G)$ may be of the form $x_i \geq K$, if $\mathcal{LP}(G)$ if unbounded and $x^*(G)$ does not exist; they are replaced in $b(G)$ by the smallest-numbered elements of G that are not above $\underline{x}^*(G)$.

An i-basis is defined as follows. For $B \in \binom{H}{d}$, note that $b(B) = B$, and define

$$I_B \equiv \{h \in H \mid b(B \cup \{h\}) \neq B\}.$$

That is, an element $h_j \in I_B$ is either above $\underline{x}^*(B)$, or there is some $h_k \in B$ with $j < k$ and

$$\underline{x}^*(B \setminus \{h_k\} \cup \{h_j\}) = \underline{x}^*(B),$$

so a lexicographically smaller subset with the same minimum can be obtained. If I_B has i members, call B an i-basis. Note that every i-minimum has a corresponding i-basis. We will count the i-minima by counting the i-bases.

Let $g_i(H)$ denote the number of i-minima of H, and let $g'_i(H)$ denote the number of i-bases. We have the following theorem.

Theorem 2.2 *If $\mathcal{A}(H)$ is an arrangement of n hyperplanes in E^d, then $g_i(H) \leq g'_i(H) = \binom{i+d-1}{d-1}$.*

Proof. As discussed above, each i-minimum of $\mathcal{A}(H)$ has a corresponding i-basis, and each i-basis determines at most one i-minimum, so $g_i(H) \leq g'_i(H)$ and it suffices to count the i-bases. Consider a random $R \in \binom{H}{r}$, where $d \leq r \leq n$. Here each element of $\binom{H}{r}$ is equally likely. Any subset has exactly one basis. On the other hand, we can express the expected number of bases of R as

$$\sum_{B \in \binom{H}{d}} \text{Prob}\{B \subset R, R \subseteq H \setminus I_B\},$$

since $B \in \binom{H}{d}$ is the basis of R if and only if $B \subset R$ and no elements of I_B appear in R. If B is an i-basis, the number of subsets $R \in \binom{H}{r}$ with $b(R) = B$ is $\binom{n-i-d}{r-d}$, since B must be in R, and the remaining $r - d$ choices of elements of R must be from $H \setminus B \setminus I_B$. Therefore the probability that i-basis B is the basis of R is $\binom{n-i-d}{r-d}/\binom{n}{r}$, and we have

$$1 = \sum_{0 \leq i \leq n-d} \frac{\binom{n-i-d}{r-d}}{\binom{n}{r}} g'_i(H), \tag{1}$$

for $d \leq r \leq n$. This equation is a special case of Lemma 2.1 of [2]. Since the matrix corresponding to this system of $n - d + 1$ linear equations in $n - d + 1$ unknowns can be rearranged to be triangular with positive diagonal elements, the system can be solved, and the reader can verify that the solution is $\binom{i+d-1}{d-1}$.
□

This bound for $g_i(H)$ is not very good for large i, since there is at most one $(n-d)$-minimum, while there are $\binom{n-1}{d-1}$ $(n-d)$-bases. However, it is easy to show that a set B of d hyperplanes yields a minimum point x if and only x is a maximum point in $\cap_{h \in B}(h^- \cup h)$. Hence $g_i(H) = g_{n-d-i}(H)$, and we have the following theorem.

Theorem 2.3 *For any simple arrangement $\mathcal{A}(H)$ of n hyperplanes in E^d, the number of i-minima $g_i(H)$ satisfies $g_i \leq \min\{\binom{i+d-1}{d-1}, \binom{n-i-1}{d-1}\}$.*

3 The Upper Bound Theorem

The g-vector of a polytope. Suppose $\mathcal{P}$ is a simple d-polytope with at most n facets, and is the set of points $\{x \in E^d \mid Ax \leq b\}$, where A is an $n \times d$ matrix, x and b are an column n-vectors, and $b \geq 0$. Since all entries of b are nonnegative, the origin is in $\mathcal{P}$. We will also write the inequalities as $a_j x \leq b_j$, for $j = 1 \ldots n$. Suppose w is an *admissible* row n-vector for $\mathcal{P}$, meaning that $wv \neq wv'$ for any two distinct vertices v and v' of $\mathcal{P}$. Orient the edges of the $\mathcal{P}$ in the direction of increasing w (*upward*) and let $g_i(\mathcal{P})$ denote the number of vertices with outdegree i, so that i of their incident edges point up. If $f_k(\mathcal{P})$ is the number of k-faces of $\mathcal{P}$, then

$$f_k(\mathcal{P}) = \sum_i \binom{i}{k} g_i(\mathcal{P}), \tag{2}$$

since each k-face F has a unique bottom vertex v, with all k edges in F incident to v pointing up. To bound the quantities $f_k(\mathcal{P})$ it is enough to bound $g_i(\mathcal{P})$. (The above condenses the discussion in Brøndsted's text of McMullen's proof of the Upper Bound Theorem[6, 1].)

The LP-dual arrangement. The linear programming problem

$$\max\{wx \mid x \in \mathcal{P}\}$$

has the dual problem

$$\min\{yb \mid y \in \mathcal{P}'\},$$

where

$$\mathcal{P}' = \{y \in E^n \mid y \in \mathcal{F}, y \geq 0\},$$

and

$$\mathcal{F} = \{y \in E^n \mid yA = w\}$$

is an $(n-d)$-flat. Letting $d' = n - d$, the d'-polytope $\mathcal{P}'$ is one cell in the arrangement $\mathcal{A}(H)$ induced by the collection H of n hyperplanes $h_j \equiv \{y \mid y_j = 0\}$, $j = 1 \ldots n$, restricted to $\mathcal{F}$. (Note that while the previous section discussed arrangements in E^d, here we consider one in a d'-flat.) We can define local minima for this arrangement where we seek minima of yb. We have the following lemma. It is standard [5, §8.2], but for completeness a proof appears below (neglecting some issues of degeneracy).

Lemma 3.1 *There is a one-to-one correspondence between i-minima of $\mathcal{A}(H)$ and vertices of $\mathcal{P}$ with outdegree i, and so $g_i(\mathcal{P}) = g_i(H)$.*

Proof. If v is a vertex of $\mathcal{P}$, then v is the solution of $\hat{A}v = \hat{b}$, a subsystem of d rows of $Ax \leq b$. Suppose $v' \in \mathcal{F}$ has zero coordinates for all but those corresponding to the rows giving $\hat{A}$. Thus v' is a vertex of $\mathcal{A}(H)$: it is the intersection of d' hyperplanes of H with $\mathcal{F}$. The nonzero coordinates of v' are determined by $v'A = w$.

First observe that v' is a local minimum $x^*(G)$ for $G = \{h_j \mid v'_j = 0\}$: note that if $y \in \mathcal{F}$, so $yA = w$, then $yb - wx = yb - yAx = y(b - Ax)$. Thus $v'b - wv = v'(b - Av) = 0$ since $v'_j = 0$ if and only if $a_j v \neq 0$. (So v' and v has the same objective function values in the dual linear programming problems.) On the other hand, if $yA = w$ and $y_j \geq 0$ when $v'_j = 0$, we have $yb - wv = y(b - Av) \geq 0$ since $b - Av \geq 0$ and $a_j v = b_j$ when $v'_j \neq 0$. Thus if $y \in \mathcal{P}'(G)$ then $yb \geq v'b$. Note that the inequality is strict if $y_j > 0$ for some j with $a_j v < b_j$.

Next to show that if v has outdegree i then v' is an i-minimum. Since $v'_j < 0$ if and only if v' is below h_j, we need to show that a coordinate $v'_j \neq 0$ corresponds to an oriented edge (v, q) where $wv - wq = w(v - q)$ has the same sign as v'_j. Suppose (v, q) is an edge of $\mathcal{P}$. Then $\hat{A}v = \hat{b} \geq \hat{A}q$, with one strict inequality $a_j v = b_j > a_j q$, and with equality for the other rows of $\hat{A}$. This implies that $w(v - q) = v'A(v - q) = v'_j a_j(v - q)$, and since $a_j(v - q) > 0$, v'_j and $w(v - q)$ have the same sign. □

We have the Upper Bound Theorem, missing the proof that the given bound is tight for dual neighborly polytopes.

Theorem 3.2 *The number of k-faces of a simple polytope in E^d with n facets is at most*

$$\sum_i \binom{i}{k} \min\{\binom{i+n-d-1}{n-d-1}, \binom{n-i-1}{n-d-1}\}.$$

Proof. The bound follows by applying the previous lemma, Equation (2), and Theorem 2.3 □

4 Concluding remarks

It is curious that the $(\leq k)$-set bounds of [2] both rely on the Upper Bound Theorem and are proven using an argument like the proof of Lemma 2.2. Perhaps some more direct argument for them exists.

References

[1] A. Brøndsted. *An Introduction to Convex Polytopes.* Springer-Verlag, Berlin, 1983.

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

[3] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, New York, 1987.

[4] P. Erdős, L. Lovasz, A. Simmons, and E. G. Straus. Dissection graphs of planar point sets. In J. N. Srivastava *et al*, editor, *A Survey of Combinatorial Theory.* North-Holland, Amsterdam, 1973.

[5] P. E. Gill, W. Murray, and M. H. Wright. *Numerical linear algebra and optimization, volume 1.* Addison-Wesley, New York, 1991.

[6] P. McMullen. The maximum number of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.

[7] K. Mulmuley. Output sensitive construction of levels and Voronoi diagrams in R^d of order 1 to k. In *Proc. 22nd Annual SIGACT Symp.*, pages 322–330, 1990.

[8] J. Pach, W. Steiger, and E. Szemeredi. An upper bound on the number of planar k-sets. In *Proc. 30th IEEE Symp. on Foundations of Computer Science*, pages 72–79, 1989.