Andrews's theorem

Talk by Travis Dillon

Notes by Sanjana Das

April 26, 2024

§1 Introduction

Today we're going to talk about Andrews's theorem, which relates the number of vertices of a lattice polytope to its volume. Throughout this talk, we'll use P to refer to a lattice polytope in $\mathbb{R}^d$ — i.e., a polytope in $\mathbb{R}^d$ whose vertices are all in $\mathbb{Z}^d$. (All polytopes are assumed to be full-dimensional and convex.)

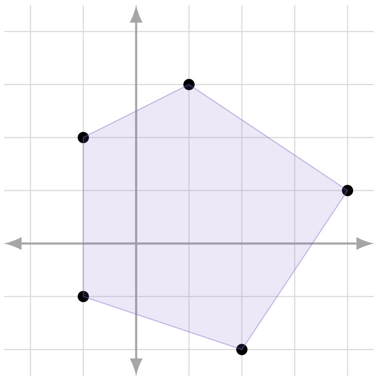

Theorem 1.1 (Andrews 1963)
If P is a lattice polytope in $\mathbb{R}^d$, then

$$\#\mathrm{vert}(P) \leq c_d \operatorname{Vol}(P)^{(d-1)/(d+1)}$$

(where $\mathrm{vert}(P)$ denotes the set of vertices of P, and c_d is a constant only depending on d).

This bound is tight, up to the constant c_d — we'll see an example soon.

Of course, it's important that P is a *lattice* polytope — otherwise we could take an arbitrary polytope with many vertices and shrink it to be tiny (so we'd have lots of vertices but tiny volume).

Remark 1.2. However, if we wanted such a statemet for a general polytope where we're only counting *lattice vertices* (i.e., we've got a polytope with some vertices not on the lattice, but we're only trying to upper-bound the number of vertices on the lattice), we can get such a statement directly from Theorem 1.1 by taking the convex hull of the lattice vertices and applying Theorem 1.1 to that subpolytope (whose volume can only be smaller than that of the original).

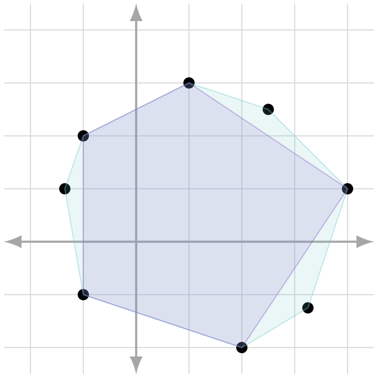

It's also important that we're looking at *vertices* of the polytope, and not just integer points — it turns out that we can also get a bound for the number of integer points, but it'll be different.

Proposition 1.3

For any lattice polytope P in $\mathbb{R}^d$, we have $\#(P \cap \mathbb{Z}^d) \leq c'_d \operatorname{Vol}(P)$ for some constant $c'_d \approx d!$.

This bound is also tight, and we'll see an example soon. Again, it's important that P is a *lattice* polytope — otherwise we could have P be extremely long and thin, so that it would contain many lattice points but have tiny volume.

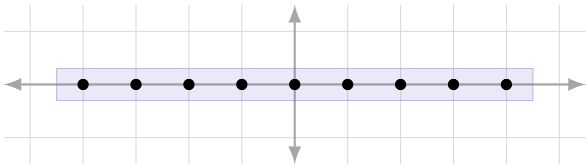

Remark 1.4. Together, Theorem 1.1 and Proposition 1.3 also give a quantitative version of the statement that a large polytope P should have significantly fewer vertices than integer points (since Theorem 1.1 has a smaller exponent of $\operatorname{Vol}(P)$ than Proposition 1.3), which is what you'd expect.

We'll first sketch the proof of Proposition 1.3, which relies on the following fact.

Fact 1.5 — Any lattice simplex in $\mathbb{R}^d$ has volume at least $1/d!$.

(This can be proven by induction on d, or by directly using the volume formula — which involves some determinant divided by $d!$.)

Proof sketch of Proposition 1.3. The main idea is to show we can decompose P into at least $\#(P \cap \mathbb{Z}^d) - d$ lattice simplices, or at least find this many disjoint lattice simplices inside P — then we're done by Fact 1.5, which tells us that each of these simplices has volume at least $1/d!$.

To get this decomposition, the idea is that if there's more than $d + 1$ vertices, then we can find some vertex we can remove while keeping P full-dimensional. Then we can remove it; this creates a new face, and we get one lattice simplex from taking that face together with the removed vertex.

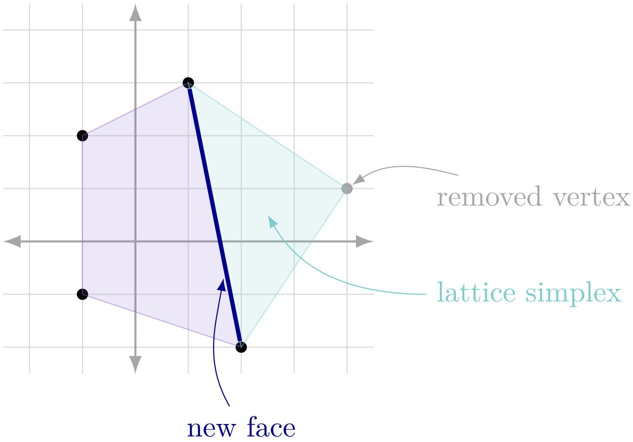

We can do this repeatedly to chop P up into (potentially large) lattice simplices; and then we chop up each of those lattice simplices further using their integer points. $\square$

§1.1 Constructions

As mentioned, both Theorem 1.1 and Proposition 1.3 are tight (up to the constants). First we'll see an example showing Proposition 1.3 is tight.

> **Example 1.6**
> The cube $[0, n]^d$ has $(n+1)^d$ integer points and volume n^d.

(Then taking $n \to \infty$ shows that Proposition 1.3 is tight.)

It's maybe more surprising that Theorem 1.1 is tight — the exponent $(d-1)/(d+1)$ looks pretty weird at first (although exponents like this actually come up not infrequently in such problems). The idea behind the construction is that we're going to take a point set X on the 'graph' of a convex function $\mathbb{R}^{d-1} \to \mathbb{R}$, and take P to be the convex hull of this point set X — this means that every point in X is a vertex of P (by the convexity of our function), so we get good control over $\#\text{vert}(P)$.

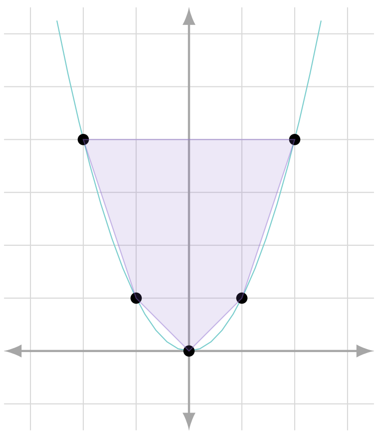

Specifically, the convex function we'll use is the function $x \mapsto \|x\|^2$ — this is convex, and if $x \in \mathbb{Z}^{d-1}$ then $\|x\|^2 \in \mathbb{Z}$ (which means the points we get will actually be lattice points).

Example 1.7

Let $X = \{(x, \|x\|^2) \mid x \in \{-n, \dots, n\}^{d-1}\}$, and let $P = \mathrm{conv}(X)$ be its convex hull.

For example, if $d = 3$, then the graph of the function $x \mapsto \|x\|^2$ is a paraboloid. And we're taking a square beneath it, considering the grid formed by all the lattice points in that square, and lifting up each of those points onto the paraboloid. And then X is the set of all the points on this paraboloid.

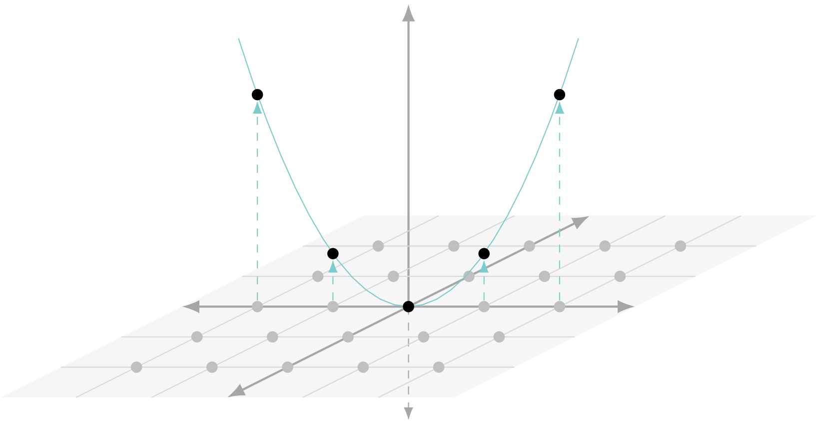

Then the vertices of P are precisely the points in X, so we have $\#\mathrm{vert}(P) = (2n+1)^{d-1}$. Meanwhile, to estimate $\mathrm{Vol}(P)$, we can imagine lifting just the corner points (i.e., the points $x \in \{-n, n\}^{d-1}$), and taking the polytope formed by these points together with the origin.

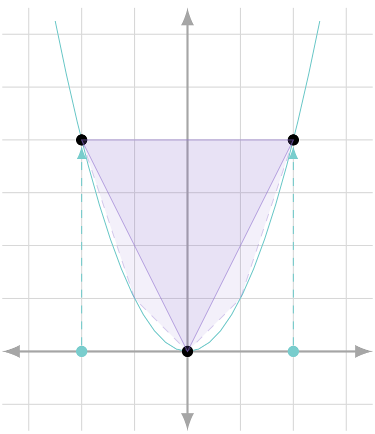

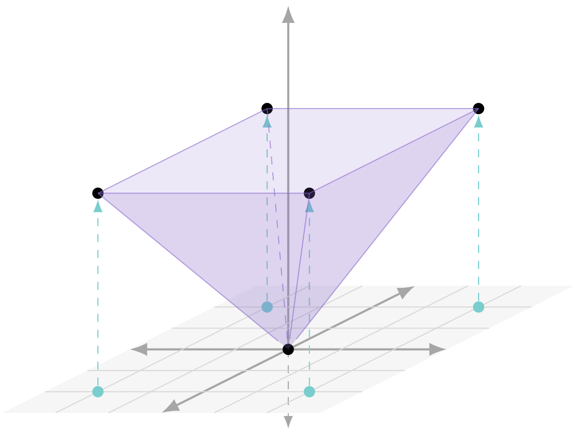

This gives us a pyramid whose base has volume $(2n+1)^{d-1}$ (since it's a lifted version of the hypercube with vertices $\{-n, n\}^{d-1}$) and whose height is $n^2(d-1)$ (the value of $\|x\|$ for all our corner points x). And this pyramid is contained in P, so (using the formula for the volume of a pyramid) we get

$$\mathrm{Vol}(P) \geq \frac{(2n+1)^{d-1} \cdot n^2 d}{d}.$$

We really want an *upper* bound on $\mathrm{Vol}(P)$ (in order to show that Theorem 1.1 is tight), but it turns out that a matching upper bound does hold, meaning that we have $\mathrm{Vol}(P) \asymp n^{d+1}$ (where the hidden constants depend on d). (Intuitively, this approximation of P by a pyramid doesn't lose too much.)

Combining these gives $\#\mathrm{vert}(P) \asymp \mathrm{Vol}(P)^{(d-1)/(d+1)}$ (as $\#\mathrm{vert}(P) \asymp n^{d-1}$ and $\mathrm{Vol}(P) \asymp n^{d+1}$). So the exponent in Theorem 1.1 is the best we could possibly hope for.

Remark 1.8. Another example showing that Theorem 1.1 is tight is $P = \operatorname{conv}(\mathbb{Z}^d \cap n\mathbb{B}^d)$ (where d is fixed and n is large). This is from a paper of Bárány from the 1980s called *The convex hull of the integer points in a large ball*, and it's maybe quite surprising — $\operatorname{Vol}(P)$ scales as n^d, but the number of vertices has this weird fractional exponent. (Note that the vertices are not just the intersection of the sphere $\mathbb{S}^{n-1}$ with the lattice — there'll be more integer points inside as well.)

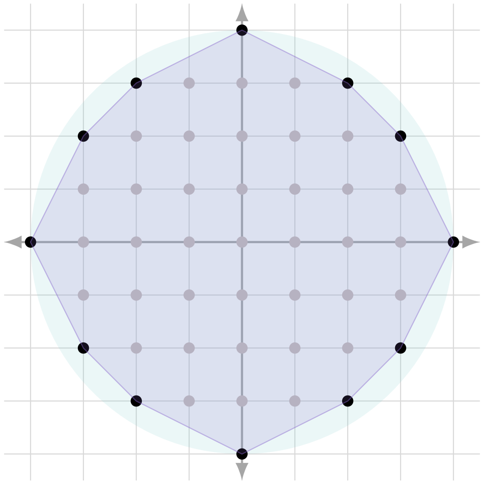

§2 The proof

There are lots of proofs of Theorem 1.1; the one we'll go through is from a paper of Konyagin and Sevastyanov. The main idea is to induct on d; the base case $d = 1$ is immediate (when $d = 1$, Theorem 1.1 just says $\#\operatorname{vert}(P)$ is bounded by a constant (since the exponent of $\operatorname{Vol}(P)$ is 0), which is true — it's at most 2).

§2.1 A first attempt

First, why is induction on d a strategy that might make sense? Imagine that we've got our polytope P in dimension d, and we've proven Theorem 1.1 for dimension $d-1$; this in particular means it applies to every facet (i.e., $(d-1)$-dimensional face) of P. So let's call these facets $F_1, \ldots, F_k$, and let's suppose they have $m_1, \ldots, m_k$ vertices (respectively). Then we can use the dimension $d-1$ case of the theorem to upper-bound m_i in terms of $\operatorname{Vol}_{d-1}(F_i)$ (where we write Vol_{d-1} to specify that we're taking the volume of a $(d-1)$-dimensional object, rather than a d-dimensional object). But we can't apply this directly, because the hyperplane that F_i lives in might intersect the d-dimensional lattice $\mathbb{Z}^d$ in a 'shifted' $(d-1)$-dimensional lattice (rather than just $\mathbb{Z}^{d-1}$), whose fundamental parallelogram might have a different volume (compared to 1). So now we need to normalize by this parallelogram — letting $\Pi_1, \ldots, \Pi_k$ be the fundamental parallelograms corresponding to each facet, we get

$$m_i \lesssim \left(\frac{\operatorname{Vol}_{d-1}(F_i)}{\operatorname{Vol}_{d-1}(\Pi_i)}\right)^{(d-2)/d} \tag{1}$$

for each i (we're going to ignore constant factors depending on d — so the implicit constants in our asymptotic notation are allowed to depend on d).

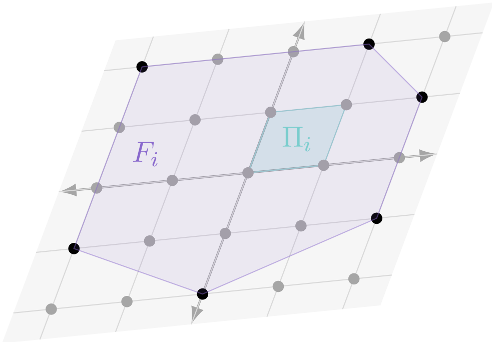

Now let m be the total number of vertices of P, so that $m \leq \sum m_i$ (this will overcount, since a vertex may be in multiple facets, but that's fine). And we have upper bounds for the m_i's, so we get

$$m \lesssim \sum \left(\frac{\mathrm{Vol}_{d-1}(F_i)}{\mathrm{Vol}_{d-1}(\Pi_i)} \right)^{(d-2)/d},$$

which is an upper bound for m (our number of vertices) in terms of the volumes of the *facets* F_i — which we can think of roughly as a bound for m in terms of the *surface area* of P. But this is where the idea hits a problem. We wanted a bound for m in terms of the *volume*. And the surface area isn't necessarily related to the volume in any reasonable way, at least in the direction we need — it's true that a shape with large volume has to have large surface area, but the reverse direction (which is what we need here) isn't true (for example, you can imagine taking a very long, thin parallelepiped).

§2.2 A fix — the reverse isoperimetric inequality

This idea is good, but we've hit a problem here. But it turns out that we can fix this problem and make the idea work using the *reverse isoperimetric inequality*.

Theorem 2.1 (Reverse isoperimetric inequality)
For any convex body K in $\mathbb{R}^d$, there is a volume-preserving affine map A such that

$$\mathrm{Vol}_{d-1}(\partial A \cdot K)^d \lesssim \mathrm{Vol}_d(A \cdot K)^{d-1}.$$

We use $A \cdot K$ to refer to the image of K under the affine map A, and $\partial A \cdot K$ to refer to its surface.

Remark 2.2. By *volume-preserving* we mean that A preserves d-dimensional volume; this does *not* mean that A preserves $(d-1)$-dimensional volume (so $\mathrm{Vol}_d(A \cdot K) = \mathrm{Vol}_d(K)$, but $\mathrm{Vol}_{d-1}(\partial A \cdot K)$ does not necessarily equal $\mathrm{Vol}_{d-1}(\partial K)$ — if it did, then A wouldn't be doing anything useful).

The (normal) isoperimetric inequality states that $\mathrm{Vol}_{d-1}(\partial K)^d \gtrsim \mathrm{Vol}_d(K)^{d-1}$. (More precisely, for a given surface area, the body with maximum volume is a ball, and $\mathrm{Vol}_{d-1}(\partial K)^d \asymp \mathrm{Vol}_d(K)^{d-1}$ if K is a ball.) The reverse inequality isn't true if we just take K with no transformation (as mentioned above). But the intuition is that we can use our affine transformation A to sort of round K out and make it look more like a ball, so that the reverse inequality *does* hold for this transformed body.

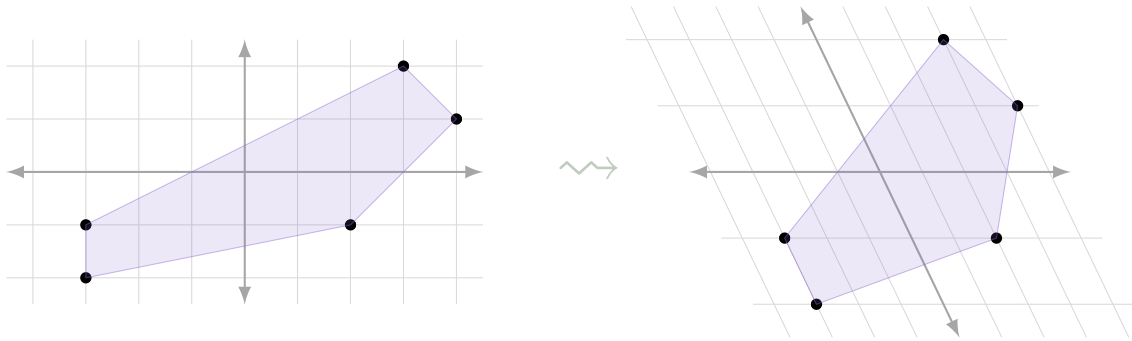

So this inequality gives us some hope — maybe we can transform our polytope first and then try to run this argument (where we use the $(d-1)$-dimensional case of the theorem to bound the number of vertices on each face by its $(d-1)$-dimensional volume). Then we'll eventually get to a point where we have a bound on the number of vertices in terms of the surface area, and we'll hopefully be able to use the reverse isoperimetric inequality to turn this into a bound in terms of volume.

§2.3 Proof outline

We're again going to start by using the inductive hypothesis on each facet F_i, as in (1). Our affine transformation A isn't necessarily going to preserve the surface measures of F_i or Π_i, but because it's affine, it *will* preserve their ratio — so we'll get

$$m_i \lesssim \left(\frac{\mathrm{Vol}_{d-1}(F_i)}{\mathrm{Vol}_{d-1}(\Pi_i)}\right)^{(d-2)/d} = \left(\frac{\mathrm{Vol}_{d-1}(A\cdot F_i)}{\mathrm{Vol}_{d-1}(A\cdot \Pi_i)}\right)^{(d-2)/d}. \tag{2}$$

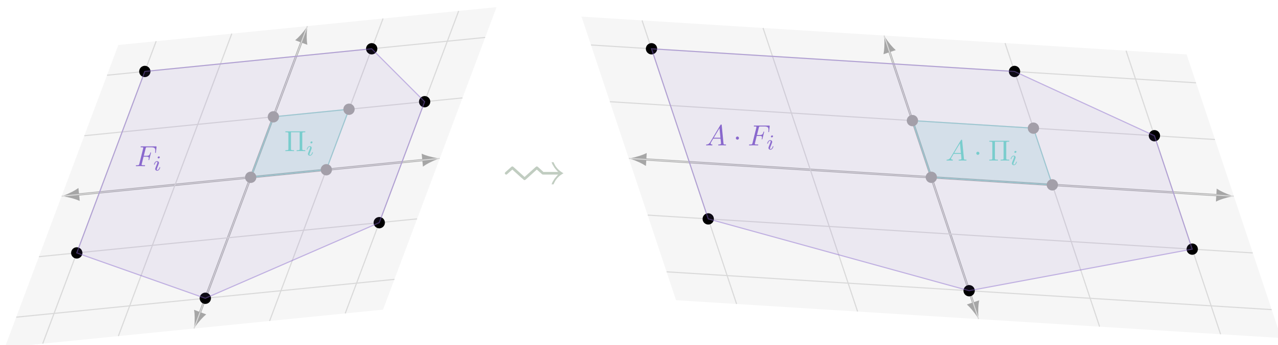

So the proof has two steps:

(1) First, we'll try to get a *lower* bound on $\mathrm{Vol}_{d-1}(A\cdot\Pi_i)$ for each i; this will give us an *upper* bound on each m_i (and therefore on m) in terms of the facet volumes of our transformed polytope $A\cdot P$.

(2) Then we'll use some general inequality techniques to turn this into a bound on m in terms of the surface area of $A\cdot P$, and we'll use the reverse isoperimetric inequality to finish.

§2.4 A simpler version of (1) — bounding the volumes of the Π_i

As mentioned above, the first step of the proof is getting a lower bound on $\mathrm{Vol}_{d-1}(A\cdot\Pi_i)$. For now, though, we're going to ignore the affine transformation A and just try to get a bound on $\mathrm{Vol}_{d-1}(\Pi_i)$. This isn't exactly what goes into the proof, but it'll illustrate several of the ideas, and we'll see later how to fix it in order to get a bound on $\mathrm{Vol}_{d-1}(A\cdot\Pi_i)$. (This is because if we tried to bound the volume of $A\cdot\Pi_i$ from the start, there'd be extra technical details that would obscure things.)

Question 2.3. How do we get a lower bound on $\mathrm{Vol}_{d-1}(\Pi_i)$?

We're going to order the facets so that $\mathrm{Vol}_{d-1}(\Pi_1) \le \cdots \le \mathrm{Vol}_{d-1}(\Pi_k)$. And now the tricky bit of the proof is that we're going to define vectors $h_i \perp \Pi_i$ (equivalently $h_i \perp F_i$) such that $|h_i| = \mathrm{Vol}_{d-1}(\Pi_i)$. We'll also say that they're pointing outwards from the polytope (this makes them uniquely defined).

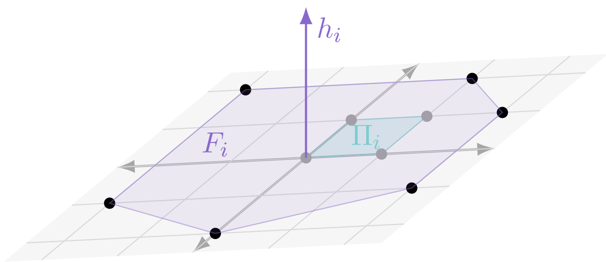

And these vectors h_i have a funny property — they turn out to have integer coordinates.

Lemma 2.4

We have $h_i \in \mathbb{Z}^d$ (for each i).

Written in this way, this fact might seem sort of crazy. But there are a few ways to see this; we'll see one by direct computation.

Proof. We're going to give an explicit formula for h_i (defining each of its coordinates); it'll be clear that the vector defined by this formula has integer coordinates, and we'll check that it satisfies all the properties h_i is supposed to (which means it's really equal to h_i).

Let $v_1, \ldots, v_{d-1}$ be the integer vectors that generate the fundamental parallelotope Π_i, and let

$$M = \begin{bmatrix} | & \cdots & | \\ v_1 & \cdots & v_{d-1} \\ | & \cdots & | \end{bmatrix}$$

be the matrix they form (so M has d rows and $d-1$ columns).

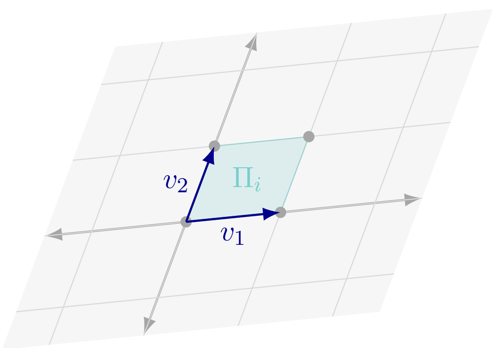

For each $1 \le j \le d$, let M_j be the $(d-1)\times(d-1)$ matrix obtained by deleting the jth row of M, and define the jth coordinate of h_i as $(h_i)_j = \det M_j$. Then it's clear that h_i has integer entries (since each M_j is an integer matrix).

To check that $h_i \perp \Pi_i$, consider any vector v in the hyperplane of Π_i (so $v \in \mathrm{Span}\{v_1, \ldots, v_{d-1}\}$). Then we have $\langle h_i, v\rangle = \pm\det(M \mid v)$ (where $(M \mid v)$ is the $d \times d$ matrix formed by M with an additional column of

v) by the cofactor formula for the determinant. But since v is in the span of the first $d-1$ columns of M, this determinant is 0, and therefore $h_i \perp v$.

Meanwhile, to check that $|h_i| = \mathrm{Vol}_{d-1}(\Pi_i)$, we'll consider the parallelepiped with base Π_i and height h_i.

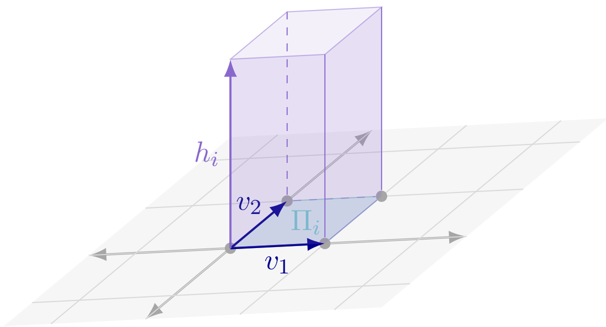

On one hand, the volume of this parallelepiped is $\mathrm{Vol}_{d-1}(\Pi_i) \cdot |h_i|$ (since in general, the volume of a prism is the volume of its base times its height). On the other hand, we know all the vectors that define it — namely, $v_1, \ldots, v_{d-1}$, and h_i — so we can calculate its volume as the determinant

$$\begin{vmatrix} | & \cdots & | & | \\ v_1 & \cdots & v_{d-1} & h_i \\ | & \cdots & | & | \end{vmatrix},$$

which can be shown to be $\mathrm{Vol}_{d-1}(\Pi_i)^2$. So we get $|h_i| = \mathrm{Vol}_{d-1}(\Pi_i)$, as desired.

So the vector h_i that we've defined (using this explicit formula) satisfies both properties in our original definition of h_i — we have $h_i \perp \Pi_i$ and $|h_i| = \mathrm{Vol}_{d-1}(\Pi_i)$ — which means it really is the correct vector (as in the original definition). And since it has integer entries (we defined its entries as determinants of integer matrices), we're done. $\square$

> **Remark 2.5.** In 3 dimensions, this formula for h_i is just the cross product $v_1 \times v_2$. So we can think of this sort of as a higher-dimensional generalization of the cross product.

Now the authors do something quite surprising. We ordered the facets such that $\mathrm{Vol}_{d-1}(\Pi_1) \leq \cdots \leq \mathrm{Vol}_{d-1}(\Pi_k)$, which means that

$$|h_1| \leq \cdots \leq |h_k|.$$

And now if we're trying to bound $\mathrm{Vol}_{d-1}(\Pi_r) = |h_r|$, we'll consider $\mathrm{conv}(h_1, \ldots, h_r)$ — the convex hull of the first r of these vectors. On one hand, all of $h_1, \ldots, h_r$ have length at most $|h_r|$, which means they're all contained in the ball of radius $|h_r|$, and so their convex hull is as well; this means

$$\mathrm{Vol}_d(\mathrm{conv}(h_1, \ldots, h_r)) \leq |h_r|^d \,\mathrm{Vol}_d(\mathbb{B}_d) \lesssim |h_r|^d$$

(where $\mathbb{B}_d$ is the unit ball in d dimensions; its volume only depends on d, which we think of as a constant). On the other hand, this convex hull has at least r integer points (namely, $h_1, \ldots, h_r$ themselves), so the argument from our proof of Proposition 1.3 gives that

$$\mathrm{Vol}_d(\mathrm{conv}(h_1, \ldots, h_r)) \geq \frac{r-d}{d!}$$

(we can decompose the polytope into at least $r-d$ lattice simplices and use Fact 1.5 to say that each has volume at least $1/d!$). Combining these gives that

$$\mathrm{Vol}_{d-1}(\Pi_r) = |h_r| \gtrsim \left(\frac{r-d}{d!}\right)^{1/d} \gtrsim r^{1/d}.$$

Remark 2.6. Alternatively, we could just say that the ball of radius $|h_r|$ has at most $|h_r|^d$ integer points (up to constant factors), since (for example) it's contained in the cube $[-|h_r|, |h_r|]^d$, and $h_1, \ldots, h_r$ are all integer points inside this ball; this also gives $|h_r|^d \gtrsim r$.

Remark 2.7. There's a detail we're brushing under the rug — we need $\operatorname{conv}(h_1, \ldots, h_r)$ to be full-dimensional for this argument to work — but this isn't too important.

§2.5 Step (2) — concluding

We've now proven the bound of $\mathrm{Vol}_{d-1}(\Pi_r) \gtrsim r^{1/d}$ for all r. What we *really* wanted was to bound $\mathrm{Vol}_{d-1}(A \cdot \Pi_r)$ (for the affine transformation A given by the reverse isoperimetric inequality applied to P); it turns out that we can get the same bound of $\mathrm{Vol}_{d-1}(A \cdot \Pi_r) \gtrsim r^{1/d}$, using the same proof as above with a few small modifications. We'll see how to do this later, but for now we'll assume this and see how to finish the proof.

Recall that we defined m as the *total* number of vertices of P and m_i as the number of vertices on F_i, so that $m \leq \sum m_i$. We'll let $\widehat{m} = \sum m_i$, and we'll actually try to bound $\widehat{m}$.

First, the inductive hypothesis gave us the bound (2) on m_i in terms of the volumes of $A \cdot F_i$ and $A \cdot \Pi_i$, and plugging in our bound $\mathrm{Vol}_{d-1}(A \cdot \Pi_i) \gtrsim i^{1/d}$ gives

$$m_i^{d/(d-2)} \lesssim \left(\frac{\mathrm{Vol}_{d-1}(A \cdot F_i)}{\mathrm{Vol}_{d-1}(A \cdot \Pi_i)}\right) \lesssim \mathrm{Vol}_{d-1}(A \cdot F_i) \cdot i^{-1/d},$$

and we can move the $i^{-1/d}$ to the other side to get

$$m_i^{d/(d-2)} \cdot i^{1/d} \lesssim \mathrm{Vol}_{d-1}(A \cdot F_i). \tag{3}$$

Now if we sum the right-hand side over all facets F_i, what we get is precisely the surface area of $A \cdot P$. So we'd like to bound $\widehat{m} = \sum m_i$ in terms of the left-hand side, and we'll do this using Hölder's inequality. In (3) we've got an exponent of $d/(d-2)$ (for m_i), and we want to choose exponents for Hölder so that this exponent becomes a 1; so the inequality we want to use is

$$\widehat{m} = \sum m_i \leq \left(\sum m_i^{d/(d-2)} i^{1/d}\right)^{(d-2)/d} \left(\sum i^{-(d-2)/2d}\right)^{2/d}. \tag{4}$$

And the first term in (4) is precisely what we had in (3), so we can bound it by

$$\sum m_i^{d/(d-2)} i^{1/d} \lesssim \sum \mathrm{Vol}_{d-1}(A \cdot F_i) = \mathrm{Vol}_{d-1}(\partial A \cdot P).$$

And now we can use the reverse isoperimetric inequality (Theorem 2.1) to go from surface area to volume — so we get

$$\sum m_i^{d/(d-2)} i^{1/d} \lesssim \mathrm{Vol}_{d-1}(\partial A \cdot P) \lesssim \mathrm{Vol}_d(A \cdot P)^{(d-1)/d} = \mathrm{Vol}_d(P)^{(d-1)/d}$$

(here we can remove A because it preserves d-dimensional volume).

Meanwhile, to bound the second term in (4), it's just a sum over i (which ranges from 1 to k, the number of facets), so we can bound it by an integral — we have

$$\sum_{i=1}^{k} i^{-(d-2)/2d} \lesssim \int_1^k x^{-(d-2)/2d}\, dx \asymp k^{(d+2)/2d}.$$

We need to get rid of k, and we'll do so using the crude bound $k \leq \sum m_i = \widehat{m}$ (since there's k terms in the sum, and each is at least 1). Then plugging these bounds into (4) gives

$$\widehat{m} \leq \mathrm{Vol}_d(P)^{(d-1)(d-2)/d^2} \cdot \widehat{m}^{(d+2)/d^2}.$$

And finally, we can move $\widehat{m}$ over to the other side to get that

$$\widehat{m} \leq \mathrm{Vol}_d(P)^{(d-1)/(d+1)},$$

which finishes the proof of Theorem 1.1 (as $m \leq \widehat{m}$).

Remark 2.8. This argument may look pretty lossy at first, so it's maybe surprising that it's tight for the construction in Example 1.7. But it turns out that the loss in these bounds is only a constant depending on the dimension for that example — in the bound $m \leq \widehat{m}$ the loss comes from vertices being in multiple facets (and therefore being overcounted in $\widehat{m} = \sum m_i$), and each vertex is in a constant number of facets (depending on d), so we only lose a constant factor here. Similarly, when we bound $k \leq \widehat{m}$, the loss comes from facets having more than one vertex, and this again corresponds to a loss of a dimension-dependent constant (since each facet has a constant number of vertices).

§2.6 Finishing step (1) — the affine transformation

Now we'll explain how to modify the argument in Subsection 2.4 to get the same bound on $\mathrm{Vol}_{d-1}(A \cdot \Pi_i)$ (instead of just $\mathrm{Vol}_{d-1}(\Pi_i)$). (Note that we can't just apply the result of Subsection 2.4 as a black box to $A \cdot P$, because $A \cdot P$ may not be a lattice polytope anymore.)

We're still going to define the vectors h_i in the same way (such that $h_i \perp \Pi_i$ and $|h_i| = \mathrm{Vol}_{d-1}(\Pi_i)$). Then if we consider the vectors $A^{-\intercal}h_i$ (where $A^{-\intercal}$ is the inverse transpose of A), we'll have $A^{-\intercal}h_i \perp A \cdot \Pi_i$ and $|A^{-\intercal}h_i| = \mathrm{Vol}_{d-1}(A \cdot \Pi_i)$.

So now we'll do essentially the same trick as in Subsection 2.4, but we'll order the facets such that $|A^{-\intercal}h_1| \leq \cdots \leq |A^{-\intercal}h_k|$ (instead of just $|h_1| \leq \cdots \leq |h_k|$). And then in order to bound $\mathrm{Vol}_{d-1}(A \cdot \Pi_r) = |A^{-\intercal}h_r|$ for some r, we'll consider $\mathrm{conv}(A^{-\intercal}h_1, \ldots, A^{-\intercal}h_r)$. On one hand, this convex hull contained in the ball of radius $|A^{-\intercal}h_r|$ (since all its vertices are contained in this ball), so we have

$$\mathrm{Vol}_d(\mathrm{conv}(A^{-\intercal}h_1, \ldots, A^{-\intercal}h_r)) \leq |A^{-\intercal}h_r|^d \, \mathrm{Vol}_d(\mathbb{B}_d) \lesssim |A^{-\intercal}h_r|^d.$$

On the other hand, since A is volume-preserving, so is $A^{-\intercal}$; this means

$$\mathrm{Vol}_d(\mathrm{conv}(A^{-\intercal}h_1, \ldots, A^{-\intercal}h_r)) = \mathrm{Vol}_d(A^{-\intercal} \cdot \mathrm{conv}(h_1, \ldots, h_r)) = \mathrm{Vol}_d(\mathrm{conv}(h_1, \ldots, h_r)).$$

And the same argument from Subsection 2.4 still gives $\mathrm{Vol}_d(\mathrm{conv}(h_1, \ldots, h_r)) \geq (r-d)/d!$ (this polytope has at least r integer points, namely $h_1, \ldots, h_r$). So we get

$$\mathrm{Vol}_{d-1}(A \cdot \Pi_r) = |A^{-\intercal h_r}| \gtrsim r^{1/d},$$

which is exactly what we wanted.

§3 An extension

This completes the proof of Theorem 1.1; to finish, we'll say one more comment about how we can extend the proof to get a stronger result.

Definition 3.1. A tower (or flag) of a polytope P is a sequence of nested faces $f_0 \subseteq f_1 \subseteq \cdots \subseteq f_{d-1} \subseteq P$ where $\dim(f_i) = i$ for each i.

This means we're starting with a single vertex f_0, and then taking an edge f_1 that contains it, and then a face f_2 containing this edge, and so on. The reason for the name is that these first three things (a vertex in an edge in a face) together look sort of like a flag.

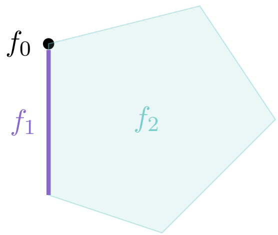

Definition 3.2. For a polytope P, we define $T(P)$ as the number of towers of P.

There's a simple recurrence for $T(P)$ — if P has facets $F_1, \ldots, F_k$, then we have $T(P) = \sum T(F_i)$ (by separating the towers based on which facet f_{d-1} is). This looks a lot like the bound $m \leq \sum m_i$ that we used in the proof of Theorem 1.1. And indeed we can go through the entire proof replacing the number of vertices with the number of towers, and nearly everything will work in the same way. So we can get the same bound as Theorem 1.1 for the number of towers.

Theorem 3.3

We have $T(P) \lesssim \mathrm{Vol}(P)^{(d-1)/(d+1)}$.

And this has a nice corollary — for any s, the number of s-dimensional faces of P is at most the number of towers (since each face can be extended in both directions to get at least one tower), giving the following bound.

Corollary 3.4

For each s, the number of s-dimensional faces of P is at most $\mathrm{Vol}(P)^{(d-1)/(d+1)}$ (up to constant factors).

The case $s = 0$ corresponds to counting the number of vertices (as in Theorem 1.1), but we can also use it to bound the number of edges or two-dimensional faces or facets or so on.

Remark 3.5. This bound should again be tight for the integer hull of a ball, as described in Remark 1.8 — the same paper by Bárány mentioned there might actually show that it's tight for all s.