Combinatorics and Geometry of Polytopes

Lecturer: Isabella Novik; transcribed, edited by Josh Swanson

May 20, 2016

Abstract

The following notes were taking during a graduate course on polytopes at the University of Washington during Spring 2016. Please send any corrections to `jps314@uw.edu`. Thanks!

Contents

March 28th, 2016: Introduction, Polytopes, Face Lattices, Graphs

1 Remark

The course web page is `www.math.washington.edu/~novik/583`. Office hours are on Tuesdays from 1:30-2:30 and Fridays from 10:30-11:30 in C-416. Textbooks include Ziegler, Grünbaum, Barvinok, and Matousek. See the class description on the course web site for more precise references.

2 Definition

We will use $\boxed{\mathbb{R}^d} := \{(x_1, \ldots, x_d) : x_i \in \mathbb{R}\}$ with the standard topology and inner product.

3 Definition

A subset $C \subset \mathbb{R}^d$ is $\boxed{\text{convex}}$ if for all $x, y \in C$, the line connecting x and y lies entirely within C, i.e.

$$\boxed{[x,y]} := \{tx + (1-t)y : 0 \leq t \leq 1\} \subset C.$$

where $[x,y]$ is the $\boxed{\text{interval}}$ between x and y.

4 Example

$\mathbb{R}^d$ itself is convex. A $\boxed{\text{hyperplane}}$ is

$$\boxed{h} := \{x : \langle a, x\rangle = b\}$$

where $a \in \mathbb{R}^d$ is non-zero, $b \in \mathbb{R}$ are fixed. The corresponding $\boxed{\text{closed half-spaces}}$ are

$$\boxed{h^-} := \{x : \langle a, x\rangle \leq b\}$$
$$\boxed{h^+} := \{x : \langle a, x\rangle \geq b\}$$

We may likewise define $\boxed{\text{open half-spaces}}$ by replacing inequality with strict inequality.

5 Remark

The arbitrary intersection of convex sets is convex, which follows immediately from the definition. We will take the convention that the empty intersection is $\mathbb{R}^d$ itself (where of course we will have fixed some dimension beforehand), which is indeed convex.

6 Definition

If $X \subset \mathbb{R}^d$, then the $\boxed{\text{convex hull}}$ of X is

$$\boxed{\text{conv}(X)} := \bigcap \text{ convex sets that contain } X,$$

which is itself convex.

7 Example

Consider six points in the plane, where four of the points are vertices of a trapezoid and two of the points are inside the trapezoid. The convex hull of those six points is the (closed) trapezoid. Note that the convex hull of finitely many points is necessarily bounded.

8 Definition

A $\boxed{V\text{-polytope}}$ is a convex hull of finitely many points.

9 Definition

An $\boxed{H\text{-polyhedron}}$ is an intersection of finitely many closed half-spaces.

10 Example

Take two half-spaces in the plane whose borders are parallel to the x-axis and which intersect in a "strip." The result is an unbounded H-polyhedron, so it cannot be a V-polytope.

11 Definition

An $\boxed{H\text{-polytope}}$ is a bounded H-polyhedron, i.e. a bounded intersection of finitely many closed half-spaces.

12 Remark

One of our first goals is to prove the following equivalence of the preceding polytope definitions, after which we will be able to just use the term $\boxed{\text{polytope}}$.

13 Theorem

A subset $X \subset \mathbb{R}^d$ is an H-polytope if and only if it is a V-polytope.

14 Definition

If $K \subset \mathbb{R}^d$ is closed and convex, then a hyperplane $h \subset \mathbb{R}^d$ is a $\boxed{\text{supporting hyperplane}}$ of K if

(i) $h \cap K \neq \varnothing$

(ii) All points of K lie on the same side of h, i.e. $h^+ \cap K = K$ or $h^- \cap K = K$.

(Condition (ii) is independent of the choice of a and b above.)

15 Definition

A $\boxed{\text{face}}$ of K is the intersection of K with a supporting hyperplane.

16 Example

Consider a semi-circle in the plane. Any tangent vector to the circular part is a supporting hyperplane. At the vertices of the semi-circle, there are many supporting hyperplanes intersecting the semi-circle only at that vertex. This shows that every point on the circular part is itself a face, and the straight part is also a face.

17 Definition

If K is convex, the $\boxed{\text{dimension}}$ of K is

$$\boxed{\dim K} := \text{dim of the smallest affine subspace that contains } K.$$

Here an $\boxed{\text{affine subspace}}$ is any translation of a linear subspace. The smallest such subspace is often called the $\boxed{\text{affine hull}}$ of K.

18 Example

The dimension of the closed unit disk in $\mathbb{R}^2$ is 2. The dimension of a line segment (not a singleton) in any $\mathbb{R}^d$ is 1.

19 Remark

A face of a convex set is a convex set. Note that any hyperplane $h = h^+ \cap h^-$ is an H-polyhedron. If P is an H-polytope, it is a bounded set of the form $P = \cap h_i^+$, and it follows that a face of P is itself an H-polytope.

20 Definition

0-dimensional faces of P are called $\boxed{\text{vertices}}$. 1-dimensional faces of P are called $\boxed{\text{edges}}$. Codimension-1 faces are called $\boxed{\text{facets}}$. Codimension-2 faces are called $\boxed{\text{ridges}}$. We call $\varnothing$ and P $\boxed{\text{improper faces}}$.

21 Remark

A preview of things to come: we'll show for a polytope P that...

- A face of a polytope is a polytope.
- The set $L(P)$ of all faces of P including $\varnothing, P$ can be partially ordered by P, giving it the structure of a graded lattice with minimum and maximum.
- A polytope has finitely many faces, i.e. $L(P)$ is finite.
- The dual $L(P)^{\text{op}}$, i.e. $L(P)$ ordered by reverse inclusion, is the face lattice of some polytope Q called the dual of P. (Q is only defined up to combinatorial type.)
- If $F \leq G$ are faces of P, the interval $[F, G]$ is also the face lattice of a polytope.

22 Example

There are many polytopes which are geometrically different but have the same face lattice. For a simple example, every (non-degenerate) quadrilateral in $\mathbb{R}^2$ has the same face lattice.

23 Definition

Two polytopes P and Q are $\boxed{\text{combinatorially isomorphic}}$ if

$$L(P) \cong L(Q)$$

as abstract lattices.

24 Remark

One may continuously deform a quadrilateral into a triangle, which clearly does not preserve the isomorphism class of the underlying face lattice.

25 Definition

The $\boxed{\text{graph}}$ of a polytope P, denoted $\boxed{G(P)}$ is a graph whose vertices are the vertices of P and whose edges are the edges of P.

26 Remark

Can we reconstruct the face lattice of a polytope from its graph? The answer is "sometimes." We have

27 Theorem (Steinitz, circa 1922)

A (finite, simple) graph is the graph of a 3-dimensional polytope if and only if it is planar and 3-connected.

Recall that planar means we can draw the graph in the plane without edges intersecting at interior points. Intuitively, if we have a polytope and "look at it" very close to one of the vertices, we'll see a planar graph. Also, a graph is 3-connected if you cannot disconnect it by removing 3 vertices. It is not too hard to show that these conditions on a graph are necessary, but sufficiency takes more work. In the three-dimensional case, the graph does indeed determine the combinatorial isomorphism type.

28 Corollary

For all 3-dimensional polytopes P, there exists a combinatorially isomorphic polytope Q such that all vertices of Q have rational coordinates.

One would naively think the analogue would be true in all dimensions by "wiggling" each vertex slightly. However, one encounters issues when faces are not simplices. Already for the cube, each face is in a sense overdetermined, so "wiggling" some of the vertices will almost never preserve the combinatorial isomorphism type.

29 Open Problem

Does every 3-polytope admit a realization with all edges having rational lengths?

The analogue of Steinitz' theorem is indeed false already in dimension 4.

30 Theorem (Richter-Gebert, 1995)

There exists a 4-dimensional polytope with 33 vertices that has no realization with rational coordinates.

(That is, there is no polytope combinatorially isomorphic to it all of whose vertices have rational coordinates.)

31 Theorem (Perles)

There exists an 8-dimensional polytope with 12 vertices that has no realization with rational coordinates.

Can we at least figure out the dimension from the graph? Not in general: for all $d \geq 4$ and all $n \geq d+1$, there exists a d-dimensional polytope P with n vertices such that $G(P) = K_n$ where K_n is the complete graph on n-vertices. Indeed, as n grows, there are exponentially many such graphs. On the other hand, we have:

32 Theorem (Blind-Blind, 1987)

The face lattice of a simple polytope can be reconstructed from its graph.

The proof is constructive/produces an algorithm. We will hopefully go through Gil Kalai's proof this quarter.

33 Definition

If P is a d-dimensional polytope and v is a vertex of P, then v lies in at least d edges. We call P a simple polytope if every vertex lies in exactly d edges.

For instance, a (closed) cube in $\mathbb{R}^3$ is a simple polytope. A square pyramid is not simple since the apex has 4 incident edges instead of 3.

34 Definition

Let P be a polytope. The face numbers of P are

$$\boxed{f_i} := \#\text{number of } i\text{-dimensional faces of } P.$$

35 Conjecture (Hirsh, 1957)

If P is a d-dimensional polytope with n facets, then the diameter of $G(P)$ is $\leq n - d$.

36 Theorem (Santos, 2010)

The conjecture is false. Further questions remain, e.g. is the diameter a polynomial?

March 30th, 2016: Convex Combinations; Theorems of Radon, Tverberg, Helly, Caratheodory

---

37 Remark

Homework 1 has been posted on our web page and is due next Friday. Some possible papers for presentations are also on the web page.

38 Remark

Today we'll discuss general facts about convex sets. Recall that a subset of $\mathbb{R}^d$ is called convex if for all $x, y \in C$, the interval $[x, y] \coloneqq \{tx + (1-t)y : 0 \le t \le 1\} \subset C$. Any intersection of convex sets is convex, and the convex hull of a subset S in $\mathbb{R}^d$ is the intersection of all convex sets containing S, which is hence convex. This description is well-defined but unwieldy, so we will first describe $\operatorname{conv}(S)$ in a "nice algebraic" way.

Note: today we will give many details, though later we will leave more details up to the reader.

39 Example

Consider a triangle with vertices x, y, z. Any point on the edge from x to z is a linear combination of x, z, and any point in the interior of the triangle lies on a line from y to a point on the edge between x and z. Writing things out formally, one finds

$$\operatorname{conv}\{x, y, z\} = \{\alpha x + \beta y + \gamma z : \alpha, \beta, \gamma \ge 0, \alpha + \beta + \gamma = 1\}.$$

40 Proposition

Given a set $S \subset \mathbb{R}^d$, we have

$$\operatorname{conv}(S) = \left\{ \sum_{i=1}^{m} \alpha_i p_i : m \in \mathbb{N}, \alpha_i \ge 0, \sum_i \alpha_i = 1, p_i \in S \right\}.$$

PROOF Call the right-hand side A. We need to show that $A \supset S$, that A is convex, and that for all convex $C \supset S$, $C \supset A$. We have:

$A \supset S$: If $p \in S$, then $p = 1 \cdot p \in A$.

A is convex: If $p, q \in A$, then we may write $p = \sum_i \alpha_i p_i$ with $\alpha_i \ge 0$, $\sum_i \alpha_i = 1$, $p_i \in S$, and $q = \sum_i \beta_i q_i$ with $\beta_i \ge 0$, $\sum_i \beta_i = 1$, $q_i \in S$. Since A allows zero coefficients, we may set $X \coloneqq \{p_i\}_{i=1}^n \cup \{q_i\}_{i=1}^m = \{x_j\}_j$ and write $p = \sum_j \alpha_j x_j$, $q = \sum_j \beta_j x_j$ with $\alpha_j, \beta_j \ge 0$ and $\sum_j \alpha_j = 1 = \sum_j \beta_j$. Then for all $0 \le t \le 1$, we have

$$tp + (1-t)q = \sum_j (t\alpha_j + (1-t)\beta_j) x_j,$$

where the coefficients in parentheses are evidently non-negative and sum to 1. Thus $tp + (1-t)q \in A$, as required.

$C \supset A$: Let $C \supset S$ be convex. Pick any $p \in A$, so $p = \sum_{i=1}^m \alpha_i p_i$, with $\alpha_i \ge 0$, $\sum_i \alpha_i = 1$, $p_i \in S$. We will show $p \in C$ by induction on m. If $m = 1$, this says $p = 1 \cdot p \in S \subset C$. If $m > 1$, then if $\alpha_m = 1$, the remaining α_i must be zero, and we again have $p \in C$, so suppose $\alpha_m < 1$. Now write

$$\begin{aligned} p &= (\alpha_1 p_1 + \cdots + \alpha_{m-1} p_{m-1}) + \alpha_m p_m \\ &= (1-\alpha_m) \sum_{i=1}^{m-1} \frac{\alpha_i}{1-\alpha_m} p_i + \alpha_m p_m \eqqcolon (1-\alpha_m) q + \alpha_m p_m. \end{aligned}$$

If $q \in C$, then since $p_m \in C$, the above shows $p \in C$ by convexity. Notice that the coefficients on q are certainly non-negative and they sum to 1 since

$$\sum_{i=1}^{m-1} \frac{\alpha_i}{1-\alpha_m} = \frac{\sum_{i=1}^{m-1} \alpha_i}{1-\alpha_m} = \frac{1-\alpha_m}{1-\alpha_m} = 1.$$

Thus q is a convex combination of $m-1$ points in C, which is inductively in A, completing the proof.

41 Remark

Points of the form $\sum_{i=1}^m \alpha_i p_i$ as above are called $\boxed{\text{convex combinations}}$. In the proof above, we essentially showed that convex combinations of convex combinations are convex combinations.

42 Remark

For the rest of the lecture, we will discuss three nice theorems due to Radon, Helly, and Caratheodory.

43 Theorem (Radon, 1913)

Let $x_1, \dots, x_m$ be $m \geq d+2$ points in $\mathbb{R}^d$. Then there is a set partition

$$\{1, 2, \dots, m\} =: [m] = S \coprod T$$

such that

$$\operatorname{conv}\{x_i : i \in S\} \bigcap \operatorname{conv}\{x_j : j \in T\} \neq \varnothing.$$

44 Example

One may think of the two sets as arising from coloring points as red or blue. In dimension $d = 1$, we have at least three points on a line. Color the outermost two points red and the interior point blue.

In dimension $d = 2$, we have at least four points. If three of the points are vertices of a triangle and the remaining point is inside the triangle, we are again done. If the four points form a non-degenerate quadrilateral, we may color opposite vertices the same color, and we are again done.

Proof We begin by adding an extra dimension, namely let $v_i := \begin{pmatrix} 1 \\ x_i \end{pmatrix} \in \mathbb{R}^{d+1}$. We have at least $m \geq d+2$ vectors in $\mathbb{R}^{d+1}$, so they must be linearly dependent. That is, there exist $\lambda_i \in \mathbb{R}$ not all zero such that $\sum_i \lambda_i v_i = 0$. In particular,

$$\sum_i \lambda_i = 0, \qquad \sum_i \lambda_i x_i = 0.$$

Set $S := \{i \in [m] : \lambda_i > 0\}$, $T := \{j \in [m] : \lambda_j \leq 0\}$. Since the sum of the λ_i is zero and not all coefficients are zero, it follows that both S and T are non-empty. Set $t := \sum_{i \in S} \lambda_i = \sum_{j \in T} -\lambda_j$. We now have

$$\sum_{i \in S} \frac{\lambda_i}{t} x_i = \sum_{j \in T} \frac{-\lambda_j}{t} x_j.$$

The coefficients have been chosen to sum to 1 and are evidently non-negative, so we are done.

45 Theorem (Tverberg, 1966)

If $p_1, \dots, p_m \in \mathbb{R}^d$ where $m \geq (r-1)(d+1)+1$, then there exists a set partition

$$[m] = S_1 \coprod \cdots \coprod S_r$$

such that

$$\bigcap_{j=1}^r \operatorname{conv}\{p_i : i \in S_j\} \neq \varnothing.$$

46 Remark

The original proof was quite long, though if we have a spare lecture we could now prove it in an hour. The $r = 2$ case is Radon's theorem. There are now "colorful" and "topological" versions of this result as well.

47 Theorem (Helly, 1921)

Let $K_1, K_2, \dots, K_n$ be convex sets in $\mathbb{R}^d$, with $n \geq d+1$ such that every $d+1$ of these sets have a point in common. Then $\cap_{i=1}^n K_i \neq \varnothing$.

48 Example

At $d = 1$, we have a collection of closed intervals in the real line. Given that any pair of them intersect non-trivially, one can quickly convince oneself that they all must intersect non-trivially using the minimum/maximum of the intervals. At $d = 2$, a Venn diagram provides an instance of this theorem: three circles arranged so that every two intersect non-trivially have a non-trivial triple overlap.

PROOF Idea: use induction on n. If $n = d + 1$, the statement is trivial. The inductive assumption for n implies that for $n - 1$, so we may assume that

$$K_1 \cap K_2 \cap \cdots \cap \widehat{K_i} \cap \cdots \cap K_n \neq \varnothing,$$

where the hat denotes we omit the ith term. We then have $p_1, \ldots, p_n \in \mathbb{R}^d$ with $n \geq d + 2$ for which we can apply Radon's theorem. The details are left to homework.

49 Remark

In the definition of $\operatorname{conv}(S)$, we allowed $m \in \mathbb{N}$ to be arbitrary. One may ask if it suffices to consider only m up to a certain size. Caratheodory's theorem says that we may choose $d + 1$ as this upper bound. Precisely:

50 Theorem (Caratheodory)

If $S \subset \mathbb{R}^d$ and $x \in \operatorname{conv}(S)$, then there exists $R \subset S$ such that $|R| \leq d + 1$ such that $x \in \operatorname{conv}(R)$.

51 Example

Consider a non-degenerate hexagon. One may triangulate the hexagon using convex hulls of triples of vertices, which is just a restatement of Caratheodory's theorem in this context.

PROOF Idea: let $x = \sum_{i=1}^{m} \lambda_i p_i$ with $p_i \in S$, $\lambda_i \geq 0$, $\sum_i \lambda_i = 1$. We may restrict to the case when $m = d+2$. As in the proof of Radon's theorem, one may use linear dependence after "lifting" to a higher dimension to eliminate one of the coefficients while the rest of the expression is still a convex combination. The remaining details are again left to homework.

52 Definition

Let $x_1, \ldots, x_m \in \mathbb{R}^d$ and consider $v_i := \begin{pmatrix} 1 \\ x_i \end{pmatrix} \in \mathbb{R}^{d+1}$. If there is a non-trivial linear dependence relation between the v_i, then we say that $x_1, \ldots, x_m$ are affinely dependent.

53 Remark

For further information on theorems along these lines, one may look at Matousek, §1.2, 1.3; or Igor Pak, Ch. 1, Ch.2.

54 Remark

Next time we will prove a separation theorem for convex sets and use that theorem to begin proving that V-polytopes and H-polytopes are equivalent.

April 1st, 2016: Separation theorem; Farkas' Lemma; Duality

---

55 Remark

Today we'll discuss the separation theorem and duality. Our goal is again laying groundwork for the equivalence of V-polytopes and H-polytopes.

56 Theorem

Let $C, D \subset \mathbb{R}^d$ be disjoint compact convex sets. Then there exists a hyperplane $h \subset \mathbb{R}^d$ such that $C \subset h^+ - h$ and $D \subset h^- - h$.

57 Example

Take two convex polygons in the plane which are disjoint. The the claim is that there is a line touching neither polygon such that each polygon is on a different side of the line.

Compactness is essential for strict separation. For instance, suppose $C \subset \mathbb{R}^2$ is the closed half-plane $y \le 0$ and D is the region above and including the right half of the hyperbola $y = 1/x$. These are convex and disjoint, but any separating hyperplane h is forced to be the x-axis, which is contained in C.

58 Remark

For a more general statement, see Matousek. The same statement holds if C is compact convex and D is merely closed convex.

Indeed, if they are both merely convex (still disjoint), it works except we can no longer guarantee that they intersect h trivially. That is, there exists a hyperplane $h \subset \mathbb{R}^d$ such that $C \subset h^+$, $D \subset h^-$.

PROOF Define $\mathrm{dist}: \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ by $(x, y) \mapsto \mathrm{dist}(x, y)$. Now $\mathrm{dist}: C \times D \to \mathbb{R}$ is a continuous map from a compact set to $\mathbb{R}$, so it achieves its minimum, say d_0. Since C, D are disjoint, $d_0 > 0$. Now, we have $x_0 \in C, y_0 \in D$ such that $\mathrm{dist}(x_0, y_0) = d_0$ and $\mathrm{dist}(x, y) \ge d_0$ for all $x \in C, y \in D$. Take h to be the hyperplane perpendicular to the line segment $[x_0, y_0]$ passing through the midpoint. Say $x_0 \in h^+ - h$, $y_0 \in h^- - h$.

We must show $C \subset h^+ - h$, $D \subset h^- - h$. Suppose to the contrary that we have some $y' \in D \cap h^+$. Since D is convex, it contains the segment $[y_0, y']$, but $\mathrm{dist}(x_0, [y', y_0]) < d_0$, a contradiction.

59 Lemma (Farkas, geometric statement)

If $P = \mathrm{conv}(V) \subset \mathbb{R}^d$ where $V \subset \mathbb{R}^d$ is finite, then either $0 \in P$ or there exists a hyperplane $h \subset \mathbb{R}^d$ such that $P \subset h^+ - h$ and $0 \in h^- - h$.

PROOF Say $0 \notin P$. Now $\{0\}$ is convex and compact. Certainly P is convex. It is the convex hull of a compact set, which is compact by homework, though we give a more direct argument as well. If $V = \{v_1, \ldots, v_r\}$, then $\mathrm{conv}(V)$ is the set of convex combinations of $v_1, \ldots, v_r$, which is evidently a "closed condition." Moreover, it is bounded: given $\sum_{i=1}^n \alpha_i v_i$ with $\alpha_i \ge 0$, $\sum_i \alpha_i = 1$, we have

$$|\sum_i \alpha_i v_i| \le \sum_i |v_i|.$$

The result now follows from the separation theorem.

60 Remark

Suppose we have a matrix with d rows, n columns,

$$A = \begin{bmatrix} v_1 & \cdots & v_n \end{bmatrix}$$

with $v_i \in \mathbb{R}^d$. Let $P := \mathrm{conv}\{v_1, \ldots, v_n\}$. Consider two cases:

$0 \in P$: We have $0 = \sum_{i=1}^n \alpha_i v_i$ with $\alpha_i \ge 0$, $\sum_i \alpha_i = 1$. That is, $Ax = 0$ has a non-negative, non-trivial solution.

$0 \notin P$: By Farkas' lemma, we have some hyperplane h such that $P \subset h^- - h$ and $0 \in h^+$. Recall that h is defined in terms of constants $a \in \mathbb{R}^d - \{0\}, b \in \mathbb{R}$, so that $0 \in h^+ - h$ says precisely that $0 = \langle 0, a \rangle > b$, so b is negative. On the other hand, $v_i \in h^- - h$ says $\langle a, v_i \rangle < b < 0$. Equivalently, $a^T A$ has all negative entries.

These considerations lead to the following version of Farkas' lemma:

61 Lemma (Farkas' lemma, algebraic version)

If A is a $d \times n$ matrix, then either $Ax = 0$ has a non-negative and non-trivial solution, or there exists some $a \in \mathbb{R}^d$ such that $a^T A$ has all negative coordinates.

62 Notation

For now we will use the term "duality" following Matousek, §5.1, 5.2. We will switch to "polarity" later when we switch sources. We will use this notion to prove the equivalence of V-polytypes and H-polytopes.

63 Remark

First, an intuitive discussion of duality. Given $X \subset \mathbb{R}^d$, we can consider the set $X^* \subset (\mathbb{R}^d)^*$ given by all linear functionals $f\colon \mathbb{R}^d \to \mathbb{R}$ whose values on X are all ≤ 1. Instead of working literally in the dual space $(\mathbb{R}^d)^*$, we will work in $\mathbb{R}^d$ by using the linear isomorphism induced by the standard inner product. Precisely, use

$$\mathbb{R}^d \xrightarrow{\sim} (\mathbb{R}^d)^*$$
$$a \in \mathbb{R}^d \mapsto \langle -, a\rangle.$$

64 Definition

The $\boxed{\text{duality transform}}$ is the mapping $\boxed{D_0}$ from non-zero points $p \in \mathbb{R}^d$ to hyperplanes $D_0(p)$ given by

$$D_0(p) := \{x \in \mathbb{R}^d : \langle x, p\rangle = 1\}.$$

65 Example

Using $P = (2,0)$, $D_0(P)$ is the line $x = 1/2$. Using $P = (1,1)$, $D_0(P)$ is the line $x + y = 1$. In general, if P has distance δ from the origin, $D_0(P)$ is the hyperplane perpendicular to the line segment $[0,P]$ at distance $1/\delta$ from the origin.

Let $\boxed{D_0(p)^-}$ denote the half-space that contains 0.

66 Definition

Given $X \subset \mathbb{R}^d$, define the $\boxed{\text{dual}}$ of X as

$$\boxed{X^*} := \bigcap_{x \in X - \{0\}} D_0(x)^-.$$

67 Example

We have $\{0\}^* = \mathbb{R}^d$ using our usual convention for empty intersections.

Now suppose X is the x-axis in $\mathbb{R}^2$. Then each non-zero point $P \in X$ has $D_0(P)^-$ given by a half-plane with boundary parallel to the y-axis containing 0. It follows that X^* here is simply the y-axis. Note that whether or not we include 0 in X does not affect the dual.

68 Remark

If $X = \text{conv}(V)$ for V a finite set, one may show that X^* is $\cap_{v \in V} \{v\}^*$. See homework.

69 Example

Let X be a triangle in the plane containing the origin as an interior point. One may check that X^* is also a triangle containing the origin.

Note that

$$X^* = \{y \in \mathbb{R}^d : \forall x \in X, \langle x, y\rangle \le 1\}.$$

From now on, we may take this as the definition.

70 Remark

Our next goal is to try to understand X^{**}. First some general remarks about X^*.

Note that X^* is the intersection of some closed half-spaces, which are closed and convex, so X^* is closed and convex. Certainly $0 \in X^*$. Also, $A \subset B$ implies $A^* \supset B^*$.

We may expect $X^{**} \supset X$, which by the above considerations implies $X^{**} \supset \overline{\text{conv}(X \cup \{0\})}$. Indeed, our next goal is to show equality holds.

71 Theorem
$X^{**} = \overline{\text{conv}(X \cup \{0\})}$.

72 Corollary
If X is convex, closed, and contains 0, then $X^{**} = X$.

Proof For the $\supset$ inclusion, as noted above it suffices to show $X^{**} \supset X$. If $x \in X$, then $\langle x, y\rangle \leq 1$ for all $y \in X^*$ by definition. Since $x \in X^{**}$ if and only if $\langle x, y\rangle \leq 1$ for all $y \in X^*$, the inclusion follows.

We will prove the $\subset$ inclusion next lecture. Equivalently, we may show that $v \notin \overline{\text{conv}(X \cup \{0\})}$ implies $v \notin X^{**}$. Equivalently, we may show that there exists $a \in X^*$ such that $\langle v, a\rangle > 1$.

73 Remark
Imagine X is a closed circle in $\mathbb{R}^3$ not intersecting the origin. We are considering the cone connecting X to the origin. For v not on that cone, by the separation theorem we have some hyperplane dividing them. We will essentially choose a above as the vector defining this hyperplane. More details next time.

April 4th, 2016: Duality, V-polytopes, and H-polytopes

---

74 Remark
Recall that if $p \in \mathbb{R}^d - \{0\}$, then $D_0(p) := \{x \in \mathbb{R}^d : \langle p, x\rangle = 1\}$, and $D_0(p)^- := \{x \in \mathbb{R}^d : \langle p, x\rangle \leq 1\}$. One may take $D_0(0)^- := \mathbb{R}^d$. If $X \subset \mathbb{R}^d$, then the dual of X is $X^* := \cap_{x\in X} D_0(x)^- = \{y \in \mathbb{R}^d : \langle x, y\rangle \leq 1, \forall x \in X\}$.

75 Theorem
For any $X \subset \mathbb{R}^d$, we have $X^{**} = \overline{\text{conv}(X \cup \{0\})}$. In particular, if $0 \in X$ and X is closed and convex, then $X^{**} = X$.

Proof Last time we proved the easy direction, $\supset$, by showing $X^{**} \supset X$. Today we will give the other containment, or really its contrapositive.

For convenience, call the right-hand side Q. Pick $v \in \mathbb{R}^d - Q$. We must show $v \notin X^{**}$. Since Q is closed and convex, and $\{v\}$ is trivially compact and convex, by the separation theorem we have some hyperplane h that strictly separates v and Q. Since $0 \in Q$, we have $0 \notin h$, so we can "normalize" h and write $h := \{x : \langle a, x\rangle = 1\}$. Since $\langle 0, a\rangle = 0 < 1$, we have $Q \subset h^- - h$ and $v \in h^+ - h$. Hence $\langle v, a\rangle > 1$ and for all $x \in Q$, $\langle x, a\rangle \leq 1$. Since $X \subset Q$, this last statement says $a \in X^*$. However, $\langle v, a\rangle > 1$ now says $v \notin X^{**}$, completing the proof.

76 Remark
During the proof, we may imagine some point v and some cone over a set X which is convex and closed, which are separated by h. After normalizing, the normal of h is in X^* but not in X^{**}.

77 Remark
Last time we proved the separation theorem for two compact convex disjoint sets, but we stated it remains true when one of the sets is merely closed instead of compact. The rough argument for proving this is to intersect the closed set with some large enough ball so that the distance is unchanged, which reduces to the "both compact" case. See Matousek for more details and the case when neither is compact.

78 Exercise
The following appear in homework:

- If $C \subset \mathbb{R}^d$ is convex, then C^* is bounded if and only if $0 \in \text{Int}(C)$.

  (Recall that the interior of a subset of a topological space is the set of points in the subset such that the subset contains an open set in the ambient space containing that point.)

- If $P = \operatorname{conv}\{v_1, \ldots, v_n\}$, then $P^* = \cap_{i=1}^n D_0(v_i)^- = \cap_{i=1}^n \{x \in \mathbb{R}^d : \langle x, v_i\rangle \le 1\}$.

79 Corollary

We have the following:

- *If P is a V-polytope, then P^* is an H-polyhedron.*
- *Moreover, if P is a V-polytope and $0 \in \operatorname{Int}(P)$, then P^* is an H-polytope.*

The corollary will essentially give us one direction of the H-polytope/V-polytope equivalence for free from the other direction, which we are now ready to prove.

80 Theorem

Each V-polytope is an H-polytope, and each H-polytope is a V-polytope.

81 Remark

This proof follows Edmonds as given by Matousek, §5.1.

Proof We begin by showing that every H-polytope is a V-polytope. We induct on d where $P \subset \mathbb{R}^d$. Begin at $d = 1$. We may assume P is in $\mathbb{R}^1$ and that P is a finite, bounded intersection of closed rays. This is either a closed interval $[a, b]$ with $a \ne b$, a point, or $\varnothing$. These are, respectively, $\operatorname{conv}\{a, b\}$, $\operatorname{conv}\{a\}$, $\operatorname{conv}(\varnothing)$, which are all V-polytopes.

Now suppose $d > 1$. Let $P = \cap_{i=1}^m h_i^-$ be a non-empty, bounded, finite intersection of closed half-spaces in $\mathbb{R}^d$. Let

$$F_i := P \cap h_i = (\cap_{j=1}^m h_j^-) \cap (h_i^- \cap h_i^+).$$

Thus $F_i \subset h_i$ is an H-polyhedron when we identify h_i with $\mathbb{R}^{d-1}$, and it inherits boundedness, so F_i is an H-polytope. By induction, F_i is a V-polytope, so we have a finite set $V_i \subset h_i$ such that $F_i = \operatorname{conv}(V_i)$. Set $V := \cup_i V_i$. It suffices to show that $P = \operatorname{conv}(V)$.

Pick any $x \in P$. Let ℓ be any line passing through x and consider $\ell \cap P$. Since P is compact, convex, we have $\ell \cap P = [y, z]$. We claim $y, z \in \cup_i h_i$. If not, say $y \notin \cup_i h_i$, then $y \in P - \cup_i h_i = \cap_i (h_i^- - h_i)$. This latter intersection is open, so there is an open neighborhood around y contained entirely in P, so we can extend the interval $[y, z]$, a contradiction. Hence $y, z \in \cup_i F_i$, so convex combinations of y, z belong to $\operatorname{conv}(V)$, giving $x \in \operatorname{conv}(V)$ and completing the proof of this direction.

We now show that every V-polytope is an H-polytope. Let $P = \operatorname{conv}(V)$ for $V \subset \mathbb{R}^d$ finite. We may assume P is full dimensional (by restricting to the affine hull of P if needed) and $0 \in \operatorname{Int}(P)$ (by translating if needed). Then by the corollary above, P^* is an H-polytope, so by the other direction, P^* is a V-polytope. Again using the corollary, P^{**} is an H-polyhedron. By the first theorem from today's lecture, $P = P^{**}$ is a bounded H-polyhedron, completing the proof.

82 Remark

As an exercise, justify the $0 \in \operatorname{Int}(P)$ step above. Alternatively, see Isabella's email.

83 Remark

Recall that a face of P is a non-empty intersection of the form $F := P \cap h$ where h is a supporting hyperplane of P, i.e. $P \subset h^+$. Again, 0-dimensional faces are vertices, 1-dimensional faces are edges, codimension-1 faces are facets, codimension-2 faces are ridges.

At present, we do not know that polytopes have finitely many faces, or that a polytope is the convex hull of its vertices. We next deal with these deficiencies.

84 Theorem

Let $P \subset \mathbb{R}^d$ be a polytope. Consider

$$\{V \subset \mathbb{R}^d : |V| < \infty, \operatorname{conv}(V) = P\},$$

which is non-empty by the preceding theorem. Let V_0 be a minimal element of this set. Then:

(a) *V_0 is the vertex set of P.*

(b) *If F is a face of P, then the vertex set of F is $V_0 \cap F$.*

(c) *Fix a face F of P. G is a face of F if and only if $G \subset F$ and G is a face of P.*

(d) *If F, F' are faces of P, then $F \cap F'$ is a face of P.*

85 Definition
If $P \subset \mathbb{R}^d$ is a polytope, write $\boxed{\mathrm{vert}(P)}$ for the set of vertices of P. This is finite by (a).

Proof (a) Suppose $v \in \mathrm{vert}(P)$. Let h be a supporting hyperplane of P such that $P \cap h = \{v\}$ and $P \subset h^-$. Then $P - \{v\} = P \cap (h^- - h)$ is convex, so $\mathrm{conv}(P - \{v\}) \subset h^- - h$, and $v \in h \Rightarrow v \notin h^- - h$, so $v \notin \mathrm{conv}(P - \{v\})$. It follows that any V above must contain v, so $v \in V_0$, giving $\mathrm{vert}(P) \subset V_0$.

For the other inclusion, pick $v \in V_0$ and set $C := \mathrm{conv}(V_0 - \{v\})$. By minimality, $C \subsetneq P = \mathrm{conv}(V_0)$, so $v \notin C$. By the separation theorem, we have a hyperplane separating $\{v\}$ and C. We will finish this argument next lecture.

April 6th, 2016: Faces, Vertices, Simplicies, and the Face Lattice Revisited

---

86 Remark
Today we'll discuss faces of a polytope. We begin by finishing the theorem from the end of last class.

Note that homework 1 is due on Friday. There are office hours on Friday from 10:30 to 11:30.

87 Theorem
Let $P \subset \mathbb{R}^d$ be a polytope, and let V_0 be an inclusion-minimal set among all sets V such that $P = \mathrm{conv}(V)$. Then

(a) *V_0 is the vertex set of P, and so P is the convex hull of its vertex set.*

(b) *If F is a face of P, then $\mathrm{vert}(F) = F \cap \mathrm{vert}(P)$.*

(c) *If F is a face of P, then G is a face of F if and only if $G \subset F$ is a face of P.*

(d) *If F, F' are faces of P, then $F \cap F'$ is a face of P.*

Proof Last time we got through half of part (a), namely $\mathrm{vert}(P) \subset V_0$. So, consider the reverse inclusion. Let $v \in V_0$ and consider $P' := \mathrm{conv}(V_0 - \{v\})$. By minimality of V_0, we have $P' \subsetneq P$, and in particular $v \notin P'$. By the separation theorem, there exists a hyperplane h that strictly separates v and P'. Let h_v be the translation of h such that $v \in h_v$. To show that v is a vertex of P, it suffices to check that h_v is a supporting hyperplane of P and that $h_v \cap P = \{v\}$. We may say $P' \subset h^+ - h$, so $P' \subset h_v^+ - h_v$. In particular, we may take as usual

$$h_v := \{x : \langle x, a \rangle = b\}$$

where $v \in h_v$ says $\langle v, a \rangle = b$ and $V_0 - \{v\} \subset h_v^+ - h_v$ says for all $x \in V_0 - \{v\}$, $\langle v, a \rangle > b$. Then for any $x \in P = \mathrm{conv}(V_0)$ we have a convex combination

$$x = \alpha_v v + \sum_{w \in V - \{v_0\}} \alpha_w w.$$

where $\alpha_w \geq 0$ and $\sum \alpha_w = 1$. Hence

$$\begin{aligned}\langle x, a\rangle &= \alpha_v \langle v, a\rangle + \sum_{w \in V_0 - \{v\}} \alpha_w \langle w, a\rangle \\ &\geq \alpha_v b + \sum_{w \in V_0 - \{v\}} \alpha_w b = b.\end{aligned}$$

where we have equality if and only if all $\alpha_w = 0$ for $w \in V_0 - \{v\}$, so if and only if $x = v$. This completes (a).

For (b), write $F \cap \operatorname{vert}(P) = \{v_1, \ldots, v_r\}$ and $V_0 - F = \{v_{r+1}, \ldots, v_n\}$. Since F is a face, we have a supporting hyperplane

$$h \coloneqq \{x \in \mathbb{R}^d : \langle x, a\rangle = b\}$$

such that $\langle x, a\rangle = b$ for all $x \in F$ and $\langle x, a\rangle > b$ for all $x \in P - F$. If $x \in F \subset P$, we have a convex combination

$$x = \sum_{w \in V_0} \alpha_w w = \sum_{i=1}^{r} \alpha_i v_i + \sum_{j=r+1}^{n} \alpha_j v_j$$

with $\alpha_w \geq 0$, $\sum_w \alpha_w = 1$. We compute

$$\begin{aligned}b = \langle x, a\rangle &= \sum_{i=1}^{r} \alpha_i \langle v_i, a\rangle + \sum_{j=r+1}^{n} \alpha_j \langle v_j, a\rangle \\ &\geq \sum_{i=1}^{r} \alpha_i b + \sum_{j=r+1}^{n} \alpha_j b \qquad\qquad \geq .\end{aligned}$$

Indeed, equality holds if and only if $\alpha_j = 0$ for all $j \in [r+1, n]$. Thus, $x \in F$ iff $x \in \operatorname{conv}\{v_1, \ldots, v_r\} = \operatorname{conv}(V_0 \cap F)$. Note that F itself is a polytope since $F = P \cap h^+ \cap h^-$, so if we can show that $V_0 \cap F$ is inclusion-minimal for finite sets whose convex hull is F, by (a) we will be done. By minimality of V_0, no vertex of $V_0 \cap F$ is a convex combination of the remaining vertices of $V_0 \cap F$, so $V_0 \cap F$ is indeed minimal.

The proof of (c) is a problem on homework 2.

For (d), if F, F' are faces, we have h, h' hyperplanes given by $a, b; a', b'$, as above, where $h \cap P = F$, $h' \cap P' = F'$, $P \subset h^+$, $P' \subset (h')^+$. Consider the hyperplane determined by $a + a'$, $b + b'$. This is a supporting hyperplane which one may check intersects P in $F \cap F'$.

88 Remark

During the proof that $V_0 \subset \operatorname{vert}(P)$, we imagine that P is a convex polygon in $\mathbb{R}^2$ and that P' is obtained by deleting one of the vertices v. Now h is a separating hyperplane of v and P', which can be translated to pass through v.

(a) and (b) together imply that any polytope has only finitely many faces, since the faces are determined by subsets of $\operatorname{vert}(P)$, of which there are finitely many.

89 Aside

Take a d-dimensional polytope P and compute the number f_i of i-dimensional faces of P. We would hope that $f_i(P) = \min\{f_0(P), f_{d-1}(P)\}$, i.e. the number vertices or facets is where the minimum of the f_i is achieved. This is an open problem, at least for d large enough (probably open for $d \geq 8$). It is true for simple and simplicial polytopes, but not for the "intermediate" cases. Note that the following relatively natural conjecture is **false**:

$$f_0 \leq f_1 \leq \cdots \leq f_p \geq f_{p+1} \geq \cdots,$$

i.e. the f_i need not be unimodal.

90 Example

We next discuss an important example of extremely well-behaved polytopes.

91 Definition

A $\boxed{\text{simplex}}$ is the convex hull of affinely independent points, i.e. $\begin{pmatrix}1\\v_i\end{pmatrix}$ are linearly independent and our simplex is $\text{conv}\{v_1,\ldots,v_d\}$.

92 Remark

Note that injective affine transformations preserve affine independence, supporting hyperplanes, faces, convexity, etc. Hence we may as well consider the simplex with $v_i := e_i$. That is

$$\boxed{\delta_{d-1}} := \text{conv}\{e_1,\ldots,e_d\} \subset \mathbb{R}^d$$

is "the" $d-1$-dimensional simplex. (One sometimes takes one of the vertices to be 0 to be more "efficient" about using the dimension of the ambient space.)

We have that δ_0 is a single point in $\mathbb{R}^1$, δ_1 is the line segment $[e_1,e_2]$ in $\mathbb{R}^2$, δ_2 is the triangle between e_1, e_2, e_3, etc. Since all subsets of affinely independent subsets are affinely independent, it follows that faces of simplicies are themselves simplicies.

93 Proposition

Every subset of $\{e_1,\ldots,e_d\}$ is the vertex set of a face of δ_{d-1}. Hence the face lattice $L(\delta_{d-1})$ is the subset lattice on $[d]$, i.e. the boolean lattice of rank d.

Proof We have

$$\text{conv}\{e_1,\ldots,e_d\} = \{(x_1,\ldots,x_d) : x_i \geq 0, \sum_i x_i = 1\}.$$

Now let $I := \{i_0,\ldots,i_k\} \subset [d]$. and say $F_I := \text{conv}\{e_i : i \in I\}$. To see that F_I is a face, we produce a supporting hyperplane. Indeed, take

$$h_I := \{x \in \mathbb{R}^d : \langle x, e_{i_1} + \cdots + e_{i_k}\rangle = 1\}.$$

Note that the $\langle x, e_{i_1} + \cdots + e_{i_k}\rangle = x_{i_1} + \cdots + x_{i_k}$. If $x \in \delta_{d-1}$, this is always ≤ 1, and this is 1 precisely when $x \in F_I$, so $h_I \cap \delta_{d-1} = F_I$.

94 Definition

Let P be a polytope. We may now formally define the $\boxed{\text{face poset}}$ or $\boxed{\text{face lattice}}$ of P, $\boxed{L(P)}$, as the set of faces of P ordered by inclusion, which includes $\varnothing$ and P. The following theorem summarizes its basic properties.

95 Theorem

Let P be a polytope.

(1) $L(P)$ is a finite poset.

(2) $L(P)$ has minimum $\varnothing =: \widehat{0}$, maximum $P =: \widehat{1}$. (This is sometimes called "bounded".)

(3) The atoms of $L(P)$ (i.e. the minimal elements of $L(P) - \{\widehat{0}\}$) are the vertices of P.

(4) $F, F' \in L(P)$ implies $F \cap F' \in L(P)$, so every two elements of $L(P)$ have a meet.

(5) Every two elements of F and F' have a join, i.e. $L(P)$ is a lattice.

(6) If $F \in L(P)$, then $L(F) = [\widehat{0}, F]$, so $[0, F]$ is the face lattice of a polytope.

(7) $L(P)$ is atomic (i.e. every face is the join of some atoms, namely its vertices).

Proof These are all straightforward or formal consequences of the preceding theorem.

April 8th, 2016: Face Lattices; Quotients; Simple and Simplicial Polytopes

---

96 Remark

Recall that if P is a polytope, $L(P)$ is the poset of all faces of P including $\varnothing$ and P ordered by inclusion. We proved last time that $L(P)$ is in fact a finite lattice with several other nice properties.

97 Theorem

We have:

(1) $L(P)$ is a graded lattice.

(2) $L(P)^{\mathrm{op}}$ is the face lattice of some polytope (namely, P^* if P is full-dimensional and $0 \in \mathrm{Int}(P)$).

(3) For all $G \subset F$ with $G, F \in L(P)$, $[G, F]$ is the face lattice of a polytope.

98 Definition

Any polytope whose face lattice is isomorphic to the interval $[G, P]$ with $G \in L(P)$ is called the $\boxed{\text{quotient}}$ of P by G, denoted by $\boxed{P/G}$. Another name for the quotient is the $\boxed{\text{link}}$ of G in P. When G is a vertex, we call this quotient the $\boxed{\text{vertex figure}}$.

At present quotients are just defined up to combinatorial isomorphism. We will eventually see how to think of them as subcomplexes in some nice cases, for instance when $G = \{v\}$ at the end of today's lecture.

A polytope Q such that $L(Q) \cong L(P)^{\mathrm{op}}$ is called a $\boxed{\text{combinatorial dual}}$ or sometimes just "dual." To avoid confusion with the geometric object P^*, we will now call P^* the $\boxed{\text{polar}}$ of P. Hence (2) asserts that the polar of P is (usually) a combinatorial dual.

Proof We assume P is full-dimensional and $0 \in \mathrm{Int}(P)$.

(1) is proved in the next homework set. The argument uses induction, the map $L(P) \to L(P^*)$ below, and the fact that given a vertex v, $\widehat{v}$ is a facet of P^*.

We now consider (2), $L(P^*) \cong L(P)^{\mathrm{op}}$. We first define a map $L(P) \to L(P^*)$ as follows. Given $F \in L(P)$, let

$$\widehat{F} := \{x \in P^* : \langle x, y \rangle = 1, \forall y \in F\}.$$

We must show (a) $\widehat{F}$ is a face of $L(P^*)$ and (b) that $F \mapsto \widehat{F}$ is an inclusion-reversing bijection $L(P) \to L(P^*)$. The rough idea for (a) is to replace the defining condition for $\widehat{F}$ which quantifies over infinitely many y with a single point in the "interior" of F.

(a) The $\boxed{\text{relative interior}}$ of F is the set

$$\boxed{\operatorname{rel\,int}(F)} := F - \bigcup_{G \in L(F) - \{F\}} G,$$

i.e. the set of points of F in none of F's proper faces. It is non-empty since the interior of a non-empty polytope is in general non-empty. Pick $x_0 \in \operatorname{rel\,int}(F)$ and let

$$F^* := \{y \in P^* : \langle y, x_0 \rangle = 1\},$$

which is the intersection of P^* and the hyperplane defined by $a = x_0$, $b = 1$. By definition of P^*, this hyperplane is a supporting hyperplane, so F^* is a face of P^*. Trivially $\widehat{F} \subset F^*$, so it suffices to prove the other inclusion, or equivalently to take $y_0 \in P^* - \widehat{F}$ and show $y_0 \notin F^*$. We have $\langle y_0, x_1 \rangle < 1$ for some $x_1 \in F$. Since $x_0 \in \operatorname{rel\,int}(F)$, there exists $x_2 \in F$ such that $x_0 \in (x_1, x_2)$. Now $\langle y_0, x_2 \rangle \leq 1$. Writing x_0 as a convex combination of x_1 and x_2 with two non-zero coefficients, it follows immediately that that $\langle y_0, x_0 \rangle < 1$, so $y_0 \notin F^*$.

(b) Order-reversal is immediate. A nice trick to show this is a bijection is to show that doing this twice yields the identity, i.e. $\widehat{\widehat{F}} = F$. This is done in the next homework.

Now consider part (3). We know $[G, F] \subset [\widehat{0}, F] \subset L(P)$ where $[\widehat{0}, F] = L(F)$. That is, we know that lower intervals are face lattices of polytopes, and we must show the same for upper intervals. This follows immediately by duality. More precisely, the interval $[G, F]$ corresponds in an order-reversing fashion to $[\widehat{0}, \widehat{G}]$ in $L(\widehat{F})$, and dualizing this using (2) gives (up to a slight abuse of notation) $[G, F] \cong [\widehat{G}^*, \widehat{F}^*] \subset L(\widehat{F}^*)$.

99 Definition

Recall that $f_i(P)$ is the number of i-dimensional faces of P. Note that the 0-dimensional faces of P correspond to the $(d-1)$-dimensional faces of P^*, and generally that the i-dimensional faces of P correspond to the $(d-1-i)$-dimensional faces of P^*. The f-vector of P is

$$f(P) \coloneqq (f_{-1}(P), f_0(P), \ldots, f_{d-1}(P), f_d(P))$$

where sometimes $f_{-1}(P)$ and/or $f_d(P)$ are left off.

100 Example

Let P be the unit cube in $\mathbb{R}^3$ centered at 0. P^* is then two square pyramids glued together at the square. The f-vector of P (cutting off the two ends) is $(1, 8, 12, 6, 1)$, while the f-vector of P^* is $(1, 6, 12, 8, 1)$.

One may imagine picking a peak vertex v of P^*, picking a supporting hyperplane, and moving it slightly inside P^*. The resulting intersection has vertices given precisely by intersecting the hyperplane and edges containing v. The result in this case is a square, which is in fact the vertex figure in this context. Doing the same operation with the cube P yields a triangle. These observations are formalized in the next remark.

101 Remark

If v is a vertex of P, then there exists a hyperplane $h = \{x \in \mathbb{R}^d : \langle x, a \rangle = b\}$ such that $P \subset h^+$ and $P \cap h = \{v\}$. Let

$$Q \coloneqq P \cap \{x \in \mathbb{R}^d : \langle x, a \rangle = b + \epsilon\}$$

for a small $\epsilon > 0$. Check: $L(Q) \cong [\{v\}, P] \subset L(P)$.

102 Remark

Recall that the lattice of a $(d-1)$-dimensional simplex is just the boolean lattice of rank d. In particular, the face numbers are just

$$f_{i-1}(\delta^{d-1}) = \binom{d}{i}.$$

103 Definition

A polytope is simplicial if every facet is a simplex.

104 Remark

The following are equivalent:

(a) P is a simplicial d-dimensional polytope.

(b) Every facet of P has exactly d vertices.

(c) Every face of P is a simplex.

(d) $[0, F]$ is a boolean lattice for every $F \in L(P) - \{P\}$.

Pictorially, all of the lower intervals are boolean lattices, in particular the ones below the coatoms, except that $L(P)$ need not itself be a boolean lattice.

105 Definition

A polytope P is simple if P^* is simplicial.

106 Remark

The following are equivalent:

(a) P is a simple d-dimensional polytope.

(b) $[v, P] \subset L(P)$ is a boolean lattice of rank d for all vertexes v of P.

(c) $[G, P] \subset L(P)$ is a boolean lattice for all $G \in L(P) - \varnothing$.

(d) Every vertex is incident with exactly d edges.

(e) Every vertex is incident with exactly d facets.

Pictorially, all of the upper intervals are boolean lattices, in particular the ones above the atoms, except that $L(P)$ need not be.

107 Example

A cube is simple and its dual is simplicial. A square pyramid is not simple because the peak vertex lies on too many edges and is not simplicial because the square has too many vertexes.

108 Remark

Note that if P is simple, then every face of P is simple. If P is simple and v is a vertex lying on edges $e_1, \ldots, e_k$, then there is a unique face of P containing those edges. As an exercise, pick your favorite condition above to see this.

Dually, if P is simplicial, then all quotients are simplicial.

April 11th, 2016: Cyclic polytopes; Neighborliness

---

109 Remark

Today we'll discuss cyclic polytopes and neighborliness. See Matusek, §5.4; Barvinok, Ch. 6.

110 Definition

The curve $q(t) := (t, t^2, t^3, \ldots, t^d) \in \mathbb{R}^d$ is called the moment curve in $\mathbb{R}^d$. Pick n and $t_1 < t_2 < \cdots < t_n \in \mathbb{R}$. Suppose $n \geq d+1$. Now $\operatorname{conv}\{q(t_1), \ldots, q(t_n)\}$ is called a cyclic polytope, written $C(d, n)$

111 Example

Consider $d = 2$. The moment curve is just the parabola (t, t^2). A cyclic polytope for $n = 4$ with, say $t_1 < t_2 < 0 < t_3 < t_4$, is a trapezoid.

112 Remark

Any $d+1$ distinct points on the moment curve are affinely independent. To see this, note that the determinant of the matrix whose columns are these points with 1 appended is

$$\begin{bmatrix} 1 & 1 & \cdots & 1 \\ t_1 & t_2 & \cdots & t_{d+1} \\ t_1^2 & t_2^2 & \cdots & t_{d+1}^2 \\ \vdots & \vdots & \ddots & \vdots \\ t_1^d & t_2^d & \cdots & t_{d+1}^d \end{bmatrix} = \prod_{1 \leq i < j \leq d+1} (t_j - t_i) \neq 0.$$

Thus $C(d, n)$ is a d-dimensional polytope and all of its proper faces are simplices. (No facet can have more than d vertices.) Hence $C(d, n)$ is a simplicial d-dimensional polytope.

113 Definition

A (simplicial) d-polytope P is called m-neighborly if every collection of m vertices of P is the vertex set of a face of P. (The homework states this definition without the "simplicial" assumption.)

114 Example

A d-dimensional simplex is $(d+1)$-neighborly. 2-neighborliness means that every two vertices form an edge, i.e. $G(P)$ is a complete graph.

115 Theorem

We have:

(a) *Every collection of $k \le \lfloor d/2 \rfloor$ points $q(t_{i_1}), \ldots, q(t_{i_k})$ is the vertex set of a face of $C(d,n)$.*

(b) *(Gale's evenness condition.) Write $v_i := q(t_i)$. A d-tuple $V_d := (v_{i_1}, \ldots, v_{i_d})$ is the vertex set of a facet of $C(d,n)$ if and only if for every two points $v_i, v_j \in V - V_d$ (say $i < j$)*

$$|V_d \cap \{v_{i+1}, \ldots, v_{j-1}\}| \text{ is even.}$$

116 Example

We illustrate (b). Take $d = 4$ with $t_1 < t_2 < \cdots < t_8$. Suppose in our 4-tuple V_4 that we've already picked t_2, t_3, t_7. If we were to include t_1 in V_4, then between t_6 and t_8 we would have oddly many elements of V_4, so we would not get a facet. The name for (b) comes from the fact that, roughly, the blocks of chosen vertices aside from the first and last must come in even sizes.

Proof We begin with (a). Take $I = \{i_1, \ldots, i_k\} \subset [n]$ such that $k \le \lfloor d/2 \rfloor$. We must show $\{v_{i_1}, \ldots, v_{i_k}\}$ forms the vertex set of a face. It suffices to exhibit an appropriate supporting hyperplane, namely we find $c = (\gamma_1, \ldots, \gamma_d) \in \mathbb{R}^d - \{0\}$ and $\alpha \in \mathbb{R}$ such that

$$\begin{aligned} \langle c, v_i \rangle &= \alpha \qquad \forall i \in I \\ \langle c, v_j \rangle &> \alpha \qquad \forall j \in [n] - I. \end{aligned}$$

Consider

$$p(t) := (t - t_1 + 1)^{d-2k} (t - t_{i_1})^2 \cdots (t - t_{i_k})^2.$$

The first factor is roughly there to ensure $\deg p(t) = d$ without affecting the sign of p on V. Notice that

$$p(t) = \gamma_d t^d + \gamma_{d-1} t^{d-1} + \cdots + \gamma_1 t + \gamma_0.$$

Further, $p(t_i) = 0$ for all $i \in I$ and $p(t_j) > 0$ for all $j \in [n] - I$. Now take $c := (\gamma_1, \gamma_2, \ldots, \gamma_d)$ and $\alpha := -\gamma_0$. Then

$$\langle q(t_i), c \rangle = \gamma_d t_i^d + \gamma_{d-1} t_i^{d-1} + \cdots + \gamma_1 t_i = p(t_i) + \alpha.$$

By construction, this is α if $i \in I$ and is $> \alpha$ if $i \in [n] - I$, proving (a).

Now consider (b). Every $d+1$ points on the moment curve are affinely independent, so every d points are affinely independent, so every d points $V_d := \{t_{i_1}, \ldots, t_{i_d}\}$ define a unique hyperplane H in $\mathbb{R}^d$. These points form a facet precisely when this hyperplane is supporting. Hence we must show that H is a supporting hyperplane if and only if the evenness condition holds. Consider the intersection of H and the moment curve; it certainly contains V_d, and since every $d+1$ points on the moment curve are affinely independent, the intersection is precisely V_d. Hence V_d divides the normal curve into $d+1$ arcs, and in fact the intersection multiplicity must be 1 at each point, so these segments alternate which side of H, H^+ or H^-, they are on. It follows that $\mathrm{conv}(V_d)$ is a facet if and only if H is a supporting hyperplane if and only if all points of $V - V_d$ are on the same side of H if and only if every 2 points of $V - V_d$ are separated on the moment curve M by an even number of points on V_d.

117 Remark

The picture we have in mind during the proof of (b) is a hyperplane with a curve "bobbing" in and out of the hyperplane at discrete points, creating "segments" between points of intersection. It is forced to "switch sides" each time it goes through the hyperplane since the roots are simple.

118 Example

Let's compute the face numbers of $C(d,n)$. If $k \le \lfloor d/2 \rfloor$, every k vertices form a face, so

$$f_{k-1}(C(d,n)) = \binom{n}{k}, \qquad \forall k \le \lfloor d/2 \rfloor.$$

One may compute the other "half" of the f-vector from the Gale evenness condition, though the expression isn't so nice.

119 Conjecture (Upper Bound Conjecture; Motzkin, 1957))

Among all polytopes of dimension d with n vertices, the cyclic polytope $C(d,n)$ simultaneously maximizes all the face numbers. Precisely, if P is a d-polytope and $f_0(P) = n$, then

$$f_{i-1}(P) \le f_{i-1}(C(d,n)), \qquad \forall i.$$

120 Remark

The preceding example shows the conjecture holds for the first "half" of all i. It took another 13 years to show the second half indeed holds (McMullen, 1970). It turns out one may prove the simplicial case using "h-vectors" and use a simple trick to get the general form, which will be the focus of later lectures.

The original conjecture included a second part saying that cyclic polytopes are the unique polytopes with this property, which turns out to be quite false, as there are exponentially many others. In particular, the maximality statement holds for all triangulations of $(d-1)$-dimensional spheres, a result due to Stanley (1975). It also holds for all odd-dimensional (homology) manifolds and all even-dimensional manifolds whose Euler characteristic is the Euler characteristic of a sphere, namely 2, a result due to Novik.

121 Remark

Our next main goal is the following theorem. Recall that P is simplicial iff P^* is simple, and P has n vertices iff P^* has n facets. In particular, $C(d,n)^*$ is simple with n facets.

122 Theorem

Among all (simple) d-dimensional polytopes with n facets, $C(d,n)^*$ simultaneously maximizes all face numbers.

123 Remark

We will get rid of the adjective "simple" in several lectures. An important tool we'll use is the concept of h-numbers. Roughly, take a simple polytope P in $\mathbb{R}^d$ and choose a linear functional $\rho: \mathbb{R}^d \to \mathbb{R}$ "generically" in the sense that ρ is injective on $\mathrm{vert}(P)$. An obvious such functional on the cube is the "height" in 3D. In the usual orientation, this will not be a generic linear functional. However, we may rotate the cube or "wiggle it slightly" to get a generic height function. Now we can orient edges using the values of the height function, from smaller heights to larger. For instance, if we've generically rotated our cube to become a "diamond" standing on its point, the edges are all directed upwards. Since P is simple, each vertex has degree d, so the in-degree of each vertex is in $[0,d]$. For our cube, these numbers are $0,1,1,2,1,2,2,3$, say. We define h-numbers by counting the number of vertices of a specified in-degree, yielding $h_0, h_1, \ldots, h_d = 1,3,3,1$ for this example. At the moment it is not clear this vector is independent of ρ. We will use several nice consequences of this fact next time. On Friday we will prove the theorem.

PROOF See next lecture; we will prove the theorem on Friday.

April 13th, 2016: h-vectors of simple polytopes

---

124 Remark

Today we'll discuss h-vectors of simple polytopes. Homework 2 is due this Friday; homework 3 is due next Friday.

125 Definition

Let P be a simple d-dimensional polytope in $\mathbb{R}^d$. (Recall that equivalently each vertex is adjacent to precisely d edges, i.e. $G(P)$ is d-regular.) Pick a $\boxed{\text{generic}}$ linear functional $\rho: \mathbb{R}^d \to \mathbb{R}$, meaning ρ is injective on $\operatorname{vert}(P)$. (Indeed, it suffices to have $\rho(v_i) \neq \rho(v_j)$ whenever v_i and v_j are connected by an edge.) Orient the edges of P using ρ by declaring $v \to w$ whenever $\rho(v) < \rho(w)$. Note that $\deg_{G(P)}(v) = d$, so $0 \leq \operatorname{indeg}^{\rho}_{G(P)}(v) \leq d$.

In this situation, we define

$$\boxed{h_i^\rho(P)} := \#\text{ of vertices of in-degree } i.$$

For instance, $h_0^\rho(P) + h_1^\rho(P) + \cdots + h_d^\rho(P) = f_0(P)$ is the number of vertices of P.

126 Theorem

Let P be simple, ρ generic on P. Then

$$\sum_{k=0}^{d} f_k(P) x^k = \sum_{i=0}^{d} h_i^\rho(P)(x+1)^i.$$

127 Remark

Generating functions are often convenient or illuminating, though we may also write things in terms of coefficients. Using $x \mapsto x-1$ and the binomial theorem, the preceding equation is equivalent to

$$h_i^\rho(P) = \sum_{k=i}^{d} (-1)^{k-i} \binom{k}{i} f_k(P),$$

and similarly

$$f_k(P) = \sum_{i=k}^{d} \binom{i}{k} h_i^\rho(P).$$

In particular, establishing lower bounds on the f_i's can be done by establishing lower bounds on the h_i.

128 Example

We have

$$\begin{aligned} h_d &= f_d = 1 \\ h_0 &= f_0 - f_1 + f_2 - \cdots + (-1)^{d-1} f_{d-1} + (-1)^d =: \chi(P) + (-1)^d \\ h_{d-1} &= f_{d-1} - d \\ h_{d-2} &= f_{d-2} - (d-1) f_{d-1} + \binom{d}{2} \end{aligned}$$

where $\boxed{\chi(P)}$ is the $\boxed{\text{Euler characteristic}}$ of P.

129 Corollary

The h_i^ρ are independent of ρ. Hence we may define

$$\boxed{h_i(P)} := h_i^\rho(P)$$

for any simple P and any generic ρ on P.

130 Corollary

If P is simple and d-dimensional, then

(1) $h_i(P) \geq 0$

(2) $h_i(P) = h_{d-i}(P)$ *for all* i

(3) We have the Euler relation*:*

$$f_0 - f_1 + f_2 - \cdots + (-1)^{d+1} f_{d-1} = \chi(P) = \begin{cases} 0 & \text{if } d \text{ is even} \\ 2 & \text{if } d \text{ is odd.} \end{cases}$$

131 Remark

(1) is silly from our definition since the h_i's are defined to count something. However, from the alternating sum of f_k's expression, this is far from obvious.

PROOF Again (1) is immediate from our definition. For (2), $h_i(P) = h_i^\rho(P)$ is the number of vertices of in-degree i using ρ. We may use $-\rho$ as well, which interchanges in-degrees and out-degrees. Hence the number of vertices of in-degree i with respect to $-\rho$ is the number of vertices of in-degree $(d-i)$ with respect to ρ, so

$$h_i^\rho(P) = h_i^{-\rho}(P) = h_{d-i}^\rho(P).$$

For (3), we've noted $h_0(P) = \chi(P) + (-1)^d$, but $h_0(P) = 1$, forcing $\chi(P) = 0$ if d is even and 2 if d is odd.

132 Proposition

Let Q be a k-dimensional simple polytope in $\mathbb{R}^d$. For instance, Q may be a face of P above. Let ρ be a generic functional $\mathbb{R}^d \to \mathbb{R}$ with respect to Q and suppose v is a vertex of Q. Assume that all edges incident with v are oriented into v, so $\rho(v) > \rho(u)$ for all neighboring vertices u of v, i.e. $\rho(v)$ is a "local maximum." Then ρ has a "global maximum" on Q at v, i.e. $\rho(v) > \rho(w)$ for all vertexes $w \neq v$.

133 Remark

If a generic linear functional is used to orient the edges of P, then the resulting oriented graph has the following properties:

- acyclic;
- has a unique sink;
- the same holds for the induced orientation of $G(F)$ for all faces F of P.

The second point is equivalent to the proposition. The first point is clear. The third is essentially a restatement of the fact that faces of a simple polytope are simple.

PROOF The idea is that "locally a simple polytope behaves like a cube." Formally, the k edges incident to a vertex yield linearly independent vectors based at that vertex, so v together with its neighbors $u_1, \ldots, u_k$ are affinely independent. Hence we may use a bijective affine transformation T taking $\operatorname{Aff}(Q) \to \mathbb{R}^k$ sending $v \mapsto 0$, $u_i \mapsto e_i$, where $T(Q)$ is a simple polytope in $\mathbb{R}^k$. Hence without loss of generality we may suppose $v = 0$, $u_i = e_i$, and $\rho(0) = 0$ with $\rho(e_i) < 0$. Since Q is simple, the edges $\{0, e_1\}$, ..., $\{0, e_{i-1}\}$ must determine a facet. The equation of the supporting hyperplane of this facet is $x_d = 0$ since it contains the given edges. Since $e_d \in Q$, Q must have non-negative dth coordinates. Repeating this argument for each coordinate, Q lives in the first orthant of $\mathbb{R}^k$, i.e. all its points have non-negative coordinates. Hence for any $(x_1, \ldots, x_d) \in P$, we have $x_i \geq 0$, and

$$\rho(x_1, \ldots, x_d) = \sum_i x_i \rho(e_i) \leq 0 = \rho(0),$$

as required.

Proof We are now ready to prove the theorem, which will by counting pairs (F, v) where F is a k-face of P and $v \in \text{vert}(F)$ where $\rho(v)$ is the maximum value of ρ on F.

- Each face has a unique such v, so we are just counting the number of k-dimensional faces, giving $f_k(P)$.
- On the other hand, pick a vertex v and consider the number of faces F where (F, v) is such a pair. Suppose v has in-degree i with corresponding neighbors $u_1, \ldots, u_i$, so v has out-degree $d - i$. Since P is simple, the interval $[v, P]$ is a boolean lattice, so any set of s edges of P that contain v define a unique s-dimensional face. For our fixed v, we are then interested precisely in size k subsets of $\{u_1, \ldots, u_i\}$, since we can only use incoming edges, giving $\binom{i}{k}$ such pairs. Hence the number of pairs is

$$\sum_{v \in \text{vert}(P)} \binom{\text{indeg}\, v}{k} = \sum_{i=0}^{d} \binom{i}{k} h_i^\rho(P).$$

Hence we've shown

$$f_k(P) = \sum_{i=k}^{d} \binom{i}{k} h_i^\rho(P).$$

We must only show this is invertible. Multiplying this by x^k and summing from $k = 0$ to d and applying the binomial theorem gives

$$\begin{aligned}\sum_{k=0}^{d} f_k(P) x^k &= \sum_{k=0}^{d} \left(\sum_{i=k}^{d} \binom{i}{k} h_i^\rho(P) \right) x^k \\ &= \sum_{i=0}^{d} h_i^\rho(P) \sum_{k=0}^{i} \binom{i}{k} x^k \\ &= \sum_{i=0}^{d} h_i^\rho(P)(x+1)^i.\end{aligned}$$

134 Remark

Our goal for next time is to prove the upper bound theorem for simple polytopes (McMullen, 1970), namely if P is a simple d-polytope with $f_{d-1}(P) = n$, then $f_i(P) \leq f_i(C(d, n)^*)$ for all i. (We could state this for simplicial polytopes without the duals.) We will actually prove the inequality holds on the h-numbers,

$$h_i(P) \leq h_i(C(d, n)^*), \forall i.$$

Since the h_i's are symmetric, it suffices to prove this for the last "half" of the i, namely for $i \geq d/2$. We know the first half of the f-vector of the cyclic polytope, so by computing its dual, we know the second half of the f-vector of its dual. Then to compute h_i in the simple world we only need to know the f_k with $k \geq i$. Hence we have an explicit formula for $h_i(C(d, n)^*)$ when $i \geq d/2$, which is part of homework 3. We'll then actually show that $h_i(P)$ is bounded by the resulting rather nice binomial coefficient.

April 15th, 2016: Proof of the Upper Bound Theorem

---

135 Remark

Today we will prove the upper bound theorem; see Barvinok, Ch. VI.7. Next week we will discuss why the "simple" adjective may be removed.

136 Theorem

If P is a (simple) d-dimensional polytope with n facets, then

$$f_i(P) \le f_i(C(d,n)^*), \qquad \forall 0 \le i < d-1.$$

137 Remark

The theorem can be interpreted in terms of linear optimization. Given a polytope P defined as the intersection of $\{x : \langle x, a_j \rangle \le b_j\}$ for $j \in [n]$, we have an absolute bound on the number of i-dimensional faces. Recall $f_0(C(d,n)^*) = f_{d-1}(C(d,n))$ and $f_0 = h_0 + \cdots + h_d$, which by the symmetry condition from last time is roughly $2(h_0 + \cdots + h_{d/2})$. Now $h_{d-i}(C(d,n)^*) = \binom{n-k+i-1}{i} \sim n^i$ for all $i \le \lfloor d/2 \rfloor$, so in all $f_0 \sim O(n^{d/2})$ by the upper bound theorem. That is, the number of vertices needed to define the solution set of our system of linear inequalities is roughly at worst polynomial in the number of inequalities to half the dimension.

138 Remark

Recall that if P is a simple d-polytope, then

$$\sum_{i=0}^{d} h_i(P) x^i = \sum_{j=0}^{d} f_j(P)(x-1)^j$$

so that

$$f_j(P) = \sum_{i=j}^{d} \binom{i}{j} h_i(P).$$

The binomial coefficients are non-negative, so to prove the upper bound theorem, it suffices to prove the following stronger result, which does require P simple:

Indeed,

$$h_i(P) \le h_i(C(d,n)^*) \qquad \forall 0 \le i < d-1.$$

139 Remark

Last time we proved the Dehn-Sommerville relations, namely that for any simple d-polytope $h_i(P) = h_{d-i}(P)$ for all i. Hence, it suffices to prove the following version of the preceding claim:

$$h_{d-i}(P) \le h_{d-i}(C(d,n)^*), \qquad \forall i \le \lfloor d/2 \rfloor.$$

Proof We begin by computing the right-hand side of the last expression. Recall that for every $i \le \lfloor d/2 \rfloor$, every set of i vertices of $C(d,n)$ yields a face. It follows that

$$f_{i-1}(C(d,n)) = \binom{n}{i} \qquad \forall i \le \lfloor d/2 \rfloor$$

so that

$$f_{d-i}(C(d,n)^*) = \binom{n}{i}, \qquad \forall i \le \lfloor d/2 \rfloor.$$

Using the relations above, we have $h_d = f_d = 1$, $h_{d-1} = n-d$, ..., and more generally h_{d-i} depends only on face numbers at index $d-i$ and *higher*. It follows that

$$h_{d-i}(C(d,n)^*) = \binom{n-d+i-1}{i}, \qquad \forall i \le \lfloor d/2 \rfloor.$$

(The remaining details of this computation are on the homework.) Hence, it suffices to show

$$h_{d-i}(P) \le \binom{n-d+i-1}{i}, \qquad i \le \lfloor d/2 \rfloor.$$

We will prove this by induction on dimension. This requires comparing h-numbers for a polytope and for its facets, which is done using the following lemmas.

140 Notation

We set $h_{-1}(F) := 0$, and more generally any h-numbers at indexes below 0 or above $d-1$ are zero.

141 Lemma

If P is a simple d-polytope and F is a facet of P, then $h_i(P) \geq h_{i-1}(F)$.

Proof Recall our geometric definition of h-numbers. We take a polytope P and a linear functional $\rho: \mathbb{R}^d \to \mathbb{R}$ which is injective on vertices. We then direct edges in $G(P)$ from vertices with smaller ρ to larger ρ, and $h_i(P)$ is the number of vertices of in-degree i.

Since F is a facet, we have a hyperplane H such that $P \cap H = F$ and $P \subset H^-$. Say $H = \langle x : \langle x, a \rangle = b \rangle$. We can't quite use $\rho(x) = \langle x, a \rangle$ since it's not generic, but note that for this ρ, $\rho(v) > \rho(u)$ for all $v \in \mathrm{vert}(F)$, $u \in \mathrm{vert}(P) - F$. It is easy to see that we may perturb ρ slightly to some $\widetilde{\rho}$ so that $\widetilde{\rho}$ is injective on vertices of P and $\widetilde{\rho}(v) > \widetilde{\rho}(u)$ for all such v, u. Now pick a vertex v in F; since P is simple, v has d neighboring vertices (that is, connected by an edge), and F is a facet so it contains precisely $d-1$ of these. For the unique neighboring vertex u of v which is not in F, the edge $\{v, u\}$ is directed from u to v, so that

$$\mathrm{indeg}_F(v) = \mathrm{indeg}_P(v) - 1.$$

Hence

$$\begin{aligned} h_i(P) &= \#\{v \in \mathrm{vert}(P) : \mathrm{indeg}_P(v) = i\} \\ &= \#\{v \in \mathrm{vert}(F) : \mathrm{indeg}_F(v) = i-1\} \\ &\qquad + \#\{v \in \mathrm{vert}(P) - \mathrm{vert}(F) : \mathrm{indeg}_P(v) = i\} \\ &= h_{i-1}(F) + (\text{something non-negative}). \end{aligned}$$

The result follows.

142 Lemma

Let P be a simple d-polytope. Then

$$\sum_F h_i(F) = (i+1)h_{i+1}(P) + (d-i)h_i(P)$$

where the sum is over facets F of P.

Proof We will again use the geometric definition of h-numbers above. Choose a linear functional $\rho: \mathbb{R}^d \to \mathbb{R}$ generic on P. Consider counting pairs (v, F) where F is a facet of P, $v \in \mathrm{vert}(F)$, and $\mathrm{indeg}_F(v) = i$. We may of course use ρ on either a face F or P.

- Each F contributes exactly $h_i(F)$ such pairs, so the total number of pairs is the left-hand side of the equation above.
- On the other hand, pick a vertex v of P and count the number of pairs which contain it. We consider two cases.
  - Suppose $\mathrm{indeg}_P(v) = i$. Then we know v has degree d in $G(P)$ and i edges are directed into d, so $d-i$ edges are directed outwards. Also, there are precisely d facets containing d, which are defined by deciding which edge to drop. In that facet, the in-degree of v is again i if and only if we've dropped one of the $d-i$ outwardly directed edges. That is, if P is simple, each facet of P through v is defined uniquely by an edge $\{v, u_j\}$ it does not contain. It follows that the number of pairs that v contributes is $d-i$. This corresponds precisely to the second term on the right-hand side of the equation above.

– Suppose $\text{indeg}_P(v) = i + 1$. The same reasoning as in the previous case applies except in reverse, namely we must drop one of the incoming edges, of which there are $i+1$, resulting in the first term on the right-hand side of the equation above.

– The same argument shows that the remaining $\text{indeg}_P(v)$'s do not contribute.

Continuing the proof of the theorem, we now induct on i. First note that for all P,

$$h_d = 1 = \binom{n-d+0-1}{0}, \qquad h_{d-1} = n - d = \binom{n-d+1-1}{1}.$$

Hence we may assume as base cases that the required inequality holds for $i = 0$ (and $i = 1$). Inductively, suppose it holds for all $i \leq r$; we must show it holds for $i = r + 1$. By the first lemma, we have for all facets F

$$h_{d-r}(P) \geq h_{d-r-1}(F) \Rightarrow n h_{d-r}(P) \geq \sum_F h_{d-r-1}(F).$$

Now using the second lemma at $i = d - r - 1$,

$$\sum_F h_{d-r-1}(F) = (d-r)h_{d-r}(P) + (r+1)h_{d-r-1}(P).$$

Combining these two observations gives

$$(n-d+r)h_{d-r}(P) \geq (r+1)h_{d-r-1}(P).$$

Applying the inductive assumption, we now have

$$\begin{aligned} h_{d-r-1}(P) &\leq \frac{n-d+r}{r+1} h_{d-r}(P) \\ &\leq \frac{n-d+r}{r+1}\binom{n-d+r-1}{r} \\ &= \binom{n-d+r}{r+1}. \end{aligned}$$

This completes the proof.

143 Remark
For other classes of polytopes, we do not have h-vectors, but the f-vector form of the preceding theorem still works. The idea behind amplifying it is that a non-simplicial polytope has vertex sets of faces which are not affinely independent. Slightly perturbing these vertices gives a simplicial polytope, and one may check this operation can only increase face numbers.

144 Remark
The above argument really only used that $C(d,n)^*$ is simple and $\lfloor d/2 \rfloor$-neighborly. There are in fact polytopes with these properties that are not combinatorially isomorphic to $C(d,n)^*$.

April 18th, 2016: Reconstructing a Simple Polytope from its Graph

---

145 Remark
Today we'll discuss reconstructing a simple polytope from its graph. See §3.4 of Günter Ziegler's book for more. Conveniently, our machinery involving h-numbers from last week will be quite useful.

146 Theorem (Conjectured by Perles; proved by Blind-Blind (1987))
If P_1 and P_2 are simple polytopes, then $G(P_1) \cong G(P_2) \Leftrightarrow L(P_1) \cong L(P_2)$.

147 Remark
That is, isomorphisms of graphs of simple polytopes implies combinatorial isomorphism of the underlying polytopes. Recall that $G(P)$ is a graph whose vertices are the vertices of P and whose edges are the edges of P.

Proof The $\Leftarrow$ implication is essentially trivial, since the graph is encoded in the rank 2 part of the face lattice. For $\Rightarrow$, we give a very expensive algorithm. This is Kalai's proof (1988).

Let $P \subset \mathbb{R}^d$ be a simple polytope. Let $\rho: \mathbb{R}^d \to \mathbb{R}$ be a linear functional generic on P. Last week we gave $G(P)$ an acyclic orientation with the following properties:

- For all faces $F \neq \varnothing$ of P, $G(F)$ with the induced orientation is acyclic with a unique sink.
- $G(P)$ is connected. (Otherwise, any orientation of $G(P)$ will have at least 2 sinks.)
- $h_i(P)$ is the number of vertices of in-degree i, which does not depend on which ρ is chosen.

Note that given any total order $<$ on $\mathrm{vert}(P)$, we can form an acyclic orientation on $G(P)$ using this total order, namely given an edge between v to u, orient it from v to u if $v < u$. Call $<$ a "good order" if for all faces F of P, the induced graph $G(F)$ has a unique sink, including for $F = P$ but not for $F = \varnothing$. It is not trivial that good orders exist, though any generic linear functional does give rise to a good order.

148 Exercise
Come up with an example of a polytope and a total order on its vertices where the induced orientation on $G(P)$ has two sinks.

149 Definition
For any total order $<$ as above, define

$$h_i^<(P) := \#\text{of vertices of in-degree } i.$$

Further define

$$F^< := h_0^< + 2h_1^< + 2^2 h_2^< + \cdots + 2^d h_d^<$$

and

$$f := \#\text{ of non-empty faces of } P = f_0 + f_1 + \cdots + f_{d-1} + f_d.$$

Note that $h_i^<(P)$ and $F^<$ are clearly properties of $G(P)$ and $<$, whereas this is not clear for f.

150 Claim
For any total order $<$ on $\mathrm{vert}(P)$,

$$F^< \geq f$$

and moreover equality holds if and only if $<$ is a good order. Hence, we can identify good orders using just $G(P)$ and $<$.

Proof We use another double counting argument. Namely, we count pairs (F, v) where $F \in L(P) - \{\varnothing\}$, $v \in \mathrm{vert}(P)$, and v is a sink of $G(F)$, where $G(F)$ is oriented according to a total order $<$ on $\mathrm{vert}(P)$. Note that $G(F)$ has at least one sink, so the number of such pairs is at least f, and it is equal to f if and only if there is a unique sink for all F, i.e. if and only if $<$ is a good order.

On the other hand, we may count these pairs vertex by vertex as follows. For a vertex v, say v has in-degree k in $G(P)$. By simplicity, v has out-degree $d - k$. Furthermore, every subset of the k incoming edges gives rise to a distinct face where v is a sink (including the empty subset), and for v to be a sink we may not choose any of the $d-k$ outgoing edges. Hence v is a sink of $G(F)$ for precisely 2^k distinct F, so the number of pairs is precisely $\sum_{k=0}^{d} 2^k h_k^<(P) = F^>$. The result follows.

151 Claim

An induced subgraph H of $G(P)$ is a graph of a k-face of P if and only if H is connected, k-regular, and $\text{vert}(H)$ form an initial segment of some good order on $\text{vert}(P)$.

152 Remark

Recall the definition of induced subgraph. Given a graph G, an induced subgraph H of G is obtained by restricting the vertex set of G to some subset where there is an edge between u_1 and u_2 in H if and only if there is an edge between u_1 and u_2 in G. For instance, take a complete graph on 5 vertices. All induced subgraphs on 3 vertices must be complete. In particular, choosing 3 of the 5 vertices and only taking "outside edges of the pentagon," the result is not induced.

The "initial segment" condition means that there is a good (total) order on $\text{vert}(G)$ such that every element of $\text{vert}(H)$ is $<$ every element of $\text{vert}(G) - \text{vert}(H)$.

The two claims prove the theorem by giving an algorithm for finding k-faces of P. Namely, look at all good orders of $G(P)$, look at all initial segments, and check if the resulting induced graph H is connected and k-regular. This step is extremely computationally expensive.

Proof ($\Leftarrow$) Suppose F is a k-dimensional face, so F is a simple k-polytope. As remarked at the beginning of class, $G(F)$ is connected and k-regular. Now F is a face, so there exists a supporting hyperplane $\mathcal{H}$ such that $\mathcal{H} \cap P = F$ and $P \subset \mathcal{H}^+$. Now $\mathcal{H} = \{x : \langle a, x \rangle = b\}$; take $\rho : x \mapsto \langle a, x \rangle$ and perturb ρ slightly to make it generic while ensuring the values of ρ on $\text{vert}(F)$ remain smaller than the values of ρ on $\text{vert}(P) - \text{vert}(F)$. It follows that $\text{vert}(F)$ is an initial segment on the good total order given by ρ.

($\Rightarrow$) Suppose H is an induced subgraph of $G(P)$ which is connected, k-regular, and its vertices form an initial segment of $\text{vert}(P)$ under a good order $<$. Let v be the maximum of $\text{vert}(H)$ under $<$. Since H is k-regular, v has k neighbors in $\text{vert}(H)$, say $u_1, \ldots, u_k$. Since P is a simple polytope, the edges $\{v, u_i\}$ define a unique k-dimensional face F. By maximality, v is a sink of $G(F)$, and since $<$ is a good order, v is the unique sink of $G(F)$. That is, for all $w \in \text{vert}(F)$, $w \leq v$, and in particular $\text{vert}(F) \subset H$. Both H and $G(F)$ are connected and k-regular, where now $G(F) \subset H$, from which it follows that $G(F) = H$. This completes the claim and theorem.

153 Remark

The above (implicit) algorithm for constructing $L(P)$ from $G(P)$ is exponential in the size of the graph. A polynomial time algorithm was found by Eric Friedman (DCG, 2009); the paper is quite short and would be appropriate for anyone interested to present. The next homework gives a simple case of his algorithm.

General polytopes are *not* reconstructible from $G(P)$. For instance, $G(C(d, n)) = K_n$ already loses any information about d. On the other hand, any simplicial d-polytope P can be reconstructed from $\text{Skel}_{\lfloor d/2 \rfloor}(P)$ (proved by Perles), which can be generalized quite a bit to triangulations of spheres and manifolds with appropriate modifications:

154 Definition

The k-skeleton of P, $\text{Skel}_k(P)$ consists of all the faces of P of dimension $\leq k$.

If P is simplicial, then P^* is simple, so the vertices of P^* correspond to the facets of P, the edges of P^* correspond to ridges of P. The following notion is then natural:

155 Definition

The facet-ridge graph of a polytope P is the graph whose vertices are given by facets of P and where two vertices have an edge between them if their facets intersect in a ridge. The theorem above then says that any simplicial polytope may be reconstructed from its facet-ridge graph.

156 Open Problem

The following related problem is open. Can triangulations of spheres be reconstructed from their facet-ridge graph? The problem is open even in the "shellable" case, which we'll discuss more next week.

April 20th, 2016: Draft

---

157 Remark

Today we'll discuss centrally symmetric polytopes. Remember that homework 3 is due on Friday. There will be office hours on Friday from 10:30 to 11:30.

158 Remark

We begin with a few more remarks on cyclic polytopes. We had defined the cyclic polytope $C(d, n)$ as the convex hull $\operatorname{conv}\{\phi(t_1), \ldots, \phi(t_n)\}$ on the moment curve $\phi: \mathbb{R} \to \mathbb{R}^d$ given by $t \mapsto (t, t^2, \ldots, t^d)$. If $d = 2k$, then in fact the following trigonometric moment curve $\phi: \mathbb{R} \to \mathbb{R}^{2k}$ may be used:

$$\boxed{\phi(t)} := (\cos t, \sin t, \cos 2t, \sin 2t, \ldots, \cos kt, \sin kt) \in \mathbb{R}^{2k}.$$

We proved that $C(d, n)$ is $\lfloor d/2 \rfloor$-neighborly, meaning that any set of $k \leq \lfloor d/2 \rfloor$ vertices are the vertices of some face.

We also proved the upper bound theorem: if P is a d-dimensional simplicial polytope with n vertices, then $f_i(P) \leq f_i(C(d, n))$. We now remove the adjective "simplicial" from the previous sentence.

Proof (Sketch.) Let P be any d-dimensional polytope with n vertices $v_1, \ldots, v_n$. Taking sufficiently small neighborhoods of vertices, pick $\widetilde{v}_i$ "generic" in that neighborhood of v_i.Let $\widetilde{P} := \operatorname{conv}\{\widetilde{v}_1, \ldots, \widetilde{v}_n\}$. Then $\{\widetilde{v}_1, \ldots, \widetilde{v}_n\}$ are the vertices of $\widetilde{P}$, so (i) $f_0(\widetilde{P}) = n$. By genericity, (ii) $\widetilde{P}$ is simplicial. Furthermore, (iii) $f_i(P) \leq f_i(\widetilde{P})$. By the simplicial upper bound theorem, we have

$$f_i(P) \leq f_i(\widetilde{P}) \leq f_i(C(d, n)).$$

Claims (i) and (ii) are straightforward to verify, though (iii) requires a little care. (This argument essentially shows that simplicial polytopes are dense with respect to Hausdorff distance.)

We picture the above procedure by taking small neighborhoods around the vertices of a cube and perturbing them each slightly. Each square face gets broken up into two triangles.

159 Definition

A polytope $P \subset \mathbb{R}^d$ is centrally symmetric if $x \in P \Rightarrow -x \in P$. We will also assume P is full-dimensional, so $0 \in \operatorname{Int}(P)$.

160 Remark

If $v \in \operatorname{vert}(P)$, then $-v \in \operatorname{vert}(P)$, so $0 \in [v, -v]$ is an interior point, so $[v, -v]$ cannot be an edge.

161 Example

Consider two square pyramids glued at their square bases. This is centrally symmetric. It is dual to the three-dimensional cube (centered at 0). More generally, let $\boxed{C_d}$ denote the d-dimensional cube centered at 0. Then $\boxed{(C_d)^*} = \{e_1, \ldots, e_d, -e_1, \ldots, -e_d\}$ is the d-dimensional cross polytope.

In light of the remark, we modify our definition of k-neighborly in the case of centrally symmetric poltyopes.

162 Definition
A centrally symmetric d-dimensional polytope P with vertex set V is k-neighborly if for all $I \subset V$, $|I| \leq k$ such that I doesn't contain antipodal vertices, then I forms the vertex set of a face.

Indeed, $(C_d)^*$ is d-neighborly in the above sense. Cross polytopes are a natural analogue of simplices in the centrally symmetric setting.

163 Question

- How neighborly can a centrally symmetric polytope be? For instance, for cyclic polytopes, it goes up to $\lfloor d/2 \rfloor$.
- Is there a centrally symmetric analogue of $C(d,n)$?
- What is the largest number of i-dimensional faces a centrally symmetric d-polytope with n vertices can have?

Roughly, the answers for the second and third points are unknown, though we'll give some asymptotics on the first one.

164 Theorem ((a) due to Linial-Novik, 2006; (b) due to Barvinok-Novik, 2008)

(a) *A centrally symmetric d-dimensional polytope with more than 2^d vertices cannot be even 2-neighborly.*

(b) *If P is a centrally symmetric d dimensional polytope with n vertices, then $f_1(P) \leq \frac{n^2}{2}\left(1 - \frac{1}{2^d}\right)$.*

165 Remark
If P is centrally symmetric and 2-neighborly with n vertices, then $f_1(P) = \binom{n}{2} - n/2 = \frac{n^2}{2}(1 - o(1))$.

Proof We sketch some of the ideas to give the flavor of the arguments. For (a), assume P is d-dimensional, centrally-symmetric, and 2-neighborly. We must show that $|V| \leq 2^d$. The main trick is to look at a family of translates of P and not just P itself. More specifically, consider $\{P_v := P + v | v \in V\}$. Claim: if P is centrally symmetric and 2-neighborly, then all P_v have pairwise disjoint interiors. How does this help? It is an easy exercise to check that convexity of P implies $P_v \subset 2P$ where $2P$ denotes doubling each coordinate. Now we compute

$$\mathrm{vol}(\cup_{v \in V} P_v) = \sum_{v \in V} \mathrm{vol}(P_v) = \sum_{v \in V} \mathrm{vol}(P) = |V| \,\mathrm{vol}(P)$$

On the other hand, the left-hand side satisfies

$$2^d \,\mathrm{vol}(P) = \mathrm{vol}(2P) \geq \mathrm{vol}(\cup_{v \in V} P_v).$$

Since the volume is non-zero, it follows that $|V| \leq 2^d$. As for the claim, assume there exist $v \neq w$ with $v, w \in V$ such that $\mathrm{Int}(P_v) \cap \mathrm{Int}(P_w) \neq \varnothing$. Then we have $x, y \in \mathrm{Int}(P)$ such that $x + v = y + w$, so $(x - y)/2 = (w - v)/2$. Since P is centrally symmetric, if $y \in \mathrm{Int}(P)$, then $-y \in \mathrm{Int}(P)$, and it follows that $[x, -y] \subset \mathrm{Int}(P)$, so that $(x - y)/2 = (w - v)/2 \in \mathrm{Int}(P)$. Consider two cases. If $w = -v$, then $(w - v)/2 = -v$ is a vertex, which certainly is not an interior point. Hence $w \in V - \{v, -v\}$. Since P is 2-neighborly, we have an edge from w to $-v$, but then its midpoint $(w - v)/2$ is not an interior point, a contradiction.

166 Remark

During the argument for (a), we imagine P as the square in $\mathbb{R}^2$ and we compute translates of the square centered on each of the four vertices. This rough trick goes back at least to Minkowski and recurs from time to time.

The proof of the claim shows that if P is centrally symmetric and $v, w \in V$, then $\operatorname{Int} P_V \cap \operatorname{Int} P_w = \varnothing$ if and only if $\{-v, w\}$ are not the vertex set of an edge.

We turn to (b) and give fewer details. Say $|V| = n$ and again consider $P_v := P + v \subset 2P$. Hence, we have n polytopes of volume $\operatorname{vol}(P)$ sitting inside something of volume $2^d \operatorname{vol}(P)$. It follows that "on average" an interior point of $2P$ is covered by $\frac{n}{2^d}$ translates of P. Hence, for "an average" vertex v of P, the interior of P_v intersects with the interiors of at least $\frac{n}{2^d} - 1$ other polytopes P_w. Thus the average degree of a vertex of P (in the graph of P) is at most

$$(n-1) - \left(\frac{n}{2^d} - 1\right) = n\left(1 - \frac{1}{2^d}\right).$$

Multiplying by $n/2$ converts this upper bound on the average degree to the stated upper bound on the number of edges.

(One may make this rigorous by introducing indicator functions of sets and using Hölder's inequality.)

167 Theorem (Barvinok-Lee-Novik, 2013)

(a) *For all $m \geq 2$, there is a centrally symmetric polytope of dimension $2(m+1)$ which is 2-neighborly with $2(3^m - 1)$ vertices.*

(b) *For all $m, s \geq 2$, there is a centrally symmetric polytope of dimension $2(m+1)$ that has $N = 2s(3^m - 1)$ vertices and $\geq \binom{N}{2}\left(1 - \frac{1}{3^m}\right)$.*

168 Remark

(a) is roughly saying there exist centrally symmetric d-polytopes that are 2-neighborly and have roughly $\sqrt{3}^d$ vertices. (b) roughly says there exist centrally symmetric d-dimensional polytopes with $N \gg 0$ vertices and $\approx \frac{N^2}{2}\left(1 - \frac{1}{\sqrt{3}^d}\right)$ edges.

Proof We again give sketches. Define a variation on the trigonometric moment curve,

$$\Phi_m(t) := (\cos t, \sin t, \cos(3t), \sin(3t), \cos(3^2 t), \sin(3^2 t), \ldots, \cos(3^m t), \sin(3^m t)).$$

Note that $\Phi_m(t) = -\Phi_m(\pi + t)$. It is thus convenient to think of Φ_m as defined on the circle $S^1 \subset \mathbb{C}$. Claim: pick $2(3^m - 1)$ equally spaced points on S^1, namely

$$A_m := \left\{a_j := \frac{\pi(j-1)}{3^m - 1} | j = 1, 2, \ldots, 2(3^m - 1)\right\}.$$

Then $\operatorname{conv}\{\Phi_m(a_j) : j = 1, 2, \ldots, 2(3^m - 1)\}$ is a $2(m+1)$-dimensional polygon; it is centrally symmetric; it has $2(3^m - 1)$ vertices; and it is 2-neighborly.

169 Remark

The $m = 1$ case of the curve above was considered by Smilanski (1985). Cutting off the first two coordinates roughly gives $\Phi_{m-1}(3t)$. Using these two "projections" gives a way to analyze the structure of the whole polytope.

April 25th, 2016: Draft

---

170 Remark

Today we'll show that polytopes are shellable. Recall that if $\mathcal{C}$ is a pure k-dimensional polytopal complex, then a shelling of $\mathcal{C}$ is a linear ordering $F_1, F_2, \ldots, F_s$ of facets of $\mathcal{C}$ such that

(i) Either $\mathcal{C}$ is 0-dimensional or...

(ii) ...for all $1 < j \leq s$, $F_j \cap (F_1 \cup \cdots \cup F_{j-1}) \subset \partial F_j$ is non-empty and of the form $G_1 \cup \cdots \cup G_r$ for facets $G_1, \ldots, G_r$ of ∂F_j is a beginning of a shelling of $\mathcal{C}(\partial F_j)$.

Last time we saw some examples where the condition (ii) failed for different reasons, where for instance ∂F_j is not pure for a square. To be extremely precise, we would need to replace each F_i and G_i with $\mathcal{C}(F_i)$ and $\mathcal{C}(G_i)$. Indeed, we also need $\mathcal{C}(\partial F_1)$ in the "or" case above, but this is automatic from today's main theorem.

171 Definition

A polytope P is shellable if its boundary complex $\mathcal{C}(\partial P)$ is shellable.

172 Theorem (Brugesser-Mani; 1970)

Every polytope is shellable. In fact, for every d-dimensional polytope $P \subset \mathbb{R}^d$ and any point $x \in \mathbb{R}^d - P$ in "general position," there exists a shelling of P in which the facets "visible from x" come first.

173 Remark

We first make the terms in quotes precise.

174 Definition

A point $x \in \mathbb{R}^d - P$ is in general position with respect to P if x is not on the affine span of any facet (proper face) of P. Given such a point x and a facet F of P where H is the affine span of F, then F is visible from x if P and x lie on opposite sides of H.

175 Example

Let P be a regular hexagon centered on the origin with two vertxes on the horizontal axis. Letting x be a point on the vertical axis very slightly above the topmost (horizontal) facet (line segment), x is in general position and only that topmost facet is visibile from x.

176 Definition

A line $\ell \subset \mathbb{R}^d$ is in general position with respect to a polytope $P \subset \mathbb{R}^d$ if

(i) ℓ is not parallel to the affine span of any facet of P;

(ii) ℓ intersects the interior of P non-trivially;

(iii) if $H_1, \ldots, H_m$ are the affine spans of the facets of P, then the m intersection points $\{p_i\} \coloneqq \ell \cap H_i$ are all distinct;

(iv) the intersection point p_i is in general position with respect to F_i in the affine hull of F_i.

The idea of the proof is the following. We begin by picking a "generic" line passing through some facet F_1 visible from x. We travel along the line away from the polytope, imagining the polytope is a "planet" and we are in a "rocket ship." Eventually we will cross the affine span of another facet F_2, after which point F_2 will be visible from the rocket ship. After we "get to infinity" we wrap around and start coming back for a landing on the other side of the polytope. For this part, we order the facets using the order in which they become invisible from the rocket ship.

The formal proof will be by induction on the dimension. The ordering of facets obtained from this procedure is called a line shelling. A helpful observation is that reversing the order of a line shelling gives another line shelling. Hence assuming...

(i) ...our claims about line shellings hold in dimension $d-1$;

(ii) F is a $(d-1)$-dimensional polytope;

(iii) a point x is in the affine span of a facet F but not in F is in general position;

(iv) and $G_1, \ldots, G_t$ is a line shelling of F such that all facets of ∂F visible from x come first...

then $G_t, G_{t-1}, \ldots, G_1$ is also a line shelling and the facets of F not visible from x come first. The proof of this claim is just to use the same line but going in the opposite direction.

177 Remark

If P is a polytope and $F_1, \ldots, F_s$ is any shelling order (not necessarily a line shelling) of ∂F, then $F_s, F_{s-1}, \ldots, F_1$ is also a shelling. This property fails for more general objects as we saw last time. Indeed, taking five unit squares oriented along the axes centered on points $(0,0), (0,1), (0,2), (-1,2), (1,2)$, this order gives a shelling, but its reversal does not.

If P is a d-polytope in $\mathbb{R}^d$ with $0 \in \text{Int}(P)$. Suppose we have written P as a minimal intersection of hyperplanes

$$P = \cap_{i=1}^m \{x : \langle x, a_i \rangle \leq 1\},$$

(i.e. that $P = Q^*$ where $Q = \text{conv}\{a_1, \ldots, a_m\}$). Suppose $x_0 \in \mathbb{R}^d - P$ is in general position, so x_0 violates at least one inequality, say $\langle x_0, a_s \rangle > 1$. Hence the facet $F_s := P \cap \{x : \langle x, a_s \rangle = 1\}$ is visible from x_0. On the other hand, if all facets of P are visible from x_0, then $\langle x_0, a_i \rangle > 1$ for all i. But then for any $y \in \text{conv}\{a_1, \ldots, a_m\}$, $\langle x_0, y \rangle > 1$. However, $P = Q^*$, so $0 \in Q$, a contradiction since $\langle x_0, 0 \rangle = 0$. In summary, for any point in general position with respect to a polytope $P \subset \mathbb{R}^d$ with $0 \in \text{Int}(P$, there exists a facet visible from x_0 and there exists a facet invisible from x_0.

Proof We show that line shellings are shellings by induction on d.

- In the $d = 1$ case, there is very little to check, since the boundary complex is zero-dimensional.

- In the inductive step, assume that for every $(d-1)$-dimensional polytope Q and any point y in the affine hull of Q but not Q itself in general position with respect to Q, and any line $\ell' \ni y$ in general position with respect to Q, that the following holds. Note that $\ell \cap Q = [a, b]$ for some a, b; order the facets of Q by traveling from a to ∞, then from ∞ to b in the obvious way. Assume this ordering, i.e. the line shelling of Q with respect to ℓ, is in fact a shelling of Q.

  Now we have a d-polytope P, $x \in \mathbb{R}^d - P$ in general position, and a line ℓ through x in general position with respect to P. Intersecting ℓ with P and removing interior points of P yields two rays starting at some points p_1 and p_2. Say the ray starting at p_1 contains x. As before, order the facets of P by traveling along ℓ starting at p_1, going through x, and coming back around to hit p_m. Call the intersection points of the affine hulls of facets with this line in this order $p_1, p_2, \ldots, p_m$. We must show this yields a shelling of P. By genericity, it is easy to check that this is indeed a linear order on facets, say $F_1 < F_2 < \cdots < F_m$. For $j \geq 2$, we have

  $$F_j \cap (F_1 \cup \cdots \cup F_{j-1}) = \begin{cases} \text{the union of facets of } F_j \text{ visible from } p_j \\ \text{the union of facets of } F_j \text{ invisible from } p_j. \end{cases}$$

  By our earlier observation, both of these unions are non-empty. We must now argue that each of these are beginnings of shellings of F_j. But this is precisely the inductive assumption together with the observations that they are both non-empty and that the reversal of a line shelling is again a line shelling. Note that when we reach x, the facets visible from x will be precisely the facets which have been included in the shelling.

178 Corollary

Suppose P is a polytope.

(1) If F, G are any two facets of P, then there exists a shelling of ∂P starting with F and ending with G.

(2) Let v be any vertex of P. Then there exists a shelling of P such that the facets containing v come first.

PROOF For (i), pick a generic point in F and a generic point in G, and use the line between them to create a line shelling.

For (ii), if v lies in $F_1, \ldots, F_r$ and not in $F_{r+1}, \ldots, F_m$, take corresponding hyperplanes $H_1, \ldots, H_m$ where $P \subset H_i^-$. Now pick

$$w \in (H_1^+ - H_1) \cap \cdots \cap (H_r^+ - H_r) \cap (H_{r+1}^- - H_{r+1}) \cap \cdots \cap (H_m^- - H_m).$$

Take a generic line through w. The remaining details are left as an exercise.

April 27th, 2016: Draft

---

179 Remark

Today we'll say more on shellability and introduce simplicial complexes. Recall the theorem from last time, namely that every polytope is shellable. We proved a more precise result, namely if $F_1, \ldots, F_s$ is a line shelling of P, then $F_s, F_{s-1}, \ldots, F_1$ is also a line shelling of P (so both of these are shellings of ∂P). In particular,

$$F_j \cap (F_{j+1} \cup F_{j+2} \cup \cdots \cup F_s) \begin{cases} \text{is non-empty} & \text{if } j < s \\ \text{is empty} & \text{if } j = s. \end{cases}$$

180 Theorem (Euler-Poincaré Formula)

If $P \subset \mathbb{R}^d$ is a d-dimensional polytope, then

$$f_0(P) - f_1(P) + \cdots + (-1)^{d-1} f_{d-1}(P) = 1 + (-1)^{d-1}.$$

181 Remark

We had earlier proven this for simple polytopes using h-numbers. The left-hand side can be thought of as the definition of the Euler characteristic of the complex of the boundary $\mathcal{C}(\partial P)$.

182 Definition

If D is a polytopal complex of dimension d, then the reduced Euler characteristic of D is

$$\boxed{\widetilde{\chi}(D)} := -f_{-1}(D) + f_0(D) - \cdots + (-1)^d f_d(D).$$

Note that if $\mathcal{D}_1$ and $\mathcal{D}_2$ are polytopal complexes and if (miraculously) $\mathcal{D}_1 \cup \mathcal{D}_2$ is also polytopal, then

$$\widetilde{\chi}(D_1 \cup D_2) = \widetilde{\chi}(D_1) + \widetilde{\chi}(D_2) - \widetilde{\chi}(D_1 \cap D_2),$$

which is essentially a simple version of inclusion-exclusion.

PROOF We induct on d. We will prove the following: if $F_1, F_2, \ldots, F_s$ with $s := f_{d-1}(P)$ is a line shelling of P, then

$$\widetilde{\chi}(\mathcal{C}(F_1) \cup \cdots \cup \mathcal{C}(F_j)) = \begin{cases} 0 & \text{if } 1 \leq j < s \\ (-1)^{d-1} & \text{if } j = s. \end{cases}$$

The $j = s$ case translates directly to the required formula. In the base case $d = 1$, we have two facets F_1, F_2, both of which are points. The reduced Euler characteristic of either of these facets is $-1 + 1 = 0$. The Euler characteristic of the complex consisting of both of these facets is $-1 + 2 = 1 = (-1)^{1-1}$, as required.

We now give the inductive step, where we go from $d - 1$ to d. We further induct on j. At $j = 1$, we have

$$\begin{aligned}\widetilde{\chi}(\mathcal{C}(F_1)) &= \widetilde{\chi}(\mathcal{C}(\partial F_1)) + (-1)^{d-1} \cdot 1 \\ &= (-1)^{(d-1)-1} + (-1)^{d-1} = 0\end{aligned}$$

(where $1 = f_{d-1}(\mathcal{C}(F_1))$, and the second equality uses the inductive assumption). We now consider going from $j - 1$ to j. We have

$$\widetilde{\chi}(\mathcal{C}(F_1) \cup \cdots \cup \mathcal{C}(F_{j-1}) \cup \mathcal{C}(F_j)) = \widetilde{\chi}(\mathcal{C}(F_1) \cup \cdots \cup \mathcal{C}(F_{j-1})) + \widetilde{\chi}(\mathcal{C}(F_j)) - \widetilde{\chi}(\mathcal{C}(F_j) \cap (\mathcal{C}(F_1) \cup \cdots \cup \mathcal{C}(F_{j-1}))).$$

The first term on the right is zero by induction; the second term is zero by the $j = 1$ computation; and the third term is subtracting the reduced Euler characteristic of the beginning of a line shelling of F_j. By the inductive hypothesis, this become

$$= \begin{cases} 0 & \text{if } j < s \\ -\widetilde{\chi}(\mathcal{C}(\partial F_j)) & \text{if } j = s. \end{cases}$$

The $j = s$ case hence gives $-(-1)^{(d-1)-1} = (-1)^{d-1}$, completing the result.

183 Remark

We now turn to (shellable) simplicial complexes.

184 Definition

Recall that a geometric simplicial complex $\mathcal{C}$ is a finite collection of simplicies such that

(1) If $\Sigma \in \mathcal{C}$ and σ is a face of Σ, then $\sigma \in \mathcal{C}$.

(2) If $\Sigma_1, \Sigma_2 \in \mathcal{C}$, then $\Sigma_1 \cap \Sigma_2$ is a face of both Σ_1 and Σ_2.

Recall further that each such k-dimensional Σ is some $\text{conv}\{p_1, \ldots, p_{k+1}\}$ for affinely independent points $p_1, \ldots, p_{k+1}$, and moreover the facets of Σ are precisely of the form

$$\text{conv}\{p_{i_1}, \ldots, p_{i_s} : \{i_1, \ldots, i_s\} \subset [k+1]\}.$$

This observation essentially says that, at least as far as the face lattice is concerned, we can "forget the geometry." This is formalized by the next definition.

185 Definition

An abstract simplicial complex Δ on a (finite) vertex set V is a collection of subsets of V such that

- if $F \in \Delta$ and $G \subset F$, then $G \in \Delta$;
- $\{v\} \in \Delta$ for all $v \in V$

(The second axiom is sometimes omitted.)

Given a geometric simplicial complex $\mathcal{C}$, we get an abstract simplicial complex by sending each $\Sigma \in \mathcal{C}$ to its set of vertices. In particular, note that requirement (2) above is automatically satisfied since $F_1 \cap F_2 \subset F_1$ with $F_1 \in \Delta$ implies $F_1 \cap F_2 \in \Delta$. We will shortly say how to "reverse" this operation.

186 Definition

If Δ is an abstract simplicial complex, we refer to the elements of $F \in \Delta$ as faces, we say the dimension of F is $\dim F := |F| - 1$, and $\dim \Delta := \max\{\dim F : F \in \Delta\}$. In this way we can still define face numbers for Δ,

$$f_{i-1}(\Delta) := \#\{F \in \Delta : \dim F = i - 1\} = \#\{F \in \Delta : |F| = i\}$$

with corresponding f-vector

$$f(\Delta) := (f_{-1}(\Delta), \ldots, f_{\dim \Delta}(\Delta)).$$

187 Example

Suppose Δ is an abstract simplicial complex of dimension $\dim \Delta = d-1$. If $d = 0$, then $\dim \Delta = -1$ forces either $\Delta = \varnothing$ or $\Delta = \{\varnothing\}$, with $V = \varnothing$. It ends up being useful for induction to have these two complexes. In the first case, the f-vector is just (0), whereas in the second case, the f-vector is (1). In particular, note that $f_{-1}(\Delta)$ is not simply defined to be 1 as we had done for polytopes.

For a more complicated example, consider the convex hull of three vertices in $\mathbb{R}^1$ labeled $3, 1, 2$ and appearing in this order. The polytope thus produced "forgets" about the vertex 1, but we can make a simplicial complex which remembers 1 using $V = \{1, 2, 3\}$ with

$$\Delta = \{\{1,2\}, \{1,3\}, \{1\}, \{2\}, \{3\}, \varnothing\}.$$

This is the abstract simplicial complex associated to the geometric simplicial complex consisting of the convex hull of vertices 3 and 1, together with the convex hull of vertices 1 and 2.

188 Definition

A geometric realization of an abstract simplicial complex Δ is defined as follows. If Δ has vertex set $V = \{v_1, \ldots, v_n\}$, the geometric realization of Δ is a geometric simplicial complex in $\mathbb{R}^n$ defined in the naive way, namely

$$F = \{v_{i_1}, \ldots, v_{i_k}\} \subset V \mapsto \operatorname{conv}\{v_{i_1}, \ldots, v_{i_k}\} =: \boxed{||F||}$$

$$\Delta \mapsto \cup_{F \in \Delta} ||F|| =: \boxed{||\Delta||}.$$

189 Example

Using Δ on three vertices from the preceding example gives the pair of line segments in $\mathbb{R}^3$ given by $[e_3, e_1]$ and $[e_1, e_2]$, which is clearly "the same" as the complex we started with.

190 Remark

While one must technically check $||\Delta||$ is indeed a geometric simplicial complex, this is easy. Using the geometric realization, we can think of an abstract simplicial complex Δ as a topological space using the topology induced from $\mathbb{R}^n$ on $||\Delta||$. This motivates the following definition.

191 Definition

Let Δ be an abstract simplicial complex. If $||\Delta||$ is homeomorphic to the $d-1$-dimensional sphere $\mathbb{S}^{d-1}$, then we call Δ a simplicial sphere.

192 Definition

Let Δ be an abstract simplicial complex. Δ is pure of dimension $d-1$ if all of its maximal faces have the same dimension $d-1$.

193 Example

Take the boundary of a simplex with some line segments "hanging off." The geometric realization is homotopy equivalent to a sphere, but it is not homeomorphic to a sphere, so the complex is not a simplicial sphere. There are several variations on this definition which we may or may not get to, including homology spheres.

194 Remark

If Δ is d-dimensional, then it can always be embedded in $\mathbb{R}^{2d+1}$ as follows. Let $v_1, \ldots, v_n$ be the vertices of Δ and pick n "generic" points of $\mathbb{R}^{2d+1}$ and repeat the previous construction. The point is that if the points have been chosen sufficiently generically, d-dimensional convex hulls will never "accidentally intersect."

Note that some d-dimensional simplicial complexes are not embeddable in $\mathbb{R}^{2d}$. For instance, if $d = 1$, we find that the complex graph K_5 is not embeddable in $\mathbb{R}^2$, and the bipartite graph (double check) $K_{3,3}$ is not embeddable in $\mathbb{R}^2$.

195 Example

Set $V := [2d+3]$. Set

$$\Delta := \{F \subset V : |F| \leq d+1\}.$$

Then Δ is d-dimensional and it has been shown that Δ is not embeddable in $\mathbb{R}^{2d}$. If $d = 1$, this is just the K_5 example above.

For an analogue of $K_{3,3}$, take $d+1$ triples of vertices, and form faces by picking at most one vertex from each triple. This is the "join" of three vertices $d+1$ times, which we will talk about later. The result is not embeddable in $\mathbb{R}^{2d}$.

These examples begin to motivate the following wide open conjecture:

196 Conjecture

If Δ is a d-dimensional complex embeddable in $\mathbb{R}^{2d}$, then

$$f_d \leq (d+2) f_{d-1}.$$

The conjecture is known for $d = 1$ when it says $f_1 \leq 3f_0$, but this is easy.

197 Example

If P is a simplicial polytope, then $\mathcal{C}(\partial P)$ is a simplicial sphere. Indeed, any 2-dimensional simplicial sphere is the boundary complex of some 3-dimensional polytope, though this fails in dimension ≥ 3.

April 29th, 2016: Draft

---

198 Remark

Recall that we are discussing simplicial complexes, where a simplicial complex Δ on a vertex set V is a collection of subsets of V such that

- If G is a face of $F \in \Delta$, then $G \in \Delta$
- $\{v\} \in \Delta$ for all $v \in V$.

Further recall that we defined the geometric realization $||\Delta||$ of Δ. We may define the f-vector for Δ using the geometric realization or using Δ directly,

$$f(\Delta) = (f_{-1}(\Delta), f_0(\Delta), \ldots, f_{\dim(\Delta)}(|Delta)).$$

Warning: the simplicial complex Δ of the boundary of a d-polytope has $\dim(\Delta) = d - 1$.

(Homework 5 will be due next Friday; homework 6 will be due *two* Friday's from then.)

199 Definition

If Δ is a $(d-1)$-dimensional simplicial complex, then the $\boxed{h\text{-vector}}$ of Δ

$$\boxed{h(\Delta)} := (h_0, h_1, \ldots, h_d)$$

defined by

$$\sum_{i=0}^{d} h_i(\Delta) \cdot x^{d-i} = \sum_{j=0}^{d} f_{j-1}(\Delta) \cdot (x-1)^{d-i}.$$

Warning: While the h-numbers are always defined by the above formula, sometimes the h-polynomial is $\sum_{i=0}^{d} h_i(\Delta) \cdot x^i$ instead of the left-hand side of the preceding expression. Of course, when the Dehn-Sommerville relations apply, both definitions agree, but in general they can be different.

200 Remark

Note that the above condition is equivalent to

$$h_i(\Delta) = \sum_{j=0}^{i} (-1)^{i-j} \binom{d-j}{d-i} f_{j-1}(\Delta).$$

In particular, $h_0 = f_{-1} = 1$, and

$$\begin{aligned} h_d &= \sum_{j-0}^{d} (-1)^{d-j} f_{j-1}(\Delta) \\ &= f_{d-1} - f_{d-2} + \cdots \pm f_{-1} \\ &= (-1)^{d-1} \widetilde{\chi}(\Delta). \end{aligned}$$

In particular, note that h_d (and other h-numbers) might be negative! For instance, consider the “bowtie” composed of two triangles touching at a single point. The f-vector is $(1,5,6,2)$, and the h-vector is $(1,2,-1,0)$.

201 Definition

The $\boxed{\text{face poset}}$ $\boxed{\mathcal{F}(\Delta)}$ consists of the faces of Δ ordered by inclusion. This is not a lattice in general essentially because faces are not generally closed under unions in any sense. One sometimes formally adds a maximum and writes

$$\boxed{\widehat{\mathcal{F}}(\Delta)} := \mathcal{F}(\Delta) \coprod \{\widehat{1}\}.$$

202 Notation

If F is a finite set, write $\boxed{\overline{F}}$ as the collection of all subsets of F ordered by inclusion, which is an honest boolean lattice. Note that if F is the vertex set of a simplex, then $\overline{F}$ agrees with $\mathcal{C}(\operatorname{conv} F)$. Also write $\boxed{\partial\overline{F}} := \overline{F} - \{F\}$.

203 Remark

In this notation, a simplicial complex Δ is shellable if there exists an ordering $F_1, F_2, \ldots, F_m$ of facets of Δ such that

$$\overline{F}_j \cap (\overline{F}_1 \cup \cdots \cup \overline{F}_{j-1})$$

a non-empty, pure $(d-2)$-dimensional subcomplex of $\partial\overline{F}_j$ for all $j \geq 2$. (The initial segment condition is trivial for simplices.)

204 Example

Consider the six points

$$P_1 = (0,0), P_2 = (1,0), P_3 = (2,0), P_4 = (2,1), P_5 = (3,1), P_6 = (3,0)$$

with edges

$$F_1 = \{P_1, P_2\}, F_2 = \{P_2, P_3\}, F_3 = \{P_3, P_4\},$$
$$F_4 = \{P_3, P_5\}, F_5 = \{P_4, P_5\}, F_6 = \{P_5, P_6\},$$
$$F_7 = \{P_5, P_6\}$$

The shelling condition at $j = 5$ works out because $\overline{F}_5 = \{\{P_4, P_5\}, \{P_4\}, \{P_5\}, \varnothing\}$ intersects the preceding elements in $\{\{P_4\}, \{P_5\}, \varnothing\}$, which is pure of the correct dimension.

205 Definition

If $F_1, F_2, \ldots, F_m$ is a shelling of a simplicial complex Δ, write

$$\overline{F}_j \cap (\overline{F}_1 \cup \cdots \cup \overline{F}_{j-1}) = \overline{G}_1 \cup \cdots \cup \overline{G}_r.$$

Let $\{v_j^i\} \coloneqq F_j - G_i$. In this notation, let $\boxed{R_j} \coloneqq \{v_j^1, \ldots, v_j^r\}$.

206 Example

Using the preceding example, we find $R_1 = \varnothing$, $R_2 = \{3\}$, $R_3 = \{4\}$, $R_4 = \{5\}$, $R_5 = \{4,5\}$. Note that G_i by assumption has precisely one fewer vertices than F_j, so the v_j^i are well-defined.

207 Remark

Given a shelling as above, we are building up Δ successively as $\overline{F}_1 \cup \cdots \cup \overline{F}_{j-1}$. What is the collection of new faces added when adding $\overline{F}_j$ to this collection? Since $F_j - G_i = \{v_j^i\}$, we have $v_j^i \notin G_i$, so $R_j = \{v_j^1, \ldots, v_j^r\} \notin G_i$, and $R_j \notin \overline{F}_1 \cup \cdots \cup \overline{F}_{j-1}$. Hence R_j is a new face, so every element of $[R_j, F_j]$ is new as well. Might there be additional new faces outside of this interval?

Suppose $R_j \notin G \subset F_j$, so we have some i such that $v_j^i \notin G$, so $G \subset G_i$ since $F_j - G_i = \{v_j^i\}$. In particular, $[R_j, F_j]$ is precisely the set of new faces.

To summarize: if $F_1, \ldots, F_m$ is a shelling of a simplicial complex Δ, then R_j is the unique minimal face added at shelling step j. (The converse happens to hold as well.) That is, we have a partition of the face poset into disjoint boolean intervals

$$\mathcal{F}(\Delta) = \coprod_j [R_j, F_j].$$

208 Example

Consider $\partial\{a, b, c, d\}$. Take

$$F_1 = \{a, b, c\}, F_2 = \{a, b, d\}, F_3 = \{a, c, d\}, F_4 = \{b, c, d\}.$$

Compute the R_j. We have $R_1 = \varnothing$, $R_2 = \{d\}$, $R_3 = \{c, d\}$, $R_4 = \{b, c, d\}$. Geometrically, this action occurs in the boundary complex of a tetrahedron; it is a good, easy exercise to draw it out yourself.

At the first step, we added the interval $[\varnothing, abc] = \{\varnothing, a, b, c, ab, ac, bc, abc\}$. At the second step, we added the interval $[d, abd] = \{d, bd, ad, abd\}$. At the third step, we added $[cd, acd] = \{cd, acd\}$, and at the final step we added $[bcd, bcd] = \{bcd\}$.

209 Remark

The preceding decomposition $\mathcal{F}(\Delta) = \coprod_j [R_j, F_j]$ allows us to easily compute f-numbers. In particular, for all $\ell = 0, 1, \ldots, d$,

$$\begin{aligned} f_{\ell-1}(\Delta) &= \sum_{j=1}^m f_{\ell_1}([R_j, F_j]) \\ &= \sum_{j=1}^m \binom{d - |R_j|}{\ell - |R_j|} \\ &= \sum_{i=0}^d \binom{d-i}{\ell-i} \#\{j : |R_j| = i\} \end{aligned}$$

Comparing this last expression to the defining relation for h-numbers shows that $h_i(\Delta) = \#\{j : |R_j| = i\}$. We now summarize this discussion.

210 Definition
Given a shelling order for a simplicial complex Δ in the notation above, if $|R_j| = i$ then we say that the jth shelling step is of type i.

211 Corollary
If Δ is a shellable simplicial complex, then $h_i(\Delta)$ is the number of steps of type i.

212 Corollary
If Δ is a shellable simplicial complex, then $h_i(\Delta) \geq 0$ for all i.

213 Proposition
If Δ is a shellable $(d-1)$-dimensional simplicial sphere, then $h_i(\Delta) = h_{d-i}(\Delta)$ for all i.

Proof We deduce this as an immediate consequence of a nice property about restriction faces of reversed shellings in this context. Precisely,

214 Lemma
If $F_1, \ldots, F_m$ is a shelling of a simplicial sphere, then $F_m, \ldots, F_2, F_1$ is also a shelling, and $R(F_j) = F_j - R'(F_j)$, where $R(F_j)$ denotes the restriction face with respect to the original shelling and $R'(F_j)$ uses the reversed shelling.

Proof If Δ is a simplicial $(d-1)$-dimensional sphere, then $\widetilde{\chi}(\Delta) = (-1)^{d-1}$, so $h_d = (-1)^{d-1}\widetilde{\chi}(\Delta) = 1$. Hence, the total number of steps of type d is h_d, which is 1. Hence $\overline{F}_j \cap (\overline{F}_1 \cup \cdots \cup \overline{F}_{j-1}) = \partial \overline{F}_j$ for precisely one j. Since the intersection $j = m$ clearly has this property, it follows that for all $j < m$, this intersection is a proper subset of $\partial\overline{F_m}$. Consequently, $\overline{F}_j \cap (\overline{F}_{j+1} \cup \cdots \overline{F}_m) \neq \varnothing$. Since Δ is a sphere, every $(d-2)$-dimensional face of Δ is in exactly two facets. It follows that the increasing intersection above is the rest of the facets of F_j, from which the result follows.

May 2nd, 2016: Draft

---

215 Remark
Today we will be characterizing the face numbers of simplicial complexes. Recall that for a simplicial complex Δ, if Δ is shellable, then $h_i(\Delta) \geq 0$ for all i. Moreover, if Δ is a shellable $(d-1)$-dimensional sphere, then $h_i(\Delta) = h_{d-i}(\Delta)$ for all i. In fact,

216 Theorem
If Δ is any $(d-1)$-dimensional simplicial sphere (not necessarily shellable), then

- *(Klee, 1964.)* $h_i = h_{d-i}$
- *(Stanley, 1975.)* $h_i \geq 0$

Klee's $h_i = h_{d-i}$ result are also referred to as the Dehn-Sommerville relations. These results suggest questions along the lines of the following: which integer vectors are the f-vectors of (simplicial, shellable, etc.) complexes?

217 Lemma
For any pair (a, i) of positive integers, then there exists a unique expression for a of the following form:

$$a = \binom{n_i}{i} + \binom{n_{i-1}}{i-1} + \binom{n_{i-2}}{i-2} + \cdots + \binom{n_j}{j}$$

where $n_i > n_{i-1} > \cdots > n_j \geq j \geq 1$.

218 Example

Let $a = 40$, $i = 3$. One may check that $\binom{7}{3} = 35$ is the largest binomial coefficient with 3 on bottom which is not greater than 40. Then $\binom{3}{2} = 3$ is the largest binomial coefficient with 2 on bottom which is not greater than 5. The next step gives $\binom{2}{1}$, so

$$40 = \binom{7}{3} + \binom{3}{2} + \binom{2}{1}.$$

Proof We prove existence; uniqueness is left as an exercise, though both proofs are similar. We induct on a and i. Our bases case are $a \le i$, so that

$$a = \binom{i}{i} + \binom{i-1}{i-1} + \cdots + \binom{i-(a-1)}{i-(a-1)},$$

or $i = 1$, so that $a = \binom{a}{1}$. Now assume the statement holds for all (a', i') with$a' < a$ and $i' \le i$. Consider (a, i) and find the maximum n such that $\binom{n}{i} \le a$. Take $n_i := n$. Note that $\binom{n}{i} \le a < \binom{n+1}{i}$, so $m := a - \binom{n}{i}$ satisfies

$$0 \le m < \binom{n+1}{i} - \binom{n}{i} = \binom{n}{i-1}.$$

If $m = 0$, then $a = \binom{n_i}{i}$ and we are done. Otherwise, we can apply our inductive assumption to $(m, i-1)$, so

$$m = \binom{n_{i-1}}{i-1} + \binom{n_{i-2}}{i-2} + \cdots + \binom{n_j}{j}$$

where $n_{i-1} > n_{i-2} > \cdots > n_j \ge j \ge 1$. Adding $\binom{n_i}{i}$ to this expression gives a in the appropriate form so long as $n_i > n_{i-1}$. Since $m < \binom{n_{i-1}}{i-1}$, we indeed have $n_i > n_{i-1}$ since $\binom{n_i}{i-1}$ is too big to be included in the expression for m.

219 Definition

If $(a, i) \in \mathbb{Z}_{>0}$ the expression

$$a = \binom{n_i}{i} + \binom{n_{i-1}}{i-1} + \cdots + \binom{n_j}{j}$$

with $n_i > \cdots > n_j \ge j \ge 1$ is called the ith canonical representation of a. We further define "bumping" operations

$$\boxed{a^{(i)}} := \binom{n_i}{i+1} + \binom{n_{i-1}}{i} + \cdots + \binom{n_j}{j+1}$$
$$\boxed{a^{\langle i \rangle}} := \binom{n_i+1}{i+1} + \binom{n_{i-1}+1}{i} + \cdots + \binom{n_j+1}{j+1}$$

where $\binom{k}{\ell} = 0$ if $k < \ell$. For convenience we also define $0^{(i)} := 0 =: 0^{\langle i \rangle}$.

220 Theorem (Kruskal-Katona, circa 1961)

A positive integer vector

$$F = (F_{-1}, F_0, F_1, \ldots, F_{d-1})$$

is the f-vector of some $(d-1)$-dimensional simplicial complex if and only if $F_{-1} = 1$ and for all $i \ge 1$, $F_i \le F_{i-1}^{(i)}$.

221 Remark

Schützenberger proved this earlier, but nobody noticed; Stanley calls it Schützenberger's theorem, but Kruskal-Katona is overwhelmingly common. We will discuss existence today. Next time we will give a few remarks on necessity.

Proof (sketch). Write $\mathbb{N} := \{1, 2, \ldots\}$ and $\binom{\mathbb{N}}{i} := \{S \subset \mathbb{N} : |S| = i\}$.

222 Definition

The reverse lexicographic order on $\binom{\mathbb{N}}{i}$ (rev-lex; less commonly, co-lex or anti-lex) is defined as follows. If $F, G \in \binom{\mathbb{N}}{i}$, compute the symmetric difference $F \Delta G := (F - G) \cup (G - F)$. Say $F <_{\text{rev-lex}} G$ if $\max F\Delta G \in G$.

223 Remark

We run into problems if we try to use lexicographic order immediately. For instance, if $i = 2$, we begin with

$$\{1,2\} < \{1,3\} < \{1,4\} < \cdots,$$

and we never get to $\{2,3\}$. In lexicographic order, we compare left-to-right:

$$F := \{3,5,6,9,10\} >_{\text{lex}} \{3,4,7,9,10\} =: G.$$

In rev-lex order, we compare right-to-left, where here $6 < 7$ is used as the tie-breaker to give $F <_{\text{rev-lex}} G$. Here the symmetric difference is $F\Delta G = \{4,5,6,7\}$, the maximum of which is $7 \in G$.

224 Example

We have

$$\{1,2,3\} < \{1,2,4\} < \{1,3,4\} < \{2,3,4\} < \{1,2,5\} < \{1,3,5\} < \{2,3,5\} < \cdots.$$

One may ask for a given $F \in \binom{\mathbb{N}}{i}$ what position F has in the rev-lex total order.

225 Lemma

Let $F = \{k_1 + 1, k_2 + 1, \ldots, k_i + 1\} \in \binom{\mathbb{N}}{i}$. Then the number of elements of $\binom{\mathbb{N}}{i}$ smaller than F in rev-lex order is precisely

$$\binom{k_i}{i} + \binom{k_{i-1}}{i-1} + \cdots + \binom{k_1}{1}$$

such sets.

Proof Suppose $G < F$ in rev-lex order. In the first case, suppose the ith element (the largest element) of G is smaller than $k_i + 1$, i.e. $G \subset [k_i]$. There are precisely $\binom{k_i}{i}$ such sets. In the second case, suppose the ith element of G is $k_i + 1$, i.e. $G = H \cup \{k_i + 1\}$ where $H \in \binom{k_{i-1}}{i-1}$. There are $\binom{k_{i-1}}{i-1}$ such sets. Continuing in this way gives the expression.

Given a simplicial complex Δ, we know that its $(i-1)$-dimensional faces form an initial segment of $\binom{\mathbb{N}}{i}$ with respect to rev-lex order.

226 Example

Say $\binom{14}{6} + \binom{7}{5} + \binom{5}{4}$ is the number of 5-dimensional faces. How many 6-dimensional faces can we have? Notice that all 6-element subsets of $[14]$ are accounted for by the first term. All 6-element subsets of the form $G \cup \{15\}$ where $G \in \binom{[7]}{5}$ are accounted for by the second term. All 6-element subsets of the form $H \cup \{8, 15\}$ where $H \in \binom{[5]}{4}$ are accounted for by the third term. Hence the 6-dimensional faces arise from 7-element subsets of $[14]$, etc. More on this next time.

227 Theorem (Stanley, 1975)

An integer vector $H = (H_0, H_1, \ldots, H_d)$ is the h-vector of a shellable $(d-1)$-dimensional simplicial complex if and only if $H_0 = 1$, $H_1 \geq 0$, and for all $i \geq 1$, $0 \leq H_{i+1} \leq H_i^{\langle i \rangle}$.

228 Remark

If Δ is shellable, then $h_1 = n - d = \binom{n-d}{1}$ with $n = f_0$. By the theorem, $h_2 \leq h_1^{\langle 1 \rangle} = \binom{n-d+1}{2}$, and similarly $h_3 \leq h_2^{\langle} 2\rangle = \binom{n-d+2}{3}$, etc. These expressions appear in the upper bound theorem. We also have $h_i = h_{d-i}$ for all i. Hence we have the following corollary:

229 Corollary (Stanley, 1975)

If Δ is a $(d-1)$-dimensional shellable sphere, then

$$h_i(\Delta) \le h_i(C(d, f_0(\Delta))).$$

("Shellable" can be replaced by "simplicial" above.)

230 Theorem (The g-theorem)

An integer vector $h = (h_0, h_1, \dots, h_d)$ is the h-vector of some d-dimensional simplicial polytope if and only if

- $h_i = h_{d-i}$
- $1 = h_0 \le h_1 \le h_2 \le \cdots \le h_{\lfloor d/2 \rfloor}$
- $g_{i+1} := h_{i+1} - h_i \le (h_i - h_{i-1}^{\langle i \rangle}) = g_i^{\langle i \rangle}$ *for all* $i \le \lfloor d/2 \rfloor - 1$.

231 Remark

The whole statement was conjectured by McMullen in 1970. Billera-Lee proved sufficiency in 1980. Stanley proved necessity in 1980. (The story is that Stanley was so motivated by Billera-Lee's result that he proved it quite quickly.) Stanley's proof used quite a bit of commutative algebra and algebraic geometry though is only two pages. McMullen wanted an elementary proof of such an elementary statement, which he gave over a series of papers totaling around 100 pages.

232 Conjecture (The g-conjecture)

The same result holds for $(d-1)$-dimensional simplicial spheres.

233 Remark

The Billera-Lee construction continues to give sufficiency. The statement that $h_1 \le h_2$ is known to hold even for simplicial spheres (and more generally), a result which goes by the name of the lower bound theorem. We will discuss this over the coming week. The next conjectured inequality, $h_2 \le h_3$, is still wide open for simplicial spheres.

May 4th, 2016: Draft

---

234 Remark

We were discussing the Kruskal-Katona theorem. Recall that if $f = (1, f_0, f_1, \dots, f_{d-1})$, this is the f-vector of some $(d-1)$-dimensional simplicial complex if and only if $0 < f_i \le f_{i-1}^{(i)}$ for all $1 \le i \le d-1$. Last time we began discussing the $\Leftarrow$ direction, i.e. constructing simplicial complexes if the inequalities are satisfied. We introduced rev-lex order on $\binom{\mathbb{N}}{i}$. If $F = \{n_1 + 1 < n_2 + 1 < \cdots < n_i + 1\} \in \binom{\mathbb{N}}{i}$, we showed that there are exactly

$$\binom{n_i}{i} + \binom{n_{i-1}}{i-1} + \cdots + \binom{n_1}{1}$$

rev-lex smaller elements of $\binom{\mathbb{N}}{i}$. Note that some of the n_i for i small may be zero; in our canonical form, we ignored them. Our reasoning was actually more precise and described the initial segment ending at F explicitly as follows.

If $a = \binom{n_i}{i} + \binom{n_{i-1}}{i-1} + \cdots + \binom{n_j}{j}$ where $n_i > n_{i-1} > \cdots > n_j \ge j \ge 1$, then the first a elements of $\binom{\mathbb{N}}{i}$ in rev-lex order are as follows:

$$\binom{[n_i]}{i} \coprod \{H \cup \{n_i + 1\} : H \in \binom{[n_{i-1}]}{i-1}\}$$
$$\coprod \{H \cup \{n_{i-1} + 1, n_i + 1\} : H \in \binom{[n_{i-2}]}{i-2}\}$$
$$\cdots \coprod \{H \cup \{n_{j+1} + 1, \dots, n_1 + 1\} : H \in \binom{[n_j]}{j}\}.$$

235 Example

We began the following example last time; there was a minor mistake (which has been corrected above). Let $a = \binom{14}{6} + \binom{7}{5} + \binom{5}{4}$. The first a elements of $\binom{\mathbb{N}}{6}$ are then

$$\mathcal{S} := \binom{[14]}{6} \coprod \{H \cup \{15\} : H \in \binom{[7]}{5}\} \coprod \{H \cup \{8, 15\} : H \in \binom{[5]}{4}\}$$

By definition,

$$a^{(6)} = \binom{14}{7} + \binom{7}{6} + \binom{5}{5},$$

and the corresponding initial $a^{(6)}$ elements of $\binom{\mathbb{N}}{7}$ are

$$\mathcal{T} := \binom{[14]}{7} \coprod \{G \cup \{15\} : G \in \binom{[7]}{6}\} \coprod \{G \cup \{8, 15\} : G \in \binom{[5]}{5}\}$$

Recall that Δ is a simplicial complex if $(T \in \Delta,\ S \subset T)$ implies $S \in \Delta$.

236 Corollary

Let $(1, f_0, f_1, \ldots, f_{d-1})$ be a positive integer vector. Define Δ_i to be the first f_{i-1} elements of $\binom{\mathbb{N}}{i}$ in rev-lex order, with $\Delta_0 := \varnothing$. Define $\Delta := \Delta_0 \coprod \Delta_1 \coprod \cdots \coprod \Delta_{d-1}$. Then Δ is a simplicial complex if and only if $f_i \leq f_{i-1}^{(i)}$.

Proof (Sketch.) Using the preceding example, every 6-element subset of $\mathcal{T}$ is contained in $\mathcal{S}$, and if we added another 7-element subset to $\mathcal{T}$, some 6-element subset of it would not belong to $\mathcal{S}$. The proof follows quickly by using this sort of reasoning.

The corollary finishes sufficiency (i.e. existence). A complex as in the corollary is called a compressed complex (or sometimes a rev-lex complex). For necessity, one approach going back to Erdös-Ko-Rado is combinatorial shifting. This is an operation that transforms any simplicial complex Δ into a compressed simplicial complex with the same face numbers. The second direction of the corollary then gives necessity. Combinatorial shifting is accomplished by replacing single elements of subsets $S \in \Delta$ with other, smaller element in an appropriate way. This certainly does not change the face number; there are many proofs that this operation can indeed be defined and give a compressed simplicial complex. It would take a lecture or two to give a proof, so we will not take the time.

237 Example

The inequalities $0 < f_i \leq f_{i-1}^{(i)}$ for all $1 \leq i \leq d-1$ at $i = 1$ using $f_0 = n = \binom{n}{1}$ gives $f_1 \leq \binom{n}{2}$, which is just the fact that a graph with n vertices cannot have more than $\binom{n}{2}$ edges. In some sense, the Kruskal-Katona theorem is generalizing this observation to higher dimensions.

238 Definition

A d-dimensional polytope $P(n, d)$ with n vertices is a stacked polytope if

- it is a simplex $(n = d - 1)$, or
- $n > d + 1$ and $P(n, d)$ is obtained from some $P(n-1, d)$ by "building a shallow pyramid on one of the facets."

239 Remark

In contrast with cyclic polytopes, $P(n, d)$ is *not* well-defined up to combinatorial isomorphism. That is, two stacked d-dimensional polytopes with n vertices may not be combinatorially isomorphic. Nonetheless, as we showed on the homework, the face numbers are the same for all $P(n, d)$.

240 Example

Starting with a regular triangle in the plane gives $P(3,2)$. Gluing a second regular triangle onto one edge and erasing the common edge yields a quadrilateral $P(4,2)$. Doing this one more time gives a pentagon $P(5,2)$.

241 Remark

During the recursive step, we find

$$f_j(P(n,d)) = f_j(P(n-1,d)) + \begin{cases} \binom{d+1}{j+1} - \binom{d}{j+1} & \text{if } j \le d-2 \\ d-1 & j = d-1. \end{cases}$$

$$= \cdots = \begin{cases} \binom{d+1}{j+1} + (n-d-1)\binom{d}{j} & \text{if } j \le d-2 \\ (d+1) + (n-d-1)(d-1) & \text{if } j = d-1 \end{cases}$$

As we showed in the homework, the h-vector is

$$h(P(n,d)) = (1, n-d, n-d, \dots, n-d, 1).$$

242 Theorem (Lower Bound Theorem)

If P is a simplicial d-dimensional polytope with n vertices, then

$$f_j(P) \ge f_j(P(n,d))$$

for all $1 \le j \le d-1$.

243 Remark

In contrast to the upper bound theorem, this theorem really only works for simplicial polytopes, and it is easy to construct examples illustrating this.

The theorem was first proved by Barnette (1973) using shellability. We will go through Kalai's proof which uses "rigidity." The first step of the proof is due to McMullen-Perles-Walkup, which reduces theorem to the $j = 1$ case for *all* P and is sometimes called MPW reduction.

If $j = 1, d \ge 3$, the condition is $f_1(P) \ge f_1(P(n,3)) = \binom{d+1}{2} + (n-d-1)d = df_0 - \binom{d+1}{2}$, or equivalently $f_1 - df_0 + \binom{d+1}{2} \ge 0$. Recall that

$$f_1 = h_2 + (d-1)h_1 + \binom{d}{2} df_0 = dh_1 + d^2$$

so that the condition is $h_2 - h_1 \ge 0$.

Proof (Outline.)

Step 1: Show that if $h_2(Q) \ge h_1(Q)$ for all simplicial polytopes of dimension $d-1$, then $f_j(P) \ge f_j(P(n,d))$ for all $j \ge 1$.

Step 2: Show that $h_2 \ge h_1$ using rigidity theory.

We will fill in this outline for the remainder of this lecture and in the next lecture.

244 Remark

We describe MPW reduction.

245 Notation

Write $\phi_k(n,d) := f_k(P(n,d))$. Given any $(d-1)$-dimensional simplicial complex $\mathcal{C}$, write

$$\gamma(C) := f_1(\mathcal{C}) - \phi_1(n,d) = \begin{cases} h_2(\mathcal{C}) - h_1(\mathcal{C}) & \text{if } d \ge 3 \\ f_1(\mathcal{C}) - f_0(\mathcal{C}) & \text{if } d = 2. \end{cases}$$

For instance, if $\mathcal{C}$ is the boundary of a polygon, this is zero.

If P is a polytope and F is a face of P, we defined a quotient polytope P/F up to combinatorial isomorphism by $L(P/F) \cong [F, P] \subset L(P)$. If F is a vertex, we noted this can be done by intersecting P with a hyperplane "near" the vertex.

The key claim in MPW reduction is the following. If P is a d-dimensional simplicial polytope, $f_0(P) = n$, and $1 \le k \le d-1$, then

$$f_k(P) - \phi_k(n,d) = \sum_F w(F)\gamma(P/F)$$

where the sum is over faces $F \in \mathcal{C}(\partial P)$ such that $-1 \le \dim F < k-1$ where the weights $w(F)$ are all strictly positive.

We note that $L(P/F)$ is indeed a simplicial polytope, which follows essentially by noting that intervals in boolean lattices are intervals.

246 Corollary

If $\gamma(Q) \ge 0$ for all simplicial Q, then the lower bound theorem holds. Moreover, if $f_k(P) = \phi_k(n,d)$ for a single value of k, then $\gamma(P) = 0$.

May 6th, 2016: Draft

---

247 Remark

Today we will discuss the lower bound theorem, motion, and rigidity.

Recall that the lower bound theorem states the following. If P is a simplicial $(d-1)$-dimensional polytope with n vertices, then $f_{i-1}(P) \ge f_{i-1}(P(n,d))$, where $P(n,d)$ is any d-dimensional stacked polytope with n vertices.

We began discussing "step 1" of the proof last time. We had defined

$$\begin{aligned}\gamma(P) &:= f_1(P) - f_1(P(n,d)) \\ &= \begin{cases} h_2(P) - h_1(P) =: g_2(P) & \text{if } d \ge 3 \\ f_1(P) - f_2(P) & \text{if } d = 2 \end{cases}\end{aligned}$$

We discussed the MPW reduction. If P is a simplicial $(d-1)$-dimensional polytope with n vertices and $1 \le k \le d-1$, then the claim is

$$f_k(P) - \phi_k(n,d) = \sum_F w(F) \cdot \gamma(P/F)$$

for some *positively* weighted sum over $F \in \mathcal{C}(\partial P)$ with $-1 \le \dim F < k-1$, where we had set $\phi_k(n,d) := f_k(P(n,d))$. An immediate corollary of this claim is that if $\gamma(P/F) \ge 0$ for all F, then $f_k(P) \ge f_k(P(n,d))$ for all k.

248 Aside

It turns out that the statement of MPW reduction holds if we replace P with any pure $(d-1)$-dimensional complex $\mathcal{C}$ and P/F with the link of F in $\mathcal{C}$. (In this generality, we are not claiming that the γ's are non-negative, but we are claiming that the weights are non-negative.)

249 Definition

If $\mathcal{C}$ is a simplicial complex and $F \in \mathcal{C}$, then the link of F in $\mathcal{C}$ is

$$\boxed{\operatorname{lk}_{\mathcal{C}} F} := \{G - F : F \subset G \in \mathcal{C}\} = \{H : H \in \mathcal{C}, H \cap F = \varnothing, H \cup F \in \mathcal{C}\}.$$

250 Example

Consider a simplicial complex arising from a triangulation of a manifold. One would naively think the link of a vertex would be topologically a punctured ball, i.e. a sphere. This is not quite true, but the intuition is nonetheless that by considering vertex links we can use results for spheres in studying manifolds or complexes.

251 Remark

We have:

(1) $\mathrm{lk}_{\mathcal{C}} F$ is a simplicial complex. If $\mathcal{C}$ is pure and $(d-1)$-dimensional, then $\mathrm{lk}_{\mathcal{C}} F$ is also pure of dimension $d-1-|F|$.

(2) Consider the modified face poset $\widehat{\mathcal{F}}(\mathcal{C}) := \mathcal{F}(\mathcal{C}) \coprod \{\widehat{1}\}$. Then

$$\widehat{F}(\mathrm{lk}_{\mathcal{C}} F) \cong [F, \widehat{1}] \subset \widehat{F}(\mathcal{C}).$$

In particular, if $\mathcal{C} = C(\partial P)$, then $\widehat{F}(\mathcal{C}) = L(P)$, so $L(P/F) = [F, P] \cong \widehat{\mathcal{F}(\mathrm{lk}_{\mathcal{C}} \mathcal{F})}$. That is, for the boundary complex of a polytope, the link is really giving us the quotient.

In particular, the statement of MPW reduction is the following. If $\mathcal{C}$ is a pure $(d-1)$-dimensional simplicial complex with n vertices and $1 \le k \le d-1$, then

$$f_k(\mathcal{C}) - \phi_k(n,d) = \sum_{\substack{F \in \mathcal{C} \\ -1 \le \dim F < k-1}} w(F) \cdot \gamma(\mathrm{lk}_{\mathcal{C}} F)$$

where $w(F) > 0$ for all F.

(3) If $F = \{v_1, v_2\} \in \mathcal{C}$, then

$$\mathrm{lk}_{\mathcal{C}} F = \mathrm{lk}_{\mathrm{lk}_{\mathcal{C}} v_1}(v_2)$$

and we can iterate this for $|F| > 2$. This is often useful for inductive proofs.

(4) We have

$$\begin{aligned}(k+1) \cdot f_k(\mathcal{C}) &= \#\text{ pairs } (k\text{-dimensional faces } F, \text{ vertices } v \text{ of } F) \\ &= \sum_{v \in \mathrm{vert}(\mathcal{C})} f_{k-1}(\mathrm{lk}_{\mathcal{C}} v)\end{aligned}$$

(This was almost surely used in the homework.)

Proof (of MPW reduction; sketch) Say the vertices of $\mathcal{C}$ are $v_1, v_2, \ldots, v_n$ and their degrees are $m_1, m_2, \ldots, m_n$, respectively. Note that $m_i = f_0(\mathrm{lk}_{\mathcal{C}} v_i)$. Skipping some computations, one writes

$$\phi_k(n,d) = a_k(d) \cdot n + b_k(d)$$

for some $a_k(d), b_k(d)$ where $a_k(d) > 0$. One further shows

$$(k+1)(f_k(\mathcal{C}) - \phi_k(n,d)) = \cdots = 2a_{k-1}(d-1) \cdot \gamma(\mathcal{C}) + \sum_{i=1}^{n} (f_{k-1}(\mathrm{lk}\, v_i) - \phi_{k-1}(m_i, d-1)).$$

Here $\gamma(\mathcal{C}) = \gamma(\mathrm{lk}\, \varnothing)$, and induction finishes the proof.

The details of this argument are in the posted notes; they are not terribly difficult, but they are also not terribly enlightening.

252 Remark

We now turn to Step 2 of the lower bound theorem proof. Interestingly, we will need convexity only for three-dimensional complexes, which begins to explain why generalizations of the theorem hold for simplicial spheres.

More precisely, we must show that if $d \ge 3$ and P is a d-dimensional polytope, then $\gamma(P) = h_2(P) - h_1(P)$ is non-negative. The idea is to use coordinates of P to define a matrix R such that $\dim \ker R = \gamma(P)$, which is thus non-negative. This will take a bit of work.

253 Remark

We first discuss framework rigidity.

254 Definition

A d-dimensional framework G (or a "framework in $\mathbb{R}^d$") is a *finite* collection of vertices (called joints) $v_i \in \mathbb{R}^d$ for $i \in [n]$ and a collection of edges (called bars) $v_i v_j = \operatorname{conv}\{v_i, v_j\}$ for ij in some edge set $E \subset \binom{[n]}{2}$.

255 Remark

Thus, a d-dimensional framework is a graph $G = ([n], E)$ together with a map $\psi\colon [n] \to \mathbb{R}^d$. (Note that ψ need not be an *embedding*, i.e. ψ could map everything to a point, or two distinct edges could overlap, etc.)

256 Example

Let $G = K_3$. We can build frameworks for G in several ways:

- ψ could be constant, so the framework "is a point";
- ψ could send the vertices to the vertices of an equilateral triangle, in which case it is an embedding;
- ψ could send the vertices to colinear points, say with $v_1 < v_2 < v_3$, in which case the edge from 13 overlaps the edges 12 and 23.

257 Definition

A motion of a d-dimensional framework is a (smooth) function

$$\begin{aligned} [0,1] \times [n] &\to \mathbb{R}^d \\ (t, i) &\mapsto v_i(t) \end{aligned}$$

such that

- $v_i(0) = v_i$ for all i;
- $\|v_i(t) - v_j(t)\| = \|v_i - v_j\|$ for all $ij \in E$ and all t.

We think of $v_i(t)$ the position of the ith vertex at time t.

258 Example

Starting with a 2-dimensional framework which is a parallelogram, a trivial motion is translation. A less trivial motion is "squishing it to become a rectangle" while preserving edge lengths.

Any Euclidean motion of $\mathbb{R}^d$ always gives rise to such a motion on any $G \subset \mathbb{R}^d$. These include translations and rotations but not reflections. We call these trivial motions. More precisely, any

$$\psi\colon [0,1] \times \mathbb{R}^d \to \mathbb{R}^d$$

such that $\|\psi(t,v) - \psi(t,\overline{v}\| = \|v - \overline{v}\|$ for all $t \in [0,1]$ and all $v, \overline{v} \in \mathbb{R}^d$, and $\psi(0,v) = v$ for all $v \in \mathbb{R}^d$.

259 Definition

A d-dimensional framework is rigid if it allows only trivial motions. That is, for any motion of the framework, it extends to a rigid motion of the whole Euclidean space.

260 Example

We have:

- An equilateral triangle in any $\mathbb{R}^n$ with $n \geq 2$ is rigid.

- A parallelogram in $\mathbb{R}^2$ is not rigid. However, we can rigidify it by adding a vertex at the intersection of the two diagonals with four new bars and using the result as the base of a pyramid in $\mathbb{R}^3$. As an exercise, find a different vertex to a parallelogram in $\mathbb{R}^3$ such that the result is rigid.

261 Remark

If $(v_1(t), v_2(t), \ldots, v_n(t))$ for $t \in [0,1]$ is a motion of G, then $\|v_i(t) - v_j(t)\|^2$ is a constant depending on i, j for all $ij \in E$. Differentiating,

$$0 = 2\langle v_i(t) - v_j(t), v_i'(t) - v_j'(t)\rangle.$$

At $t = 0$, writing $u_i := v_i'(0)$, we get

$$\langle v_i - v_j, u_i - u_j\rangle = 0, \qquad \forall ij \in E.$$

262 Definition

An infinitesimal motion of $G \subset \mathbb{R}^d$ is a collection of vectors $u_1, \ldots, u_n \in \mathbb{R}^d$ such that

$$\langle v_i - v_j, u_i - u_j = 0, \qquad \forall ij \in E.$$

Equivalently, we require $\operatorname{proj}_{v_i - v_j} u_i = \operatorname{proj}_{v_i - v_j} u_j$ for all $ij \in E$.

263 Example

A trivial infinitesmal motion arises by taking $u_1 = \cdots = u_n$. This will eventually correspond to translation.

Return to the pyramid of the preceding example. Let v_6 be the intersection of the diagonals of the original parallelogram. Set $u_1 = \cdots = u_5 = 0$ with u_6 perpendicular to the base. This example shows that while rigid motions always induce infinitesimal motions, the converse does not hold.

May 9th, 2016: Draft

(Missed.)

May 11th, 2016: Draft

---

264 Remark

Homework 6 is on the web; it is due *next* Friday. Annie Raymond will sub on Friday and Monday; she will discuss matroidal polytopes and their intersections.

265 Remark

Recall our setup. We had a building $G = (V, E) \subset \mathbb{R}^d$ with $V = \{v_1, \ldots, v_n\}$. An infinitesimal motion of G is $u_1, \ldots, u_n \in \mathbb{R}^d$ such that $\langle v_i - v_j, u_i - u_j\rangle = 0$ for all $ij \in E$. We denote by $M(G)$ the space of all infinitesimal motions of G, and we can do the same for $M(\mathbb{R}^d)$; these are vector spaces. Each element of $M(\mathbb{R}^d)$ induces an infinitesimal motion of $M(G)$, and infinitesimal rigidity is precisely the claim that this induced map $M(\mathbb{R}^d) \to M(G)$ is surjective.

Last time we proved the following theorem:

266 Theorem

$\dim M(\mathbb{R}^d) = \binom{d+1}{2}$. *(Indeed, $M(\mathbb{R}^d) = M(\Sigma^d)$ where Σ^d is a d-dimensional simplex.)*

267 Corollary

A framework $G = (V,E)$ in $\mathbb{R}^d$ (with $\operatorname{Aff}(V) = \mathbb{R}^d$) is infinitesimally rigid if and only if $\dim M(G) = \binom{d+1}{2}$. In this case, every infinitesimal motion of G (or $\mathbb{R}^d$) is determined by its restriction to any d affinely independent joints of G.

268 Corollary (The Gluing Lemma)

If G and G' are two d-dimensional infinitesimally rigid frameworks that have d affinely independent joints in common, then $G_1 \cup G_2$ is also infinitesimally rigid.

269 Remark

We next express the space $M(G)$ as the kernel of an appropriate matrix, which offers a convenient way to compute the space. Recall that our goal is to find u_i such that $\langle v_i - v_j, u_i - u_j\rangle = 0$ for all $ij \in E$ with $G = (V,E) \subset \mathbb{R}^d$.

Form a matrix whose rows are indexed by edges of G, of which there are f_1. Group columns into blocks of size d, where each block corresponds to a vertex of G, so there are $d \cdot f_0$ columns. For a particular row indexed by ij, put zeros in each block corresponding to vertices different from v_i and v_j. For the block corresponding to v_i, place $(v_j - v_i)^T$ in it, and for the block corresponding to v_j, place $(v_i - v_j)^T$ in it. This matrix doesn't have a common name, though it's sometimes referred to as the rigidity matrix of G or the stress matrix of G, which we'll write as $\mathcal{R}(G,d)$ or perhaps just $\mathcal{R}$ if G is understood.

Note that if $u_1, \ldots, u_n$ is an infinitesimal motion of G, then $\langle v_i - v_j, u_i\rangle = \langle v_i - v_j, u_j\rangle$ for all $ij \in E$, and conversely. Consider multiplying the column vector

$$\begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ u_n \end{pmatrix}$$

on the left by $\mathcal{R}$. The result is simply a column vector with rows indexed by edges where the $ij \in E$ entry is $\langle v_j - v_i, u_i\rangle + \langle v_i - v_j, u_j\rangle$. We have just proven:

270 Theorem

$M(G) = \ker \mathcal{R}(G,d)$.

We may also consider the left kernel of $\mathcal{R}$ (or up to transpose the right kernel of $\mathcal{R}^T$). Pick a row vector λ_{ij} with columns indexed by $ij \in E$. Multiplying this on the right by $\mathcal{R}$ gives $\sum_{i:ik\in E} \lambda_{ik}(v_i - v_k)^T$.

Hence an element λ of the left null space of $\mathcal{R}$ is an assignment of a number λ_{ij} for each edge ij such that the following holds. We require $\sum_{i:ik\in E} \lambda_{ij} = 0$ for all $k = 1, \ldots, n$.

271 Definition

The left null space of $\mathcal{R} = \mathcal{R}(G,d)$ is called the stress space of G. Its elements are called stresses on G.

272 Example

Pick a regular pentagon in the plane. Number vertices from 1 to 5 clockwise (say). The edges are then $12, 23, 34, 45, 51$, and stresses are choices of numbers for each edge. At vertex 1, we require $\lambda_{12}(v_2 - v_1) + \lambda_{15}(v_5 - v_1) = 0$. Pictorially, these vectors are just the vectors pointing from 1 to 2 and from 1 to 5, which clearly form a basis, so we require $\lambda_{12} = \lambda_{15}$. In this way, we have only trivial stress in this case.

273 Example

Consider a square in $\mathbb{R}^2$ where we've added the two diagonals as additional edges. We can form a stress using edge weights of 1 for the outer edges and -1 for the diagonals. The matrix itself is straightforward to write in this case as well.

274 Theorem

Let $G = (V, E)$ be a framework in $\mathbb{R}^d$, where $\operatorname{Aff}(V) = \mathbb{R}^d$, with f_1 edges and f_0 vertices. The following are equivalent:

(1) G is infinitesimally rigid

(2) $\dim M(G) = \binom{d+1}{2}$

(3) $\operatorname{rank} \mathcal{R}(G, d) = df_0 - \binom{d+1}{2}$

(4) the dimension of the stress space is $f_1 - df_0 + \binom{d+1}{2}$

Proof The equivalence of (1) and (2) was shown above. (3) and (4) follows from rank-nullity.

275 Remark

Recall that in the case when G comes from a simplicial polytope, the number in (4) is really $h_2 - h_1$.

276 Corollary

Let P be a d-dimensional simplicial polytope (or in fact any $(d-1)$-dimensional simplicial complex). If P is infinitesimal (as a framework in $\mathbb{R}^d$), then

$$h_2(P) - h_1(P) = f_1 - df_0 + \binom{d+1}{2} \geq 0.$$

Proof From the previous theorem, this quantity is the dimension of a vector space, which is non-negative.

277 Remark

To finish the proof of the lower bound theorem, we just need to show simplicial polytopes are infinitesimally rigid. Note that this fails for more general polytopes (or complexes).

278 Theorem (Whiteley, 1984)

For $d \geq 3$, d-dimensional simplicial polytopes in $\mathbb{R}^d$ are infinitesimally rigid.

Proof We give the sketch now; we will give the details sometimes next week.

The strategy is to induct on d. At the $d = 3$ base case, by the Dehn-Sommerville relations, $h_2(P) - h_1(P) = 0$. Hence we must show the stress space is trivial, which is precisely the following older theorem:

279 Theorem (Dehn, 1916)

If P is a simplicial 3-dimensional polytope, then P admits only the trivial stress.

For the inductive step, take $d > 3$. Pick a vertex $v_0 \in P$ and consider the link of v_0 in P. This is essentially only defined combinatorially in general, but we noted that it can also be identified with the vertex figure of v_0. That is, $Q := P/v_0$ is P intersected with a hyperplane "close" to v_0. We know that Q is infinitesimally rigid in $\mathbb{R}^{d-1}$ by induction.

Now define G_0 to be the graph of the star of v_0, meaning the vertices and edges of G_0 are the vertices and edges of the facets of P that contain v_0. We'll show that Q being infinitesimally rigid in $(d-1)$ dimensions implies G_0 is infinitesimally rigid in d dimensions. Now use the gluing lemma to show that the union of all such stars is also infinitesimally rigid. This union is precisely the graph of P, which completes the proof. We expand on the gluing lemma step while it is still fresh.

Take v_0, v_1 adjacent with $G_0 = G(\operatorname{st} v_0)$, $G = G(\operatorname{st} v_1)$. Then $G_0 \cap G_1$ contains all d vertices of any facet that contains the edge v_0v_1, and this is non-empty since the polytope is pure. By the gluing lemma, $G_0 \cup G_1$ is infinitesimally rigid. This procedure can be inductively repeated.

May 13th, 2016: Draft

280 Remark

Annie Raymond is lecturing today.

281 Definition

Take a pair $M = (E(M), \mathcal{F}(M))$ where $E(M)$ is some ground set of elements and $\mathcal{F}(M)$ is a family of subsets of $E(M)$. Call the elements of $\mathcal{F}(M)$ independent sets. The pair M is a matroid if

(1) $S \in \mathcal{F}(M)$ and $T \subset S$, then $T \in \mathcal{F}(M)$

(2) $S_1, S_2 \in \mathcal{F}(M)$ and $|S_2| > |S_1|$, then there exists $e \in S_2 - S_1$ such that $S_1 \cup \{e\} \in \mathcal{F}(M)$.

282 Example

A linear matroid is a matroid obtained in the following way. Given a matrix A, let $E(M)$ be the indices of the columns of A and let $\mathcal{F}(M)$ be the sets of indices whose columns are linearly independent.

For instance, given

$$\begin{pmatrix} 1 & 0 & 1 & 2 & -1 & 0 \\ 1 & 1 & 2 & 3 & -2 & 0 \\ 0 & 1 & 1 & 1 & -1 & 1 \end{pmatrix}$$

then our ground set is $E = [6]$ and our independent sets include $\varnothing$, $\{1\}4, \{2\}, \ldots, \{6\}, \{1,2\}$, but not $\{1,2,3\}$. Since A is a rank three matrix, the largest independent sets here are size 3. The second condition on independent sets is immediate from elementary linear algebra.

283 Example

A graphic matroid is a matroid obtained from a graph $G = (V, E)$ in the following way. Let $E(M) := E$ and $\mathcal{F}(M)$ be the edge sets of forests in G. (A forest is an acyclic graph.)

For instance, if G is the complete graph on 5 vertices, then independent sets include the edges arising from the empty graph, or the boundary of a pentagon with one edge removed, etc.

284 Remark

We next discuss matroid maximization. Let $M = (E, \mathcal{F})$ be a matroid and let $c : E \to \mathbb{R}$ be a "cost function." The goal is to find an independent set $S \in \mathcal{F}$ which maximizes $c(S) := \sum_{e \in S} c_e$.

Indeed, consider the following greedy algorithm. Order E such that $c(e_1) \geq c(e_2) \geq \cdots \geq c(e_{|E|})$. Begin at $S = \varnothing$. For j from 1 to $|E|$, if $S \cup \{e_j\} \in \mathcal{F}$, then replace S with $S \cup \{e_j\}$ so long as the overall cost actually increases.

Let S_k be the set obtained in the algorithm after k elements have been added. We claim S_k maximizes $c(S)$ over all $S \in \mathcal{F}$ with $|S| = k$.

Proof Suppose $S_k = \{s_1, \ldots, s_k\}$ and consider $T = \{t_1, \ldots, t_k\} \in \mathcal{F}$ such that $c(t_1) \geq \cdots \geq c(t_k)$. Further suppose to the contrary that $c(T) > c(S_k)$, so there exists i such that $c(t_i) > c(s_i)$ and let p be the smallest such i. Now define $A := \{t_1, \ldots, t_p\}$ and $S_{p-1} := \{s_1, \ldots, s_{p-1}\}$, which are independent sets by the first property. By the second property, we may move some element of A over to S_{p-1} and get an independent set, or formally there exists t_j such that $1 \leq j \leq p$ such that $S_{p-1} \cup \{t_j\} \in \mathcal{F}$ and $t_j \notin S_{p-1}$. We now have $c(t_j) \geq c(t_p) > c(s_p)$. However, the algorithm would have added t_p before s_p, which is a contradiction.

285 Remark

Kruskal's algorithm is (nearly? double check) the preceding algorithm applied to a graphic matroid. Indeed, the above property completely characterizes matroids, though we won't make this precise or prove it.

286 Definition

The rank of a matroid is the function

$$\boxed{r}: 2^{|E|} \to \mathbb{Z}_+$$
$$S \subset E \mapsto r(S) := \max_{\substack{T \subset S \\ T \in \mathcal{F}}} |T|.$$

287 Example

In the graphic matroid of K_5, any subset of three vertices have rank 2, since the largest forest on those vertices has two edges.

288 Proposition

The rank function satisfies the following:

(1) $0 \le r(S) \le |S|$ *for all* $S \subset E$.

(2) If $S \subset T$, *then* $r(S) \le r(T)$.

(3) The submodularity *condition: for all* $S, T \subset E$,

$$r(S) + r(T) \ge r(S \cup T) + r(S \cap T)$$

Proof (1) and (2) are immediate. (3) is not hard, though we skip it for the sake of time (and interest).

289 Definition

Let $M = (E, \mathcal{F})$ be a matroid. The corresponding matroid polytope is given by

$$P := \operatorname{conv}\{\chi^S : S \in \mathcal{F}\}$$

where χ^S is the indicator function on the subset S of E. This is a polytope in $\mathbb{R}^{|E|}$.

290 Example

Let G be the complete graph on three vertices numbered $1, 2, 3$. The corresponding graphic matroid has independent sets $\varnothing, \{12\}, \{13\}, \{23\}, \{12, 13\}, \{12, 23\}, \{13, 23\}$, which have corresponding indicator vectors $(0,0,0), (1,0,0), (0,1,0), (0,0,1), (1,1,0), (1,0,1), (0,1,1)$. The corresponding matroid polytope is the convex hull of these indicator vectors.

291 Remark

One is often interested in optimizing some linear function over some discrete set, which one can sometimes encode as optimizing the function over the corresponding matroidal polytope. However, the number of vertices can sometimes be prohibitively enormous, and one may instead prefer to maximize using the half-plane description of the polytope.

292 Definition

Continuing the notation of the preceding definition, define

$$Q_I := \operatorname{conv}\{x \in \mathbb{R}^{|E|} : \forall S \subset E, \sum_{e \in S} x_e \le r(S); \forall e \in E, x_e \in \{0, 1\}\}.$$

The right-hand side is an *integer* program, which in general is very hard to solve. We have $Q_I = P$ essentially immediately. We may avoid an integer program by relaxing the final constraint as in the following:

$$Q := \{x \in \mathbb{R}^{|E|} : \forall S \subset E, \sum_{e} x_e \le r(S); \forall e \in E, 0 \le x_e (\le 1)\}.$$

One has $P \subset Q$ immediately. Q_I is the convex hull of the integer points inside of Q.

May 16th, 2016: Draft

---

293 Remark

Recall the end of last lecture. We had a matroid $M = (E, \mathcal{F})$, and we had a corresponding matroid polytope $P := \text{conv}(\chi^S : S \in F)$. We had a continuous analogue of the matroid polytope $Q := \text{conv}\{x \in \mathbb{R}^{|E|} : \forall S \subset E, \sum_{e \in S} x_e \leq r(S); \forall e \in E, x_e \geq 0\}$. It is clear that $P \subset Q$. Perhaps surprisingly, the reverse containment also holds. We will give several proofs of this.

294 Proposition

In the notation above, $P = Q$.

Proof (First proof of proposition.) Intuitively, we begin by asking if we are missing any equalities for Q, i.e. is $\dim P = \dim Q$? We may assume every element $e \in E$ is independent, i.e. $\{e\} \in \mathcal{F}$, since otherwise $x_e = 0$ is forced in both P and Q. Hence P contains $\text{conv}\{0, e_1, \ldots, e_{|E|}\}$, so P is full dimensional, so Q must be as well:

$$|E| \leq \dim P \leq \dim Q \leq |E|.$$

Now we ask if we are missing any inequalities for Q. Let $\alpha^T x \leq \beta$ be a valid inequality for P, i.e. there exists some $v \in P$ such that $\alpha^T v = \beta$ and P lies on one side of $\alpha^T x = \beta$. Then

$$\begin{aligned}\beta &= \max_{x \in P} \alpha^T x \\ &= \max_{S \in \mathcal{F}} \alpha(S)\end{aligned}$$

where $\alpha(S) := \sum_{e \in S} \alpha_e$. Let $F := \{x^* \in P : \alpha^T x^* = \beta\}$ be the face induced by $\alpha^T x \leq \beta$ on P. Further let $S := \{e \in E : \alpha_e = \max_{f \in E}\{\alpha_f\}\}$.

Claim: for any $T \in \mathcal{F}$ such that $\alpha(T) = \beta$, we have that $|S \cap T| = r(S)$.

To prove the claim, assume not. $S \cap T \in \mathcal{F}$ lies in S, so $|S \cap T| < r(S)$. Extend $S \cap T$ to some set X such that $X \in \mathcal{F}$, $X \subset S$, and $|X| = r(S)$ using the second property of matroids repeatedly. Similarly extend X to a set Y such that $Y \in \mathcal{F}$, $Y \subset T \cup X$, $|Y| = |T|$. Note that

$$\begin{aligned}\alpha(Y) - \alpha(T) &= \sum_{e \in Y} \alpha_e - \sum_{e \in T} \alpha_e \\ &= \sum_{e \in Y - T} \alpha_e - \sum_{e \in T - Y} \alpha_e.\end{aligned}$$

Now $Y - T \subset S$ and $T - Y \subset E - S$ and $|Y - T| = |Y| - |T \cap Y| = |T| - |T \cap Y| = |T - Y|$, forcing this sum to be non-negative. Hence $\alpha(Y) > \alpha(T) = \beta$, so $\alpha^T x \leq \beta$ is not valid, being violated by $Y \in \mathcal{F}$, a contradiction.

Given the claim, for any $T \in \mathcal{F}$ such that $\alpha(T) = \beta$, we have $|T \cap S| = r(S)$, so taking $\sum_{e \in S} x_e \leq r(S)$, F is contained in the face induced by $\sum_{e \in S} x_e \leq r(S)$. (Note to self: write this out a little more carefully.)

Proof (Second proof of proposition.) The above proof focused on half-spaces. Another way to proceed is to focus on vertices instead, which we'll get to if we have time. We'll give a third way here, which focuses on duality.

For any cost function c, $\max_{x \in P} c^T x \leq \max_{x \in Q} c^T x$ since $P \subset Q$. It we can show equality always holds, it will follow that $P = Q$. More precisely, we are maximizing $c^T x$ over $\sum_{e \in S} x_e \leq r(S)$ for all $S \subset E$; $x_e \geq 0$ for all $e \in E$. By LP-duality (which we will not discuss), this is equal to the minimum of $\sum_{S \subset E} r(S) y_S$ such that $\sum_{S : e \in S} y_S \geq c_e$ for all $e \in E$; $y_S \geq 0$ for all $S \subset E$.

Goal: show that for any c, we can find $S^* \in \mathcal{F}$ such and a dual feasible solution y^* that have the same weight. For S^*, we can use the greedy algorithm from last lecture. Recall that it ordered elements as weakly decreasing according to cost, iterated through those, and added elements whenever they both preserved independence and increased the cost. Hence we use the greedy algorithm on all elements e such that $c(e) \geq 0$, say $e_1, \ldots, e_q$. Suppose the greedy algorithm returns the set $S_k = \{s_1, \ldots, s_k\}$.

Build a dual solution as follows. For any index $j \le k$, let $S_j := \{s_1, \ldots, s_j\}$ and let $U_j := \{e_1, e_2, \ldots, e_\ell\}$ where $e_{\ell+1} = s_{j+1}$.

Note that $r(U_j) = r(S_j) = j$ since the greedy algorithm is greedy; this is why we restricted the s's to be non-negative. Let $y_S := 0$ for all $S \subset E$ except for $y_{U_j} = c(s_j) - c(s_{j+1})$ for all $1 \le j \le k$, where $c(s_{k+1}) := 0$. We claim this is valid for the dual program of the same weight as the solution for the primal program. We certainly have $y_S \ge 0$ for all S since the s's are weakly decreasing, so we must check that $\sum_{S \ni e} y_S \ge c_e$. For fixed e, let t be the least index such that $e \in U_t$. We have

$$\begin{aligned}\sum_{S \ni e} y_S &= \sum_{j=t}^{k} y_{U_j} \\ &= c(s_t) - c(s_{t+1}) + c(s_{t+1}) - c(s_{t+2}) + \cdots \\ &= c(s_t).\end{aligned}$$

Now $c(s_t) \ge c_e$ since s_t occurred before e in the ordering by the greedy algorithm. Hence y is a feasible dual solution. We now compute its weight:

$$\begin{aligned}\sum_{S} r(S) y_S &= \sum_{j=1}^{k} r(U_j) y_{U_j} \\ &= \sum_{j=1}^{k} j \cdot (c(s_j) - c(s_{j+1}) \\ &= \sum_{j=1}^{k} (j - (j-1)) c(s_j) \\ &= \sum_{j=1}^{k} c(s_j) = c(S_k).\end{aligned}$$

May 18th 2016: Draft

---

295 Remark

Isabella is lecturing again today. The plan for the rest of the quarter is as follows. Today we'll finish the proof of the lower bound theorem. Next time we'll discuss a generalization of the lower bound theorem. On Monday Sean will be lecturing on flow polytopes. Then we'll head towards the Hirsch conjecture.

296 Remark

Recall where we were in the proof of the lower bound theorem. We had reduced it to Dehn's theorem and Whiteley's theorem. We discussed infinitesimal rigidity in d-dimensions, which we'll sometimes refer to as $\boxed{d\text{-rigidity}}$ for clarity.

Recall that if $G = (V, E) \subset \mathbb{R}^d$ is a framework in $\mathbb{R}^d$ with $\mathrm{Aff}(V) = \mathbb{R}^d$, a stress on G is a function $\lambda \colon E \to \mathbb{R}$ by $ij \mapsto \lambda_{ij}$ such that $\sum_{j : ij \in E} \lambda_{ij}(v_i - v_j) = 0$ for all i. We had the following theorem in this context: G is infinitesimally d-rigid if and only if the dimension of the stress space is $f_1 - df_0 + \binom{d+1}{2}$, which is $h_2 - h_1$ assuming G is the graph of some $(d-1)$-dimensional complex and $d \ge 3$.

We begin by proving Whiteley's theorem.

297 Theorem (Whiteley, 1984)

For all $d \ge 3$, any d-dimensional simplicial polytope is infinitesimally d-rigid.

Proof To prove this, we just need to check the base case $d = 3$ and induct. The base case was proven much earlier:

298 Theorem (Dehn, 1916)
If P is a simplicial 3-dimensional polytope, then P admits only the trivial stress $\lambda = 0$.

Proof Fix P. Assume there exists a non-trivial stress λ. Label each edge $ij \in E$ by $+, -$, or 0 depending on the sign of λ_{ij}. If there exists $v \in V$ such that all edges incident with v are labeled 0, remove v and consider the convex hull of the remaining vertices.

Can we get a 2-dimensional complex? If so, the stress was supported on a polygon, and we showed that polygons admit only trivial stresses, which would imply $\lambda = 0$, contrary to our assumption, so the result is still 3-dimensional. Can we get more edges? Yes; just label them with 0.

Can we get a non-simplicial polytope? Yes. However, we may triangulate each 2-face by adding diagonals and label them with 0's. The result may still be non-simplicial, which will not bother us. Note that we did not introduce more vertices, so we may repeat this procedure to find a polytope Q whose vertices are a subset of the vertices of P with the property that no vertex has all zero edges incident with it, all non-zero edges are edges of P, and every 2-face is triangulated.

Now for each triangle Δ of Q in some 2-face, label the corners of Δ with 0, 1, or 1/2 as follows. The edgets of the triangle have signs $\epsilon_1, \epsilon_2, \epsilon_3$. Label the vertex on the edges labeled by both ϵ_1 and ϵ_2 as follows. If $\epsilon_1 = \epsilon_2$, use label 0. If $\epsilon_1 = -\epsilon_2 (\neq 0)$, use label 1. If $\epsilon_1 = 0, \epsilon_2 \neq 0$ (or $\epsilon_1 \neq 0, \epsilon_2 = 0$), use label 1/2.

Claim 1: consider a vertex v of Q. The sum over all triangles Δ containing v of the corner label at v is ≥ 4.

Claim 2: for all triangles Δ, the sum of the corner labels is ≤ 2.

Assuming these claims for the moment, let S denote the sum of all corner labels. By claim 1, $S \geq 4f_0(Q)$. By claim 2, $S \leq 2f_2(Q)$ where $f_2(Q)$ here really means the number of triangles. Since Q is 3-dimensional, $f_0 - f_1 + f_2 = 2$, and $3f_2 = 2f_1$, we have $f_2 = 2f_0 - 4$. (Indeed, $f_1 = 3f_0 - 6$.) But then

$$4f_0 \leq S \leq 2f_2 = 2(2f_0) - 4) = 4f_0 - 8,$$

a contradiction.

We turn to the proof of the claims. Begin with claim 2. There are only a handful of cases to check by symmetry. If all edge labels are positive, all vertex labels are 0, so the sum is 0. If there are two +'s and one −, the vertex labels are $1, 1, 0$, which sum to 2. If there are two +'s and one 0, the vertex labels are $1/2, 1/2, 0$, which sum to 1. We won't write out the other cases, but they work out.

For claim 1, we'll show that there are ≥ 4 as we "circle" along the edges incident with v. Here we imagine taking a hyperplane a small distance away from v and we intersect it with Q, and we order the edges incident with v using the resulting polygon. By the stress condition, we need $\sum_{i:iv} \lambda_{iv}(v - v_i) = 0$. Consider a supporting hyperplane of v in Q. The edges incident with v all lie on the same side of the hyperplane, i.e. $v - v_i$ all point in the same direction relative to the normal of the hyperplane, so the λ_{iv} cannot possibly be all of the same sign. Hence there is at least 1 sign change, and there must be an even number of sign changes.

We must now argue that there cannot be exactly 2 sign changes, so suppose to the contrary that there are exactly 2 sign changes. Recall that no 3 non-zero edges incident with v are coplanar since we began with a 3-polytope. Hence we can strictly separate the +'s from the −'s, i.e. there exists a hyperplane H such that all edges incident with v with label + on one side of H and all edges labeled − on the other side. Now project these edges $v - v_i$ onto the normal n of H, say pointing in the direction of the + edges. The contribution of the + edges to $\sum_{i:iv} \lambda_{iv}(v - v_i)$ is (positively) in the direction of n, and the contribution of the − edges is *also* (positively) in the direction of n, a contradiction.

The inductive step uses the following theorem:

299 Theorem

If P is a d-dimensional simplicial polytope ($d \geq 4$), then the vertex stars of P are infinitesimally d-rigid.

Proof Pick a vertex v_0 of P, and assume without loss of generality $v_0 = 0$. Assume without loss of generality that the neighbors of v_0 are vertices with coordinates $(v_1, a_1), \ldots, (v_m, a_m)$ where $v_i \in \mathbb{R}^{d-1}$ and $a_1, \ldots, a_m \in \mathbb{R}$ with $a_i > 1$. Let $Q := P/v_0 = P \cap \{(x_1, \ldots, x_d) : x_d = 1\}$ be the corresponding vertex figure. The vertices of Q have coordinates $(v_1/a_1, 1), \ldots, (v_m/a_m, 1)$, since they arise from the intersections of edges of P with the above hyperplane.

By the inductive assumption, Q is infinitesimally $(d-1)$-rigid. Let G_0 be the graph of the star of v_0. Claim: G_0 is infinitesimally d-rigid. Up to combinatorial isomorphism, G_0 is the graph of $(\overline{v_0} * \operatorname{lk}_{\partial P} v_0)$. We showed in homework

$$h(\overline{v_0} * \operatorname{lk}_{\partial P} v_0; x) = h(\overline{v_0}; x) \cdot h(\operatorname{lk}_{\partial P} v_0; x) = h(\operatorname{lk}_{\partial P} v_0; x).$$

Now

$$h_2(\operatorname{st} v_0) - h_1(\operatorname{st} v_0) = h_2(\operatorname{lk} v_0) - h_1(\operatorname{lk} v_0) = \dim(\text{stress space of } Q)$$

where the second equality uses the fact that Q is infinitesimally rigid. To prove the claim, it then suffices to show that the dimension of the stress space of G_0 is the same as the dimension of the stress space of Q. We will complete this next time.

May 20th, 2016: Draft

---

300 Remark

Today we will discuss generalizations of the lower bound theorem, and finish our proof from last time.

Proof (Continued from last time.) Recall that we had a d-dimensional polytope P with $d \geq 4$, vertices $v_1, \ldots, v_m \in \mathbb{R}^{d-1}$ with $\mathbb{R} \ni a_i > 1$, $v_0 = 0 \in \mathbb{R}^d$, where $(v_1, a_1), \ldots (v_m, a_m)$ are the neighbors of v_0. If Q is the intersection of $x_d = 1$ and P, we know that Q is infinitesimally $(d-1)$-rigid. We need to show that the vertex figure of P at v_0, call it G_0, is infinitesimally rigid.

Equivalently, we need to show that the dimension of the stress d-space of G_0 is the same as the dimension of the stress $(d-1)$-space of Q. If (λ_{ij}) is a stress of G_0, we have for each $i = 1, \ldots, m$

$$\sum_{\substack{ij \in E(G_0) \\ j \neq 0}} \lambda_{ij} \begin{pmatrix} v_i - v_j \\ a_i - a_j \end{pmatrix} + \lambda_{i0} \begin{pmatrix} v_i \\ a_i \end{pmatrix} = 0$$

$$\sum_{j=1}^{m} \lambda_{0j} \begin{pmatrix} v_j \\ a_j \end{pmatrix} = 0.$$

The first equation imply

$$\sum \lambda_{ij}(v_i - v_j) + \lambda_{i0} v_i = 0$$
$$\sum \lambda_{ij}(a_i - a_j) + \lambda_{i0} a_i = 0$$
$$\Rightarrow \lambda_{i0} - \frac{1}{a_i}\sum \lambda_{ij}(a_i - a_j).$$

Indeed, the second equation is then redundant. Moreover, setting $\overline{\lambda}_{ij} := a_i a_j \lambda_{ij}$, one checks that if (λ_{ij}) is a stress on G_0, then $\overline{\lambda}_{ij}$ is a stress on Q. In the other direction, given a stress α_{ij} on Q, one can form a stress on G_0 $\widetilde{\alpha}_{ij} := \frac{1}{a_i a_j}\alpha_{ij}$ with α_{i0} determined by the formula above. The result then follows, completing the proof of the lower bound theorem.

301 Remark

We turn to generalziations of the lower bound theorem.

302 Theorem (Whiteley)

Let P be any d-dimensional polytope (not necessarily simplicial), $d \geq 3$. Triangulate every 2-face of P by adding non-intersecting diagonals. Let $\widetilde{G}(P)$ be the corresponding graph. Then $\widetilde{G}(P)$ is infinitesimally rigid.

303 Corollary (Kalai, 1987)

Let P be a d-dimensional polytope (not necessarily simplicial), $d \geq 3$. Let $\boxed{f_2^k(P)}$ denote the number of 2-faces of P that are k-gons. Then

$$f_1(P) + \sum_{k\geq 3}(k-3)f_2^k(P) \geq df_0 - \binom{d+1}{2}.$$

Proof (Sketch.) The proof is roughly the same as in the above proof of the lower bound theorem. The left-hand side is the number of edges of $\widetilde{G}(P)$, so the left-hand side minus the right-hand side is precisely the dimension of the stress space of $\widetilde{G}$ by Whiteley's theorem, which is non-negative.

304 Aside

Consider

$$\begin{aligned}\sum_{k\geq 1}(k-3)f_2^k(P)& \\ &= \sum k f_2^k(P) - 3\sum f_2^k(P)\\ &= \#\{\text{pairs (2-face } F \text{ of } P, \text{ vertex } v \text{ of } F)\} + f_2(P)\\ &=: f_{02} - 3f_2.\end{aligned}$$

Here f_{02} is an example of a $\boxed{\text{flag } f\text{-number}}$, which we will not presently discuss further.

305 Definition

If $G = (V, E)$ is an abstract graph, we say it is $\boxed{\text{generically } d\text{-rigid}}$ if almost all embeddings (in the sense above; these need not be injective) of G in $\mathbb{R}^d$ are infinitesimally rigid.

306 Remark

Homework problem number 3 in this terminology says that graphs of simplicial 3-polytopes are generically 3-rigid. We have a low-dimensional miracle: every 2-dimensional simplicial sphere is the boundary complex of some 3-dimensional polytope. We can then restate the result of the homework problem as saying graphs of simplicial 2-dimensional spheres are generically 3-rigid.

307 Lemma (The Cone Lemma)

If $G = (V, E)$, the $\boxed{\text{cone}}$ over G is

$$\boxed{CG} := \{(V \coprod \{v_0\}, E \coprod \{v_0 v : v \in V\}\}.$$

Then G is generically $(d-1)$-rigid if and only if CG is generically d-rigid.

308 Lemma (The Gluing Lemma)
Let $G_1 = (V_1, E)$, $G_2 = (V_2, E_2)$ such that

- *G_1 and G_2 are generically d-rigid,*
- *$|V_1 \cap V_2| \geq d$.*

Then $\boxed{G_1 \cup G_2} = (V_1 \cup V_2, E_1 \cup E_2)$ *is generically d-rigid.*

309 Theorem (Kalai, 1987)
The graph of any connected $(d-1)$-dimensional simplicial manifold Δ (with $d \geq 3$) is generically rigid, and so $h_2(\Delta) - h_1(\Delta) \geq 0$.

Proof (Sketch.) The proof is again essentially the same as our proof of the lower bound theorem. The base case is the restated homework problem above. One uses a version of Whiteley's theorem for generic rigidity together with the above generic versions of the cone lemma and the gluing lemma. One difficulty is that if Δ is $(d-1)$-dimensional, all vertex links of Δ are only *homology* $(d-2)$-spheres. By the cone lemma, graphs of vertex links are generically $(d-1)$-rigid, and graphs of vertex starts are generically d-rigid. Indeed, Kalai showed equality holds precisely for boundaries of stacked polytopes.

Gromov independently proved this result using his own notion of combinatorial rigidity, also in 1987.

310 Conjecture (Kalai, 1987)
Let Δ be a $(d-1)$-dimensional simplicial (connected) manifold, $d \geq 4$. Then

$$h_2(\Delta) - h_1(\Delta) \geq \binom{d+1}{2} \beta_1(\Delta)$$

where β_1 is the dimension of the first homology of δ (computed with coefficients in some field; part of the conjecture is that the bound holds for all fields).

311 Example
For the torus $S^1 \times S^1$, $\beta_1 = 2$, which is independent of the field used. For the two-sphere S^2, $\beta_1 = 0$ for all fields. However,

$$\beta_1(\mathbb{R}P^2; k) = \begin{cases} 0 & \text{if } \operatorname{char} k \neq 2 \\ 1 & \text{if } \operatorname{char} k = 2. \end{cases}$$

312 Conjecture
We have

$$h_2 - h_1 \geq \binom{d+1}{2} m(\Delta)$$

where $m(\Delta)$ is the minimum number of generators of $\pi_1(\Delta)$.

313 Example
We have $m(S^1 \times S^1) = 2$ since $\pi_1(S^1 \times S^1) \cong \mathbb{Z} \times \mathbb{Z}$. Also, $\pi_1(S^2) = 0$ so $m(S^2) = 0$. Furthermore, $m(\mathbb{R}P^2) = 1$. It is a standard fact that $m(\Delta) \geq \beta_1(\Delta)$, so the second conjecture is stronger than the first. Indeed, $m(\Delta) \geq \beta_1(\Delta)$ can be arbitrarily poor. Poincaré spheres are getting at this issue; all homology agrees with spheres, but the fundamental group is non-zero. Suspending the Poincaré sphere twice gives a 5-dimensional manifold.

314 Theorem (Novik-Swartz, 2009; Murai, 2015)
The first conjecture holds.

315 Theorem (Novik-Murai, two weeks ago)
The second conjecture holds (very likely; very new).

316 Definition

A simplicial complex is "flag" roughly if it is defined by its graph. More precisely, a simplicial complex Δ is flag if it is a clique complex of its graph.

317 Example

Barycentric subdivisions are always flag. For instance, take the boundary of a triangle, which is not flag. However, subdividing it adds vertices at the three midpoints and divides each edge in two, and the result is flag. The boundary complex of the three-dimensional cross polytope is also flag.

318 Conjecture (Gal)

If Δ is a flag complex such that $\|\Delta\|$ is a $(d-1)$-dimensional sphere, $d \geq 4$, then $f_1 - (2d-3)f_0 + 2d(d-2) \geq 0$.

319 Remark

By comparison, the lower bound theorem says

$$f_1 - df_0 + \binom{d+1}{2} \geq 0.$$

The bound is known to hold for 3- and 4-spheres, though it is wide open for 5-spheres. One could also be ambitious and replace "sphere" with "manifold."

List of Symbols

$a^{\langle i \rangle}$	Bumped canonical representation, page 41
$f_2^k(P)$	Count of k-gon 2-faces of P, page 58
f_i	Face number, page 5
h	Hyperplane, page 2
$h(\Delta)$	h-vector of a simplicial complex, page 38
h^+	Closed half-space, page 2
h^-	Closed half-space, page 2
$h_i(P)$	h-number of P, page 21
$h_i^\rho(P)$	h-number of P using ρ, page 21
r	Rank of a matroid, page 53

Index