Triangulations Of Point Sets

Applications,

Structures,

Algorithms.

Jesús A. De Loera Jörg Rambau Francisco Santos

MSRI Summer school July 21–31, 2003

(Book under construction)

Triangulations Of Point Sets

Applications,

Structures,

Algorithms.

Jesús Loera Jörg Rambau Francisco Santos

MSRI Summer school July 21–31, 2003

(Book under construction)

Triangulations Of Point Sets

Applications,

Structures,

Algorithms.

Jesús A. De Loera   Jörg Rambau   Francisco Santos

MSRI Summer school July 21–31, 2003

(Book under construction)

Outline of the book/course

July 21: Motivation and fundamental notions (FS)

July 22: Life in two dimensions (JdL)

July 23: Regular triangulations and secondary polytopes (JR)

July 24: Non-regular triangulations (FS)

July 25: A friendly space of triangulations: cyclic polytopes (JR)

July 28: Unfriendly spaces of triangulations (FS)

July 30: Enumeration (JdL)

July 29: Optimization (JR)

July 31: Further selected topics (JdL, FS, JR)

Triangulations

A **polytope** is the convex hull of finitely many points

$$\mathrm{conv}(p_1, \ldots, p_n) := \{\sum \alpha_i p_i : \alpha_i \geq 0 \ \forall i = 1, \ldots, n, \sum \alpha_i = 1\}$$

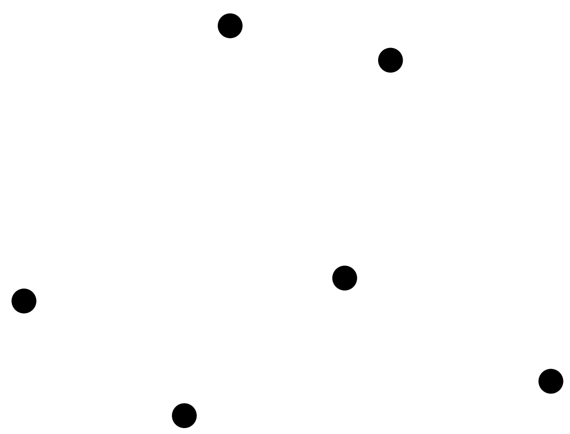

A finite point set

Triangulations

A **polytope** is the convex hull of finitely many points

$$\mathrm{conv}(p_1, \ldots, p_n) := \{\sum \alpha_i p_i : \alpha_i \geq 0 \ \forall i = 1, \ldots, n, \sum \alpha_i = 1\}$$

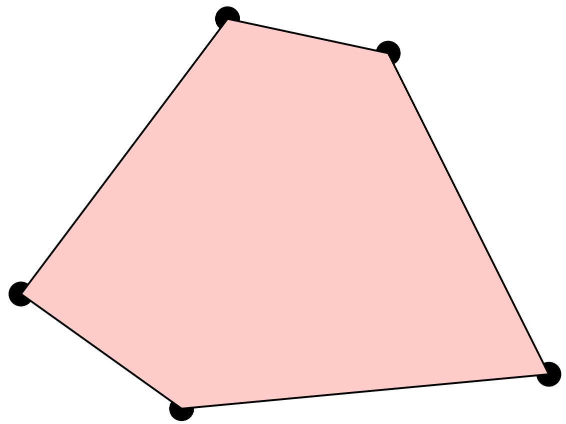

Its convex hull

Triangulations

A **triangulation** is a partition of the convex hull into **simplices** such that

The union of all these simplices equals $\mathrm{conv}(A)$. (Union Property.)

Any pair of them intersects in a (possibly empty) common face. (Intersec. Prop.)

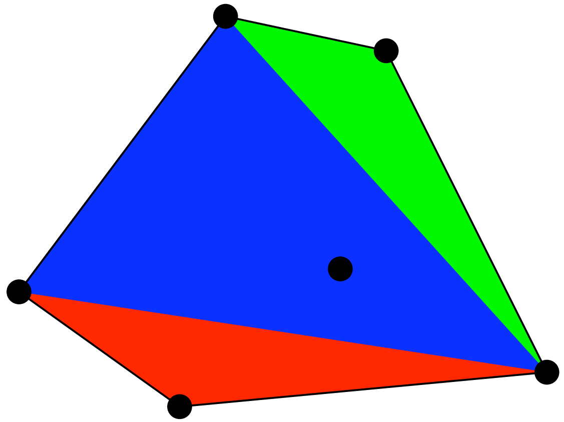

A triangulation of P

Triangulations

The following are **not** triangulations:

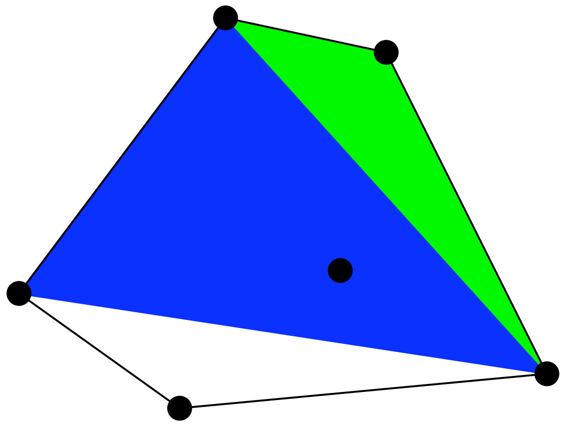

The union is not the whole convex hull

Triangulations

The following are **not** triangulations:

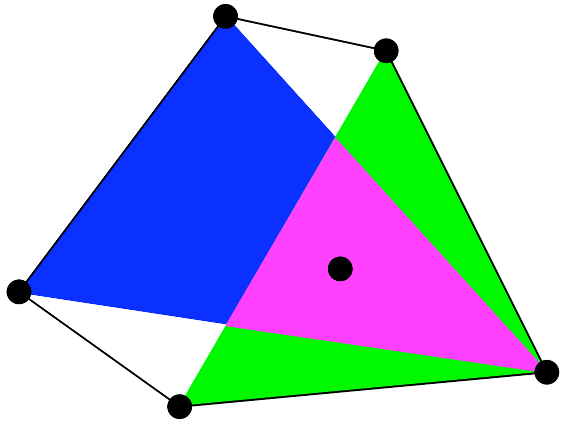

The intersection is not okay

Triangulations

Recall that a **simplex** is the convex hull of any set of **affinely independent points**. Equivalently, it is any polytope of dimension d with $d+1$ vertices.

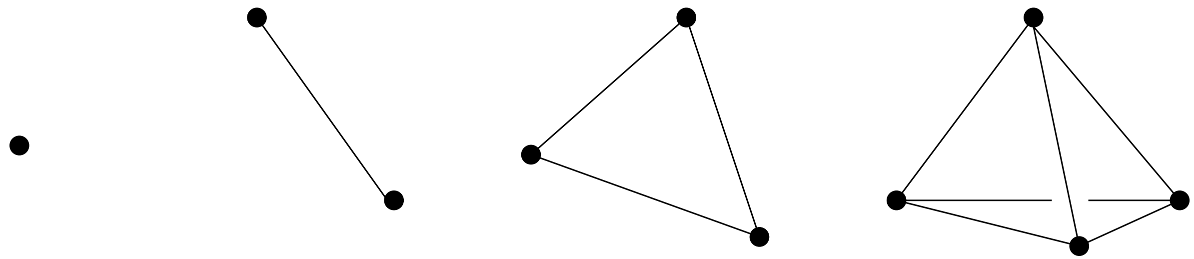

0–simplex 1–simplex 2–simplex 3–simplex

The simplest simplices

A d-simplex has exactly $\binom{d+1}{i+1}$ faces of dimension i, $(i = -1, 0, \ldots, d)$, which are themselves i-simplices.

Triangulations of a point configuration

A **point configuration** is a finite set of points in $\mathbb{R}^d$, possibly with repetitions.

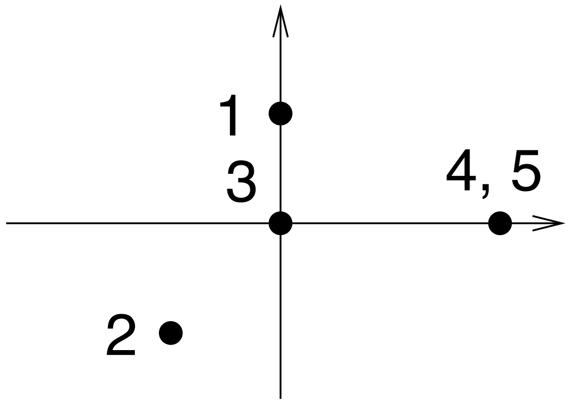

$a_1 = (0, 1)$

$a_2 = (-1, -1)$

$a_3 = (0, 0)$

$a_4 = a_5 = (2, 0)$

A point set with repetitions

Triangulations of a point configuration

A triangulation of a point set $\mathcal{A}$ is a triangulation of $\operatorname{conv}\mathcal{A}$ with **vertex set contained in** $\mathcal{A}$.

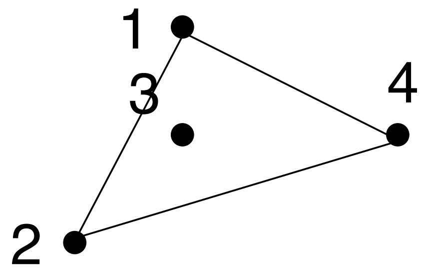

124

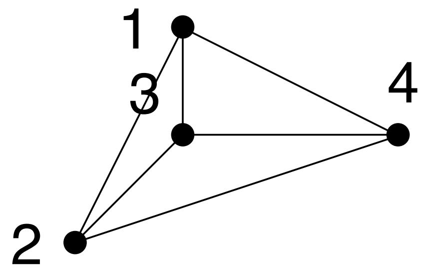

123, 134, 234

The two triangulations of $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$

Remark: Don't need to use all points

Triangulations of a point configuration

A triangulation of a point set $\mathcal{A}$ is a triangulation of $\operatorname{conv} \mathcal{A}$ with **vertex set contained in** $\mathcal{A}$.

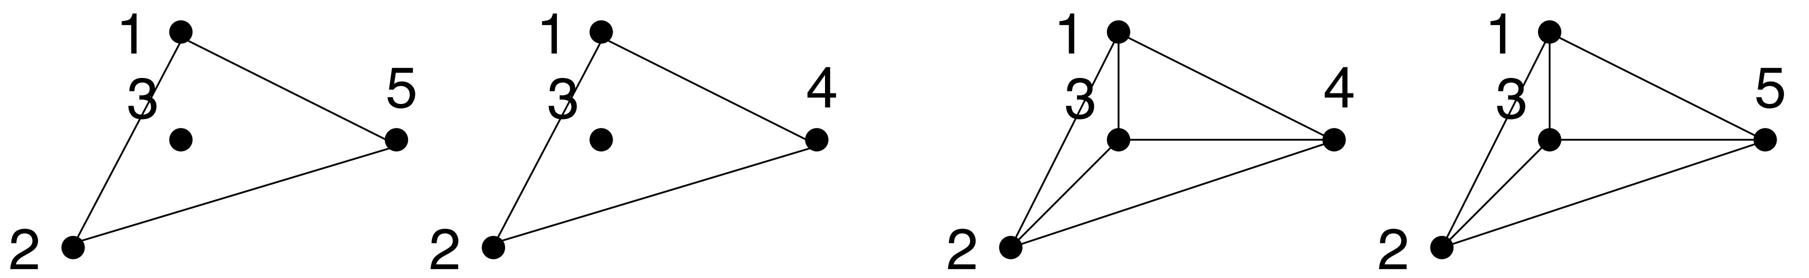

125 | 124 | 123, 134, 234 | 123, 135, 235

The four triangulations of $\mathcal{A} = \{a_1, a_2, a_3, a_4, a_5\}$

Remark: Repeated points count!

Why study triangulations?

Why study triangulations?

. . . well, we expect to have you convinced by the end of the course.

Why study triangulations?

The main intuition is that subdividing a geometric object into simple pieces allows to do (several types of) things "piece by piece".

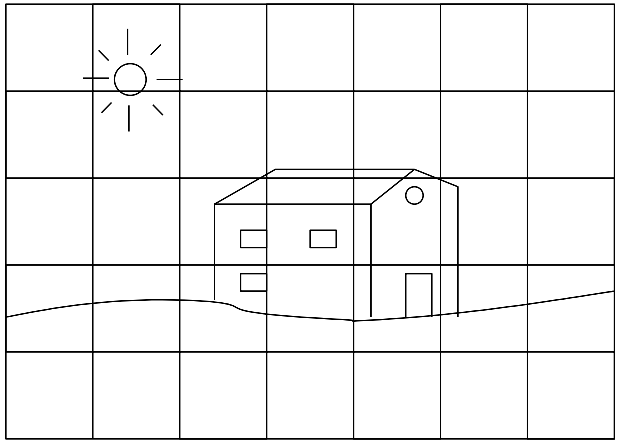 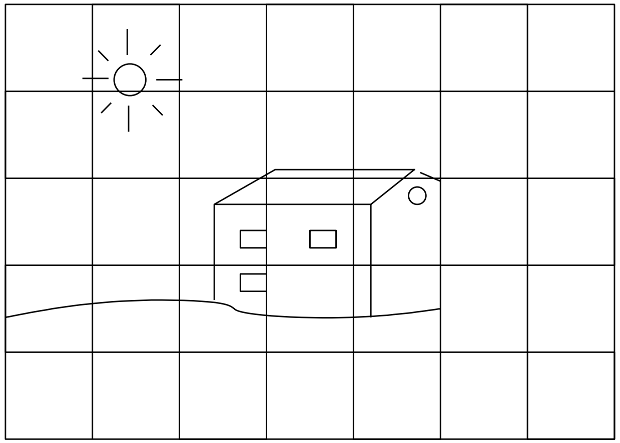

And, actually, triangulations arise in several parts of mathematics.

Triangulations and combinatorics

Triangulations and combinatorics

1. Triangulations of a convex polygon

2. The order polytope of a poset

Triangulations of a convex n-gon

Triangulations of a convex n-gon

To triangulate the n-gon, you just need to insert $n-3$ non-crossing diagonals:

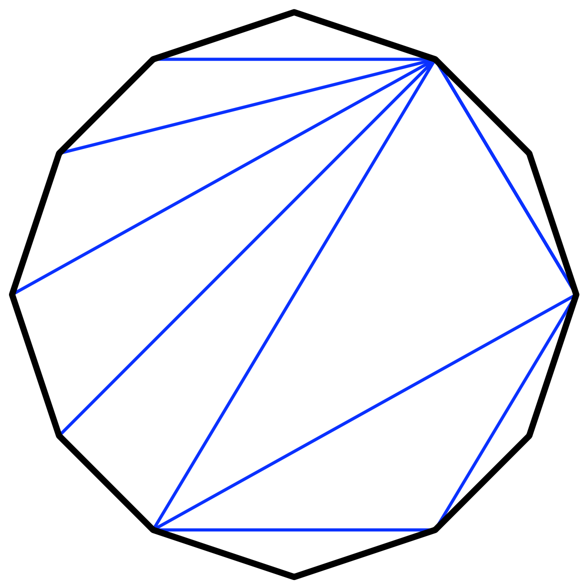

A triangulation of the 12-gon

Triangulations of a convex n-gon

To triangulate the n-gon, you just need to insert $n-3$ non-crossing diagonals:

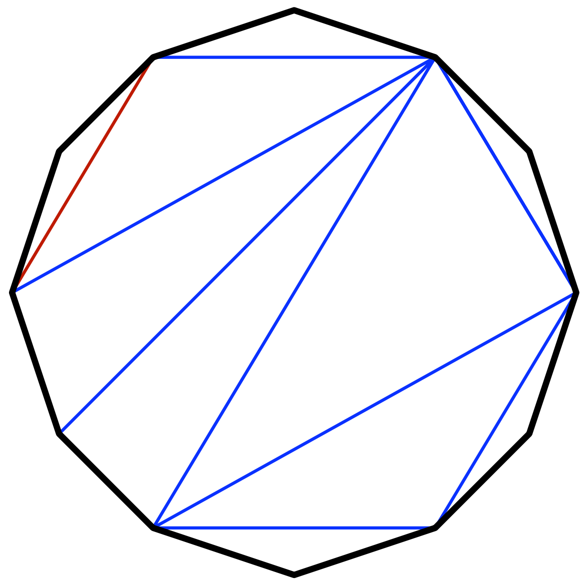

Another triangulation of the 12-gon, obtained by **flipping** an edge

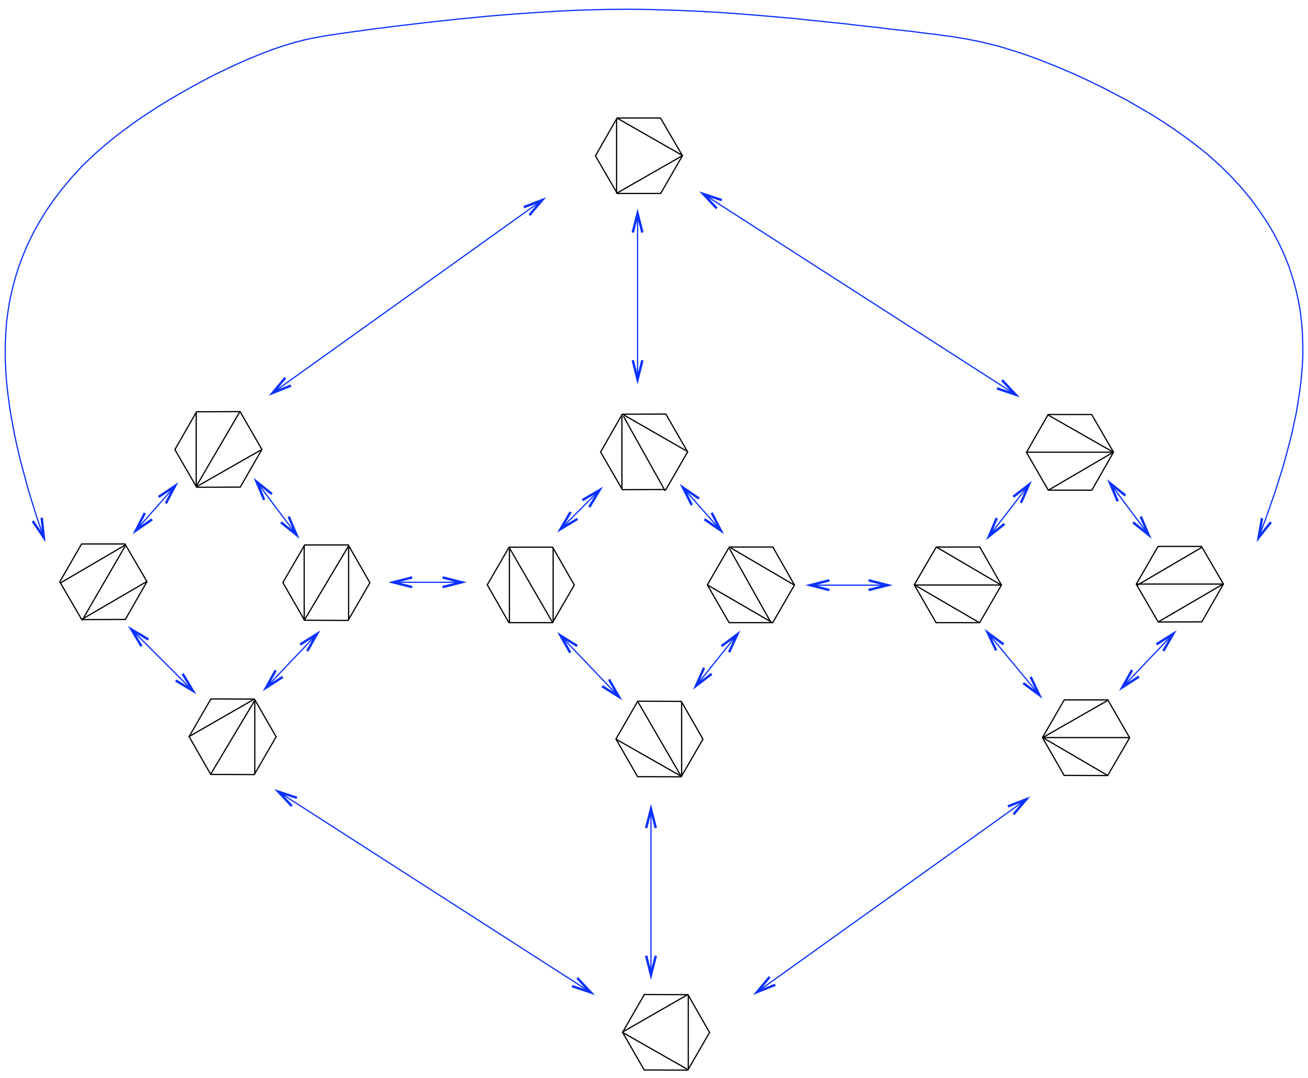

The Graph of flips for a hexagon

Some obvious properties of triangulations and flips of an n-gon

- The graph is regular of degree $n - 3$.
- The graph has dihedral symmetry.

Some non-obvious properties of triangulations and flips of an n-gon

- It is the graph of a polytope of dimension $n-3$, called **the associahedron**. (and we'll see why on Wednesday).

- The graph has diameter bounded above by $2n-10$ for all n (and equal to that number for large n).

- There are exactly $\frac{1}{n-1}\binom{2n-4}{n-2}$ triangulations (we'll see why tomorrow).
  That is to say, the **Catalan number** C_{n-2}:

  $C_n := \frac{1}{n+1}\binom{2n}{n}$,

n	0	1	2	3	4	5	6
C_n	1	1	2	5	14	42	132

The Catalan number C_n not only counts the triangulations of a $n+2$-gon:

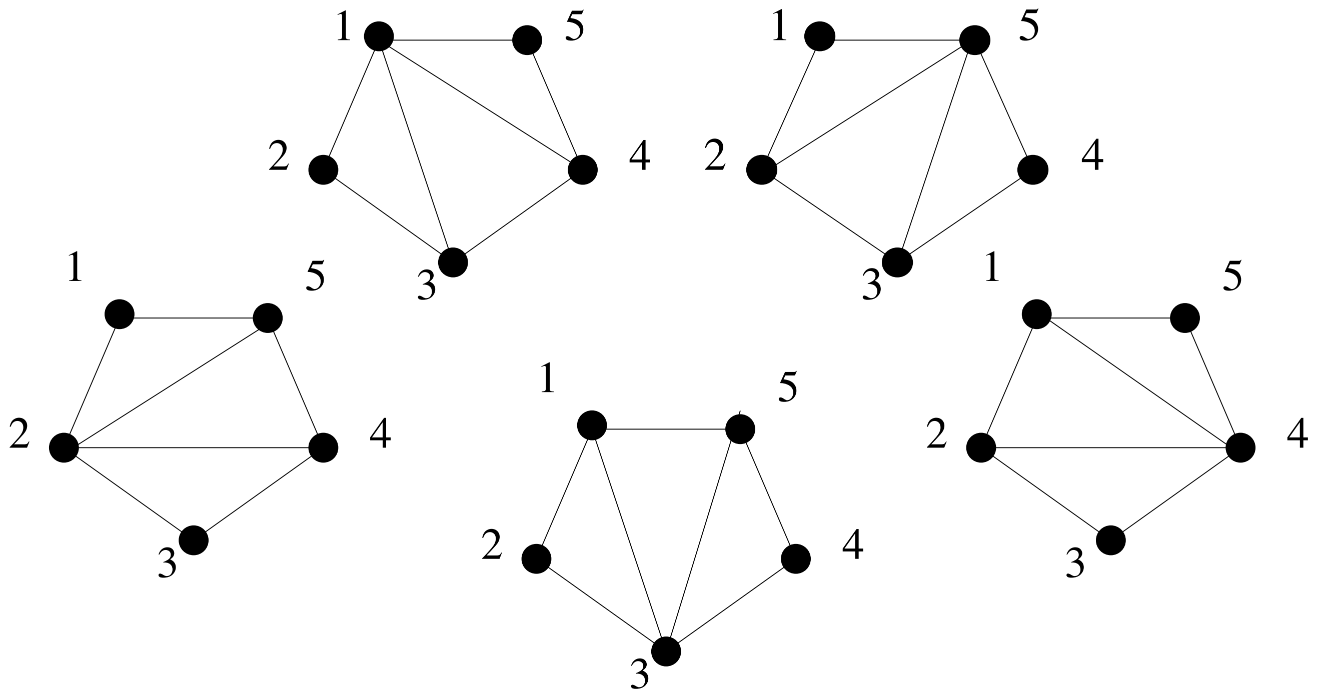

It also counts. . .

1. Binary trees on n-nodes.

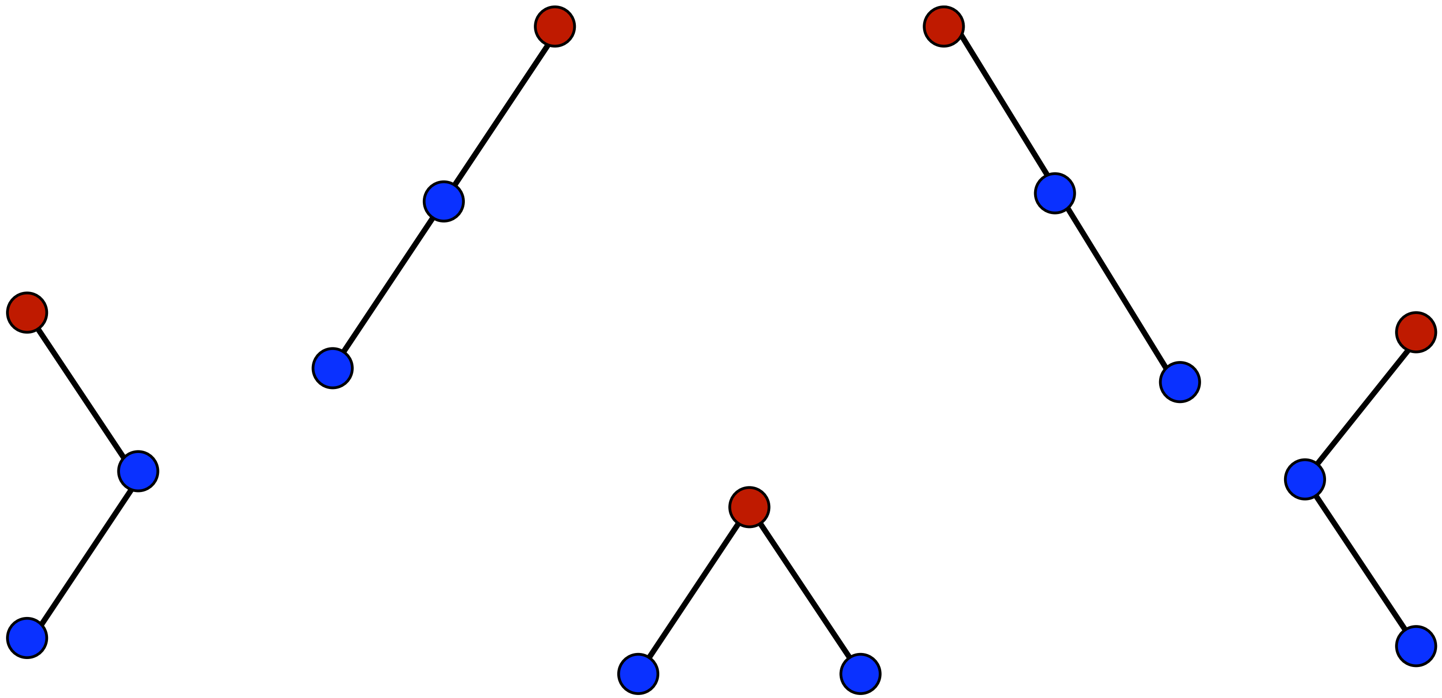

1. Binary trees on n-nodes.

2. Monotone lower-diagonal lattice (integer) paths from $(0,0)$ to (n,n).

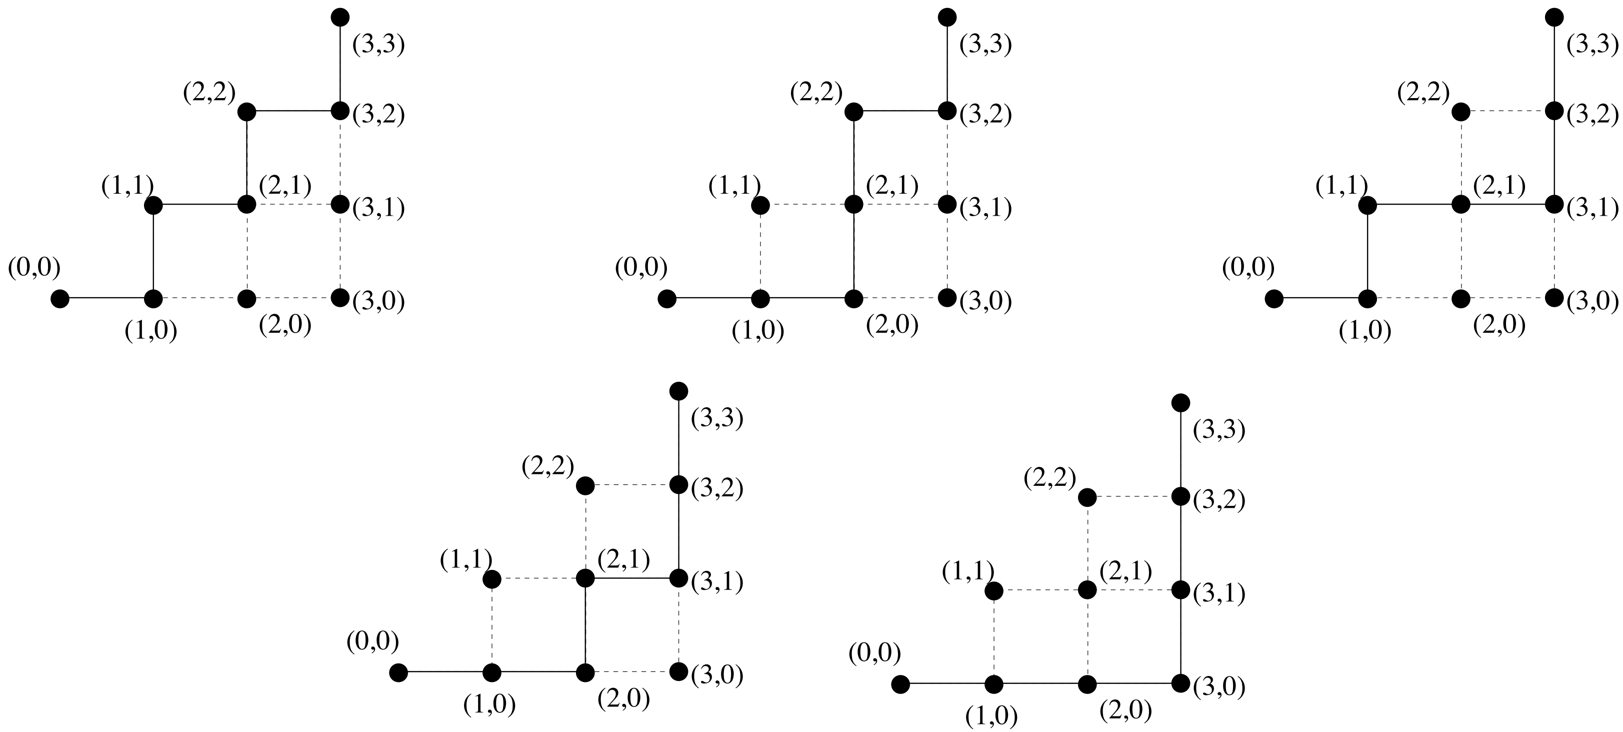

1. Binary trees on n-nodes.

2. Monotone lower-diagonal lattice (integer) paths from $(0, 0)$ to (n, n).

3. Sequences of $2n$ signs with exactly n of each and with more $+$'s than $-$'s in every initial segment.

$$+ - + - + - \qquad + + - - + - \qquad + - + + -$$

$$+ - + - + - \qquad + + + - - -$$

1. Binary trees on n-nodes.

2. Monotone lower-diagonal lattice (integer) paths from $(0, 0)$ to (n, n).

3. Sequences of $2n$ signs with exactly n of each and with more $+$'s than $-$'s in every initial segment.

. . . and **some other 60 combinatorial structures**,

according to Exercise 6.19 in

R. Stanley, *Enumerative combinatorics*, Cambridge University Press, 1999.

Triangulations and partially ordered sets

Triangulations and partially ordered sets

1. A partially ordered set (or poset) is a finite set P with a relation $<$ that is reflexive, antisymmetric, and transitive.

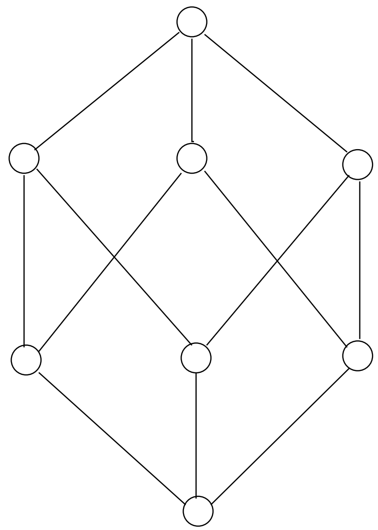

The Hasse diagram of the Boolean poset on 3 elements

1. A partially ordered set (or poset) is a finite set P with a relation $<$ that is reflexive, antisymmetric, and transitive.

2. A linear extension of a poset on n vertices is a bijection λ from the set of vertices of P to $\{1, \ldots, n\}$ such that $\lambda(x) < \lambda(y)$ whenever $x < y$ in P. In other words, it is a total order compatible with $<$.

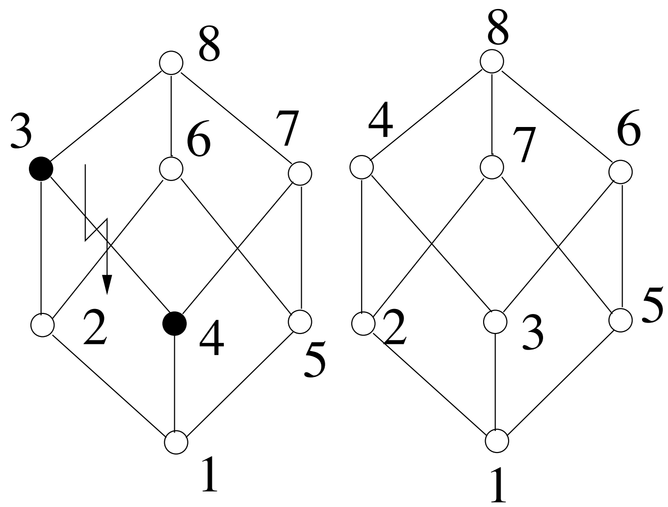

1. A partially ordered set (or poset) is a finite set P with a relation $<$ that is reflexive, antisymmetric, and transitive.

2. A linear extension of a poset on n vertices is a bijection λ from the set of vertices of P to $\{1, \ldots, n\}$ such that $\lambda(x) < \lambda(y)$ whenever $x < y$ in P.

3. A (lower) order ideal is a subset such that if $x \in I$ and $y < x$ then $y \in I$.

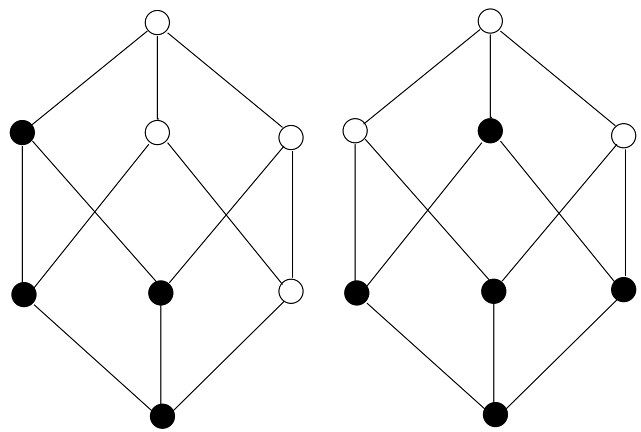

The order polytope

Given a poset P with elements $p_1, \ldots, p_n$, we define its order polytope $O(P)$ in $\mathbb{R}^n$ by the following linear constraints:

$$O(P) = \{\mathbf{x} \in [0,1]^n : \ x_i \geq x_j \text{ if } p_i > p_j \text{ in } P\}.$$

The order polytope

$$O(P) = \{\mathbf{x} \in [0,1]^n : \; x_i \geq x_j \text{ if } p_i > p_j \text{ in } P\}.$$

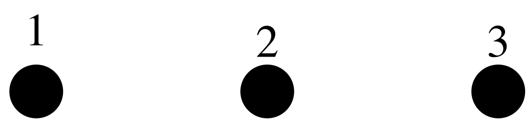

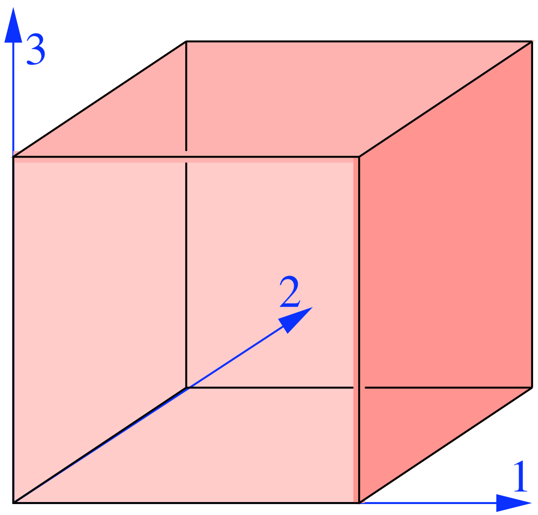

The order polytope of an antichain with n elements is the whole n-cube $[0,1]^n$

The order polytope

$$O(P) = \{\mathbf{x} \in [0,1]^n : \; x_i \geq x_j \text{ if } p_i > p_j \text{ in } P\}.$$

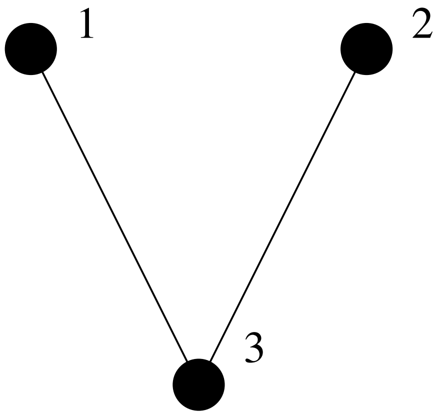

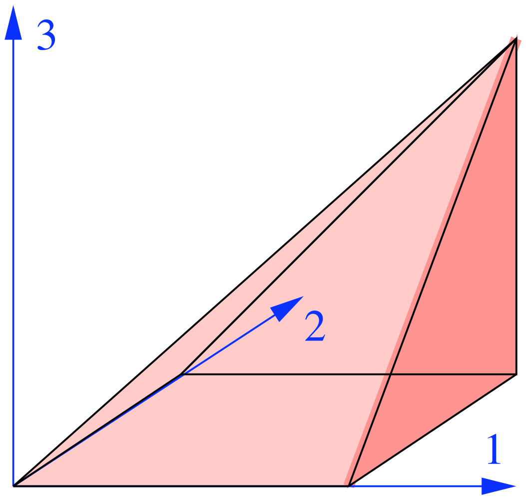

The order polytope of this is a pyramid

The order polytope

$$O(P) = \{\mathbf{x} \in [0,1]^n : \ x_i \geq x_j \text{ if } p_i > p_j \text{ in } P\}.$$

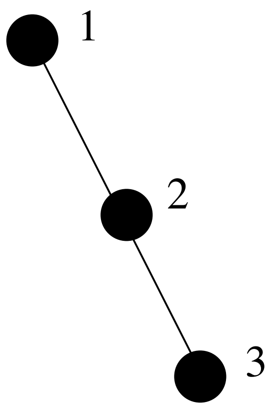

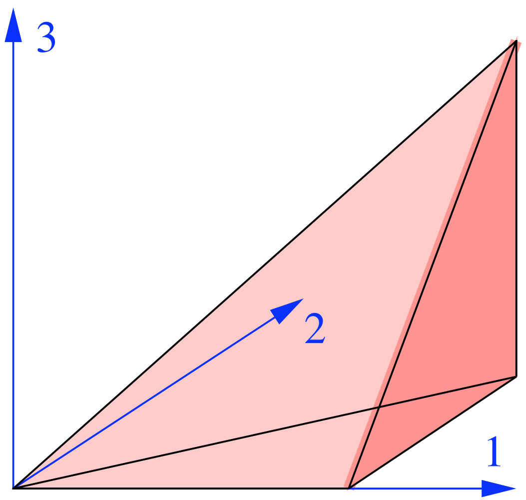

The order polytope of a chain of n elements is an n-simplex
(and a unimodular one)

The order polytope

$$O(P) = \{\mathbf{x} \in [0,1]^n :\ x_i \geq x_j \text{ if } p_i > p_j \text{ in } P\}.$$

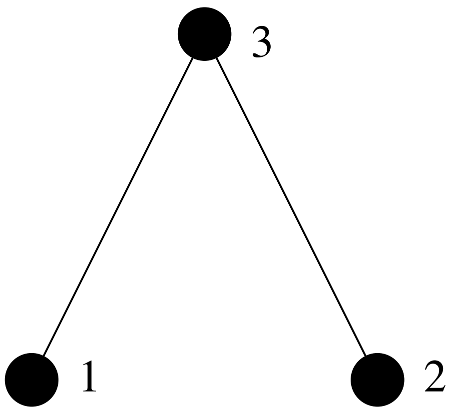

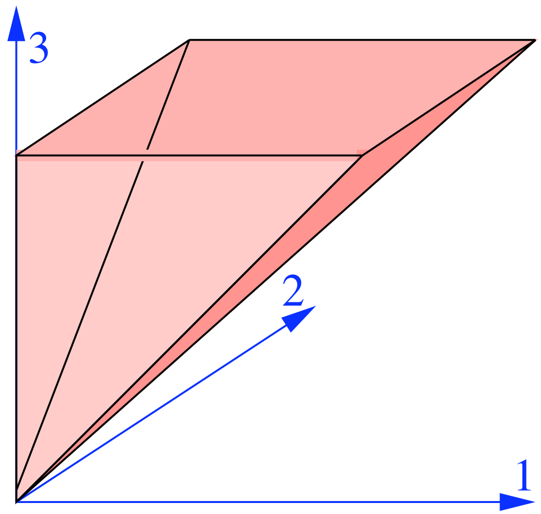

The order polytope of this is a pyramid upside-down

The order polytope

Theorem: The following hold for the order polytope $O(P)$ of a poset P:

1. The vertices of $O(P)$ are **integer** and are in bijection with the **order ideals** of the poset P ($x_i = 0 \Leftrightarrow p_i$ is in the ideal).

2. The order polytopes (simplices) of the different linear extensions of P **form a (unimodular) triangulation** of the polytope $O(P)$.

The order polytope

Theorem: The following hold for the order polytope $O(P)$ of a poset P:

1. The vertices of $O(P)$ are **integer** and are in bijection with the **order ideals** of the poset P ($x_i = 0 \Leftrightarrow p_i$ is in the ideal).

2. The order polytopes (simplices) of the different linear extensions of P **form a (unimodular) triangulation** of the polytope $O(P)$.

Corollary: The **volume** of the order polytope $O(P)$ is $1/n!$ times the number of linear extensions of the poset P.

The order polytope

Corollary: It is $\#P$-hard to compute the volume of a polytope given by its facets. (Even if we assume that its vertices have all coordinates 0 or 1).

Proof: Brightwell and Winkler proved in 1992 that it is $\#P$-hard to compute the number of linear extensions of a poset.

Triangulations and optimization

Triangulations and optimization

1. Parametric linear programming

2. Sperner Lemma

Triangulations and optimization

1. Parametric linear programming

2. Sperner Lemma

3. (Voronoi diagrams and Delaunay triangulations)

Triangulations of vector sets

Triangulations of vector sets

Let $A = \{a_1, \dots, a_n\}$ be a finite set of real **vectors** (a vector configuration).

The cone of A is $\quad \mathrm{cone}(A) := \{\sum \lambda_i a_i : \alpha_i \geq 0, \forall i = 1, \dots, n\}$

Two vector configurations, and their cones

A simplicial cone is one generated by linearly independent vectors.

Triangulations of vector sets

A triangulation of a vector configuration A is a partition of $\text{cone}(A)$ into simplicial cones with generators contained in A and such that:

(UP) The union of all these simplices equals $\text{conv}(A)$. (Union Property.)

(IP) Any pair of them intersects in a common face (Intersection Property.)

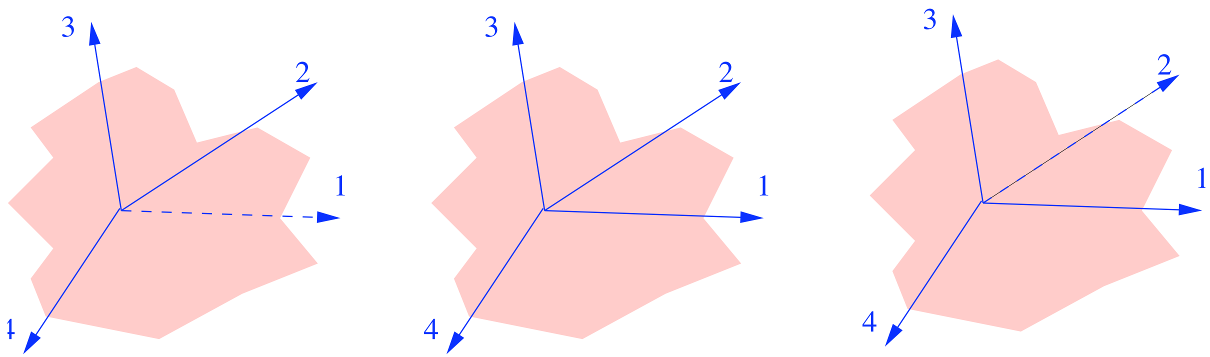

The three triangulations of the first configuration

A cone is pointed if it is contained (except for the origin) in an **open** half-space. If this happens for $\mathrm{cone}(A)$, then A is called acyclic.

Remark: Triangulations of a $\{$pointed/acyclic$\}$ $\{$cone/vector set$\}$ of dimension d are the same as the triangulations of the $\{$polytope/point$\}$ set of dimension $d-1$ obtained cutting by an affine hyperplane:

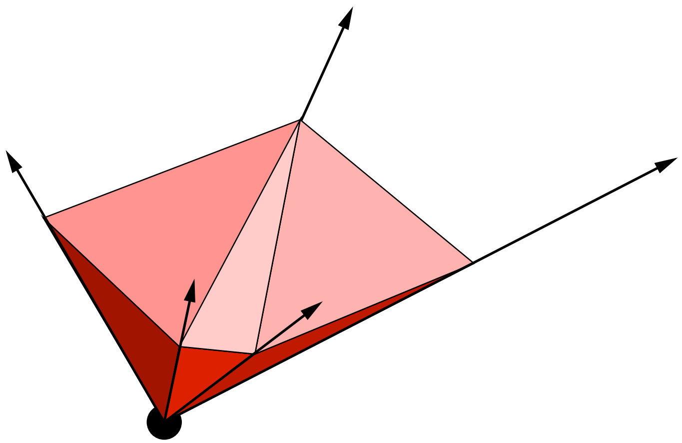

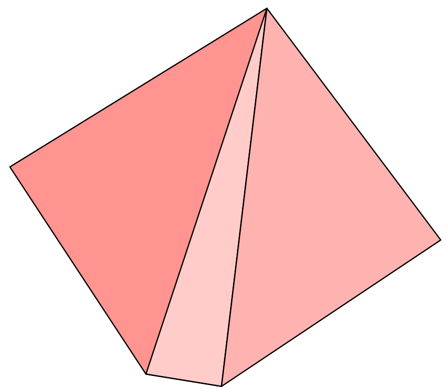

Linear programming

Linear programming

Let $A = (a_1, \ldots, a_n) \in \mathbb{R}^{d \times n}$ be a matrix. Let $b \in \mathbb{R}^d$ and $c \in \mathbb{R}^n$. To this data one associates the linear programming problem $LP_{A,c}(b) := \min\{c(x) : Ax = b,\ x \geq 0\}$:

Linear programming

Let $A = (a_1, \ldots, a_n) \in \mathbb{R}^{d \times n}$ be a matrix. Let $b \in \mathbb{R}^d$ and $c \in \mathbb{R}^n$. To this data one associates the linear programming problem $LP_{A,c}(b) := \min\{c(x) : Ax = b,\ x \geq 0\}$:

"minimize the cost function $c(x)$

subject to $Ax = b$ and $x \geq 0$"

We say that the linear program $LP_{A,c}(b)$ is feasible if $\{x \in \mathbb{R}^n : Ax = b\}$ is not empty. It is bounded if c has a lower bound in $\{x \in \mathbb{R}^n : Ax = b\}$.

Linear programming

Example: $A = \begin{pmatrix} 0 & 3 & 1 \\ 2 & 1 & 0 \end{pmatrix}$, $b = \begin{pmatrix} 6 \\ 4 \end{pmatrix}$. Then,

$$Ax = b \quad \Leftrightarrow \quad x = (1, 2, 0) + \lambda(1, -2, 6)$$

The linear program is feasible ($(1, 2, 0)$ is a feasible solution). It is bounded for every c, because big and small values of λ will make some coordinate of x negative.

Remarks:

- feasible $\quad \Leftrightarrow \quad b \in \text{cone}(A) := \text{cone}(\{a_1, \ldots, a_n\}).$

Remarks:

- feasible $\Leftrightarrow$ $b \in \text{cone}(A) := \text{cone}(\{a_1, \ldots, a_n\})$.
- bounded for every c $\Leftrightarrow$ $\ker(A) \cap \mathbb{R}^n_{\geq 0} = \{0\}$ $\Leftrightarrow$ $\text{cone}(A)$ is pointed $\Leftrightarrow$ bounded for $c = (-1, \ldots, -1)$.

Remarks:

- feasible $\quad\Leftrightarrow\quad b \in \mathrm{cone}(A) := \mathrm{cone}(\{a_1, \ldots, a_n\})$.
- bounded for every $c \quad\Leftrightarrow\quad \ker(A) \cap \mathbb{R}^n_{\geq 0} = \{0\} \quad\Leftrightarrow\quad \mathrm{cone}(A)$ is pointed $\Leftrightarrow$ bounded for $c = (-1, \ldots, -1)$.
- if b and c are generic, there is (at most) one optimal solution. In this case, the optimal solution has d non-zero coordinates and the corresponding columns of A form a basis of $\mathrm{cone}(A)$. They are called the optimal basis of $LP_{A,c}(b)$.

Remarks:

- feasible $\Leftrightarrow$ $b \in \text{cone}(A) := \text{cone}(\{a_1, \dots, a_n\})$.
- bounded for every c $\Leftrightarrow$ $\ker(A) \cap \mathbb{R}^n_{\geq 0} = \{0\}$ $\Leftrightarrow$ $\text{cone}(A)$ is pointed $\Leftrightarrow$ bounded for $c = (-1, \dots, -1)$.
- if b and c are generic, there is (at most) one optimal solution. In this case, the optimal solution has d non-zero coordinates and the corresponding columns of A form a basis of $\text{cone}(A)$. They are called the optimal basis of $LP_{A,c}(b)$.
- if we knew the optimal basis σ, we could find the optimal solution by just solving a linear system of equations:

$$Ax = b, \qquad \text{and} \qquad x_i = 0 \quad \forall i \notin \sigma.$$

Parametric linear programming

Parametric linear programming

Let us study how the previous linear program depends on the right hand side b. That is, study the family of linear programs

$$LP_{A,c} = \{LP_{A,c}(b) : b \in \mathrm{cone}(A)\}$$

Question: How does the optimal basis depend on b?

Parametric linear programming

Theorem (Walkup-Wets 1969) Let $LP_{A,c}(b)$ denote the linear program

$$\min\{cx : Ax = b, x \geq 0\},$$

where c and A are fixed.

Then, there exists a triangulation T of $\mathrm{cone}(A)$ such that the optimal basis of $LP_{A,c}(b)$ for each $b \in \mathrm{cone}(A)$ is precisely the (generators of) the simplicial cone $\mathrm{cone}(\sigma)$ with $\sigma \in T$ and $b \in \mathrm{cone}(\sigma)$.

Idea of proof: Consider the lifted vector configuration $\tilde{A} = \begin{pmatrix} a_1 & \cdots & a_n \\ c_1 & \cdots & c_n \end{pmatrix} \subset \mathbb{R}^{d+1}$. The triangulation of A in question is the lower envelope of $\mathrm{cone}(\tilde{A})$.

Parametric linear programming

Example: $A = \begin{pmatrix} 0 & 3 & 1 \\ 2 & 1 & 0 \end{pmatrix}, \qquad b = \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}, \qquad c = \begin{pmatrix} c_1 & c_2 & c_3 \end{pmatrix}.$

$$Ax = b \quad \Leftrightarrow \quad x = x_0 + \lambda(1, -2, 6)$$

Then:

• if $c_1 - 2c_2 + 6c_3 \geq 0$, then the optimal basis is $(*, *, 0)$ or $(0, *, *)$, and this happens depending on whether $b \in \mathrm{cone}(a_1, a_2)$ or $b \in \mathrm{cone}(a_2, a_3)$

• if $c_1 - 2c_2 + 6c_3 \leq 0$, then the optimal basis is $(*, 0, *)$ for every $b \in \mathrm{cone}(a_1, a_2, a_3)$

Regular triangulations

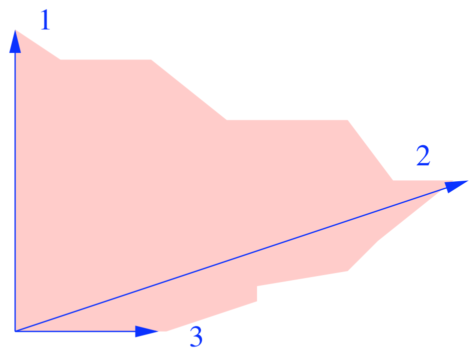

$c_2 \leq \frac{c_1+6c_3}{2}$

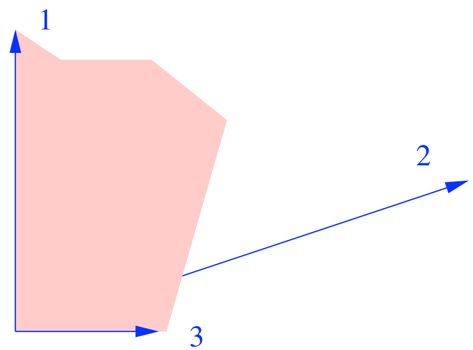

$c_2 \geq \frac{c_1+6c_3}{2}$

Remark: Different c's may provide different triangulations. But, for some A's, **not all triangulations can be obtained in this way.**

Regular triangulations

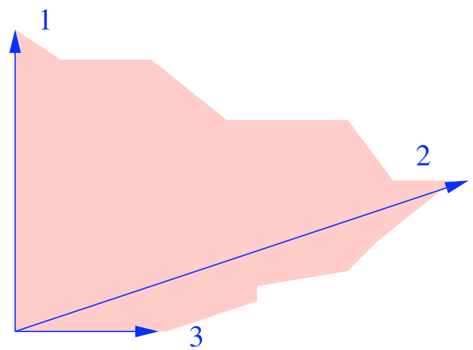

$c_2 \leq \frac{c_1+6c_3}{2}$

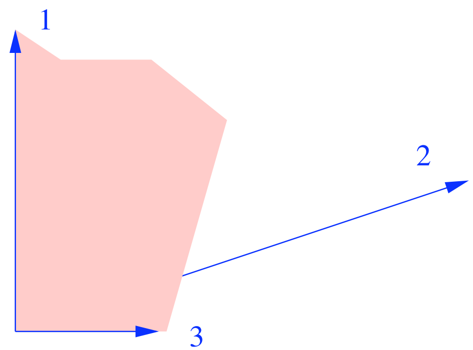

$c_2 \geq \frac{c_1+6c_3}{2}$

Remark: Different c's may provide different triangulations. But, for some A's, **not all triangulations can be obtained in this way**.

The triangulations that **can** be obtained like this are called regular.

$$c = (\ 0,\ 0,\ 0, -5, -4, -3), \qquad A = \begin{bmatrix} 4 & 0 & 0 & 2 & 1 & 1 \\ 0 & 4 & 0 & 1 & 2 & 1 \\ 0 & 0 & 4 & 1 & 1 & 2 \end{bmatrix},$$

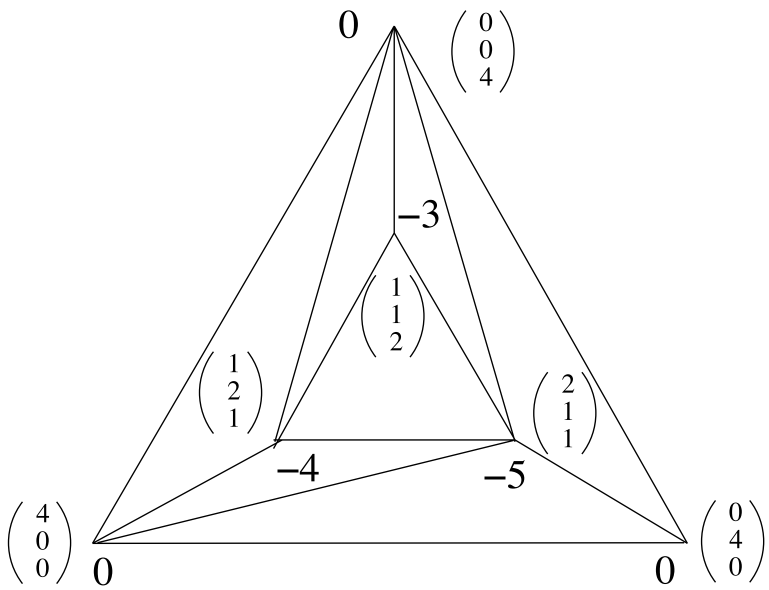

The cone triangulation associated with the cost vector c. This shows a two-dimensional slice of the cone.

$$c = (\ ?,\ ?,\ \ ?,\ \ ?,\ \ ?,\ \ ?), \qquad A = \begin{bmatrix} 4 & 0 & 0 & 2 & 1 & 1 \\ 0 & 4 & 0 & 1 & 2 & 1 \\ 0 & 0 & 4 & 1 & 1 & 2 \end{bmatrix},$$

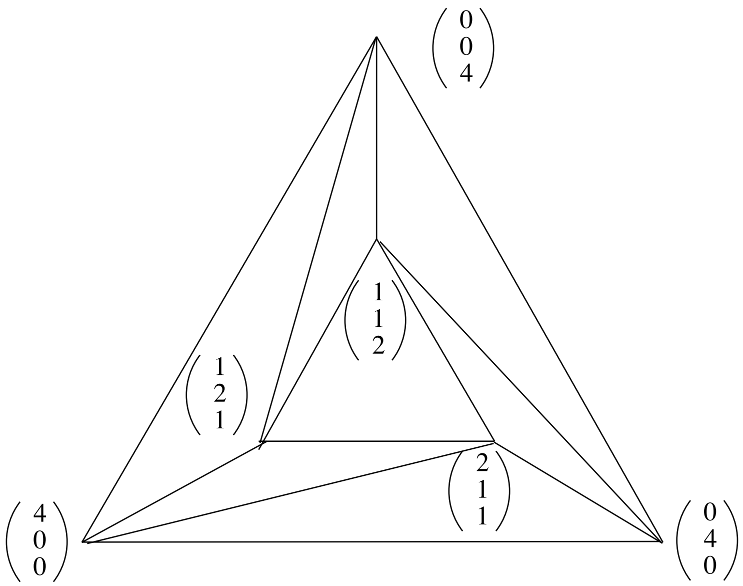

A triangulation not associated with any cost vector c. That is to say, a non-regular triangulation.

Parametric linear programming (cont.)

Let us get back to the linear programs $LP_{A,c}(b)$, for a fixed matrix A. But suppose that now c varies, too. By the previous theorem, each value of c will provide a different triangulation of $\mathrm{cone}(A)$.

Question: What values $b, b' \in \mathrm{cone}(A)$ are guaranteed to provide the same optimal solution of $LP_{A,c}(b)$ no matter what c?

Answer: clearly, those which are contained in exactly the same bases of A. That is to say, those in the same chamber of the chamber complex of A.

The chamber complex

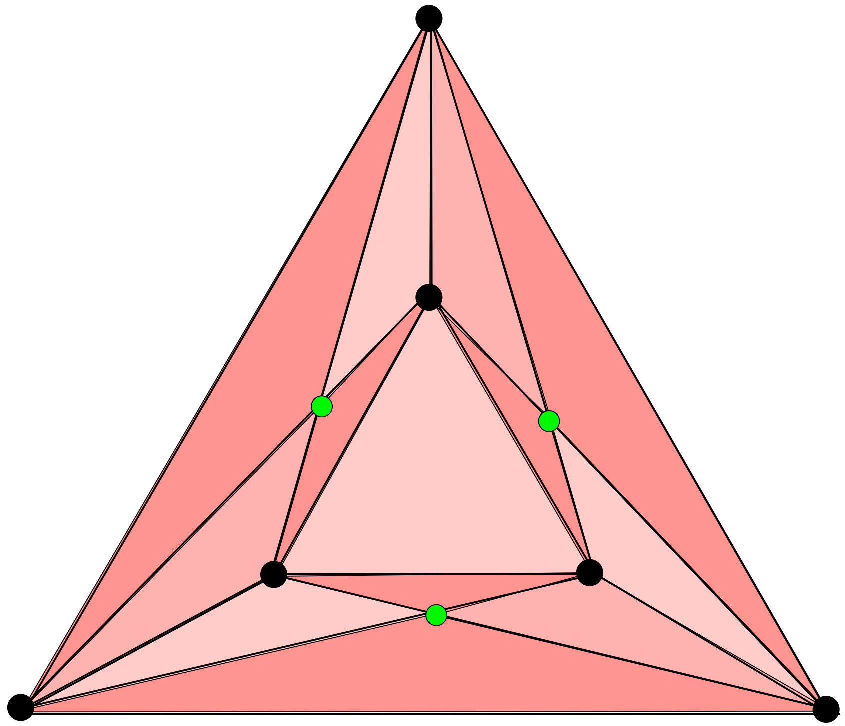

The chamber complex of $\mathrm{cone}(A)$.

. . . curiously enough:

Theorem (Billera, Filliman, Sturmfels 1990) For any vector configuration A there is another vector configuration A^* (its Gale transform) such that the chambers of A correspond to regular triangulations of A^* and viceversa

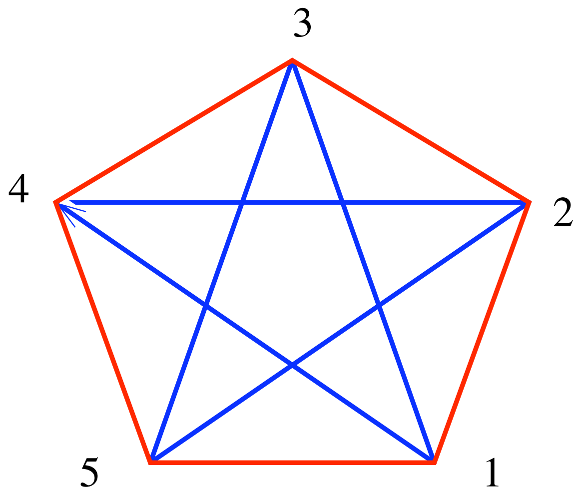

5 regular triangulations
11 chambers

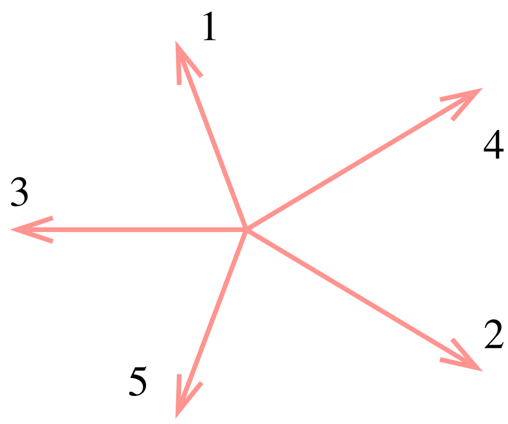

11 regular triangulations
5 chambers

Sperner's lemma and fixed points

Sperner's lemma and fixed points

Lemma (Sperner) Let A be a point configuration whose convex hull is a d-dimensional simplex Δ and let T be a triangulation of A. Let $\Delta_1, \ldots, \Delta_{d+1}$ denote the $d+1$ facets $\Delta_1, \ldots, \Delta_{d+1}$ in the simplex Δ.

Label all the vertices of T using the numbers $1, 2, \ldots, d+1$ in such a way that no vertex that lies on the facet Δ_i receives the label i.

Then there is a simplex in T whose vertices carry all the different $d+1$ labels.

Sperner's lemma and fixed points

Proof By induction on the dimension: start with a fully labeled simplex of one dimension less in the boundary; then dive into the big simplex until you find a fully labeled simplex in the triangulation.

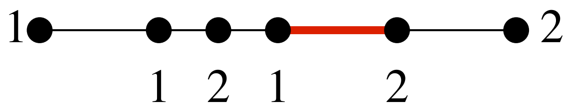

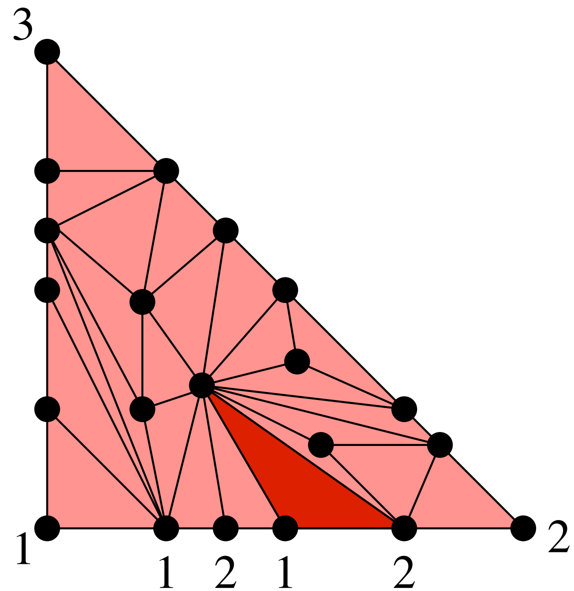

Sperner's lemma and fixed points

Corollary (Brower's fixed point theorem) If C is a topological d-dimensional ball and $f : C \to C$ is a continuous map, then there is a point in C such that $f(x) = x$.

Proof: For any given triangulation T, Sperner Lemma allows you to find a simplex in which the i-th barycentric coordinate of the i-th vertex does does not increase. Doing this for finer and finer triangulations, converges to a fixed point.

Sperner's lemma and fixed points

The algorithmic performance of Sperner Lemma depends heavily on the size (number of simplices) of your triangulations. This raises the question of what is the smallest size of a triangulation. Unfortunately, this is a hard problem:

Theorem (Below, de Loera, Richter-Gebert, 2000) It is NP-coplete to compute the smallest size triangulation of a polytope, even in dimension 3.

Remark Even for the d dimensional cube I^d, the smallest size triangulation has only been computed up to $d = 7$, and the asymptotics of the minimum size of a triangulations is not known.

Triangulations and algebra

Triangulations and algebra

1. Real algebraic varieties and Viro's Theorem

2. Computing and counting zeroes via triangulations

Hilbert's sixteenth problem (1900)

Hilbert's sixteenth problem (1900)

"What are the possible (topological) types of non-singular real algebraic curves of a given degree d?"

Observation: Each connected component is either a pseudo-line or an oval. A curve contains one or zero pseudo-lines depending in its parity.

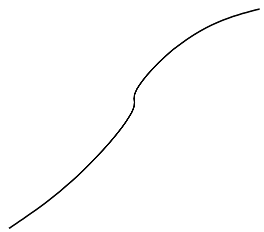

A pseudoline. Its complement has one component, homeomorphic to an open circle. The picture only shows the "affine part"; think the two ends as meeting at infinity.

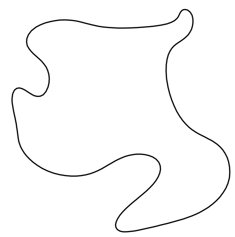

An oval. Its interior is a (topological) circle and and its exterior is a Möbius band.

Partial answers:

Bezout's Theorem: A curve of degree d cuts every line in at most d points. In particular, there cannot be nestings of depth greater than $\lfloor d/2 \rfloor$

Harnack's Theorem: A curve of degree d cannot have more than $\binom{d-1}{2} + 1$ connected components (recall that $\binom{d-1}{2}$= genus)

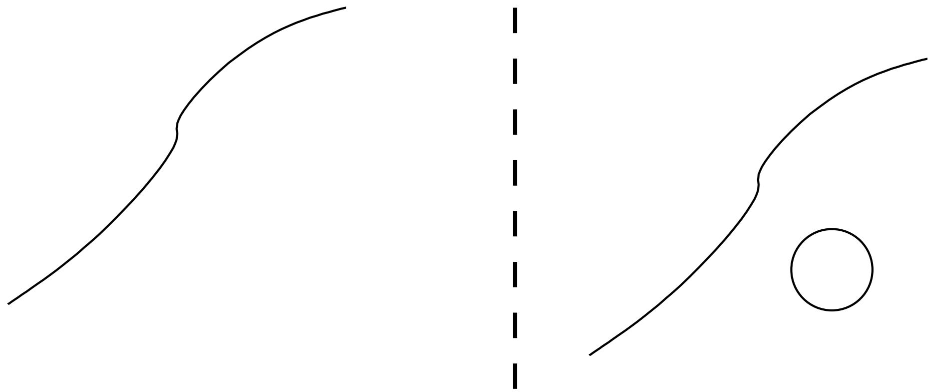

Two configurations are possible in degree 3

Partial answers:

Bezout's Theorem: A curve of degree d cuts every line in at most d points. In particular, there cannot be nestings of depth greater than $\lfloor d/2 \rfloor$

Harnack's Theorem: A curve of degree d cannot have more than $\binom{d-1}{2} + 1$ connected components (recall that $\binom{d-1}{2} =$ genus)

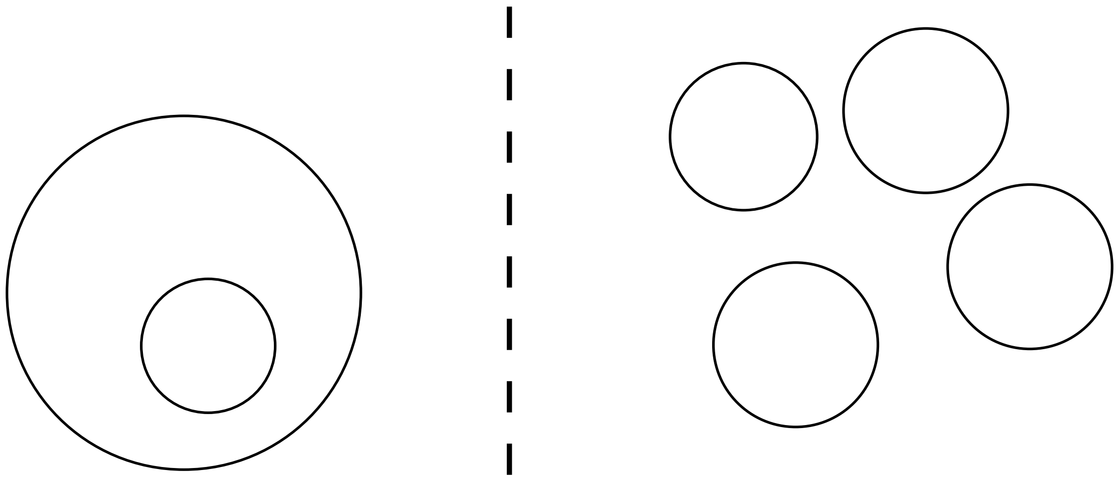

Six configurations are possible in degree 4. Only the maximal ones are shown.

Partial answers:

Bezout's Theorem: A curve of degree d cuts every line in at most d points. In particular, there cannot be nestings of depth greater than $\lfloor d/2 \rfloor$

Harnack's Theorem: A curve of degree d cannot have more than $\binom{d-1}{2} + 1$ connected components (recall that $\binom{d-1}{2} =$ genus)

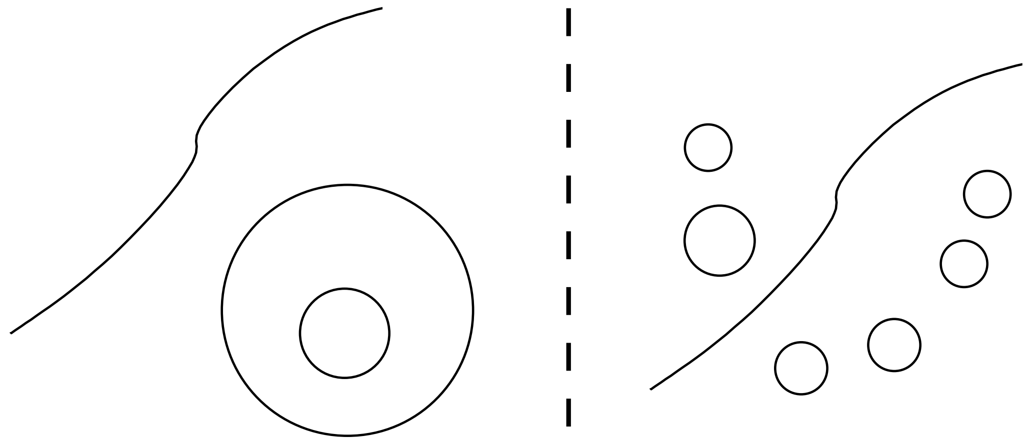

Eight configurations are possible in degree 5. Only the maximal ones are shown.

Partial answers:

Bezout's Theorem: A curve of degree d cuts every line in at most d points. In particular, there cannot be nestings of depth greater than $\lfloor d/2 \rfloor$

Harnack's Theorem: A curve of degree d cannot have more than $\binom{d-1}{2} + 1$ connected components (recall that $\binom{d-1}{2} =$ genus)

In degree six, the possibilities are:

- A single nest with three ovals.
- A number of zero to eleven unnested ovals.
- An oval having i ovals inside (unnested to one another) and j ovals outside, with $i + j \leq 10$.

. . . but not all of them occur. For example:

Theorem (Petrovskii, 193?) Let po and io be the numbers of even and odd ovals (that is to say, ovals nested in an even or odd number of other ovals, respectively) of a nonsingular curve of even degree $d = 2k$. Then:

$$-3/2(k^2 - k) \leq po - io \leq 3/2(k^2 - k) + 1.$$

In particular, a curve of degree six with 11 components cannot have all ovals unnested. Other restrictions were found and in the 60's, **Gudkov** completed the classification of non-singular real algebraic curves of degree six.

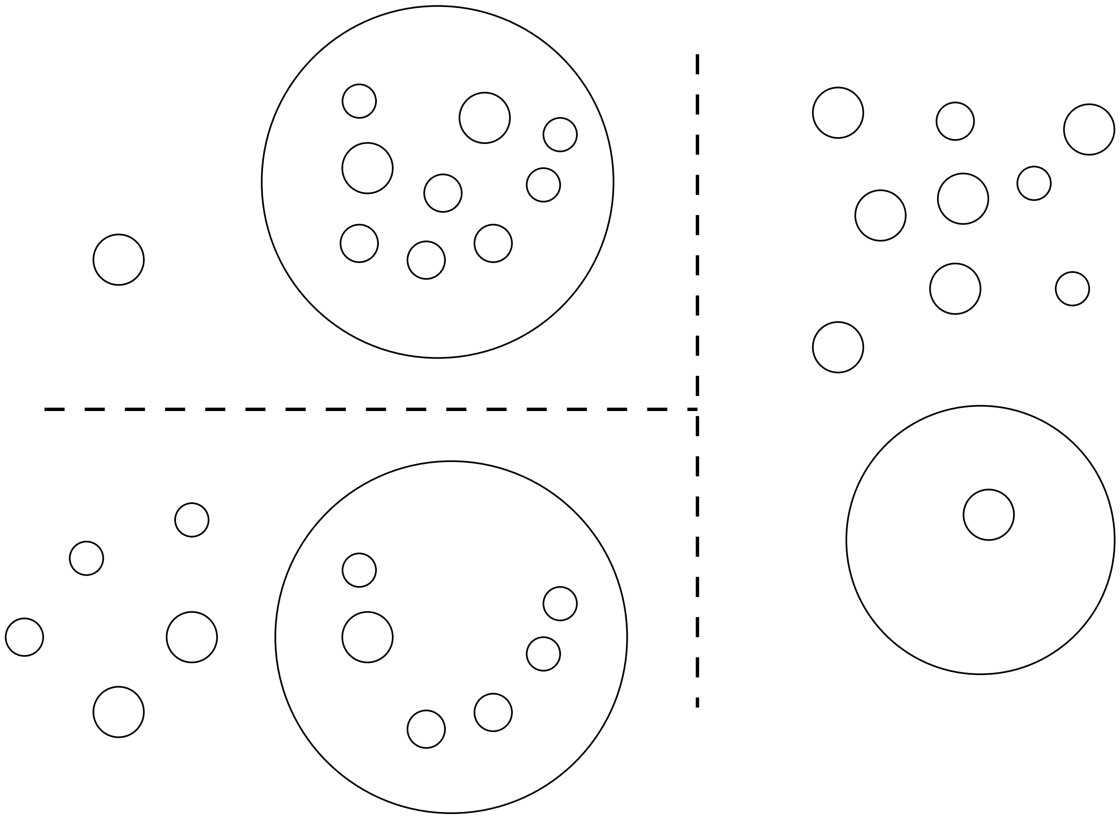

The three curves of degree six with eleven ovals. There are 56 types in total, six of them “maximal”

What about dimension 7? It has only recently been solved (**Viro, 1984**) with a method that involves triangulations.

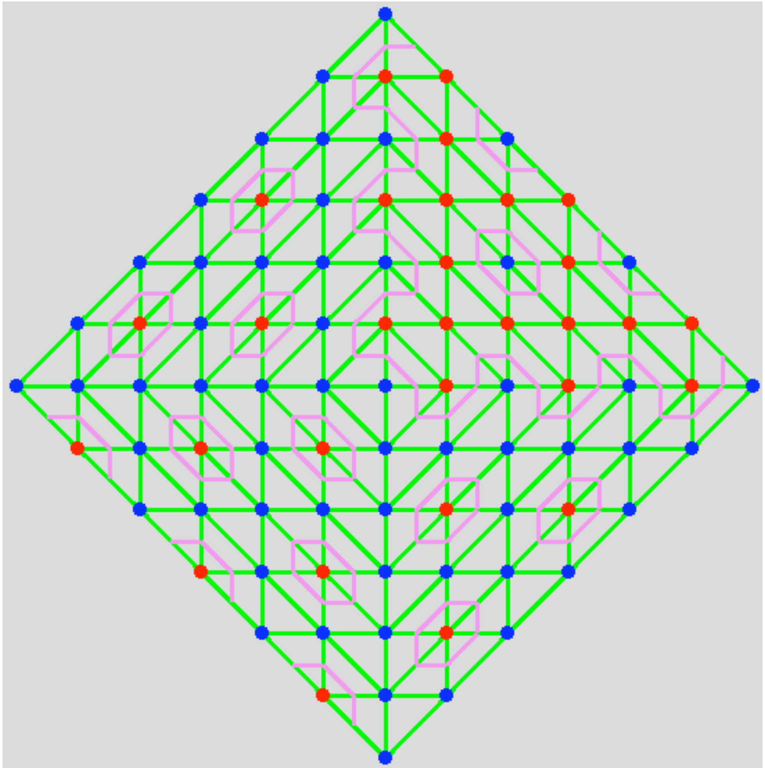

A curve of degree 6 constructed using Viro's method

Viro's method:

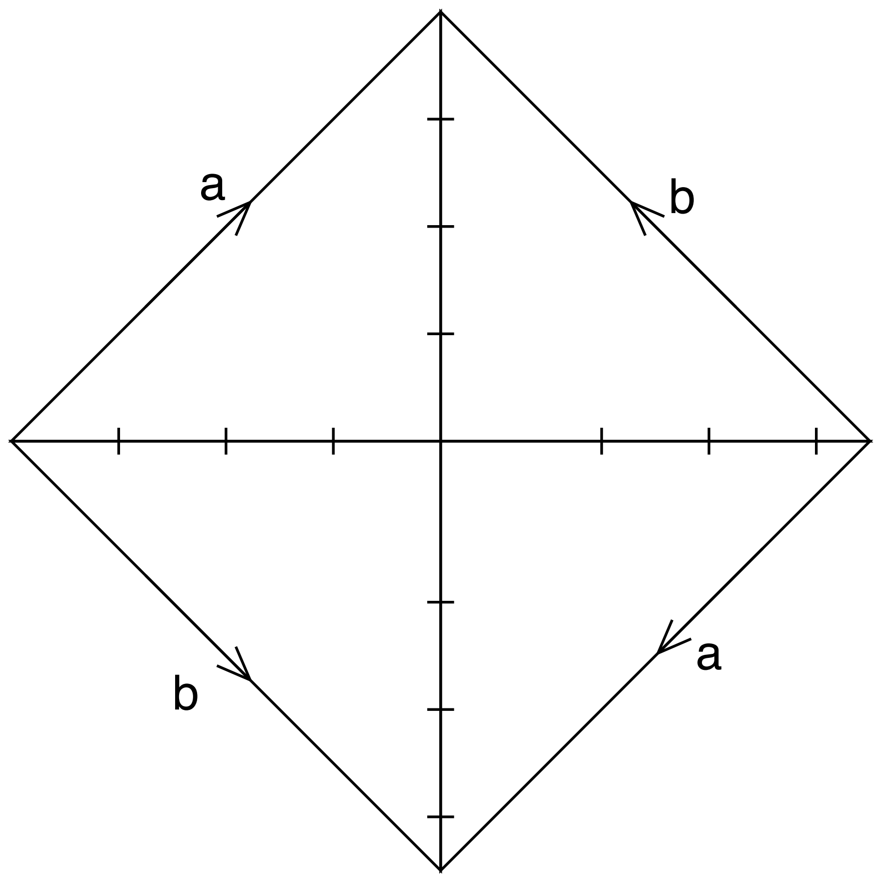

For any given d, construct a topological model of the projective plane by gluing the triangle $(0,0),(d,0),(0,d)$ and its symmetric copies in the other quadrants:

Viro's method:

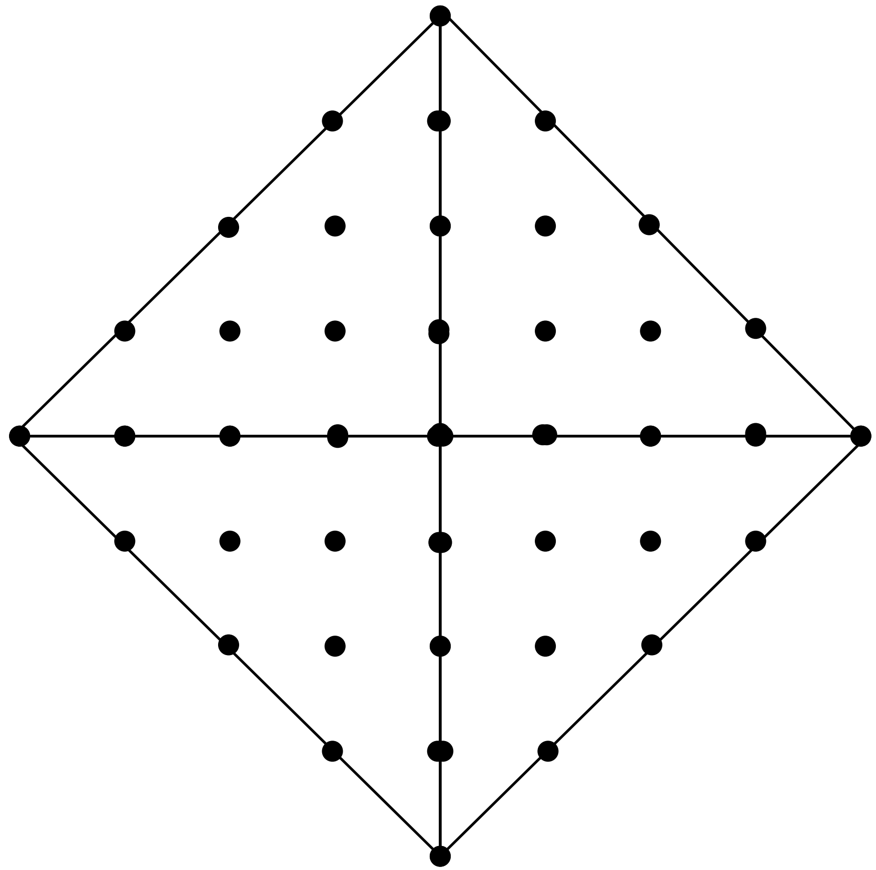

Consider as point set all the integer points in your rhombus (remark: those in a particular orthant are related to the possible homogeneous monomials of degree d in three variables).

Viro's method:

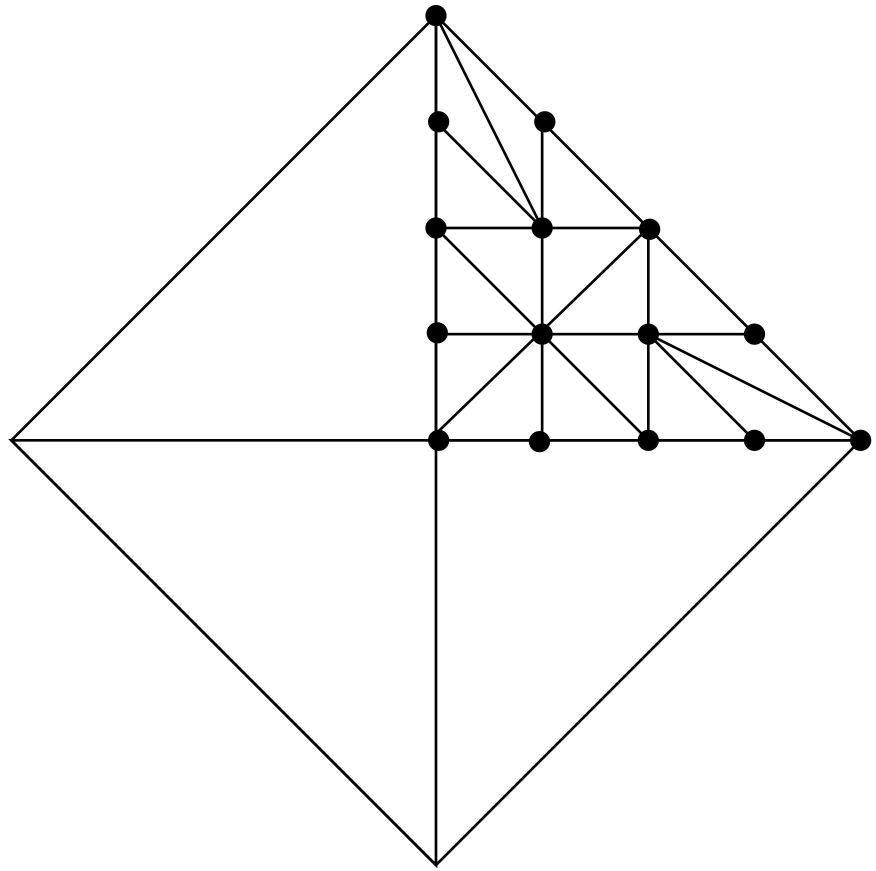

Triangulate the positive orthant arbitrarily . . .

Viro's method:

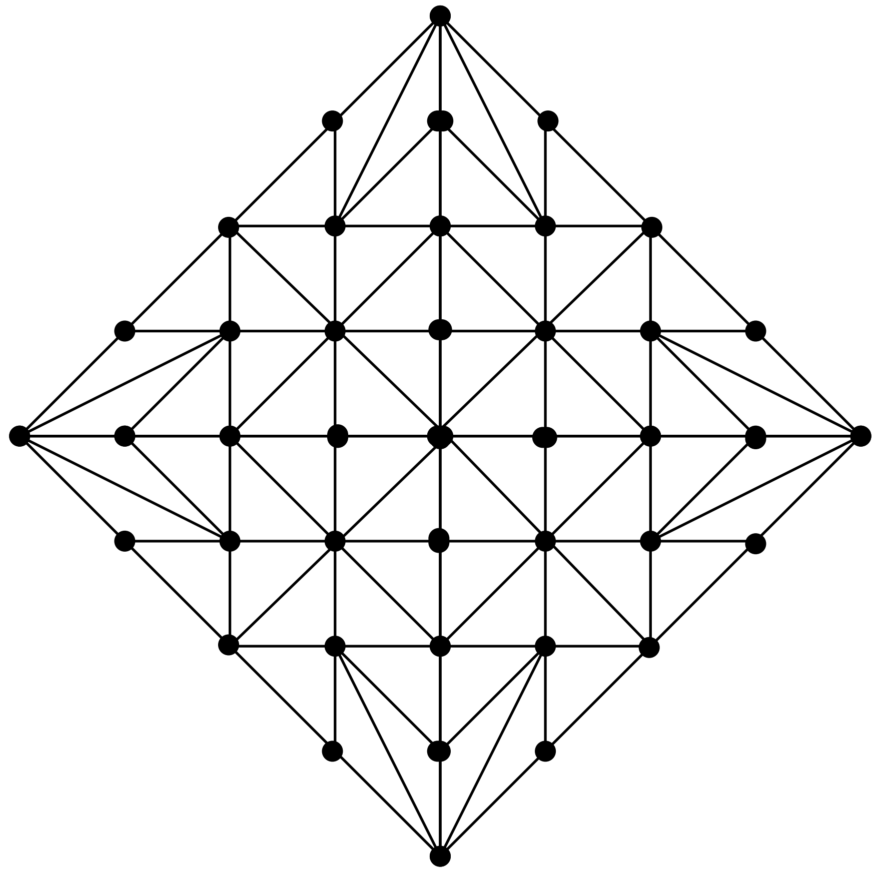

Triangulate the positive quadrant arbitrarily . . .

. . . and replicate the triangulation to the other three quadrants by reflection on the axes.

Viro's method:

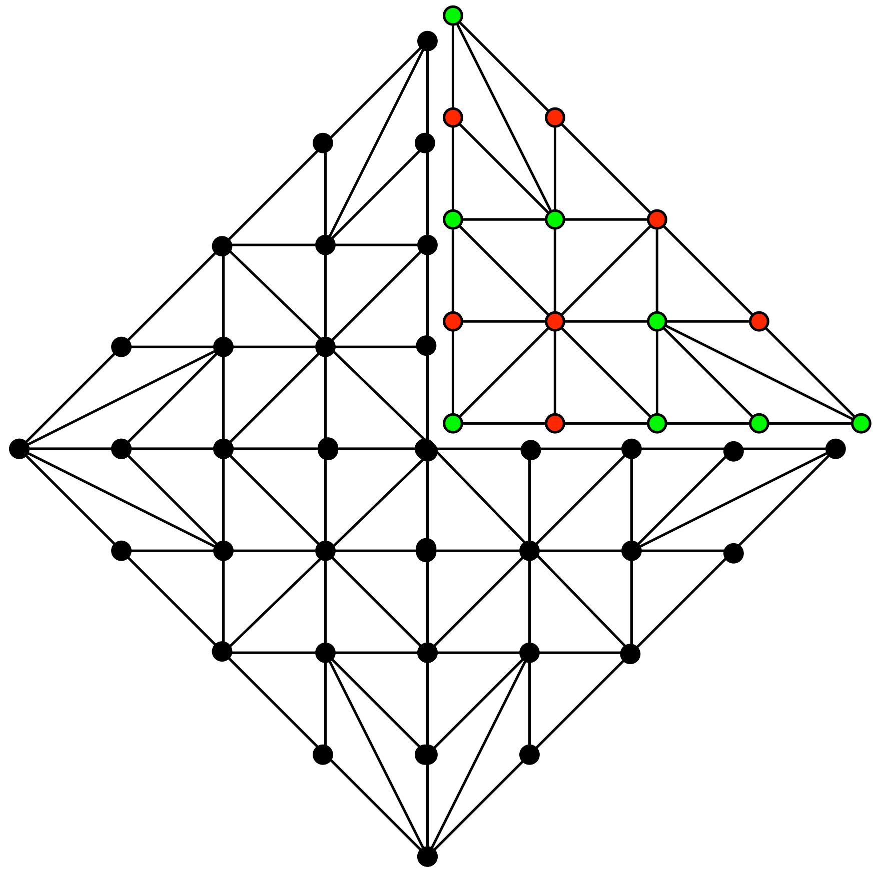

Choose arbitrary signs for the points in the first quadrant

Viro's method:

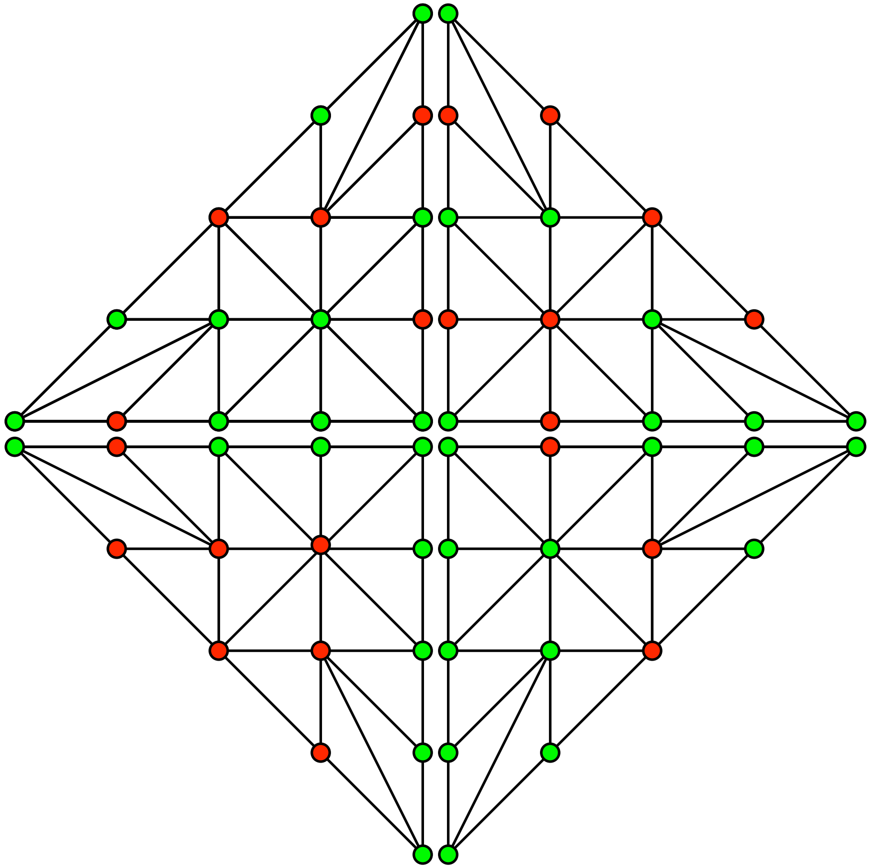

Choose arbitrary signs for the points in the first quadrant . . . and replicate them to the other three quadrants, taking parity of the corresponding coordinate into account.

Viro's method:

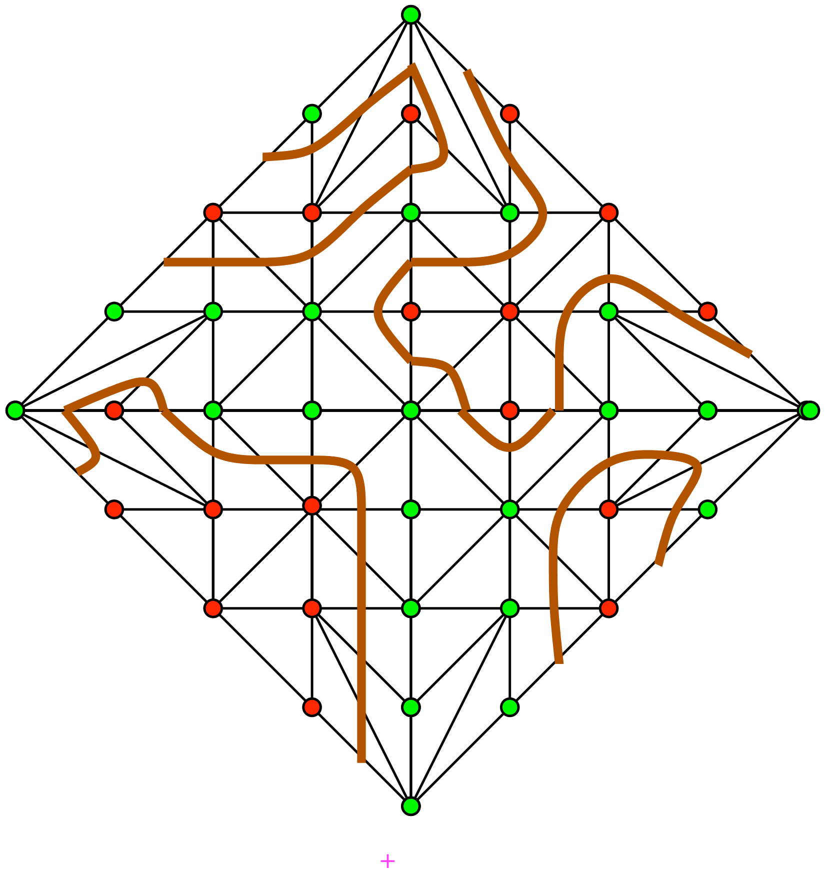

Finally draw your curve in such a way that it separates positive from negative points.

Viro's Theorem

Theorem (Viro, 1987) If the triangulation T chosen for the first quadrant is regular then there is a real algebraic non-singular projective curve f of degree d realizing exactly that topology.

Viro's Theorem

Theorem (Viro, 1987) If the triangulation T chosen for the first quadrant is regular then there is a real algebraic non-singular projective curve f of degree d realizing exactly that topology.

More precisely, let $w_{i,j}$ $(0 \leq i \leq i + j \leq d)$ denote "weights" (↔cost vector↔lifting function) producing your triangulation and let $c_{i,j}$ be any real numbers of the sign you've given to the point (i, j).

Then, the polynomial

$$f_t(x, y) = \sum c_{i,j} x^i y^j z^{d-i-j} t^{w(i,j)}$$

for any positive and sufficiently small t gives the curve you're looking for.

Viro's Theorem

• The method works exactly the same in higher dimension (and produces smooth real algebraic projective hypersurfaces).

Viro's Theorem

- The method works exactly the same in higher dimension (and produces smooth real algebraic projective hypersurfaces).

- It was extended to varieties of higher codimension by Sturmfels.

Viro's Theorem

• The method works exactly the same in higher dimension (and produces smooth real algebraic projective hypersurfaces).

• It was extended to varieties of higher codimension by Sturmfels.

• It was used by **I. Itenberg** in 1993 to disprove Ragsdale's conjecture, dating from 1906!

Viro's Theorem

• The method works exactly the same in higher dimension (and produces smooth real algebraic projective hypersurfaces).

• It was extended to varieties of higher codimension by Sturmfels.

• It was used by **I. Itenberg** in 1993 to disprove Ragsdale's conjecture, dating from 1906!

• What happens if we do the construction with a non-regular triangulation?

Viro's Theorem

• The method works exactly the same in higher dimension (and produces smooth real algebraic projective hypersurfaces).

• It was extended to varieties of higher codimension by Sturmfels.

• It was used by **I. Itenberg** in 1993 to disprove Ragsdale's conjecture, dating from 1906!

• What happens if we do the construction with a non-regular triangulation?
Well, then the formula in the theorem cannot be applied (there is no possible choice of weights). But nobody knows an explicit case in which the curve given by the combinatorial procedure does not have the type of a curve of the corresponding degree.

Counting solutions of sparse polynomial systems

Counting solutions of sparse polynomial systems

Let f and g be two polynomials in two unknowns, of degrees d and d'. Bezout's Theorem says that if the number of (complex, projective) solutions of $f(x,y) = g(x,y) = 0$ is finite, then it is bounded above by dd'.

Question: Can we get better bounds if we know that most of the possible monomials in f and g have zero coefficient? Observe that in one dimension:

• The number of distinct roots of a polynomial is at most equal to twice the number of monomials minus one (Descartes rule of signs)

• The number of non-zero roots, counted with multiplicity, cannot exceed the difference between the highest and lowest degree monomials in the polynomial (as follows from "Bezout in one dimension").

What we want is a generalization of the second statement. For this:

- To every possible monomial $x^i y^j$ we associate its corresponding integer point (i, j) (as in Viro's Theorem).

- To a polynomial $f(x, y) = \sum c_{i,j} x^i y^j$ we associate the corresponding integer point set, and call the convex hull of it the Newton polytope of f.

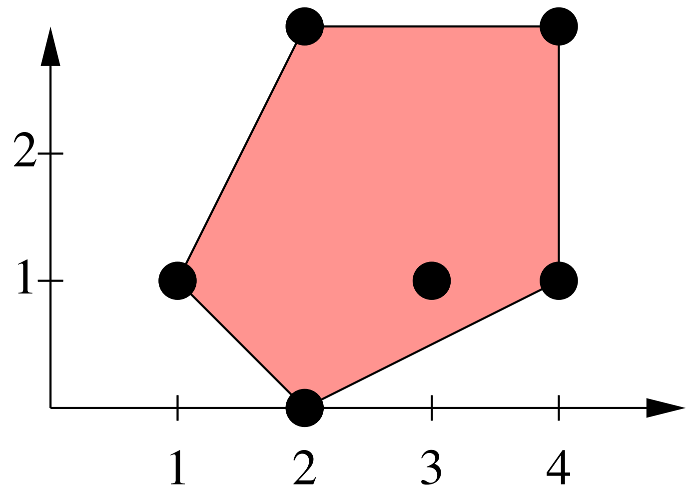

The Newton polytope for the polynomial $x^2 + xy + x^3y + x^4y + x^2y^3 + x^4y^3$

What we want is a generalization of the second statement. For this:

• To every possible monomial $x^i y^j$ we associate its corresponding integer point (i, j) (as in Viro's Theorem).

• To a polynomial $f(x, y) = \sum c_{i,j} x^i y^j$ we associate the corresponding integer point set, and call the convex hull of it the Newton polytope of f.

The question then is:

Can we bound the number of common zeroes of f and g, counted with multiplicities, in terms of $N(f)$ and $N(g)$? (For example, in terms of their areas...)

What we want is a generalization of the second statement. For this:

- To every possible monomial $x^i y^j$ we associate its corresponding integer point (i,j) (as in Viro's Theorem).

- To a polynomial $f(x,y) = \sum c_{i,j} x^i y^j$ we associate the corresponding integer point set, and call the convex hull of it the Newton polytope of f.

The question then is:

Can we bound the number of common zeroes of f and g, counted with multiplicities, in terms of $N(f)$ and $N(g)$? (For example, in terms of their areas...)

Remark: sparse polynomial $\sim$ polynomial with a fixed set of allowed monomials.

YES !

YES !

Theorem (Bernstein, 1975) The number of common zeroes of f and g in $(\mathbb{R}^*)^2$ (that is, out of the coordinate axes) is bounded above by the mixed area of the two polygons $N(f)$ and $N(g)$.

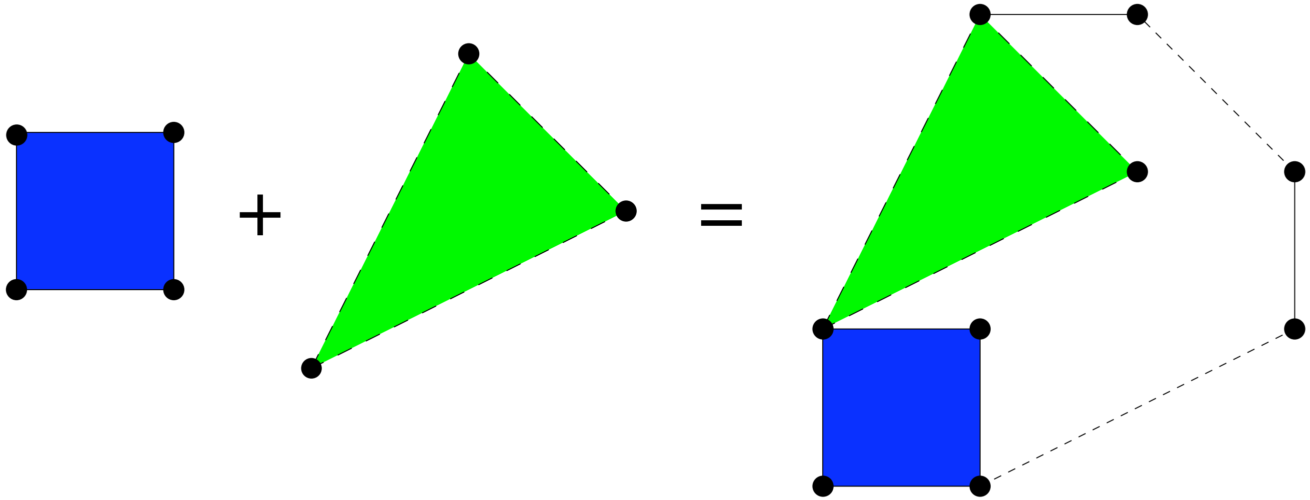

Mixed area of a triangle and a rectangle.

YES !

Theorem (Bernstein, 1975) The number of common zeroes of f and g in $(\mathbb{R}^*)^2$ (that is, out of the coordinate axes) is bounded above by the mixed area of the two polygons $N(f)$ and $N(g)$.

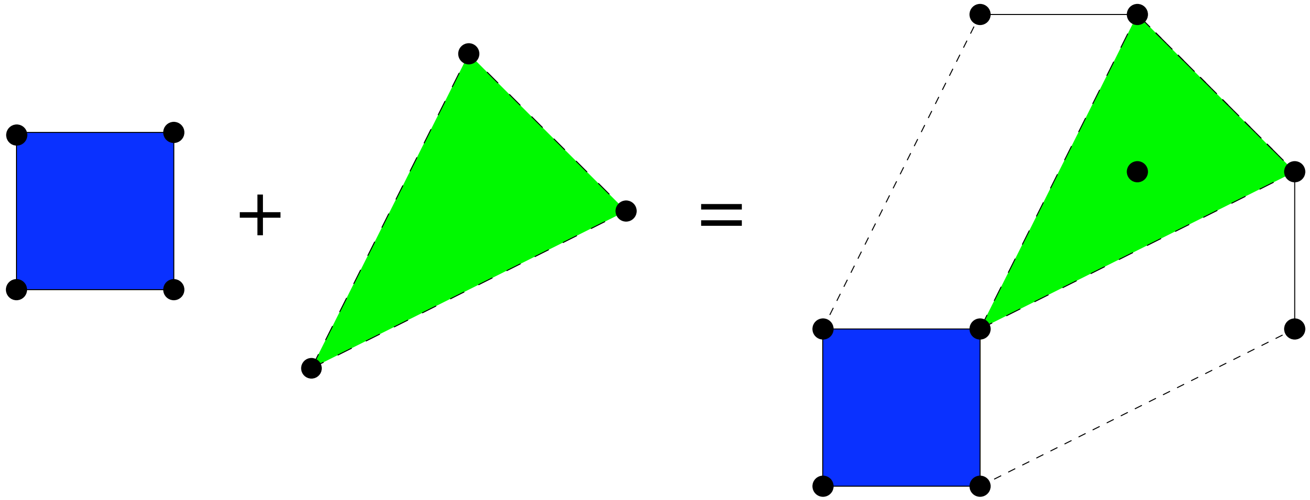

Mixed area of a triangle and a rectangle.

YES !

Theorem (Bernstein, 1975) The number of common zeroes of f and g in $(\mathbb{R}^*)^2$ (that is, out of the coordinate axes) is bounded above by the mixed area of the two polygons $N(f)$ and $N(g)$.

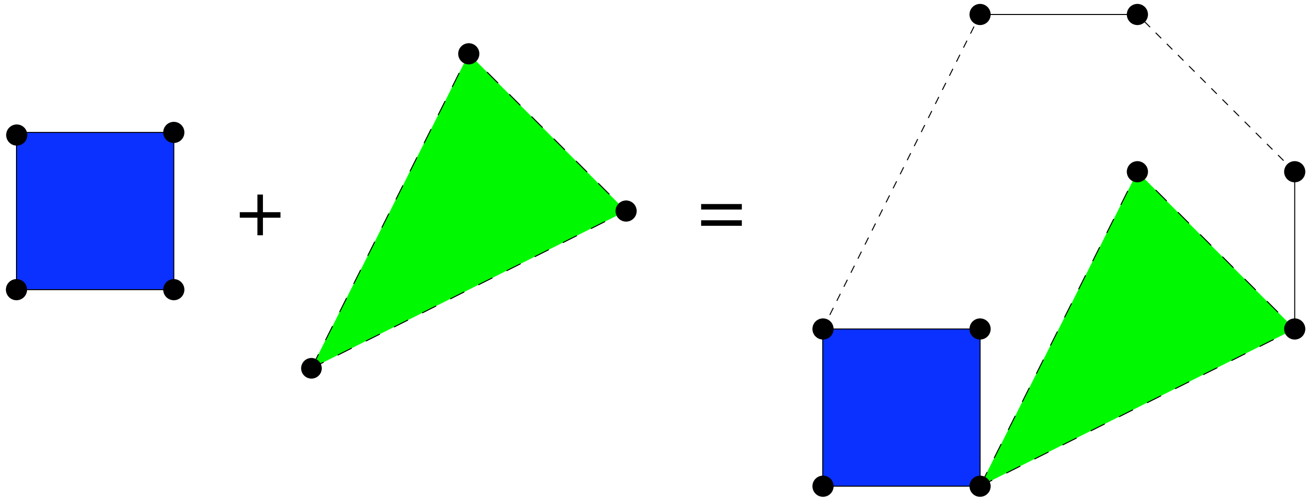

Mixed area of a triangle and a rectangle.

YES !

Theorem (Bernstein, 1975) The number of common zeroes of f and g in $(\mathbb{R}^*)^2$ (that is, out of the coordinate axes) is bounded above by the mixed area of the two polygons $N(f)$ and $N(g)$.

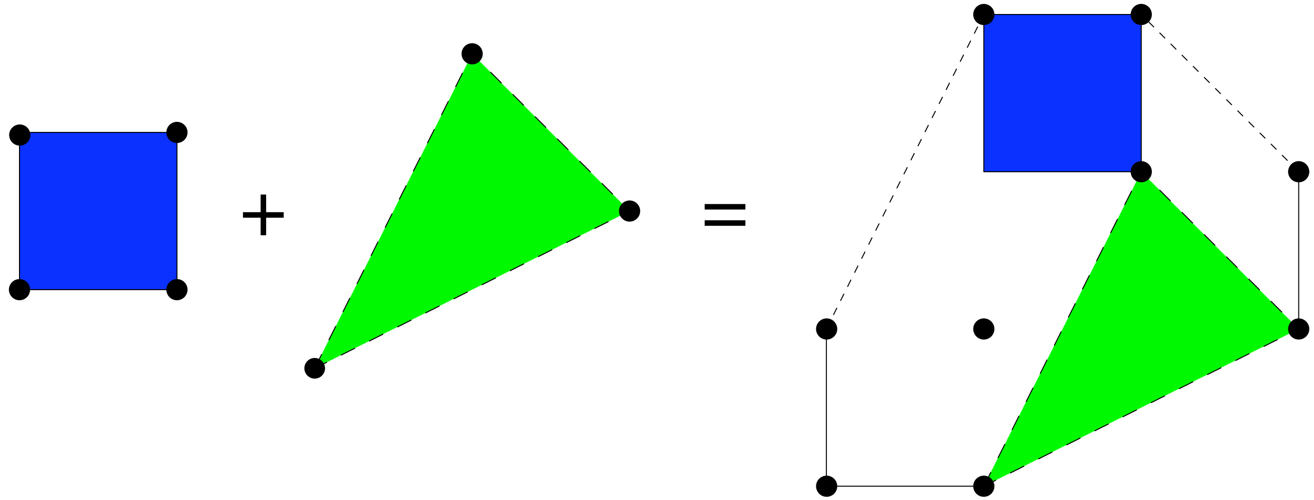

Mixed area of a triangle and a rectangle.

YES !

Theorem (Bernstein, 1975) The number of common zeroes of f and g in $(\mathbb{R}^*)^2$ (that is, out of the coordinate axes) is bounded above by the mixed area of the two polygons $N(f)$ and $N(g)$ and it equals the mixed area for sufficiently generic choices of the coefficients.

YES !

Theorem (Bernstein, 1975) The number of common zeroes of f and g in $(\mathbb{R}^*)^2$ (that is, out of the coordinate axes) is bounded above by the mixed area of the two polygons $N(f)$ and $N(g)$ and it equals the mixed area for sufficiently generic choices of the coefficients.

The theorem is valid for n polynomials $f_1, \ldots, f_n$ in n variables, except we have to define their mixed volume.

Mixed volume

Mixed volume

Definition 1: Let $Q_1, Q_2, \ldots, Q_n$ be n polytopes in $\mathbb{R}^n$. Their mixed volume $\mu(Q_1, \ldots, Q_n)$ equals

$$\sum_{I \subset \{1,2,\ldots,n\}} (-1)^{|I|} \operatorname{vol}\left(\sum_{j \in I} Q_i \right).$$

Mixed volume

Definition 2: Let $Q_1, Q_2, \ldots, Q_n$ be n polytopes in $\mathbb{R}^n$. Their mixed volume $\mu(Q_1, \ldots, Q_n)$ equals

the coefficient of $x_1 x_2 \cdots x_n$ in the homogeneous polynomial $\mathrm{vol}(x_1 Q_1 + \cdots + x_n Q_n)$.

Mixed volume

Definition 3: Let $Q_1, Q_2, \ldots, Q_n$ be n polytopes in $\mathbb{R}^n$. Their mixed volume $\mu(Q_1, \ldots, Q_n)$ equals

the sum of the volumes of the mixed cells in any (fine) mixed subdivision of $Q_1 + \cdots + Q_n$.

Mixed volume

Definition 3: Let $Q_1, Q_2, \ldots, Q_n$ be n polytopes in $\mathbb{R}^n$. Their mixed volume $\mu(Q_1, \ldots, Q_n)$ equals

the sum of the volumes of the mixed cells in any (fine) mixed subdivision of $Q_1 + \cdots + Q_n$.

In particular, to compute the number of zeroes of a sparse system of polynomials $f_1, \ldots, f_n$ one only needs to compute a "(fine) mixed subdivision" of $N(f_1) + \cdots + N(f_n)$.

A cooking recipe for fine mixed subdivisions:

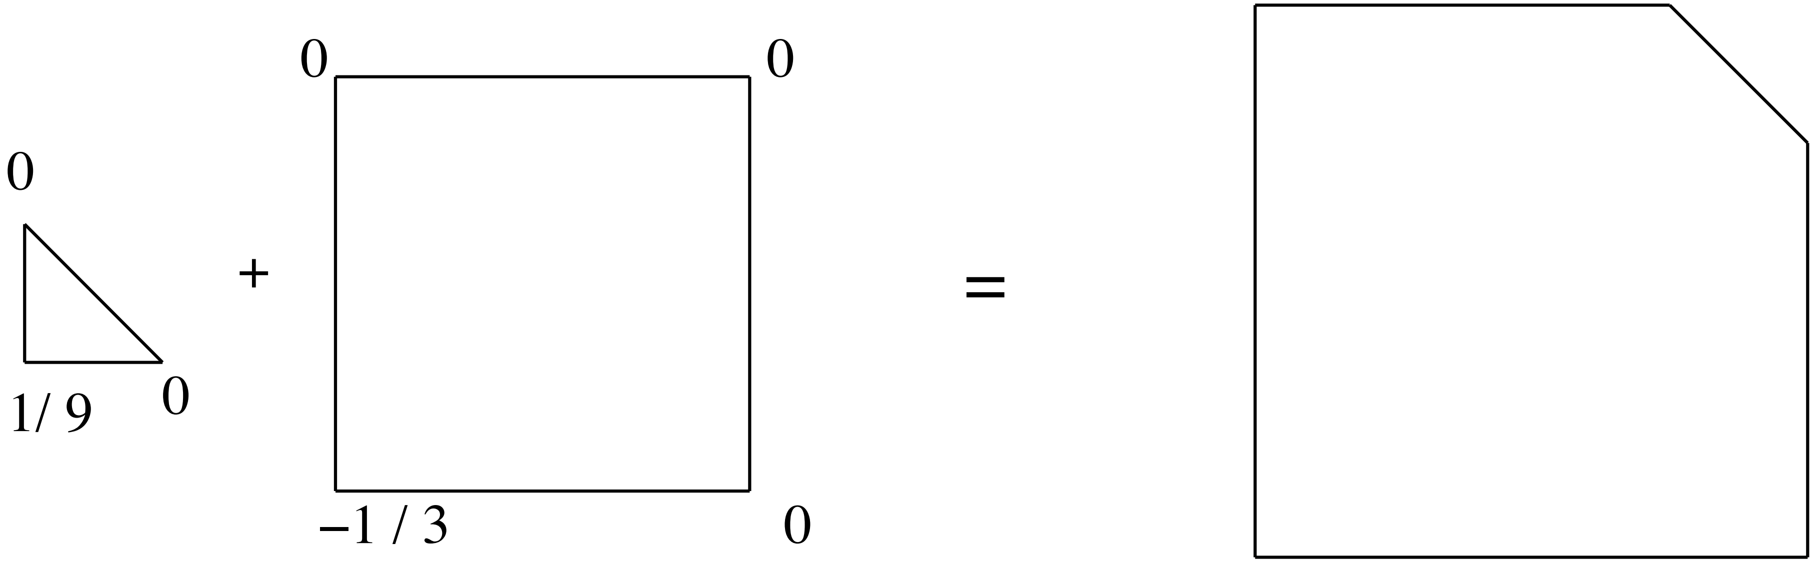

Choose sufficiently generic (e.g. random) numbers $w_a \in \mathbb{R}$, one for each a in each of the Q_i's

A cooking recipe for fine mixed subdivisions:

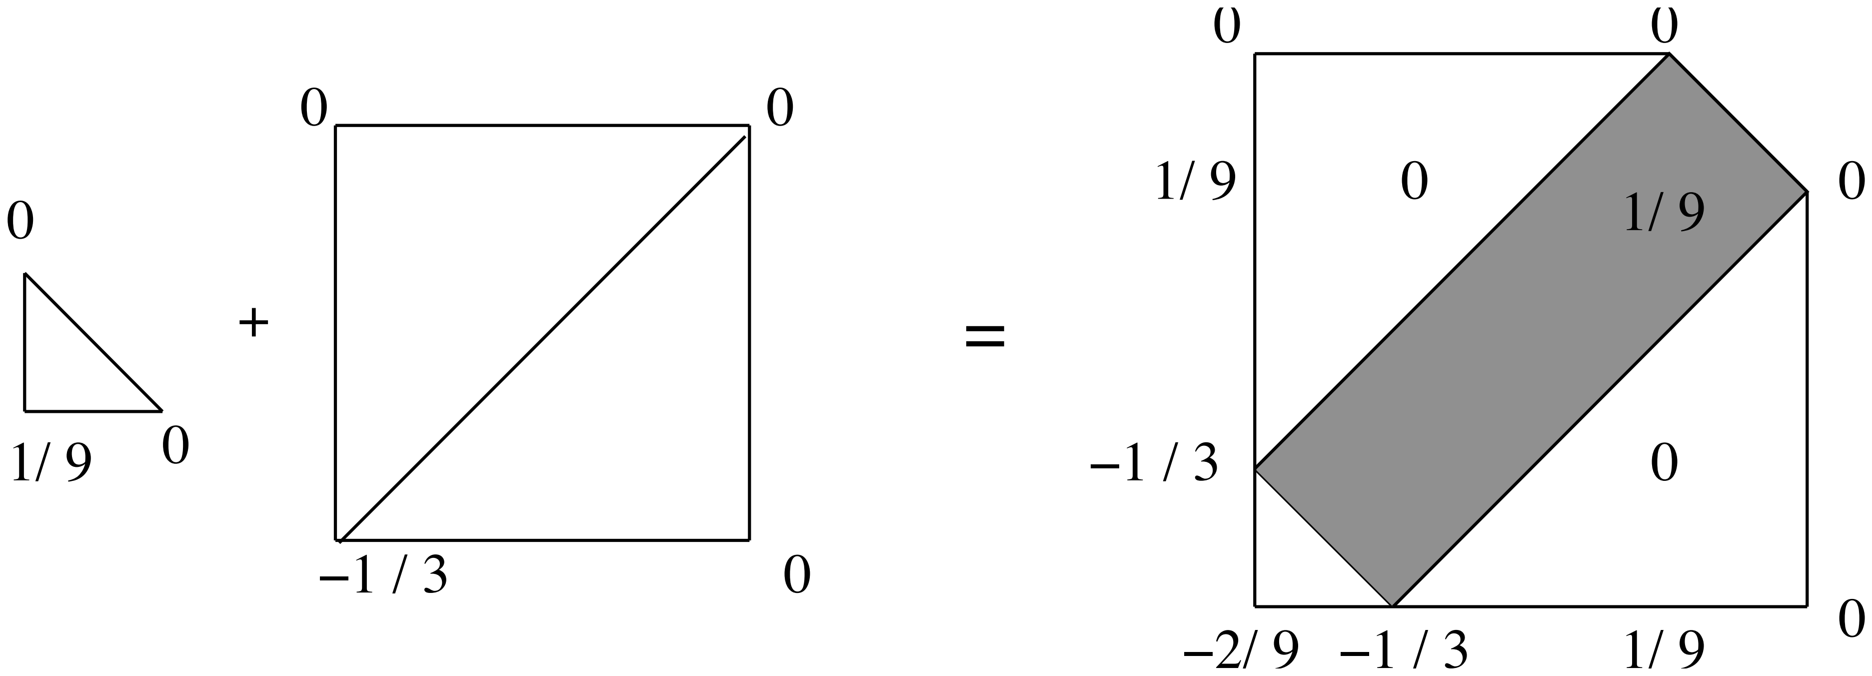

Use the numbers to lift the points of $Q_1 + \cdots + Q_n$ and compute the lower envelope of the lifted point configuration.

A cooking recipe for fine mixed subdivisions:

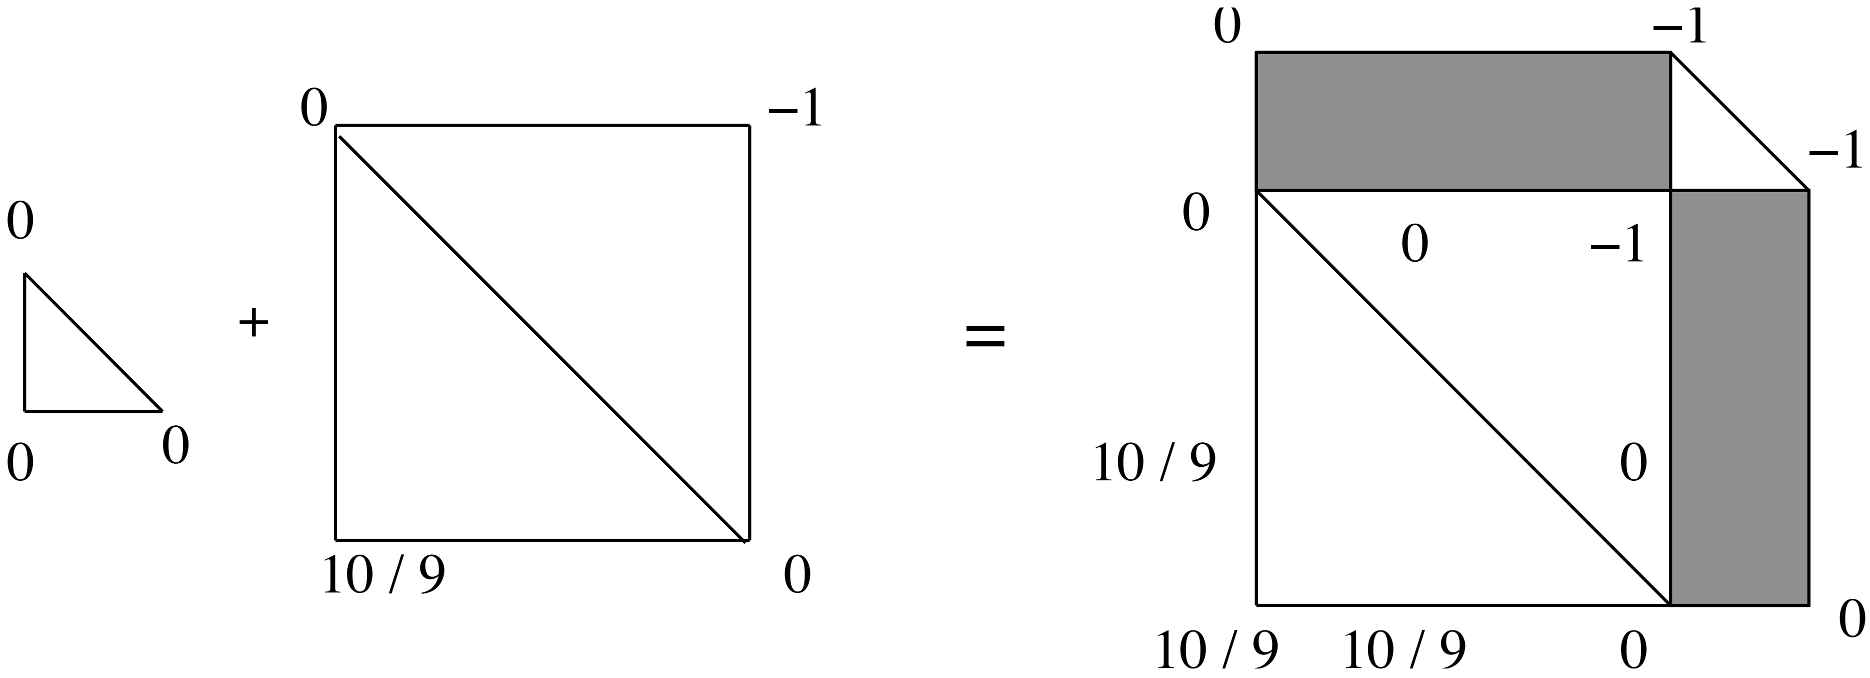

Use the numbers to lift the points of $Q_1 + \cdots + Q_n$ and compute the lower envelope of the lifted point configuration.

A cooking recipe for fine mixed subdivisions:

- The procedure is, clearly, a generalization of the notion of regular triangulation.

A cooking recipe for fine mixed subdivisions:

- The procedure is, clearly, a generalization of the notion of regular triangulation.
- ...actually it is a regular subdivision.

A cooking recipe for fine mixed subdivisions:

• The procedure is, clearly, a generalization of the notion of regular triangulation.

• ...actually it is a regular subdivision.

• But, moreover, for every family of polytopes $Q_1, \ldots, Q_n$ in $\mathbb{R}^d$ there is another polytope $\mathcal{C}(Q_1, \ldots, Q_n)$ in $\mathbb{R}^{n+d-1}$ such that

mixed subdivisions of $Q_1, \ldots, Q_n \leftrightarrow$ subdivisions of $\mathcal{C}(Q_1, \ldots, Q_n)$

fine mixed subdivisions of $Q_1, \ldots, Q_n \leftrightarrow$ triangulations of $\mathcal{C}(Q_1, \ldots, Q_n)$

A cooking recipe for fine mixed subdivisions:

• The procedure is, clearly, a generalization of the notion of regular triangulation.

• ...actually it is a regular subdivision.

• But, moreover, for every family of polytopes $Q_1, \dots, Q_n$ in $\mathbb{R}^d$ there is another polytope $\mathcal{C}(Q_1, \dots, Q_n)$ in $\mathbb{R}^{n+d-1}$ such that

$$\text{mixed subdivisions of } Q_1, \dots, Q_n \leftrightarrow \text{subdivisions of } \mathcal{C}(Q_1, \dots, Q_n)$$

$$\text{fine mixed subdivisions of } Q_1, \dots, Q_n \leftrightarrow \text{triangulations of } \mathcal{C}(Q_1, \dots, Q_n)$$

That is to say, the number of roots of a sparse system of polynomials can be computed via triangulations.

THE END

T H E   E N D

(of lecture 1)