On the Number of Corner Cuts

Uli Wagner
Institut für Theoretische Informatik, ETH Zürich
ETH Zentrum, CH-8092 Zürich, Switzerland
E-mail: uli@inf.ethz.ch
Tel. +41-1-632 71 83, Fax. +41-1-632 11 72

Abstract

A *corner cut* in dimension d is a finite subset of $\mathbf{N}_0^d$ that can be separated from its complement in $\mathbf{N}_0^d$ by an affine hyperplane disjoint from $\mathbf{N}_0^d$. Corner cuts were first investigated by Onn and Sturmfels [6], the original motivation stemming from computational commutative algebra. Let us write $\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}}$ for the set of corner cuts of cardinality k; in the computational geometer's terminology, these are the *k-sets* of $\mathbf{N}_0^d$. Among other things, Onn and Sturmfels give an upper bound of $O(k^{2d\frac{d-1}{d+1}})$ for the size of $\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}}$ when the dimension is fixed. In two dimensions, it is known (see [3]) that $\#\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}} = \Theta(k \log k)$. We will see that in general, for any fixed dimension d, the order of magnitude of $\#\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}}$ is between $k^{d-1} \log k$ and $(k \log k)^{d-1}$. (It has been communicated to me that the same bounds have been found independently by Gaël Rémond.) In fact, the elements of $\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}}$ correspond to the vertices of a certain polytope, and what our proof will show is that the above upper bound holds for the total number of flags of that polytope.

Key words: corner cut, Gröbner fan, k-set, reverse search.

1 Introduction

A *corner cut* in dimension d is a finite subset T of $\mathbf{N}_0^d$ that can be separated from its complement by an affine hyperplane h disjoint from $\mathbf{N}_0^d$. We will call h a *separating hyperplane* or *separating flat* for T. Corner cuts were first investigated by Onn and Sturmfels [6]. Their motivation stems from computational commutative algebra; for instance, they show that the corner cuts of fixed size k in dimension d are in one-to-one correspondence with the Gröbner bases of a certain kind of ideals in the polynomial ring $K[x_1, \ldots, x_d]$, K any infinite field: They prove that if I is the vanishing ideal of a generic configuration of k points in K^d, then the Gröbner fan of I is just the normal fan of the *corner cut polyhedron* $P_k^d := \operatorname{conv}\{\sum_{\boldsymbol{x}\in X} \boldsymbol{x} \mid X \subset \mathbf{N}_0^d \text{ and } |X| = k\}$. The above-mentioned correspondence then follows from the fact that the vertices of this polyhedron are exactly the points of the form $\sum T := \sum_{\boldsymbol{x}\in T} \boldsymbol{x}$ for a k-element corner cut T (see [6] for the details; see also [4] for the original definition of the so-called k-set polytope).

Furthermore, they give an upper bound of $O_d(k^{2d\frac{d-1}{d+1}})$ for the number of such corner cuts (the subscript d indicates that the constants hidden in the Landau notation depend on d).

Let us write $\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}}$ for the set of corner cuts of cardinality k, *k-cuts* for short, in dimension d. Apart from the above-mentioned relations to algebraic geometry, estimating the number $\#\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}}$ of such corner cuts seems to be of interest in its own right, since it is a rather natural special instance of the *k-set problem*, which concerns the maximal possible number of *k-sets* of a set S of n points in d-space, that is, k-element subsets $T \subseteq S$ that can be separated from $S \setminus T$ by a hyperplane disjoint from P (see [2], [8] or [7] for the recent developments); this, in turn, can be generalized to *shaped partition problems*, see, e.g., [5].

When studying the k-set problem, it is one of the first observations that one can safely assume general position of the points. Our case has a different flavor, not so much because we are dealing with an infinite point set (we will see below that we can restrict our attention to a suitable finite subset of the lattice points), but because the point configuration is so highly degenerate that we cannot afford to pass to general position (by invoking some perturbation arguments, say) lest we might increase the number of k-sets dramatically: There are examples [8] of sets of n points in general position in $\mathbf{R}^d$ which have at least $n^{d-1}e^{\Omega(\sqrt{\log k})}$ many k-sets. As we will see in Section 2, the number n of nonnegative integer points belonging to some k-element corner cut is roughly $k(\log k)^{d-1}$. Therefore, since $e^{\sqrt{\log k}}$

grows faster than any given power of $\log k$, general k-set estimates are of no avail if we want to establish an upper bound of the form $k^{d-1}\,\mathrm{polylog}(k)$.

In Section 4, we will prove that

$$\#\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}} = O_d((k\log k)^{d-1}). \tag{4}$$

In the *planar case*, Corteel et al. [3] showed that

$$\#\binom{\mathbf{N}_0^2}{k}_{\mathrm{cut}} = \Theta(k\log k)\,. \tag{1}$$

In Section 3 we will use this knowledge to derive a general lower bound

$$\#\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}} = \Omega_d(k^{d-1}\log k)\,, \tag{3}$$

which matches (4) up to a polylogarithmic factor. (It has been communicated to me that the same bounds have been found independently by Gaël Rémond.)

In fact, our proof of the upper bound immediately extends to the overall number of flags (that is, of ascending chains $F_1 \subsetneq \ldots \subsetneq F_s$ of faces) of the corner cut polyhedron P_k^d.

2 How Many Lattice Points Are Involved?

Given an oriented hyperplane h, let h^- and $\overline{h^-}$ denote the open, respectively closed, negative halfspace of h, and similarly for the positive halfspaces. Moreover, unless explicitly stated otherwise, let us assume that whenever we encounter a hyperplane h not containing $-\mathbf{1} = (-1,\ldots,-1)$, it is oriented in such a way that $-\mathbf{1} \in h^-$.

Suppose that a corner cut T of size k contains a lattice point $\boldsymbol{u} = (u_1,\ldots,u_d)$. Then the whole "lattice box" $Q = ([0,u_1]\times\ldots\times[0,u_d]) \cap \mathbf{N}_0^d$ is contained in T, and therefore $k \geq \prod_{i=1}^d (1+u_i)$. In other words, all corner cuts of cardinality k are subsets of the finite set $S_k^d := \{\boldsymbol{u} \in \mathbf{N}_0^d \mid \prod_{i=1}^d (1+u_i) \leq k\}$.

Observation 2.1. *For any real number $y \geq 1$,*

$$|S_y^d| \leq y(1+\log y)^{d-1} \tag{2}$$

Proof. We proceed by induction on d: For $d = 1$, $|S_y^1| = \lfloor y \rfloor \leq y$. Moreover, for $d > 1$ and integer $1 \leq j \leq y$, $\#\{\boldsymbol{u} \in S_y^d \mid 1 + u_d = j\} = |S_{y/j}^{d-1}| \leq (y/j)(1 + \log(y/j))^{d-2}$. Therefore,

$$\#S_y^d \leq \sum_{j=1}^{\lfloor y \rfloor} (y/j)(1 + \log(y/j))^{d-2} \leq y(1 + \log y)^{d-2} \underbrace{\sum_{j=1}^{\lfloor y \rfloor} \frac{1}{j}}_{\leq 1 + \log y}$$

□

On the other hand:

Observation 2.2. *Every k-set T of S_k^d which contains the origin is a k-cut.*

Proof. Suppose that, contrarily, $T \ni \mathbf{0}$ is a k-set of S_k^d such that, for every hyperplane h such that $T = h^- \cap S_k^d$, there exists some nonnegative lattice point $\boldsymbol{u} \notin S_k^d$ with $u \in h^-$; call such a point $\boldsymbol{u}$ a *violator* and h a *witness* for $\boldsymbol{u}$. Clearly, $d > 1$, and T is not contained in any coordinate hyperplane (else we are done by induction).

Now, consider a violator $\boldsymbol{u}$ which minimizes $\prod_{j=1}^d (1 + u_j)$ and a witness hyperplane $h = \{\boldsymbol{x} \in \mathbf{R}^d \mid \langle \boldsymbol{\nu}, \boldsymbol{x} \rangle = t\}$ for $\boldsymbol{u}$. If $Q = \{\boldsymbol{m} \in \mathbf{N}_0^d \mid m_j \leq u_j \text{ for } 1 \leq j \leq d\}$, then $Q \setminus \{\boldsymbol{u}\}$ contains at least k points since $\boldsymbol{u} \notin S_k^d$ Moreover, $Q \subset h^- = \{\boldsymbol{x} : \langle \boldsymbol{\nu}, \boldsymbol{x} \rangle < t\}$: By assumption, $\mathbf{0} \in h^-$, and so $t > 0$; moreover, all entries of $\boldsymbol{\nu}$ must be positive, for if $\nu_j \leq 0$, then all the points $i\boldsymbol{e} = (0, \ldots, i, \ldots, 0)$, $i \in \{0 \ldots k-1\}$, would belong to, and hence constitute, $h^- \cap S_k^d = T$, contradicting the fact that T is not contained in any coordinate hyperplane. Since $\boldsymbol{\nu}$ and t are positive, $\boldsymbol{u} \in h^-$ implies $Q \subset h^-$.

Therefore, $Q \setminus \{\boldsymbol{u}\} = T$. But then all $u_j \geq 1$ (otherwise T would be contained in some coordinate hyperplane) and so, for some j, h must intersect the x_j-axis at a distance greater than $u_j + 1$ from the origin. This, however, contradicts the assumption that h separates T from $S_k^d \setminus T$ because the point $(u_j + 1)\boldsymbol{e}_j$ belongs to $S_k^d \setminus T$ and lies below h. □

3 The Lower Bound

Lemma 3.1. *If $\#\binom{\mathbf{N}_0^{d-1}}{k}_{\mathrm{cut}} \geq \varphi(k)$ for some function φ then $\#\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}} \geq \sum_{i=1}^k \varphi(i)$*

Proof. We will show that for every $(d-1)$-dimensional corner cut T of size j, $1 \leq j \leq k$, there is some corner cut $\hat{T}$ of cardinality k in dimension d with $T = \mathbf{N}_0^{d-1} \cap \hat{T}$.

Take some T of size j as above with separating $(d-2)$-dimensional flat $f \subset \mathbf{R}^{d-1}$. Let us identify $\mathbf{R}^{d-1}$ with $\mathbf{R}^{d-1} \times \{0\} \subset \mathbf{R}^d$, and for a real parameter $t > 0$, consider the hyperplane h_t spanned by the flat f and the the point $(0, \ldots, 0, t)$. Observe that by choosing f in a sufficiently generic manner, we may assume that no h_t contains more than one lattice point.

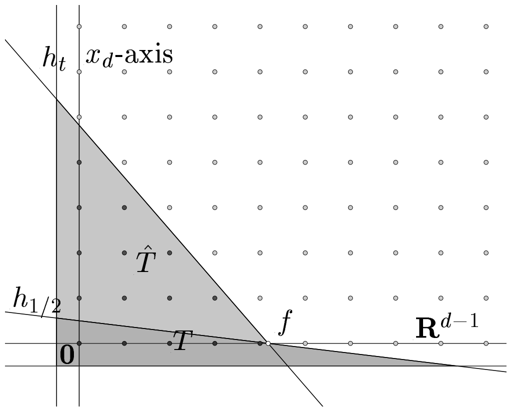

Figure 1: Lifting a corner cut.

Each of the hyperplanes h_t defines a certain d-dimensional corner cut T_t. For $t < 1$, this is exactly our original T. But as t grows, more and more lattice points will be included, one at a time by our assumption on f, until for some appropriate $\hat{t}$, we obtain a d-dimensional corner cut $\hat{T} := T_{\hat{t}}$ of size k as advertised. (Note that $\hat{T}$ depends on the choice of a separating flat f.) $\square$

Since we know that in the plane, $\#\binom{\mathbf{N}_0^2}{k}_{\text{cut}} = \Omega(k \log k)$, and since $\sum_{j=1}^{k} j^{d-2} \log j \geq \int_1^k x^{d-2} \log x \, dx = \frac{1}{d-1} k^{d-1} \log k - \int_1^k x^{d-2} \sim \frac{1}{d-1} k^{d-1} \log k$, we conclude inductively:

Theorem 3.2. *The number of k-cuts in d dimensions satisfies*

$$\#\binom{\mathbf{N}_0^d}{k}_{\text{cut}} = \Omega_d(k^{d-1} \log k). \tag{3}$$

4 The Upper Bound

Let S be a full-dimensional discrete point set in $\mathbf{R}^d$ — we will mainly be concerned with the case when S is finite, or $S = \mathbf{N}_0^d$. Consider an oriented hyperplane h, spanned by d affinely independent points from S, such that $|h^- \cap S| < k$ and $|\overline{h^-} \cap S| \geq k$; let us call h a *k-hyperplane* of S. Note that, for a given $\boldsymbol{\nu} \neq \mathbf{0}$, there is at most one k-hyperplane h of S with outer normal $\boldsymbol{\nu}$. Observe that h might be a k-hyperplane for several values of k. For instance, according to our definition, the coordinate hyperplanes are k-hyperplanes of $\mathbf{N}_0^d$ for every $k > 0$. On the other hand, if h is a k-hyperplane of $\mathbf{N}_0^d$ other than a coordinate hyperplane, then, since h is spanned by nonnegative lattice points, its outer normal vector $\boldsymbol{\nu} = (\nu_1, \ldots, \nu_d)$ must be strictly positive, i.e. all $\nu_i > 0$. In this case, h will be called a *proper* k-hyperplane of $\mathbf{N}_0^d$.

Furthermore, let us say that a k-hyperplane h of S is *incident* to a k-set T of S if $S \cap h^- \subset T \subseteq \overline{h^-}$. Note that every k-set is incident to some k-hyperplane h. Observe furthermore that if h_0 is a separating hyperplane for T then the open wedge $(h^- \cap h_0^+) \cup (h^+ \cap h_0^-)$ contains no points from S; that is, we can get from h_0 to h by a rotation (about the $(d-2)$-dimensional axis $h \cap h_0$) without traversing any points from S.

Now, consider $J := h_0^- \cap h \cap S$. This is a j-set (as witnessed by the separating flat $h_0 \cap h$) of the $(d-1)$-dimensional point set $S \cap h$, where $j = k - |h^- \cap S|$. What is more, we can recover T from h and J: Take any $(d-2)$-dimensional separating flat f for J in h; then a small rotation of h about f gives a separating hyperplane h_0 for T.

Lemma 4.1. *Let T be a k-set of a d-dimensional point set S. Set $f_d := \mathbf{R}^d$ and $k_d := k$. Then T can be uniquely represented by a sequence $(f_{d-1}, \ldots, f_1)$, where, for $1 \leq i \leq d-1$, f_i is an i-dimensional oriented flat, spanned by points from $S \cap f_{i+1}$, that forms a k_{i+1}-hyperplane within f_{i+1}, and $k_i = k_{i+1} - |S \cap f_{i+1} \cap f_i^-|$.*

Proof. Applying the above observation recursively, we see that each k-set T of S can be represented by a sequence $(h = f_{d-1}, f_{d-2}, \ldots, f_2)$ as specified (up to the last entry) together with a $j := k_2$-set J of the two-dimensional point configuration $S \cap f_2$.

But in two dimensions, every j-set J can be uniquely represented by a j-line (that is, a j-hyperplane w.r.t. the surrounding 2-dimensional space): If ℓ_0 is a separating line for J take a j-line f_1 incident to J such that the angle α of f_1 w.r.t. ℓ_0 is negative (that is, ℓ arises from ℓ_0 by a clockwise rotation).

Conversely, given f_1 and j, let $i = j - |S \cap f_1^-|$ and let $\boldsymbol{a}$ and $\boldsymbol{b}$ be the i^{th} and $(i+1)^{\text{th}}$ point of S on f_1, respectively (in the direction of f_1). Then a small counterclockwise rotation of f_1 about the midpoint of $\boldsymbol{a}$ and $\boldsymbol{b}$ gives a separating line ℓ_0 for J (see Figure 2). Observe that it is crucial to know

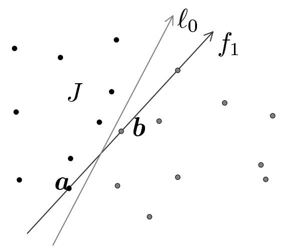

Figure 2: An example with $j = 9$ and $i = 1$.

both f_1 *and* j (we can reproduce the latter from our knowledge of k and $(f_{d-1}, \ldots, f_2)$). $\square$

Corollary 4.2. *Every k-set T of a discrete d-dimensional set S can be encoded by a sign vector $\boldsymbol{\varepsilon} \in \{+1, -1\}^{d-1}$ together with a $(d-1)$-tuple $(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d-1})$ of vectors of the following kind: There exist points $\boldsymbol{p}_0, \boldsymbol{p}_1, \ldots, \boldsymbol{p}_{d-1} \in S$ that span a k-hyperplane h incident to T and such that $\boldsymbol{v}_i = (\boldsymbol{p}_i - \boldsymbol{p}_0)$ for $i \in \{1 \ldots d-1\}$.*

Proof. Given T, represent it by $(f_{d-1}, \ldots, f_1)$ as above. Now, pick a pair $(\boldsymbol{p}_0, \boldsymbol{p}_1)$ of points that span the k_2-line f_1. Inductively, we construct a sequence of points $\boldsymbol{p}_i \in S$, $i \in \{0 \ldots d-1\}$ such that $\boldsymbol{p}_0, \boldsymbol{p}_1, \ldots, \boldsymbol{p}_i \in S$ span the flat $h_{(i)}$. Hence, if we set $\boldsymbol{v}_i := (\boldsymbol{p}_i - \boldsymbol{p}_0)$, then the vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_i$ span an i-dimensional linear flat parallel to f_i, and by chosing appropriate signs $\varepsilon_i \in \{+, -\}$ we can also record the orientation of f_i within f_{i+1}.

Moreover, this encoding is one-to-one: Given $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d-1}$ and $\varepsilon_1, \ldots, \varepsilon_{d-1}$, we get the outer normal of the k-hyperplane f_{d-1} of S, and hence f_{d-1} itself, since we know k. Then $k_{d-1} = k - S \cap f_{d-1}^-$, and by induction, we can reconstruct the sequence $(f_{d-2}, \ldots, f_1)$ from $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d-2}$, $\varepsilon_1, \ldots, \varepsilon_{d-2}$ and k_{d-1}. But once we know all f_i's, the set T is uniquely determined. $\square$

Applying this to $S = \mathbf{N}_0^d$, we see that every k-cut T can be uniquely encoded by a sign vector $\boldsymbol{\varepsilon} \in \{+, -\}^{d-1}$ and the $d \times (d-1)$-matrix $V = [v_{ij}]$ whose columns are the vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_{d-1}$ constructed above. Now, we derive some properties of V that will allow us to estimate the number of such matrices.

If h is one of the coordinate hyperplanes, then $T \subset h$ is essentially a corner cut in dimension $d-1$, and we can handle the number of these inductively. Thus, we may assume that h is a proper k-hyperplane, i.e. that its outer normal vector is strictly positive.

Observation 4.3. *A "corner simplex" $\Delta = \mathbf{R}^d_{\geq 0} \cap h^-$ which contains r lattice points has volume* $\mathrm{vol}_d(\Delta) \leq r$.

Proof. For each lattice point $\boldsymbol{u} \in \Delta$, consider the unit cube $\{\boldsymbol{x} \mid u_j \leq x_j \leq u_j + 1$ for all j. Clearly, Δ is contained in the union of these boxes, whose volume is r. ☐

Thus, the open simplex bounded by our proper k-hyperplane h and the coordinate hyperplanes has volume at most k; it follows that the same is true for its closure.

Let $B = [a_1, b_1] \times \ldots \times [a_d, b_d]$ be the *bounding box* of the points $\boldsymbol{p}_i$, $i \in \{0 \ldots d-1\}$, i.e. the smallest axis-parallel box containing them. Since the hyperplane h spanned by the $\boldsymbol{p}_i$ is not one of the coordinate hyperplanes, B is a full-dimensional box, i.e. $b_i - a_i > 0$ for all $i \in \{1 \ldots d\}$.

Observation 4.4. *Let H be any halfspace containing all $\boldsymbol{p}_i$'s. Then* $\mathrm{vol}_d(B \cap H) \geq \frac{1}{d!} \mathrm{vol}_d(B)$.

Proof. Let $H = \{\boldsymbol{x} \in \mathbf{R}^d \mid \boldsymbol{\nu} \cdot \boldsymbol{x} \leq t\}$. Suppose without loss of generality $\nu_d \geq 0$, and consider the projection of the points onto the the hyperplane $\{\boldsymbol{x} \in \mathbf{R}^d \mid x_d = a_d\}$. Then these projected points are contained in $H \cap \{x_d = a_d\}$. By induction, $P = H \cap B \cap \{x_d = a_d\}$ has $(d-1)$-dimensional volume at least $\frac{1}{(d-1)!} \prod_{i=1}^{d-1} (b_i - a_i)$. But then, the pyramid whose base is P and whose apex is any $\boldsymbol{p}_j$ maximizing the x_d-coordinate is contained in $B \cap H$ and has volume as guaranteed. ☐

For each row index $i \in \{1 \ldots d\}$ of the matrix $V = [v_{ij}]$, choose a column index $j(i) \in \{1 \ldots d-1\}$ such that $|v_{ij(i)}| = \max_j |v_{ij}|$. Then the i^{th} side of the bounding box B has lenght $b_i - a_i \geq |v_{ij(i)}| \geq 1$, whence $\prod_{i=1}^d |v_{ij(i)}| \leq \mathrm{vol}_d(B) \leq d! \, \mathrm{vol}(\overline{h^-} \cap B) \leq d! k$, by Observation 4.4.

Now, fix a sequence $(j(1), \ldots, j(d))$ and positive integers $m_1, \ldots, m_d$ such that $\prod_{i=1}^d m_i \leq d! k$, that is, $(m_1 - 1, \ldots, m_d - 1) \in S^d_{d!k}$. What is the number of integer matrices $V = [v_{ij}]$ such that $m_i = |v_{ij(i)}| = \max_j |v_{ij}|$ for all $i \in \{1 \ldots d\}$? Well, for each entry $v_{ij(i)}$ we get to choose a sign from $\{+, -\}$, while for the entries v_{ij} with $j \neq j(i)$ we may select any integer from $\{-m_i \ldots m_i\}$. Thus, we have $2^d \prod_{i=1}^d (2m_i + 1)^{d-2} = O_d(k^{d-2})$ possibilities to choose the entries when $j(1), \ldots, j(d)$ and $m_1, \ldots, m_d$ are fixed. Since there are $(d-1)^d = O_d(1)$ choices for the $j(i)$'s and, by Observation 2.1, $O_d(k(\log k)^{d-1})$ choices for the m_i's, we get a total of at most $O_d((k \log k)^{d-1})$ candidate matrices V. These, together with sign tuples $\varepsilon \in \{+,-\}^{d-1}$ suffice to encode all k-cuts for which we have picked an incident k-hyperplane h which is not one of the coordinate hyperplanes, hence there are at most $O_d((k \log k)^{d-1})$ such k-cuts. But the remaining k-cuts correspond to lower-dimensional corner cuts, and by induction, there are at most $O_d((k \log k)^{d-2})$ of those. We have proved:

Theorem 4.5. *The number of k-cuts in d dimensions satisfies*

$$\#\binom{\mathbf{N}_0^d}{k}_{\text{cut}} = O_d((k \log k)^{d-1}) \tag{4}$$

5 Concluding Remarks

(1) The above proof of the upper bound of $O_d((k \log k)^{d-1})$ immediately extends to the total number of faces of the corner cut polyhedron $P_k^d = \operatorname{conv}\{\sum X \mid X \subset \mathbf{N}_0^d \text{ and } |X| = k\}$, and indeed even to the number of *flags* of P_k^d, where a flag of a polytope P is a chain $F_1 \subsetneq \ldots \subsetneq F_s$ of faces of the polytope. In fact, the proof really is about flags. To see this, let $F = \operatorname{conv}\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m\}$ be an r-dimensional face of P_k^d, where each $\boldsymbol{v}_i = \sum T_i$ for some k-cut T_i (these are the bounded faces of P_k^d; on the other hand each unbounded r-face is contained in some r-dimensional intersection of coordinate hyperplanes, so there are only 2^d unbounded faces alltogether). Then F corresponds, in a one-to-one fashion, to an r-dimensional affine flat f which is parallel to F and spanned by points from $\mathbf{N}_0^d$, such that there exists a hyperplane h with the following properties:

1. $f \subseteq h$, and $h \cap \mathbf{N}_0^d = f \cap \mathbf{N}_0^d$,
2. $h^- \cap \mathbf{N}_0^d = T_1 \cap \ldots \cap T_m$,
3. each $T_i \subseteq h^- \cup f$, and the sets $T_i \cap f$ are precisely the j-sets of $\mathbf{N}_0^d \cap f$, where $j = k - |h^- \cap \mathbf{N}_0^d|$.

(h is an appropriate translate of the supporting hyperplane of P_k^d which defines F, and vice versa.)

Now, the total number of flags is at most 2^d times the number of maximal flags. Thus, consider such a maximal chain $\boldsymbol{v} = F_0 \subsetneq \ldots \subsetneq F_{d-1}$ of bounded faces of P_k^d. This corresponds to a chain $f_1 \subsetneq \ldots \subsetneq f_{d-1}$ of flats as above, where $\dim f_i = \dim F_i = i$ for $0 \leq i < d$. Now, f_{d-1} is a k-hyperplane of

$\mathbf{N}_0^d$ (with the additional property that $|\overline{f_{d-1}^-} \cap \mathbf{N}_0^d| > k$), and $f_{d-1}, \ldots, f_1$ is a descending sequence of nested flats like the one constructed in Lemma 4.1. But as we have seen, there are at most $O_d((k \log k)^{d-1})$ such sequences, hence at most that many flags.

The proof can also be adapted to show that there are at most $O_d((k \log k)^{d-1})$ many pairs ("incidences") $(h, \boldsymbol{p})$ where h is a proper k-hyperplane and $\boldsymbol{p} \in h \cap \mathbf{N}_0^d$.

(2) For the lower bound, we considered a corner cut $\lambda \in \binom{\mathbf{N}_0^{d-1}}{j}_{\mathrm{cut}}$, $1 \leq j \leq k$, and showed that the "fiber" $\{\mu \in \binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}} \mid \mu \cap \mathbf{N}_0^{d-1} = \lambda\}$ is non-empty. Since $\#\binom{\mathbf{N}_0^1}{j}_{\mathrm{cut}} = 1$ for all j in dimension 1, this argument would only give a linear number of corner cuts of size k in dimension 2. But this number is $k \log k$, so the "average" fiber must have size $\log k$. This might be considered as evidence that the lower bound (3) is not optimal. (To get a lower bound that matches the upper one we have, we would have to show that the average fiber is of size $\log k$ in every dimension.)

(3) On the algorithmic side, as noted in Section 2, enumerating the elements of $\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}}$ reduces to enumerating the k-sets of $S = S_k^d$ which contain the origin $\mathbf{0}$. For this purpose, we can use the k-set enumeration algorithm of Andrzejak and Fukuda [1] which is based on the paradigm of reverse search. Conceptually, the algorithm implicitly builds a subtree of the edge graph of the corner-cut polyhedron P_k^d and traverses this graph in a depth first manner. The crucial issue is how to find the neighbors of a given vertex $\boldsymbol{v} = \sum T$. As shown in [1], this reduces to solving certain linear programs. Some care is needed to ensure that the algorithm handles degenerate point configurations correctly, but on the other hand, the corner cut set-up is particularly nice in that the coefficients of the LP's will be integers in the range $\{0 \ldots k\}$. In total, the running time is (cf. [1])

$$\#\binom{\mathbf{N}_0^d}{k}_{\mathrm{cut}} kn^2 \, \mathrm{lp}_k(d, |S_k^d|)\,, \tag{5}$$

where $\mathrm{lp}_k(d, n)$ denotes the time required to solve a linear program in d variables, with n constraints and coefficients from $\{0, \ldots, k\}$. If we substitute (4) and $|S_k^d| = O_d(k(\log k)^{d-1})$ in (5), we obtain, for fixed dimension d, a time complexity of

$$k^{d+2}(\log k)^{3d-3} \, \mathrm{lp}_k(d, k(\log k)^{d-1})\,.$$

Acknowledgments. I first learned about corner cuts in a seminar talk in which Artur Andrzejak presented the results from [6]. My work was initiated by that presentation and by the discussions that followed it. I would also like to thank Komei Fukuda, Ingo Schurr and Emo Welzl for helpful comments and discussions.

References

[1] Artur Andrzejak and Komei Fukuda. Optimization over k-set polytopes and efficient k-set enumeration. In *Proceedings of 6th International Workshop on Algorithms and Data Stuctures (WADS'99)*, volume 1663 of *LNCS*, pages 1–12, 1999.

[2] Artur Andrzejak and Emo Welzl. Halving point sets. In *Proceedings of the International Congress of Mathematicians Berlin 1998, Documenta Mathematica, Extra Volume ICM 1998*, Vol. III.

[3] Sylvie Corteel, Gaël Rémond, Gilles Schaeffer, and Hugh Thomas. The number of plane corner cuts. *Advances in Applied Mathematics*, 23(1):49–53, 1999.

[4] Herbert Edelsbrunner, Pavel Valtr, and Emo Welzl. Cutting dense point sets in half. *Discrete & Computational Geometry*, 17:243–255, 1997.

[5] Frank K. Hwang, Shmuel Onn, and Uriel G. Rothblum. A polynomial time algorithm for shaped partition problems. *SIAM Journal on Optimization*, 10:70–81, 1999.

[6] Shmuel Onn and Bernd Sturmfels. Cutting corners. *Advances in Applied Mathematics*, 23(1):29–48, 1999.

[7] Micha Sharir, Shakhar Smorodinsky, and Gábor Tardos. An improved bound for k-sets in three dimensions. In *Proceedings of the Sixteenth Annual Symposium on Computational Geometry*, 2000.

[8] Géza Tóth. Point sets with many k-sets. In *Proceedings of the Sixteenth Annual Symposium on Computational Geometry*, 2000.