Triangulations Of Polytopes

and Algebraic Geometry

Francisco Santos
Universidad de Cantabria, Santander, SPAIN
http://personales.unican.es/santosf

Anogia, August 25 2005

Outline of the talk

1. Triangulations of polytopes (a brief overview).
2. A polytope with a disconnected graph of triangulations.
3. A disconnected toric Hilbert scheme.
4. Tropical polytopes and products of simplices.

1. Triangulations

Polytopes

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

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

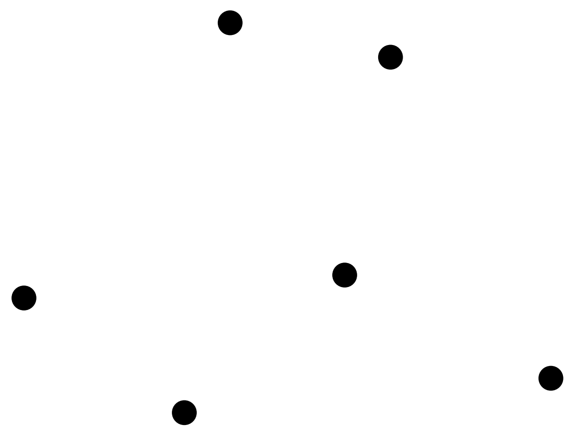

A finite point set

Polytopes

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

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

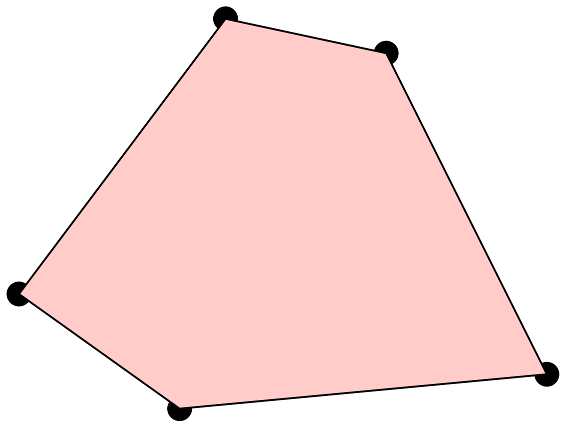

Its convex hull

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.)

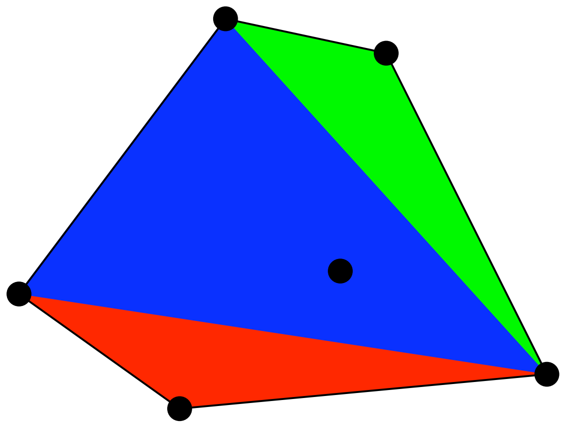

A triangulation of P

Triangulations

The following are **not** triangulations:

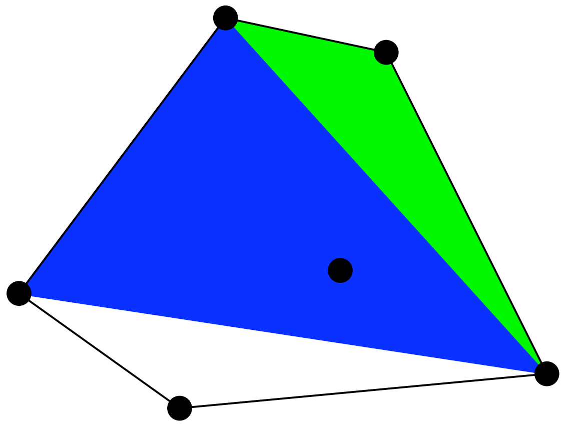

The union is not the whole convex hull

Triangulations

The following are **not** triangulations:

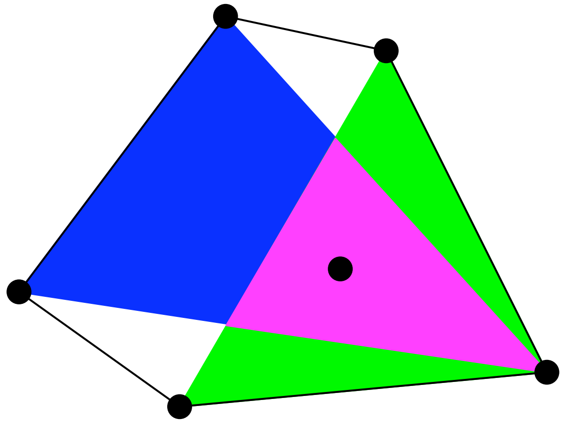

The intersection is not okay

Triangulations of a point configuration

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

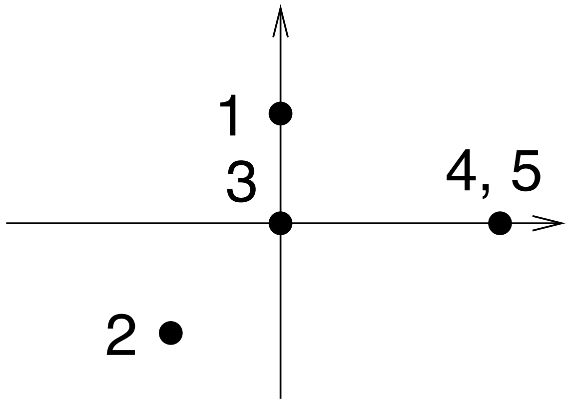

$a_1 = (0, 1)$

$a_2 = (-1, -1)$

$a_3 = (0, 0)$

$a_4 = a_5 = (2, 0)$

A point set with repetitions

Triangulations of a point configuration

A triangulation of a configuration $\mathcal{A}$ is a triangulation of $\mathrm{conv}(\mathcal{A})$ with **vertex set contained in** $\mathcal{A}$. **Remark:** Don't need to use all points

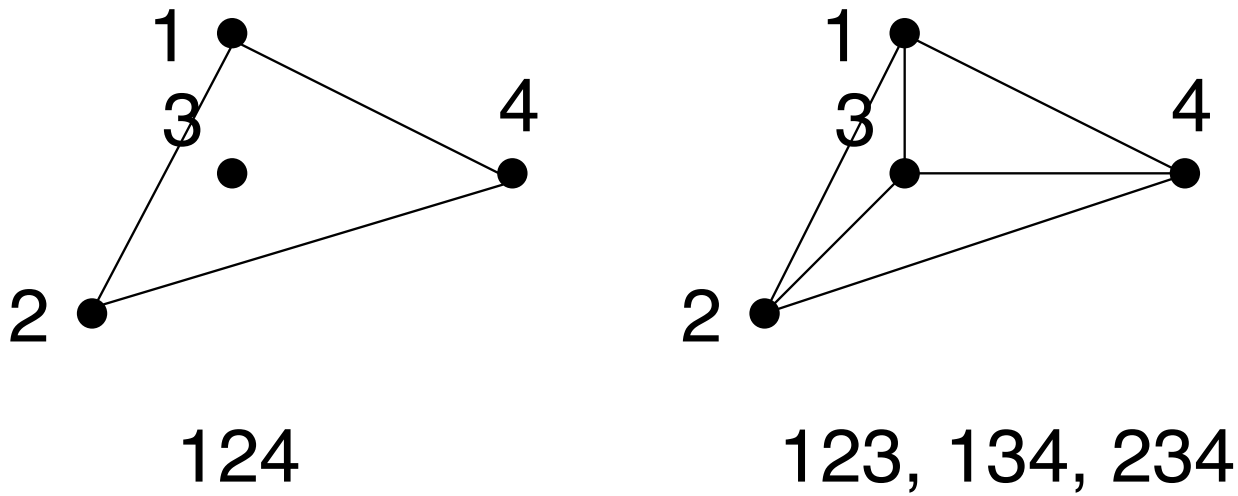

The two triangulations of $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$

Triangulations of vector sets

Triangulations of vector sets

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

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

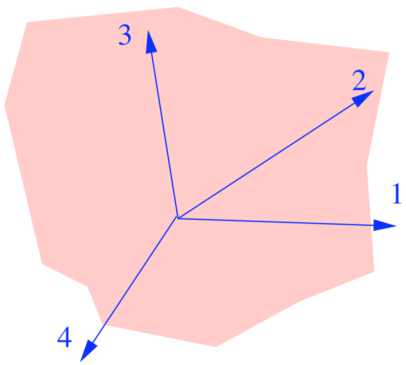

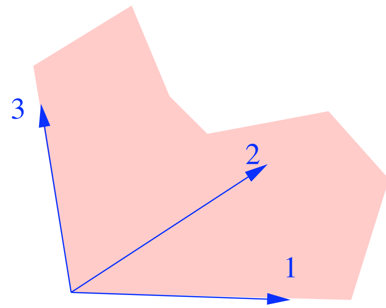

Two vector configurations, and their cones

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 $\mathrm{cone}(A)$ into simplicial cones with generators contained in A and such that:

(UP) The union of all these cones equals $\mathrm{conv}(A)$. (Union Property.)
(IP) Any pair of them intersects in a common face (Intersection Property.)

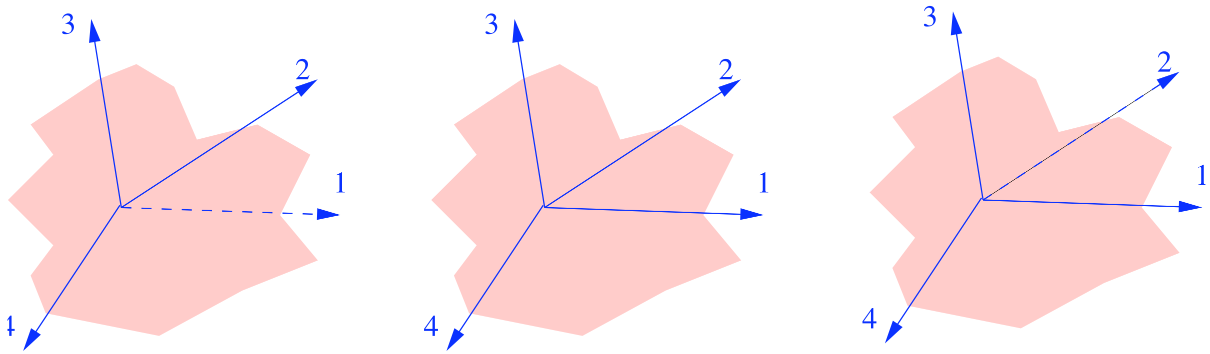

23,34,24 12,23,34,14 13,34,14

The three triangulations of the first configuration

Remark: Triangulations of a $\{$pointed cone/acyclic 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:

(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).

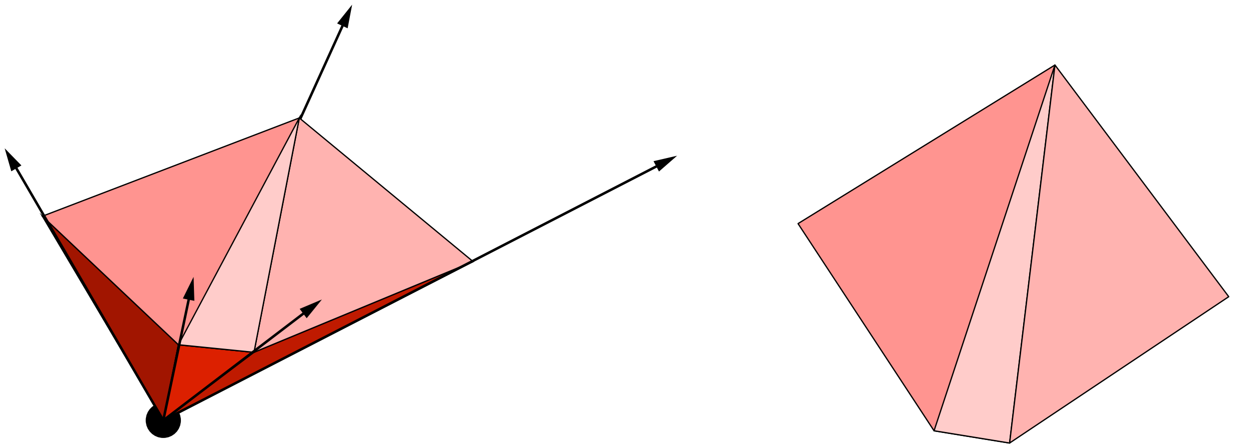

A point configuration can be considered a particular case of vector configuration

Example: Triangulations of a convex n-gon

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

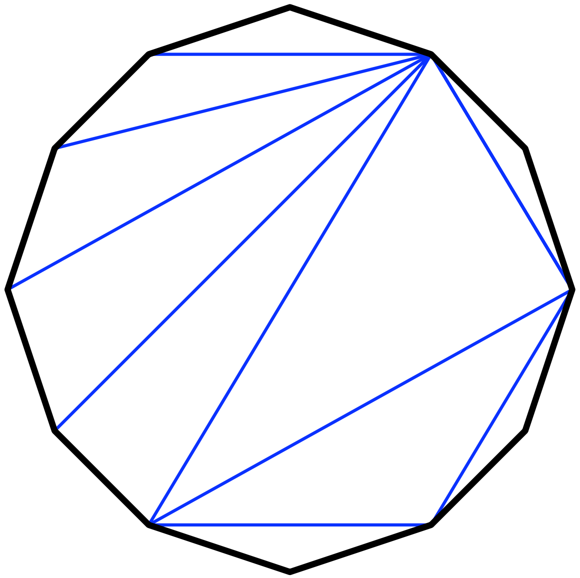

A triangulation of the 12-gon

Example: Triangulations of a convex n-gon

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

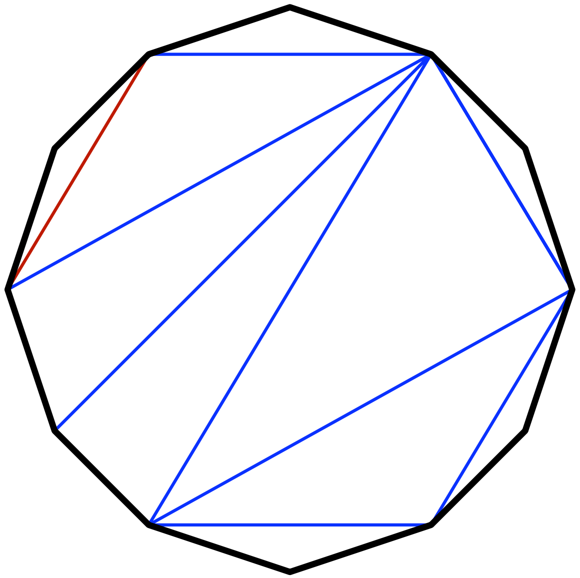

Another triangulation of the 12-gon, obtained by **flipping** an edge

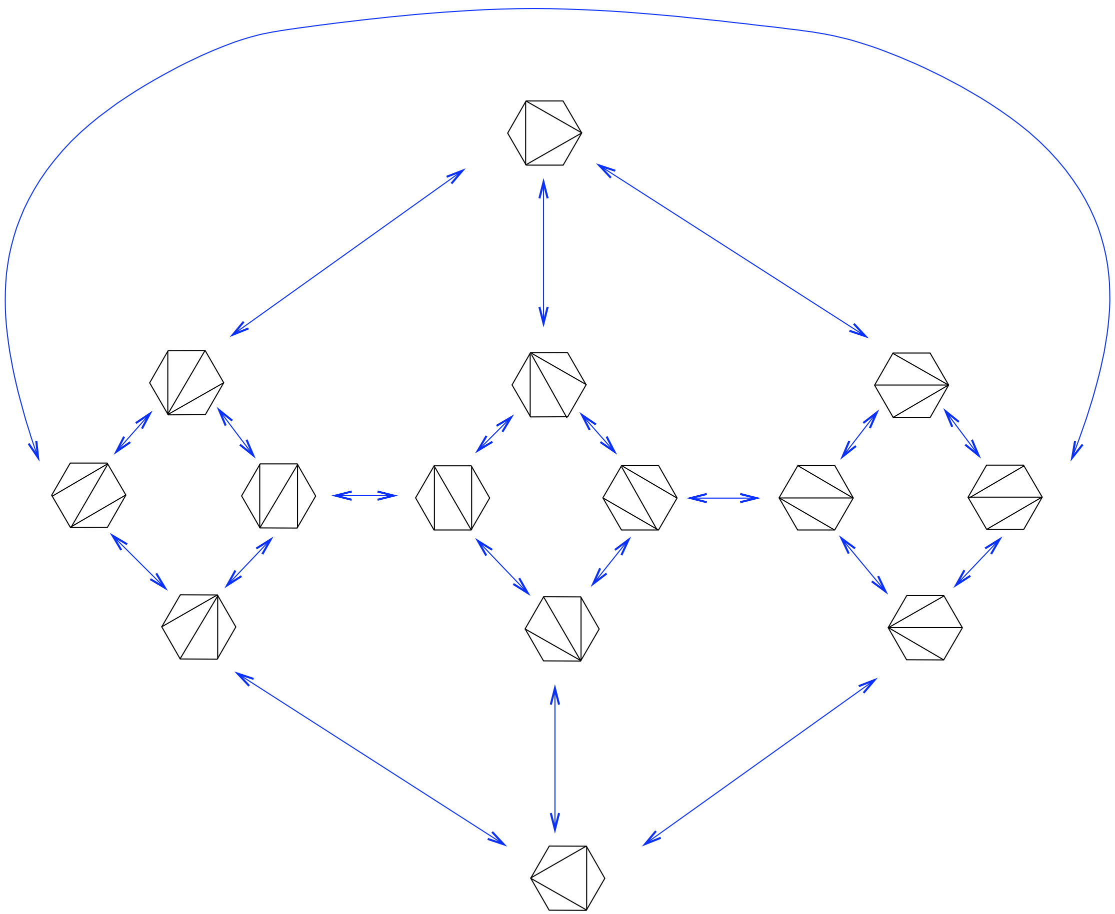

The Graph of flips in triangulations of a hexagon

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** [Stasheff 1963, Haiman 1984, Lee 1989].

- The graph has diameter bounded above by $2n-10$ for all n [easy] and equal to $2n-10$ for large n [hard, Sleator-Tarjan-Thurston, 1988].

- There are exactly $\frac{1}{n-1}\binom{2n-4}{n-2}$ triangulations.
  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:

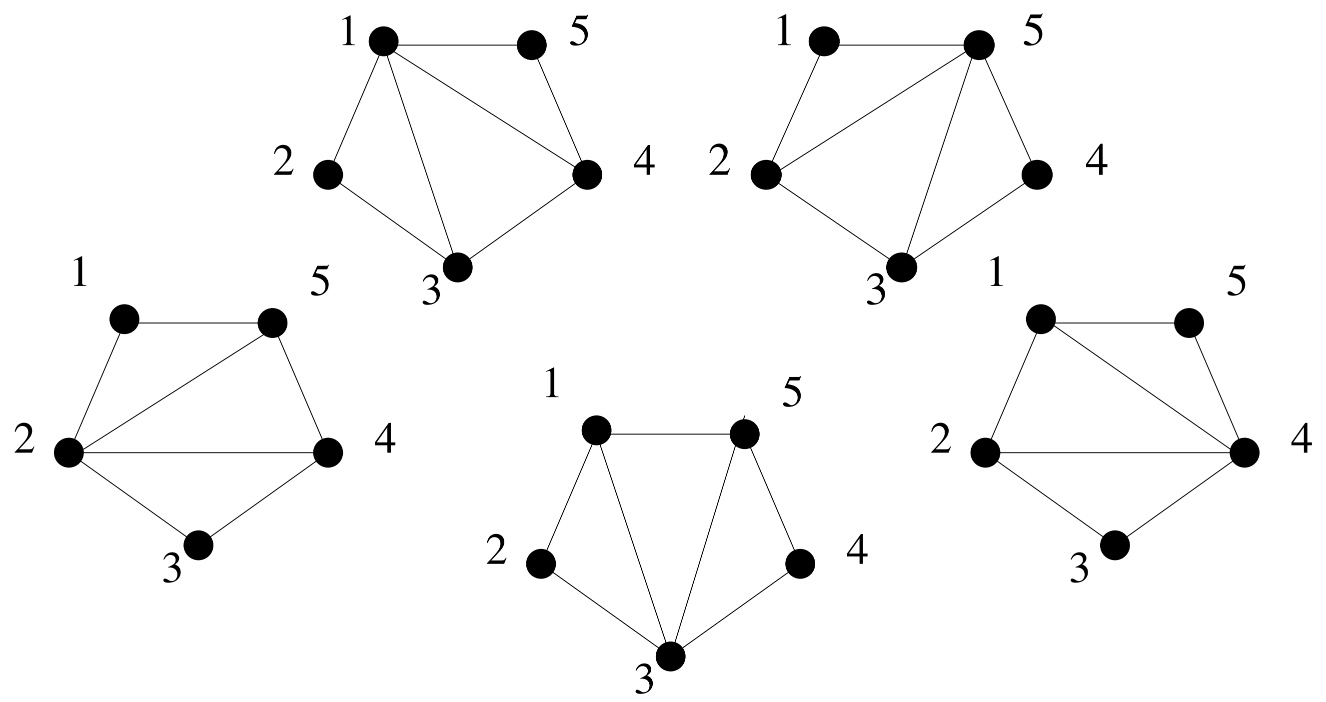

It also counts. . .

1. Binary trees on n-nodes.

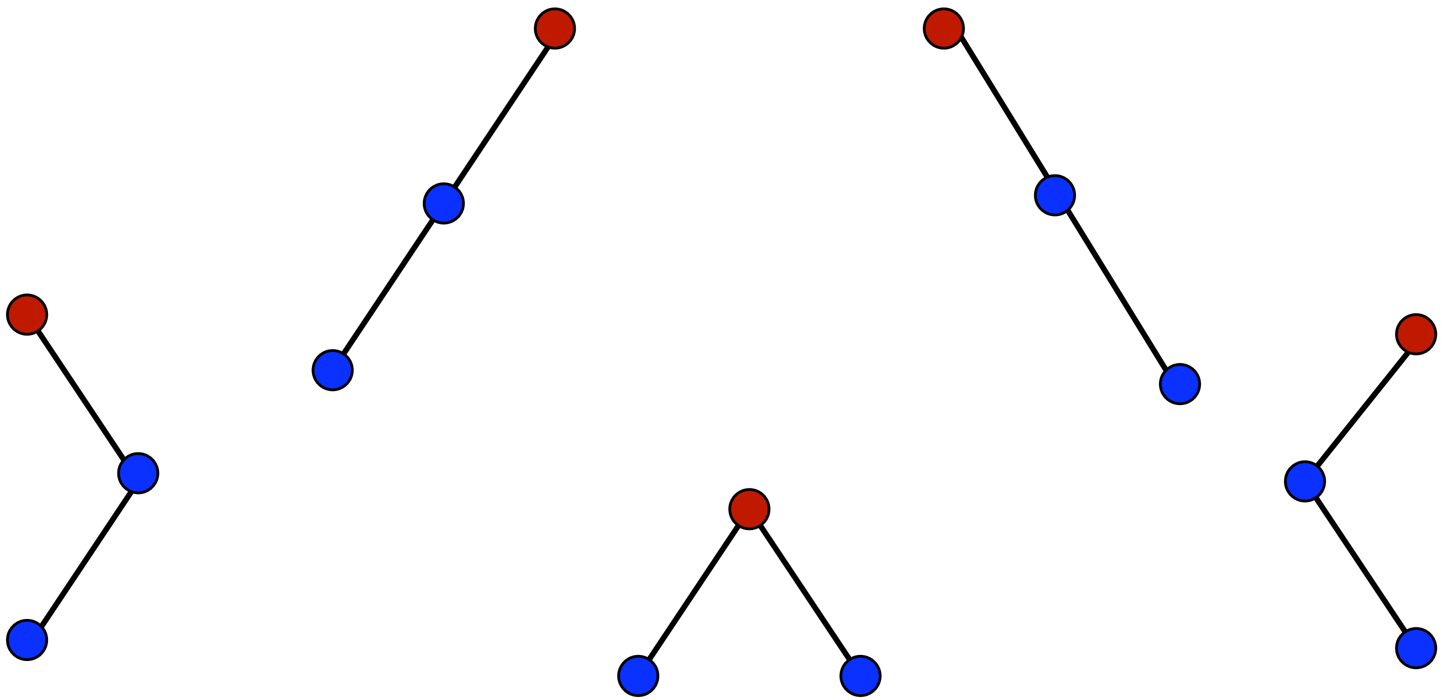

1. Binary trees on n-nodes.

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

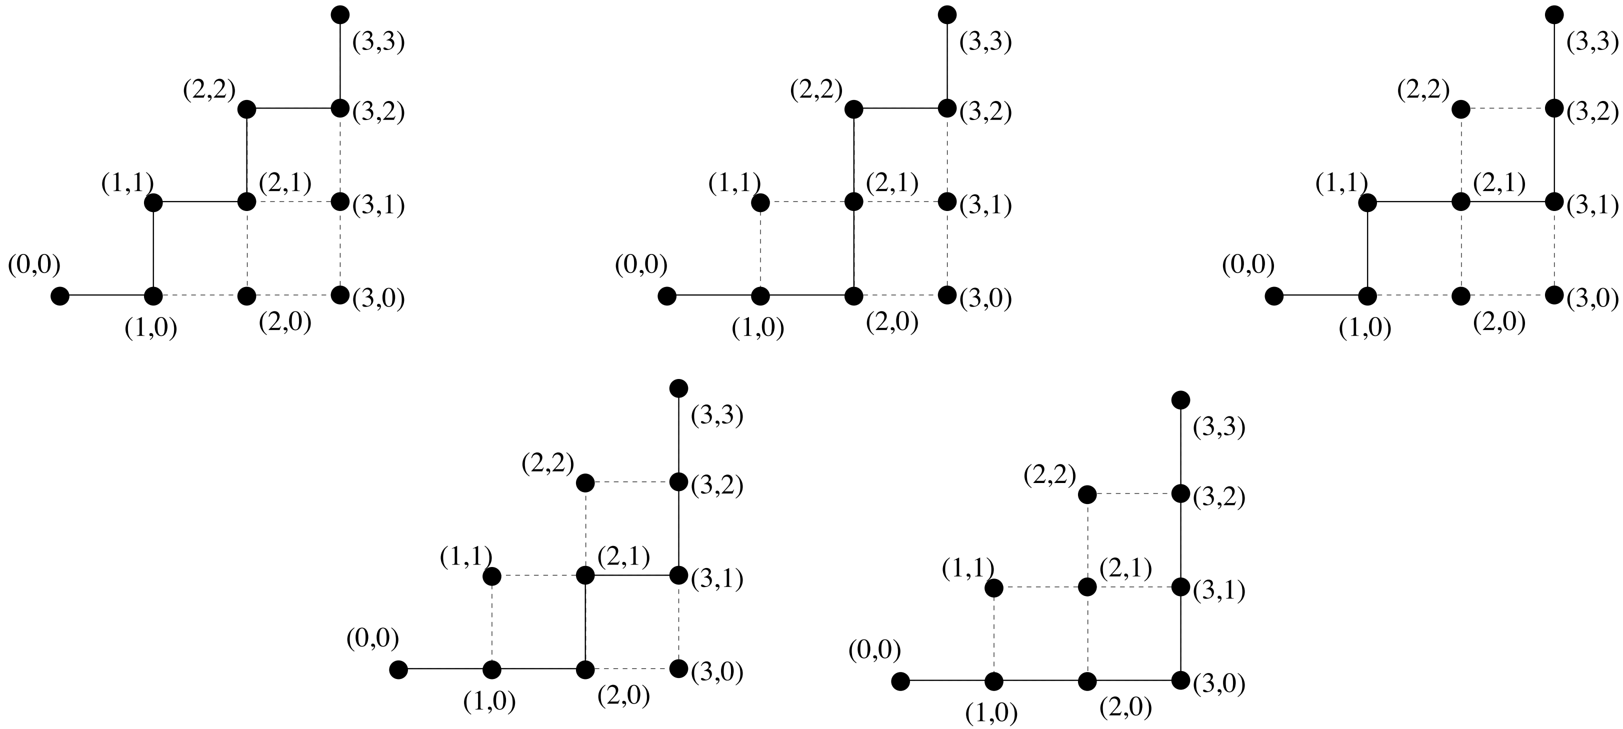

1. Binary trees on n-nodes.

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.

Regular triangulations

Let $A = \{a_, \ldots, a_n\} \subset \mathbb{R}^d$ be a vector configuration. Let $w = (w_1, \ldots, w_n) \in \mathbb{R}^n$ be a vector.

Consider the lifted vector configuration $\tilde{A} = \begin{Bmatrix} a_1 & \cdots & a_n \\ w_1 & \cdots & w_n \end{Bmatrix} \subset \mathbb{R}^{d+1}$. The lower envelope of $\operatorname{cone}(\tilde{A})$ (projected down to $\mathbb{R}^d$) forms a polyhedral subdivision of A. If w is "sufficiently generic" then it forms a triangulation.

Regular triangulations

Let $A = \{a_, \ldots, a_n\} \subset \mathbb{R}^d$ be a vector configuration. Let $w = (w_1, \ldots, w_n) \in \mathbb{R}^n$ be a vector.

Consider the lifted vector configuration $\tilde{A} = \begin{Bmatrix} a_1 & \cdots & a_n \\ w_1 & \cdots & w_n \end{Bmatrix} \subset \mathbb{R}^{d+1}$. The lower envelope of $\mathrm{cone}(\tilde{A})$ (projected down to $\mathbb{R}^d$) forms a polyhedral subdivision of A. If w is "sufficiently generic" then it forms a triangulation.

Obviously, different w's may provide different triangulations.

Regular triangulations

Let $A = \{a_, \ldots, a_n\} \subset \mathbb{R}^d$ be a vector configuration. Let $w = (w_1, \ldots, w_n) \in \mathbb{R}^n$ be a vector.

Consider the lifted vector configuration $\tilde{A} = \left\{ \begin{matrix} a_1 & \cdots & a_n \\ w_1 & \cdots & w_n \end{matrix} \right\} \subset \mathbb{R}^{d+1}$. The lower envelope of $\mathrm{cone}(\tilde{A})$ (projected down to $\mathbb{R}^d$) forms a polyhedral subdivision of A. If w is "sufficiently generic" then it forms a triangulation.

Obviously, different w's may provide different triangulations.

More interestingly, for some A's, **not all triangulations can be obtained in this way.** The triangulations that do are called regular triangulations.

Regular triangulations

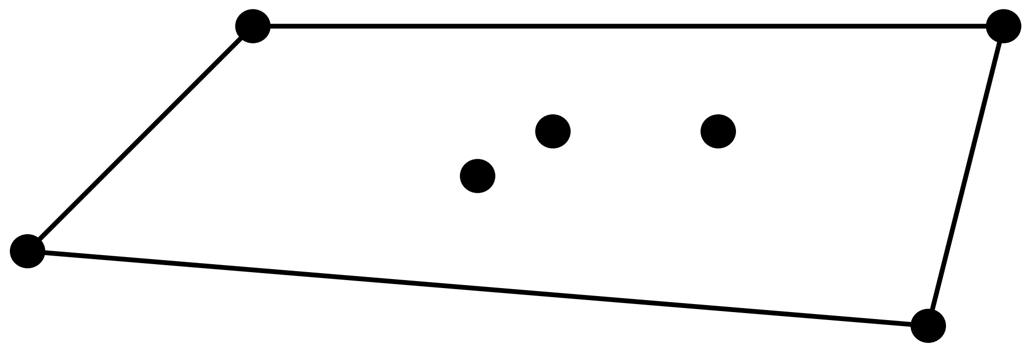

Regular triangulations

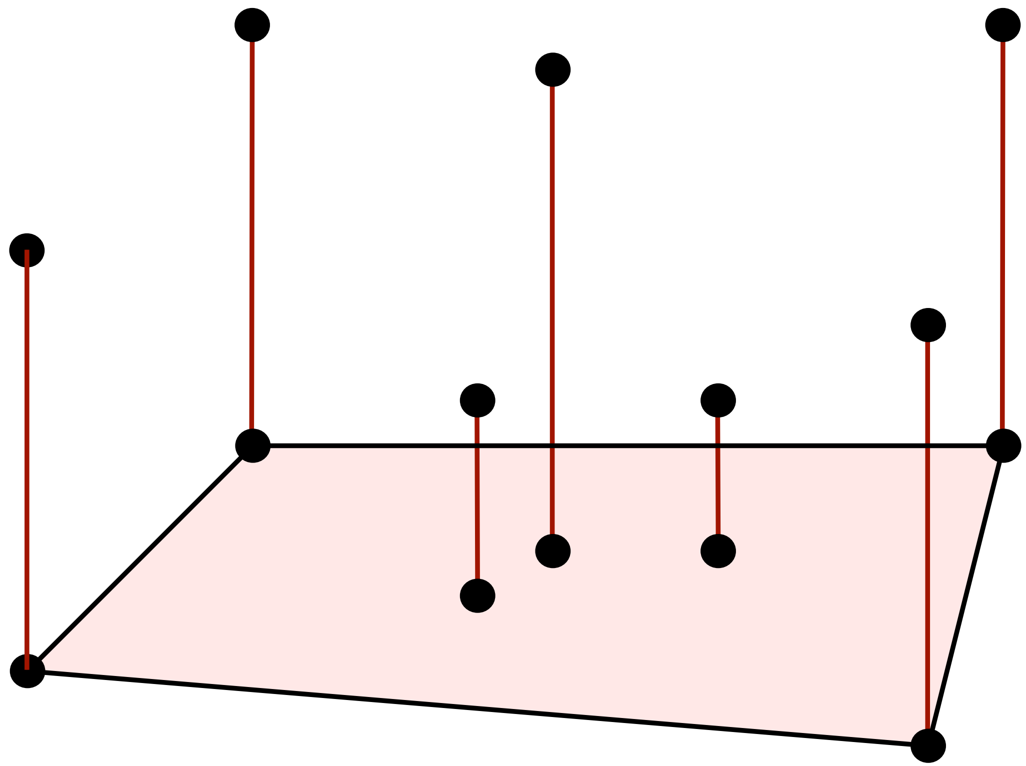

Regular triangulations

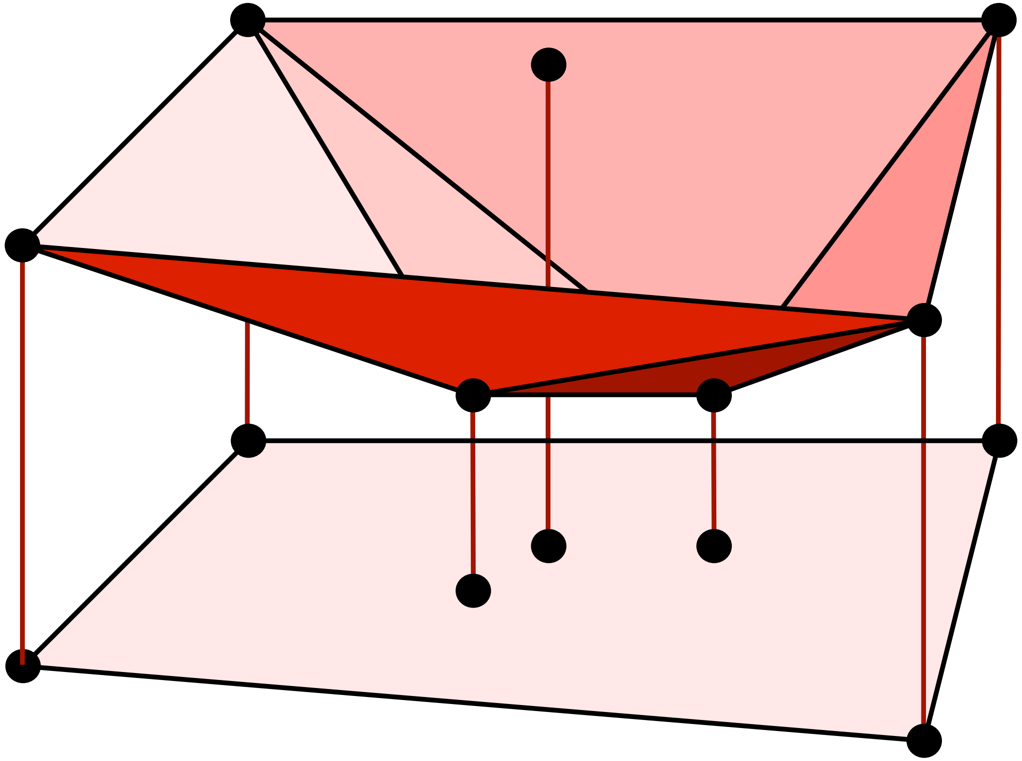

Regular triangulations

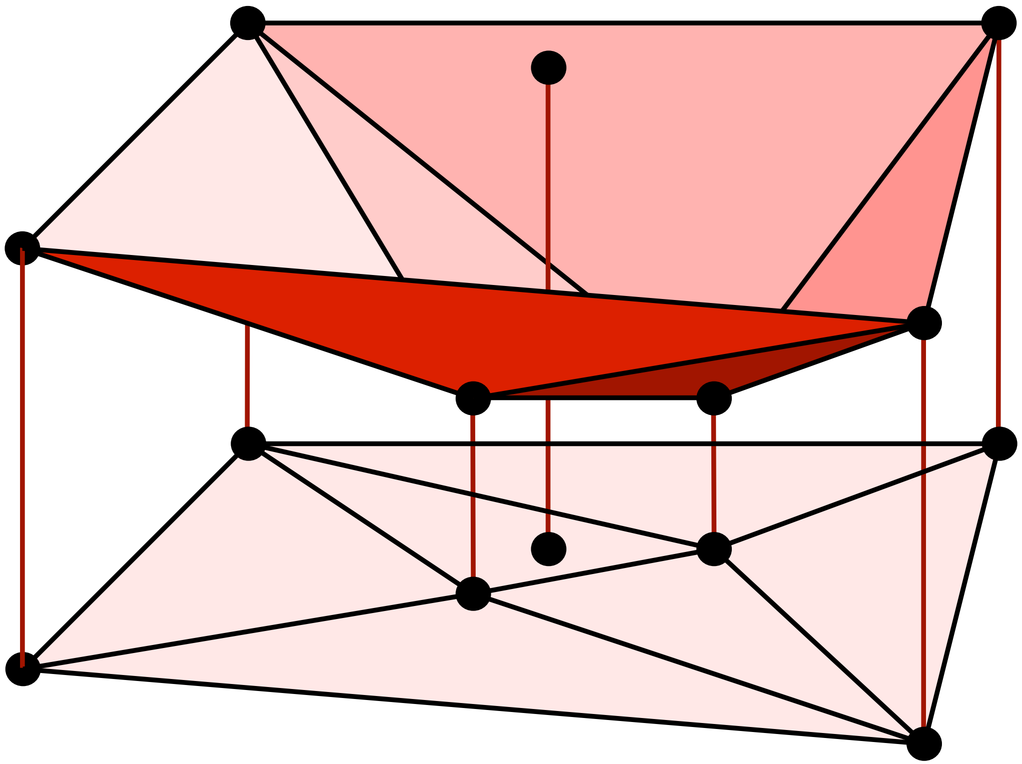

Example:

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

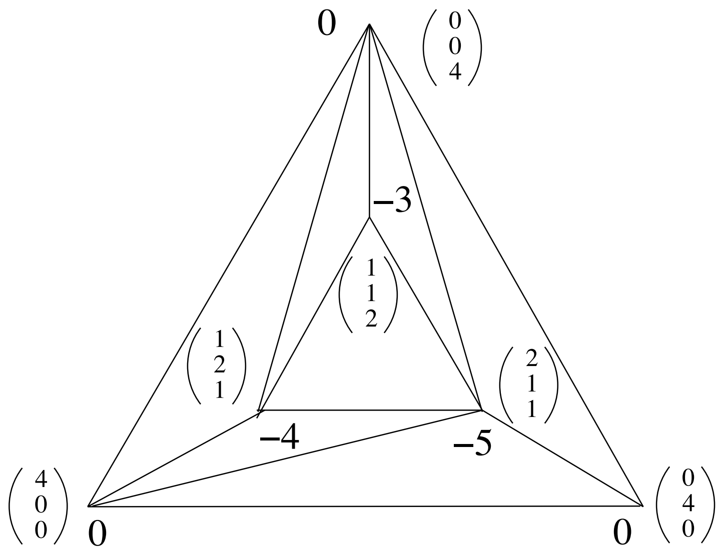

The cone triangulation associated with the lifting vector h. This shows a two-dimensional slice of the 3d-cone.

Another example:

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

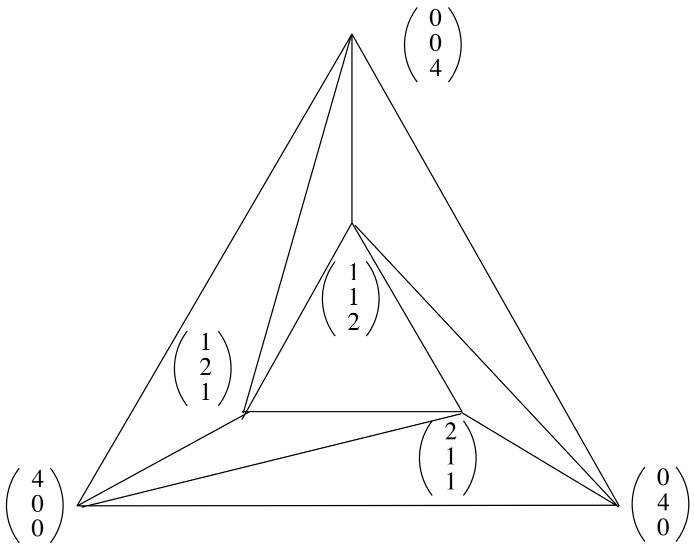

A triangulation not associated with any lifting vector h. That is to say, a non-regular triangulation.

The secondary polytope

Theorem [Gelfand-Kapranov-Zelevinskii, 1990]
The poset of **regular** polyhedral subdivisions of a point set (or acyclic vector configuration) A equals the face poset of a certain polytope of dimension $n-k$ ($n =$ number of points, $k =$ rank $=$ dimension $+1$).

This is called the secondary polytope of A.

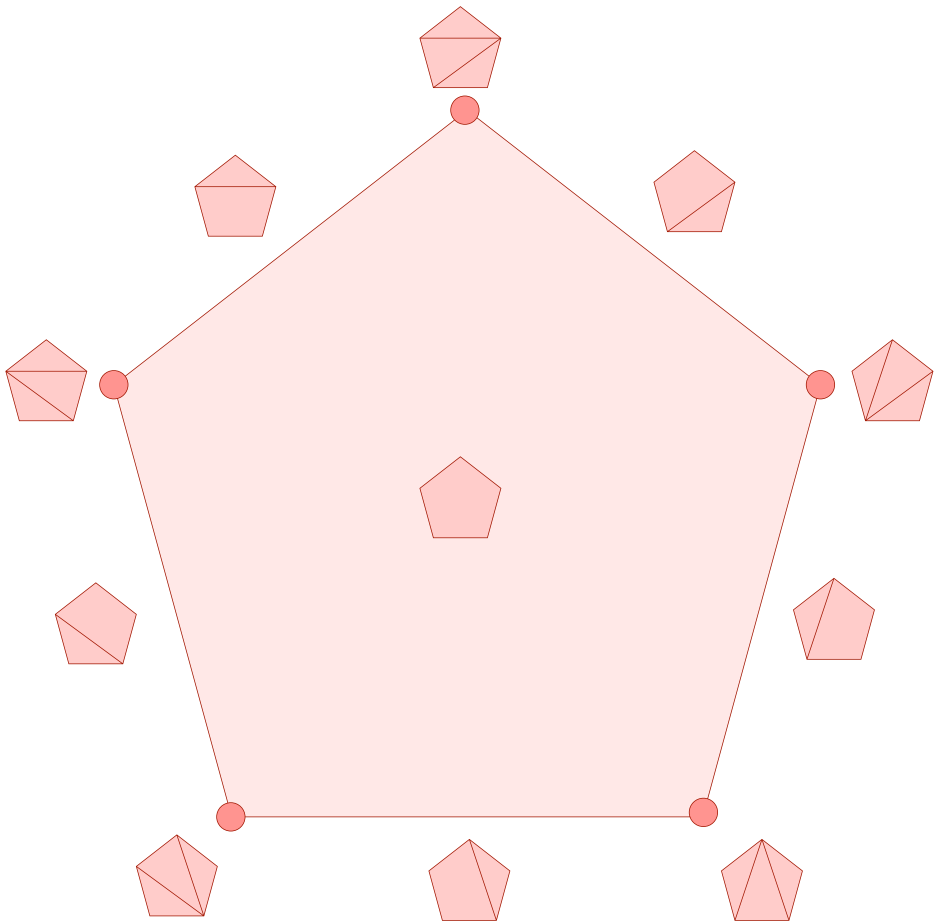

Secondary polytope of a pentagon (again a pentagon)

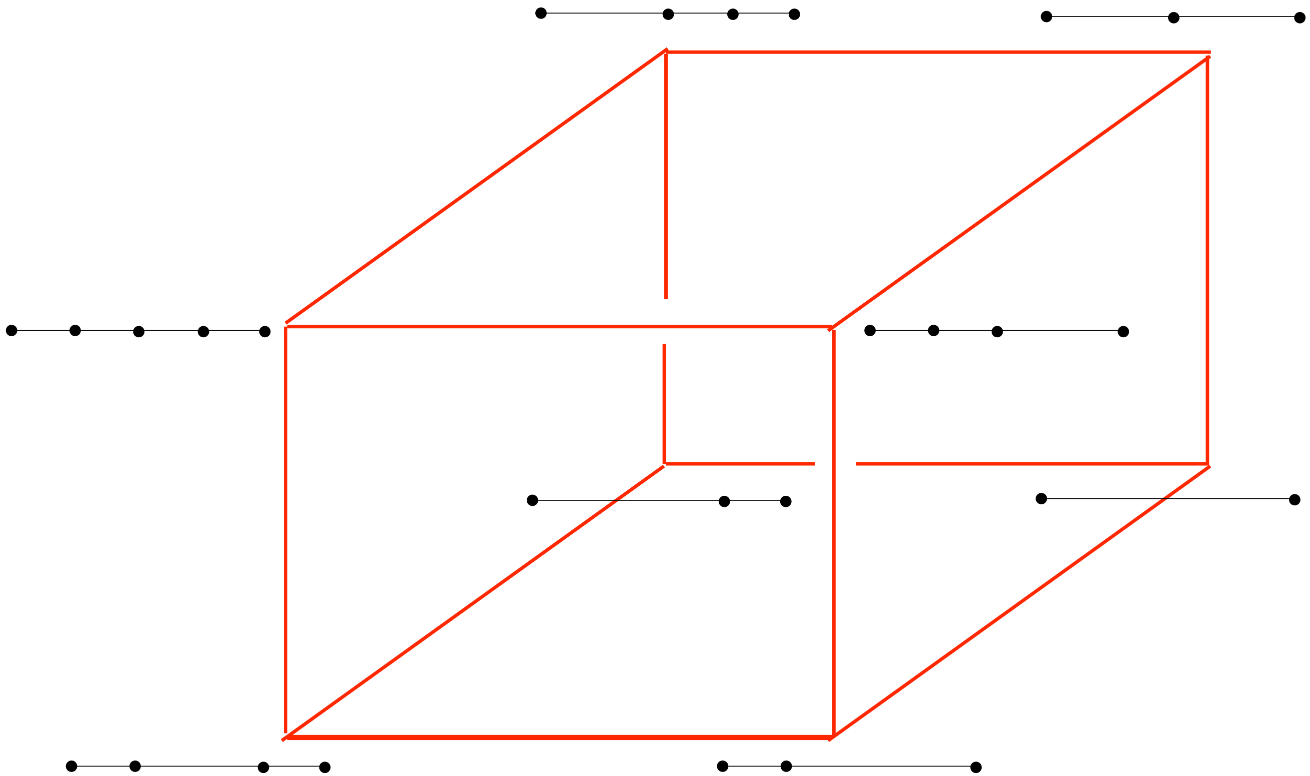

Secondary polytope of a 1-dimensional configuration (a cube)

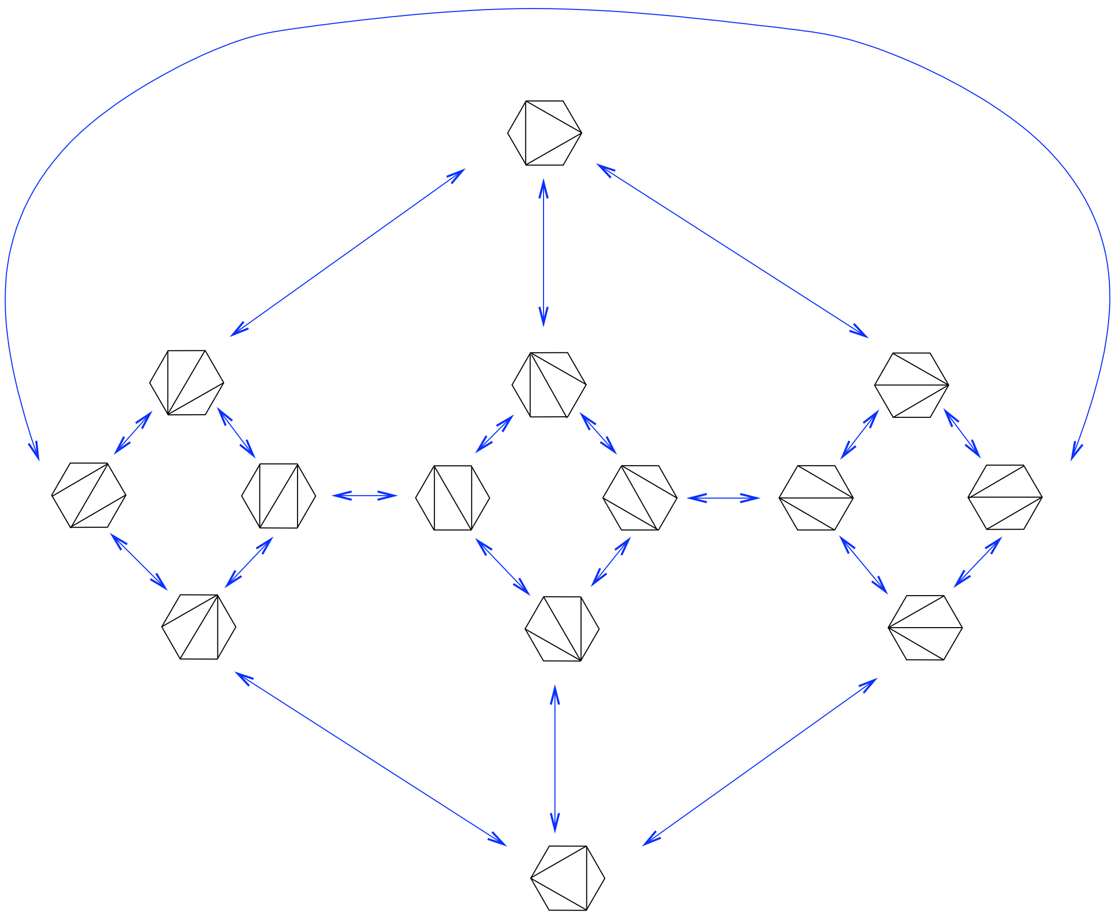

Secondary polytope of a hexagon (sort of)

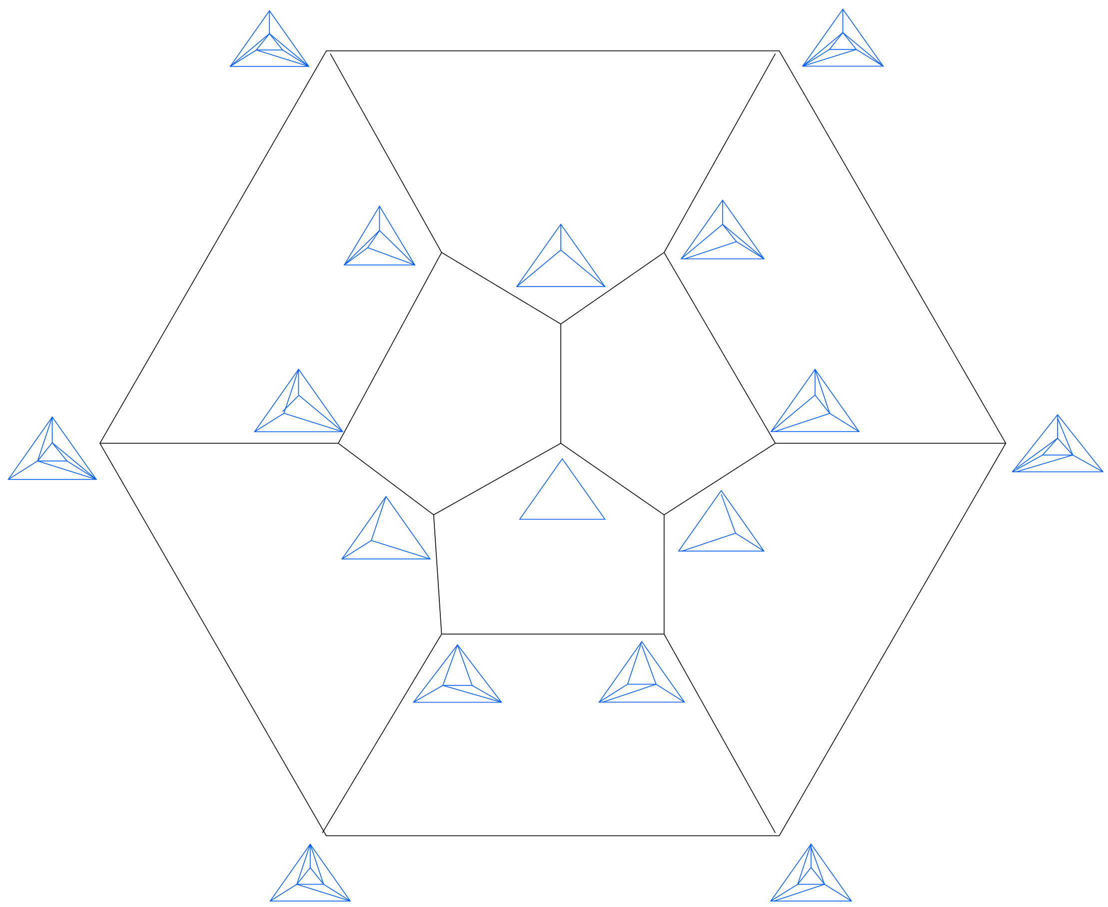

Secondary polytope of a point set with non-regular triangulations

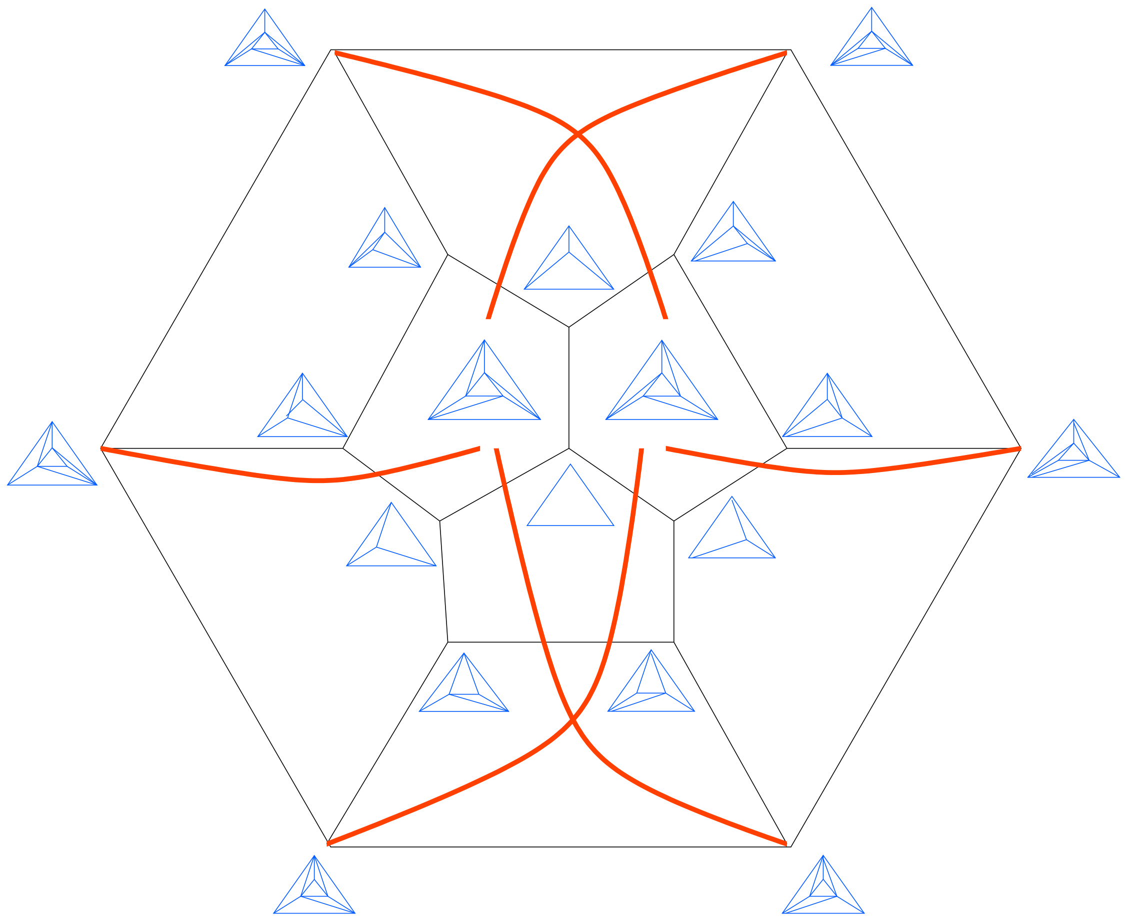

... and graph of all triangulations of the same point set

Bistellar flips

They are the "minimal possible changes" among triangulations. They correspond to edges in the secondary polytope.

Definition 1: In the poset of polyhedral subdivisions of a point set A, the minimal elements are the triangulations. It is a fact that if a subdivision is only refined by triangulations then it is refined by exactly two of them. We say these two triangulations differ by a flip.

That is to say, flips correspond to next to minimal elements in the poset of polyhedral subdivisions of A

Bistellar flips

They are the "minimal possible changes" among triangulations. They correspond to edges in the secondary polytope.

Definition 2: A circuit is a minimally (affinely/linearly) dependent set of (points/vectors). It is a fact that a circuit has exactly two triangulations. If a triangulation T of A contains one of the two triangulations of a circuit C, a flip on C consists on changing that part of T to become the other triangulation of C.

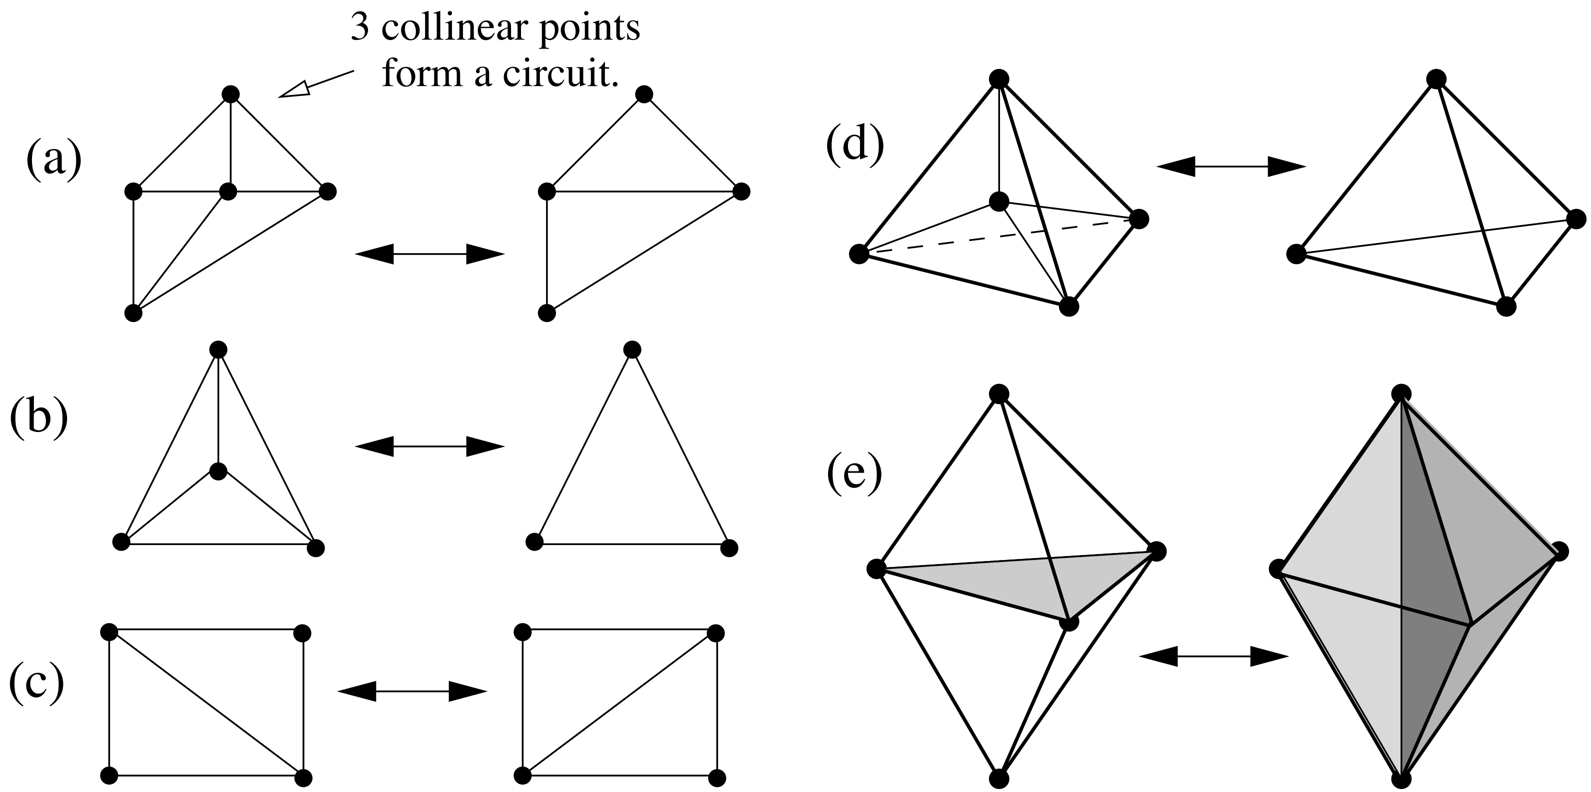

Triangulated circuits and their flips, in dimensions 2 and 3

The number of flips of a triangulation

Flips between regular triangulations correspond exactly to edges in the secondary polytope.

The number of flips of a triangulation

Flips between regular triangulations correspond exactly to edges in the secondary polytope. Hence,

Theorem: For every point set A, the graph of flips between regular triangulations is $n-d-1$-connected (in particular, every triangulation has at least $n-d-1$ flips).

The number of flips of a triangulation

Flips between regular triangulations correspond exactly to edges in the secondary polytope. Hence,

Theorem: For every point set A, the graph of flips between regular triangulations is $n-d-1$-connected (in particular, every triangulation has at least $n-d-1$ flips).

The same happens for non-regular triangulations if d or n are "small":

If $d=2$, then the graph is connceted [Lawson 1977] and every triangulation has at least $n-3$ flips [de Loera-S.-Urrutia, 1997] (but it is not known if the graph is $n-3$-connected).

The number of flips of a triangulation

Flips between regular triangulations correspond exactly to edges in the secondary polytope. Hence,

Theorem: For every point set A, the graph of flips between regular triangulations is $n-d-1$-connected (in particular, every triangulation has at least $n-d-1$ flips).

The same happens for non-regular triangulations if d or n are “small”:

If $d=3$ **and** the points are in convex position, then every triangulation has at least $n-4$ flips [de Loera-S.-Urrutia, 1997] (but it is not known if the graph is connected).

The number of flips of a triangulation

Flips between regular triangulations correspond exactly to edges in the secondary polytope. Hence,

Theorem: For every point set A, the graph of flips between regular triangulations is $n-d-1$-connected (in particular, every triangulation has at least $n-d-1$ flips).

The same happens for non-regular triangulations if d or n are "small":

If $n \leq d+4$, then every triangulation has at least 3 flips and the graph is 3-connected [Azaola-S., 2001].

The number of flips of a triangulation

But:

1. In **dimension 3**, there are triangulations with arbitrarily large n and only O($\sqrt{n}$) flips [S., 1999].

2. In **dimension 4**, there are triangulations with arbitrarily large n and only O(1) flips [S., 1999].

3. In **dimension 5**, there are polytopes with disconnected graph of flips [S., 2004].

4. In **dimension 6**, there are triangulations with arbitrarily large n and no flips [S., 2000].

2. A disconnected graph of triangulations

Point sets with a disconnected graph of triangulations

The first example known had 324 points and dimension 6 [S. 2000].

Point sets with a disconnected graph of triangulations

The first example known had 324 points and dimension 6 [S. 2000].

The smallest example known has 26 points and dimension 5 [S. 2004].

Point sets with a disconnected graph of triangulations

The first example known had 324 points and dimension 6 [S. 2000].

The smallest example known has 26 points and dimension 5 [S. 2004].

Here we are going to describe a somewhat simpler example with 50 points and dimension 5.

Point sets with a disconnected graph of triangulations

The first example known had 324 points and dimension 6 [S. 2000].

The smallest example known has 26 points and dimension 5 [S. 2004].

Here we are going to describe a somewhat simpler example with 50 points and dimension 5.

It is not known whether examples exist in dimensions 3 or 4.

Ingredient 1: Triangulations of a prism

Triangulations of a prism

We call prism of dimension d the product of a $d-1$-simplex and a segment.

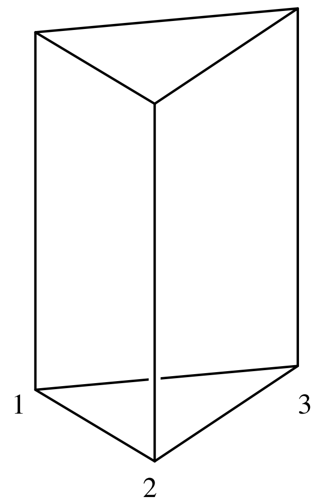

The triangular prism $\Delta^2 \times I$

Triangulations of a prism

To triangulate the d-prism, start with the top base and join to the d bottom vertices, one by one.

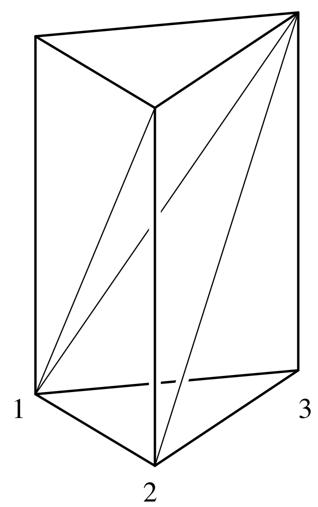

A triangulation of $\Delta^2 \times I$

All triangulations of the d-prism can be obtained in this way.

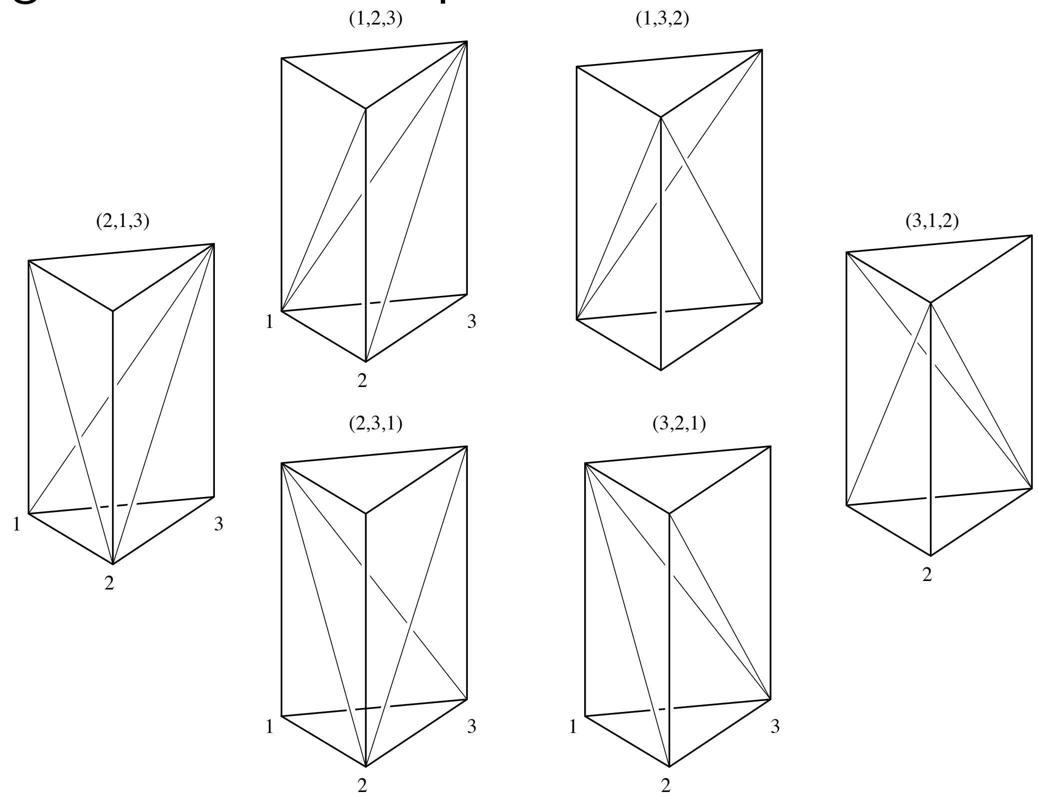

There are $d!$ of them, in bijection with the $d!$ orderings on the bottom vertices.

Triangulations of a prism

The secondary polytope of the d-prism is the permutahedron of dimension $d-1$: the convex hull of the points $(\sigma_1, \ldots, \sigma_d)$ where σ runs over all permutations of $(1, 2, \ldots, d)$.

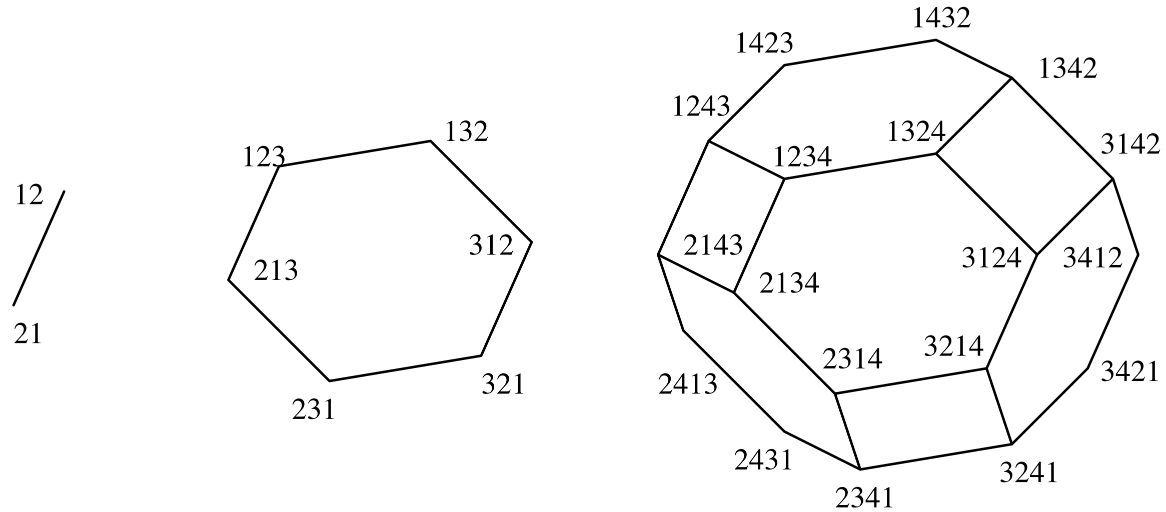

Triangulations of a prism

Put differently, triangulations of $\Delta^{d-1} \times I$ are in bijection to acyclic orientations of the complete graph K_d.

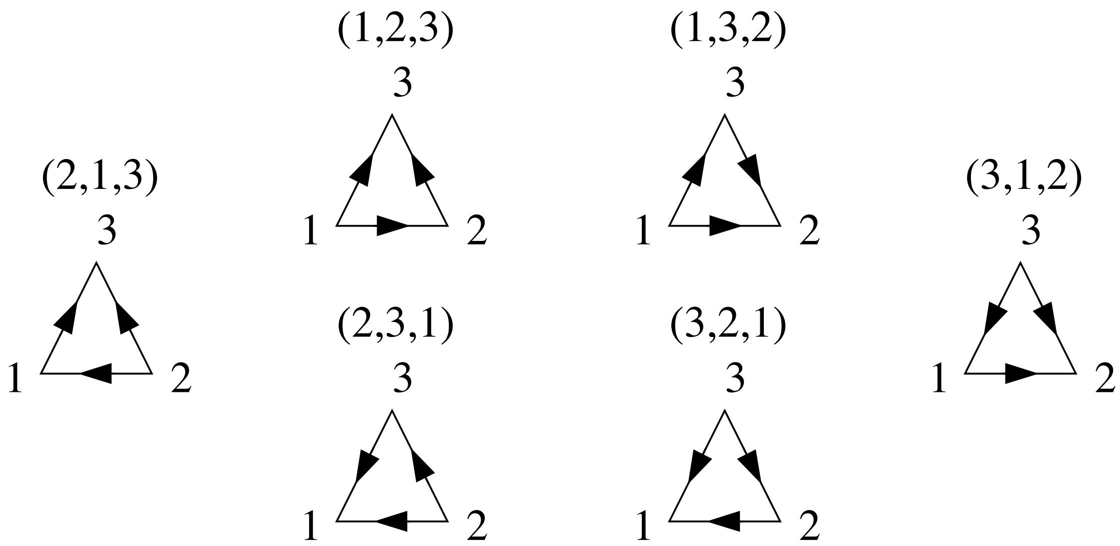

Triangulations of a prism

Flips correspond to reversals of a single edge, whenever this does not create a cycle.

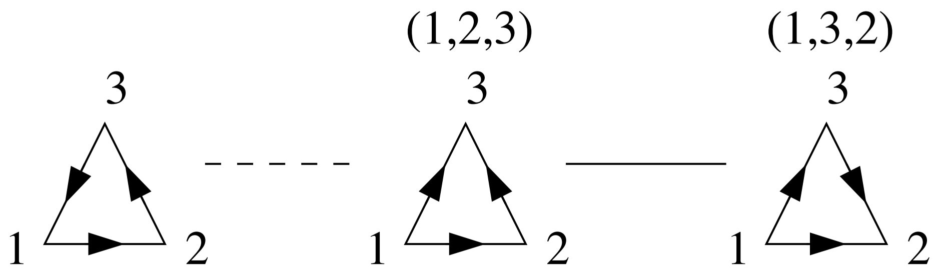

Not a flip A flip

Ingredient 2: Locally acyclic orientations

Definition: A locally acyclic orientation (l.a.o.) in a simplicial complex K is an orientation of its graph which is acyclic on every simplex.

A reversible edge in a l.a.o. is an edge whose reversal creates no local cycle. The graph of l.a.o.'s of K has as nodes all the l.a.o.'s and as edges the single-edge reversals.

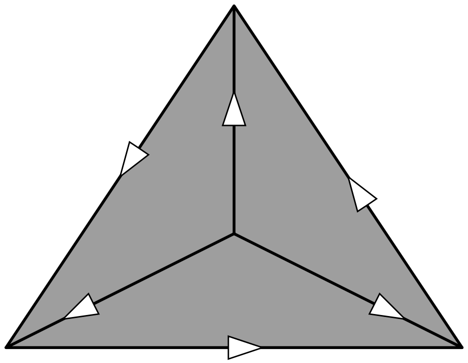

A l.a.o. with three reversible edges

Example: a simplex

If K is a simplex with k vertices, it has $k!$ l.a.o.'s. The graph is the graph of the permutahedron.

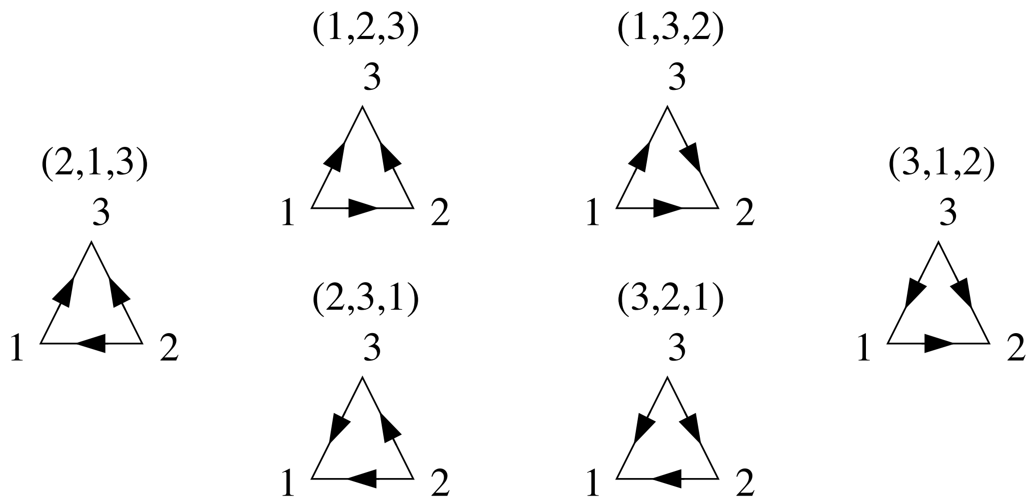

Example: a l.a.o. of the boundaty of an octahedron

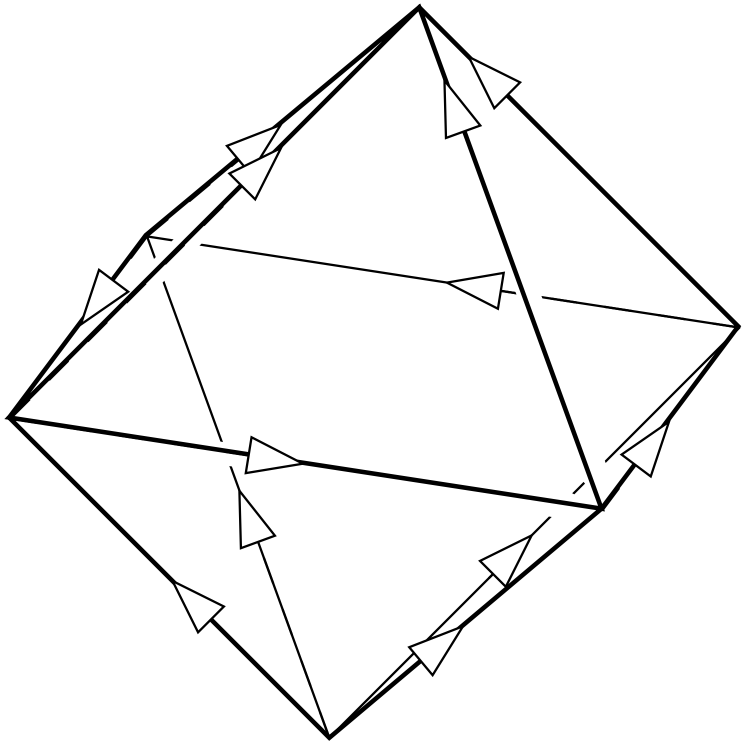

It has a cycle. The only reversible edges are the **four** edges in the cycle.

Crucial remark

There is a bijective correspondence between l.a.o. of K and triangulations that refine $K \times I$.

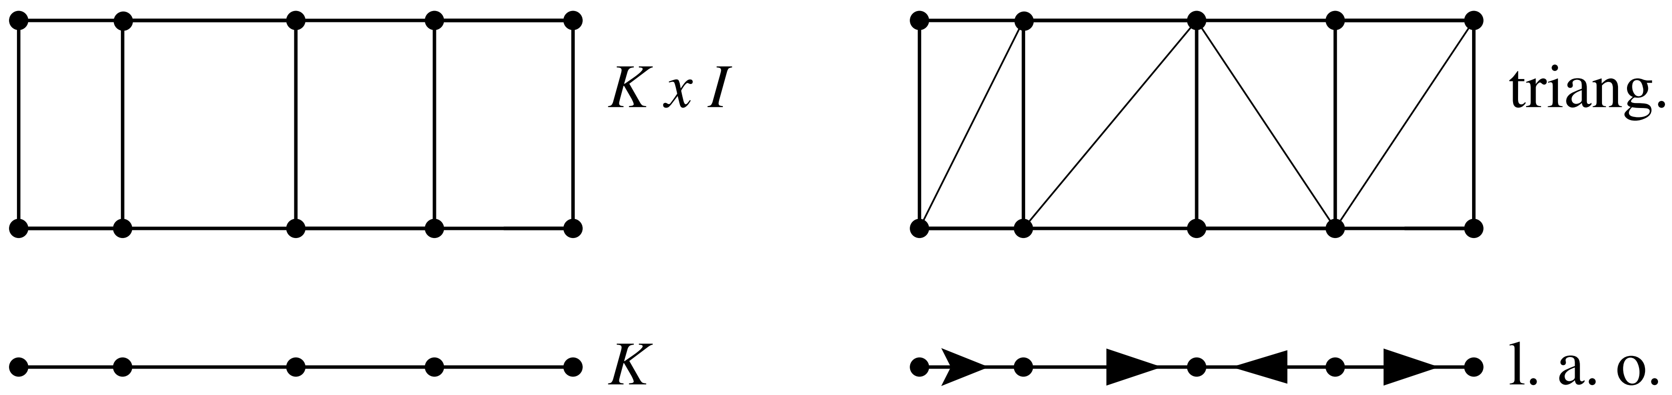

Triangulating boundary subcomplexes

Suppose now that K is a simplicial subcomplex of the face complex of a polytope P. Then, there is a bijection betweem l.a.o.'s of K and triangulations of the subcomplex $K \times I$ in $P \times I$.

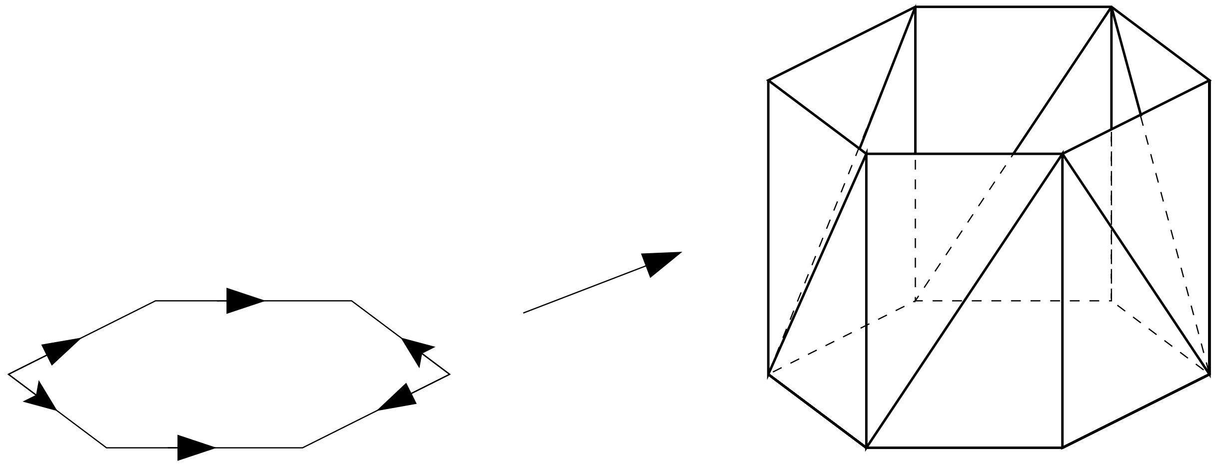

Triangulating boundary subcomplexes

Suppose now that K is a simplicial subcomplex of the face complex of a polytope P. Then, there is a bijection betweem l.a.o.'s of K and triangulations of the subcomplex $K \times I$ in $P \times I$.

Moreover:

- Every triangulation of $P \times I$ in particular triangulates $K \times I \Rightarrow$ "**there is a map** ϕ : **triangulations**$(P \times I) \rightarrow$ **l.a.o's**(K)".

Triangulating boundary subcomplexes

Suppose now that K is a simplicial subcomplex of the face complex of a polytope P. Then, there is a bijection betweem l.a.o.'s of K and triangulations of the subcomplex $K \times I$ in $P \times I$.

Moreover:

- Every triangulation of $P \times I$ in particular triangulates $K \times I \Rightarrow$ "**there is a map** ϕ : **triangulations**$(P \times I) \to$ **l.a.o's**(K)**"**.

- If two triangulations of $P \times I$ differ by a **flip**, the corresponding l.a.o.'s coincide or differ by a **single-edge reversal** $\Rightarrow$ **the map** ϕ **is continuous**, as a map between the graph of triangulations of $P \times I$ and the graph of l.a.o.'s of K.

Triangulating boundary subcomplexes

Suppose now that K is a simplicial subcomplex of the face complex of a polytope P. Then, there is a bijection betweem l.a.o.'s of K and triangulations of the subcomplex $K \times I$ in $P \times I$.

Corollary: *If the image of ϕ is a disconnected subgraph of l.a.o.'s of K, then the graph of triangulations of $P \times I$ is not connected.*

Ingredient 3: The 24-cell

The 24-cell is one of the six regular polytopes in dimension four. It is self-dual.

Its faces are 24 octahedra, 96 triangles, 96 edges and 24 vertices. There are six octahedra incident to each vertex.

One coordinatization consists of the following 24 vertices:

- The sixteen points $(\pm 1, \pm 1, \pm 1, \pm 1)$.
- The eight points $(\pm 2, 0, 0, 0)$, $(0, \pm 2, 0, 0)$, $(0, 0, \pm 2, 0)$, $(0, 0, 0, \pm 2)$.

A 24-cell from a 4-cube

It can be constructed from a 4-cube (16 vertices) by adding a point beyond each of its eight 3-cubes. Each new point "divides" a 3-cube into "six half-octahedra", and these 6×8 half octahedra are glued in pairs.

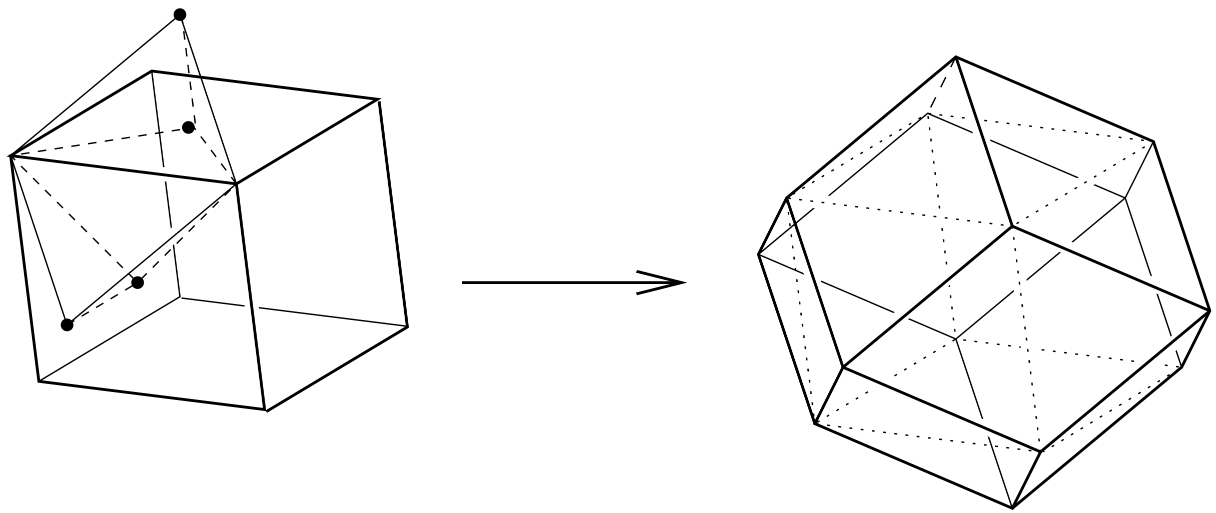

A 3d analogue of the construction of the 24-cell from a 4-cube

A l.a.o. of the 2-skeleton of the 24-cell

Let K be the complex consisting of the 96 triangles in the 24-cell (the "2-skeleton"). To define a l.a.o. in K, we consider the boundary of the 4-cube as consisting of two (oriented) cycles of four 3-cubes each (a "3-sphere obtained by gluing two solid tori along the boundary").

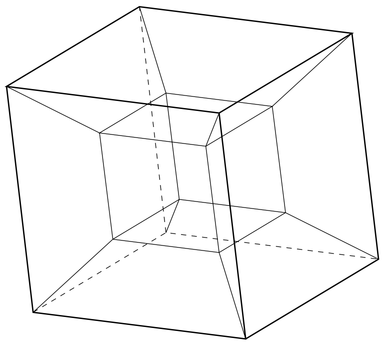

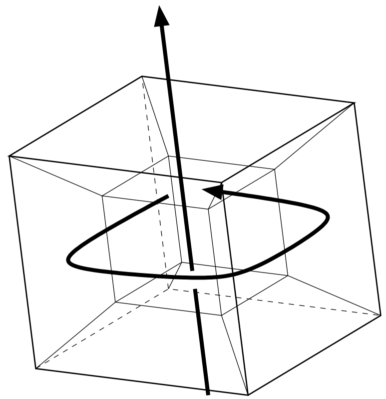

We orient each edge in the 24-cell in the way "most consistent" with the cycles:

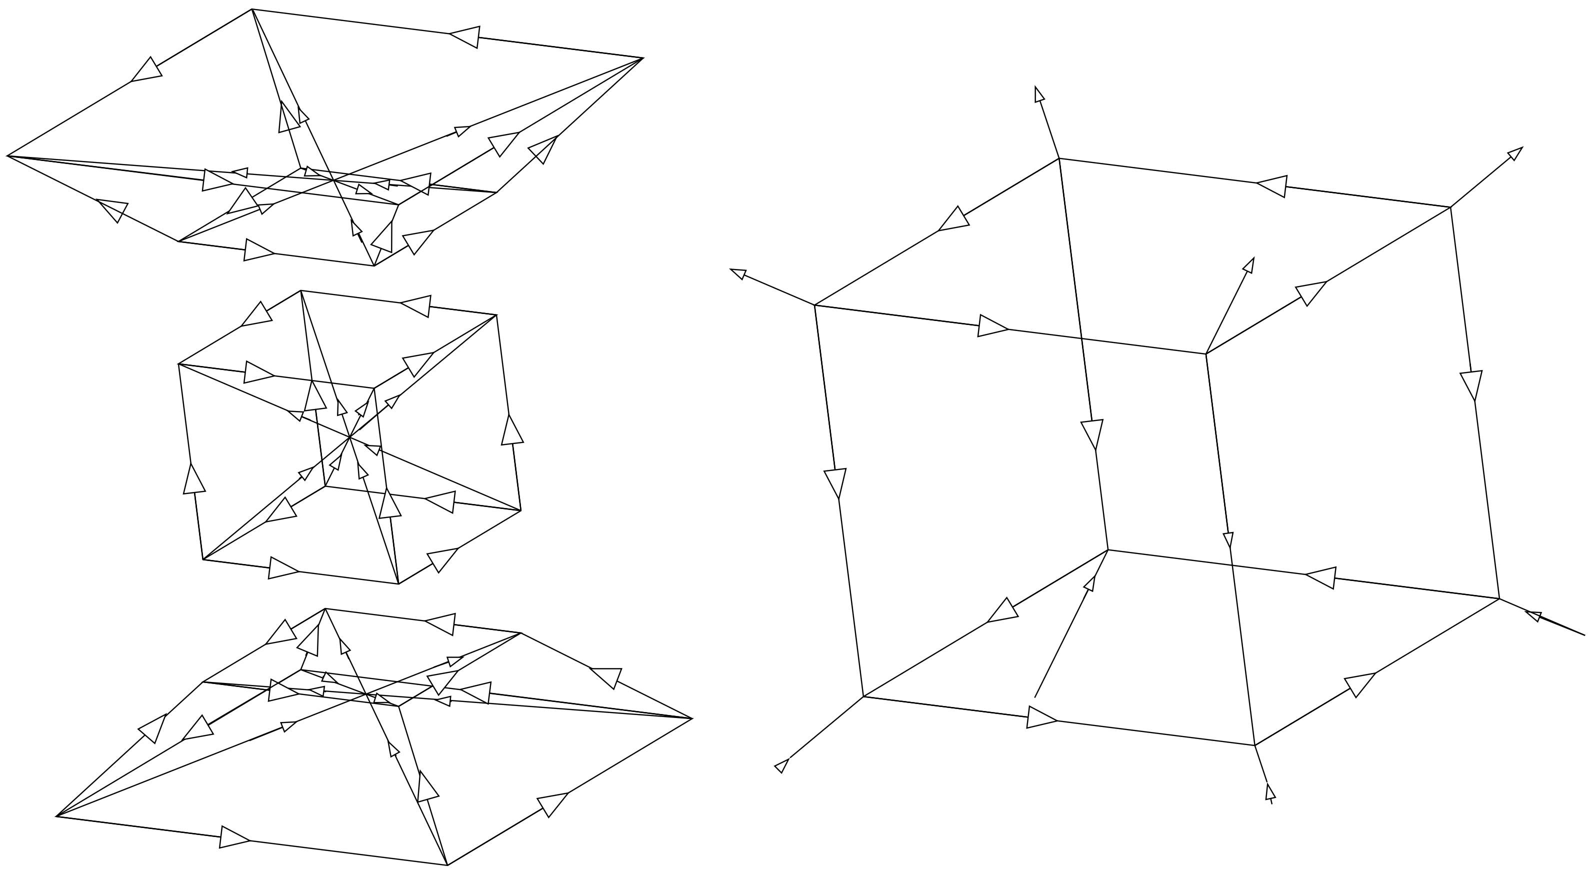

the “vertical cycle”

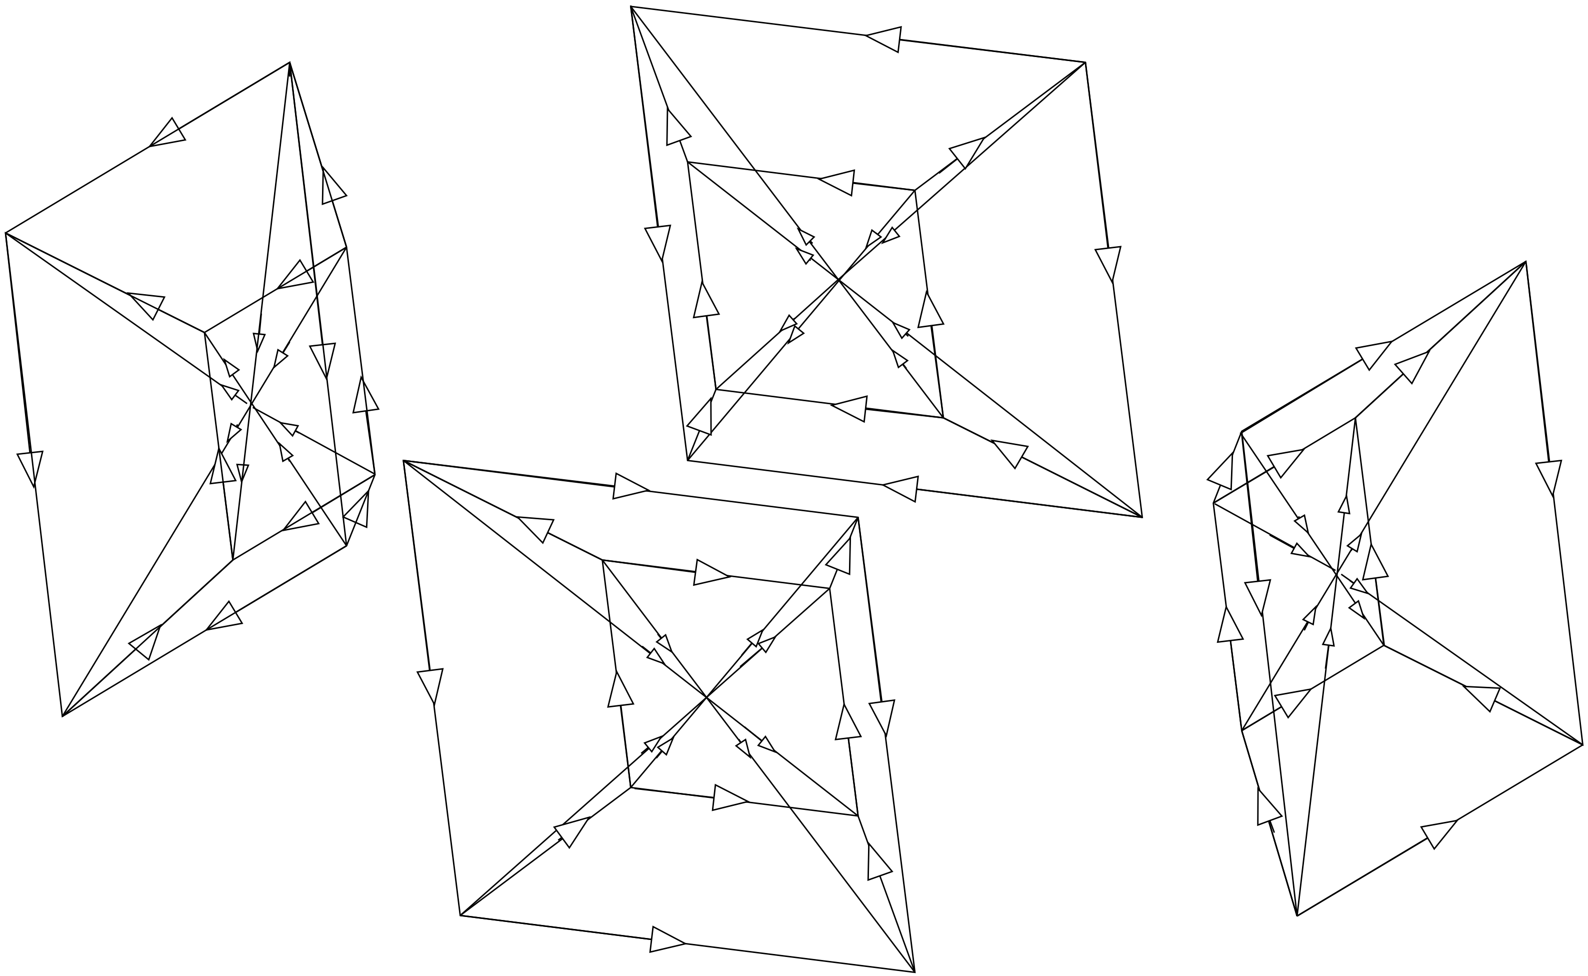

the “horizontal cycle”

The graph of l.a.o.'s of the 2-skeleton of the 24-cell

It turns out this is a locally acyclic orientation **with no reversible edges at all**. (Each of the 24 octahedra gets the l.a.o. with a cycle of reversible edges, but every edge is in the cycle of only one of the three octahedra it belongs to).

The graph of l.a.o.'s of the 2-skeleton of the 24-cell

It turns out this is a locally acyclic orientation **with no reversible edges at all**. (Each of the 24 octahedra gets the l.a.o. with a cycle of reversible edges, but every edge is in the cycle of only one of the three octahedra it belongs to).

Hence, **the graph of l.a.o.'s of the 2-skeleton of a 24-cell is not connected**. Actually, it has at least thirteen connected components, twelve of them consisting of an isolated vertex.

The graph of l.a.o.'s of the 2-skeleton of the 24-cell

It turns out this is a locally acyclic orientation **with no reversible edges at all**. (Each of the 24 octahedra gets the l.a.o. with a cycle of reversible edges, but every edge is in the cycle of only one of the three octahedra it belongs to).

Hence, **the graph of l.a.o.'s of the 2-skeleton of a 24-cell is not connected**. Actually, it has at least thirteen connected components, twelve of them consisting of an isolated vertex.

. . . but, does this imply that the graph of triangulations of 24-cell$\times I$ is not connected?

The graph of l.a.o.'s of the 2-skeleton of the 24-cell

It turns out this is a locally acyclic orientation **with no reversible edges at all**. (Each of the 24 octahedra gets the l.a.o. with a cycle of reversible edges, but every edge is in the cycle of only one of the three octahedra it belongs to).

Hence, **the graph of l.a.o.'s of the 2-skeleton of a 24-cell is not connected**. Actually, it has at least thirteen connected components, twelve of them consisting of an isolated vertex.

. . . but, does this imply that the graph of triangulations of 24-cell$\times I$ is not connected?

Remember that: "there is a map ϕ : triangulations$(P \times I) \rightarrow$ l.a.o's(K). If the image of ϕ is a disconnected subgraph of l.a.o.'s of K, then the graph of triangulations of $P \times I$ is not connected".

A point set with disconnected graph of triangulations

We still need to check that the l.a.o. we have is "in the image of ϕ". That is to say, that the triangulation of $K \times I$ it represents can be extended to a triangulation of $P \times I$ ($P =$ 24-cell).

A point set with disconnected graph of triangulations

We still need to check that the l.a.o. we have is "in the image of ϕ". That is to say, that the triangulation of $K \times I$ it represents can be extended to a triangulation of $P \times I$ ($P =$ 24-cell).

. . . But we are allowed to add points to the interior of the configuration!

A point set with disconnected graph of triangulations

We still need to check that the l.a.o. we have is "in the image of ϕ". That is to say, that the triangulation of $K \times I$ it represents can be extended to a triangulation of $P \times I$ ($P =$ 24-cell).

. . . But we are allowed to add points to the interior of the configuration!

Theorem [S. 2004] Let A consist of the 24 vertices of the 24-cell, together with the origin. Let K be the 2-skeleton of the 24-cell. Then, the triangulation of $K \times I$ represented by the above l.a.o. of K can be extended to a triangulation of $A \times I$. Hence, the graph of triangulations of $A \times I$ has at least thirteen connected components, each with at least 3^{48} triangulations.

A point set with disconnected graph of triangulations

Moreover:

- The triangulations we have constructed are unimodular, which has interesting algebro-geometric consequences (see part 3).

- The construction can be put into convex position: there is a 5-polytope with 50 vertices with a disconnected graph of triangulations.

- The construction can be iterated: for every k there is a 5-polytope with $26 + 24k$ vertices whose graph of triangulations has at least 13^k connected components.

Interlude — Viro's Theorem

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.

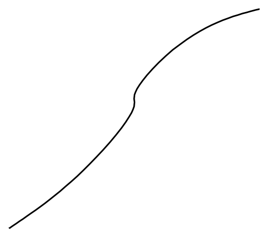

A pseudoline. Its complement has one component, homeomorphic to an open circle. The picture only shows the "affine part"; think the two ends as meeting at infinity.

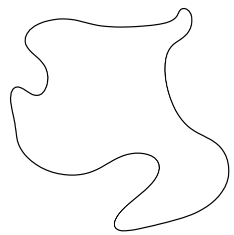

An oval. Its interior is a (topological) circle and and its exterior is a Möbius band.

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)

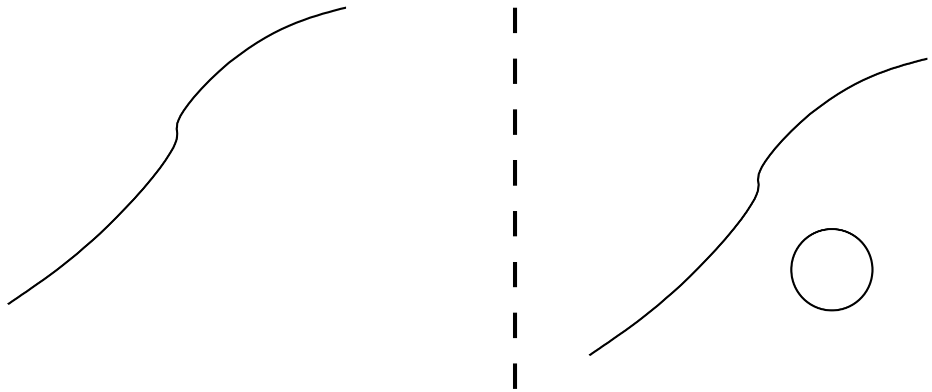

Two configurations are possible in degree 3

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)

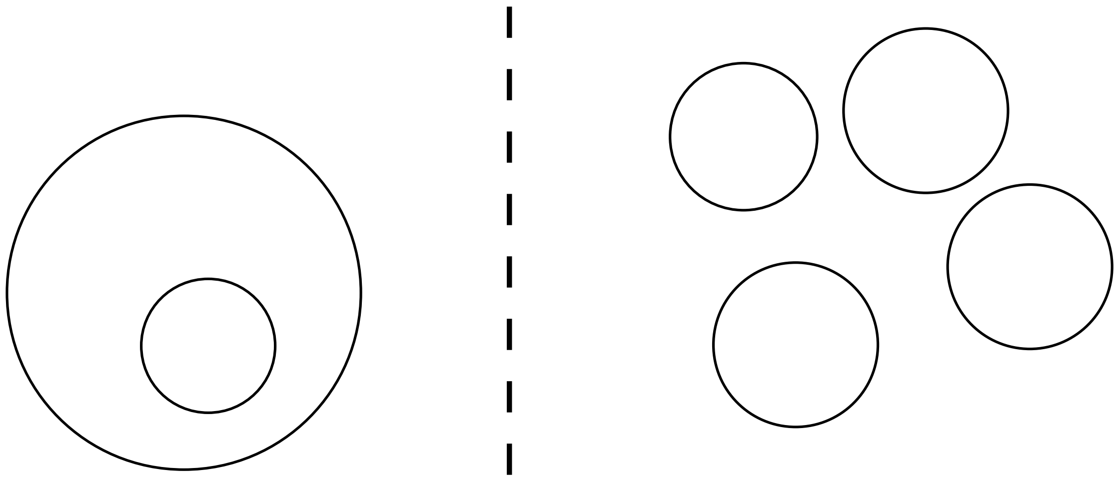

Six configurations are possible in degree 4. Only the maximal ones are shown.

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)

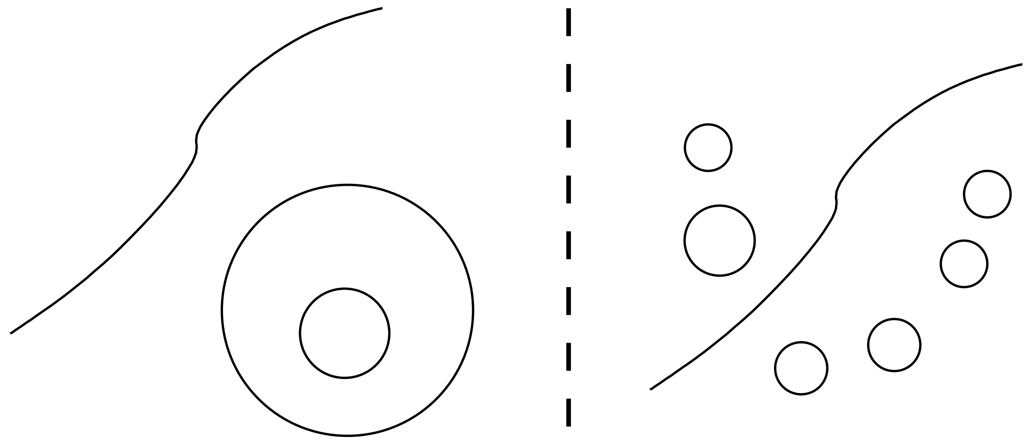

Eight configurations are possible in degree 5. Only the maximal ones are shown.

All that was known when Hilbert posed the problem, but the classification of non-singular real algebraic curves of degree six was not completed until the 1960's [Gudkov]. There are 56 types degree six curves, three with 11 ovals:

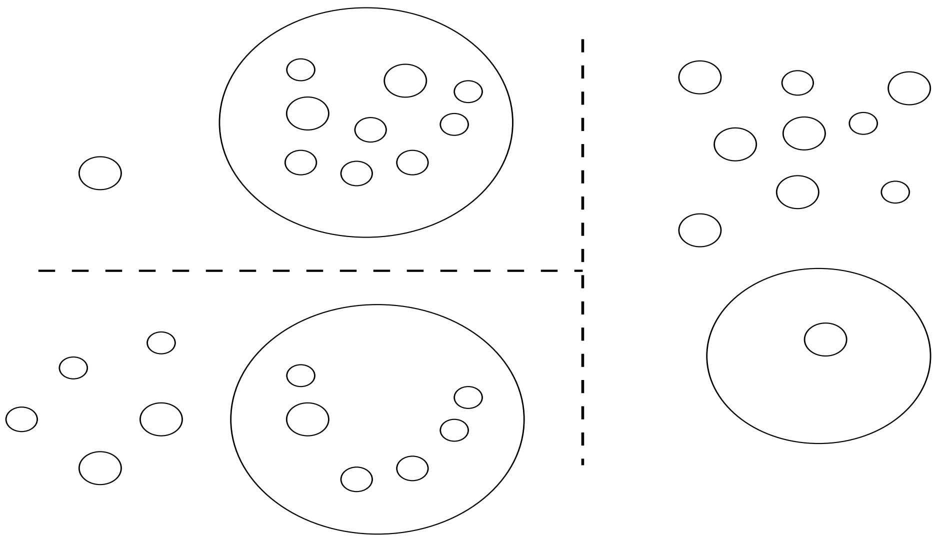

Dimension 7 was solved by **Viro**, in 1984 with a method that involves triangulations.

Viro's method:

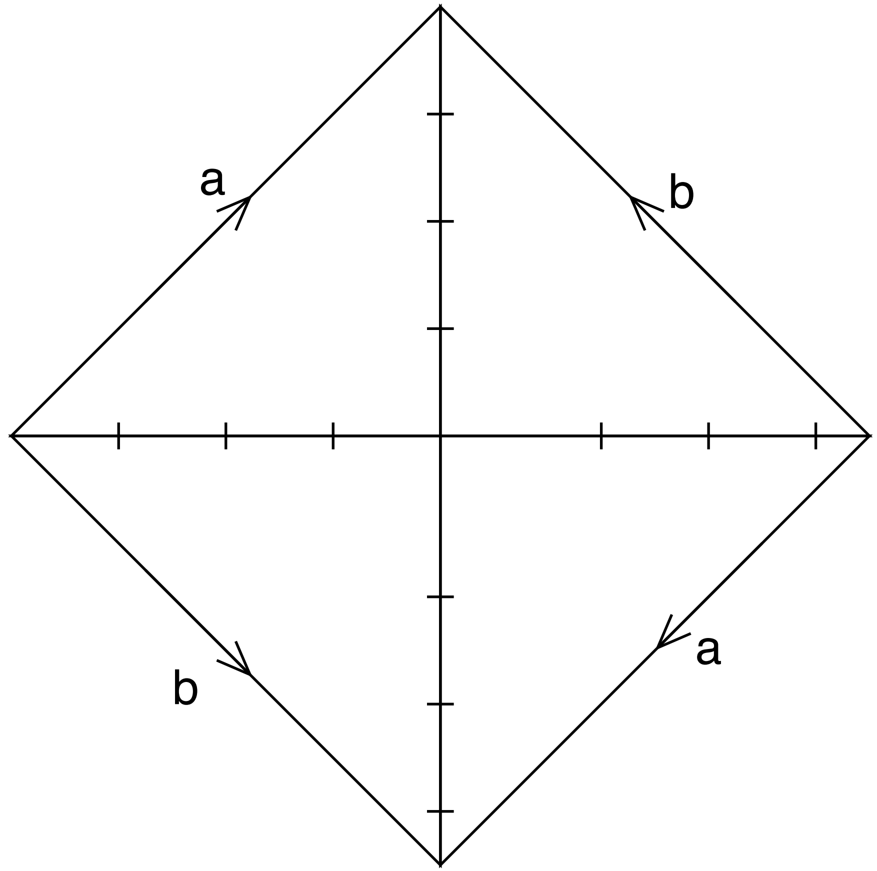

For any given d, construct a topological model of the projective plane by gluing the triangle $(0,0),(d,0),(0,d)$ and its symmetric copies in the other quadrants:

Viro's method:

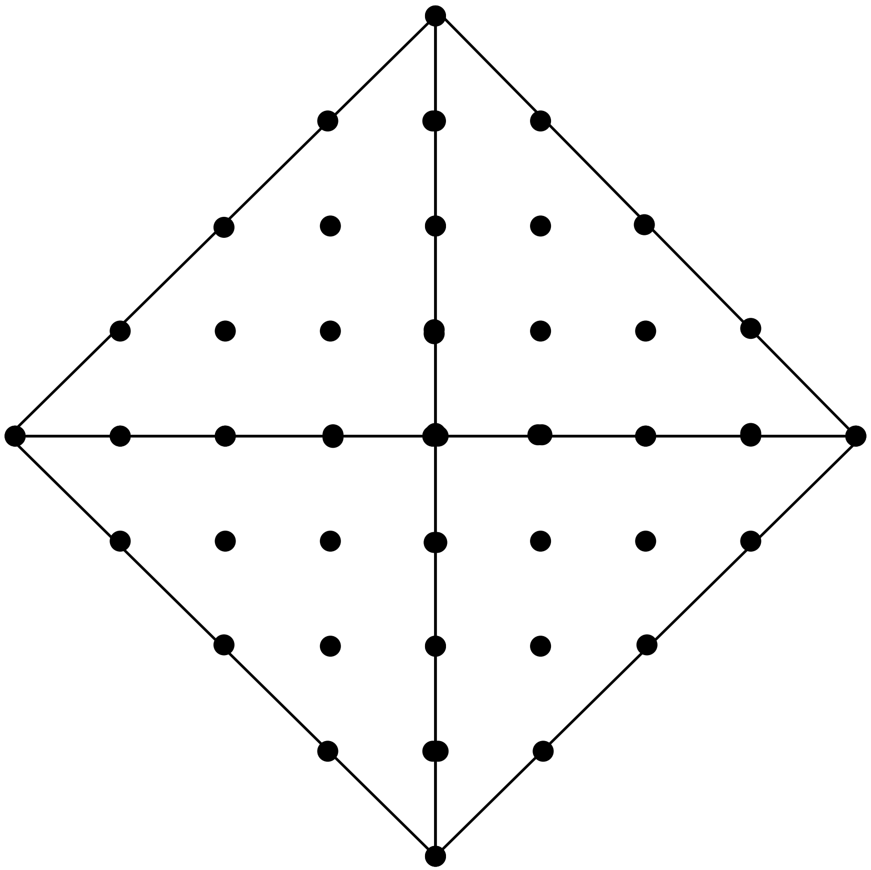

Consider as point set all the integer points in your rhombus (remark: those in a particular orthant are related to the possible homogeneous monomials of degree d in three variables).

Viro's method:

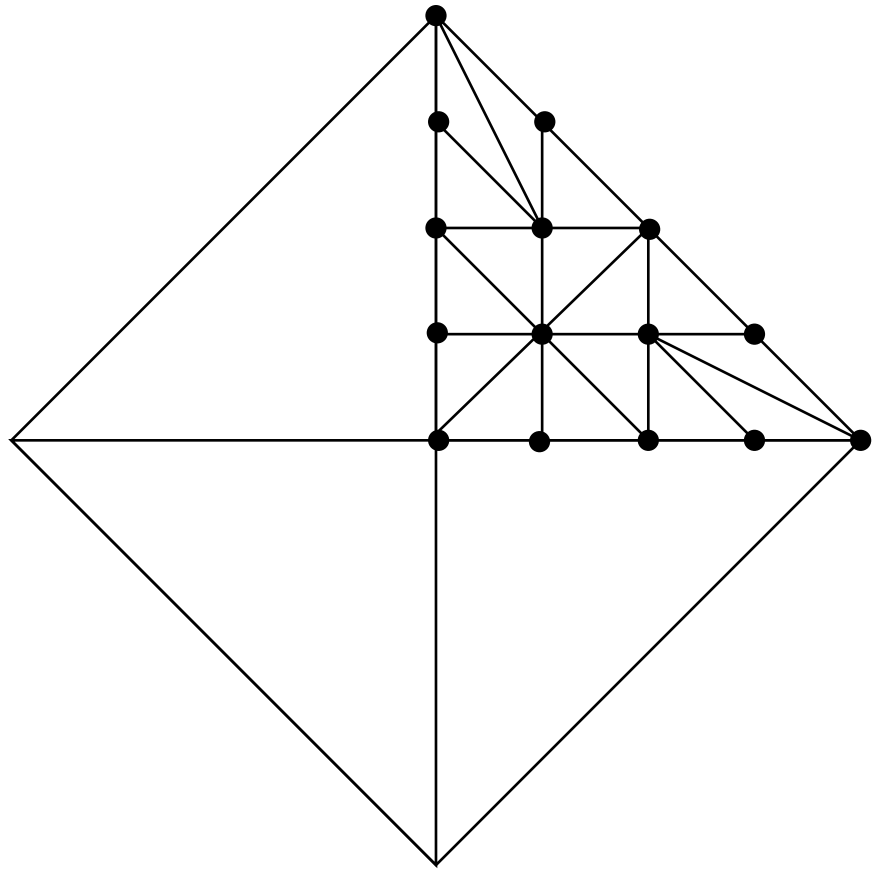

Triangulate the positive orthant arbitrarily . . .

Viro's method:

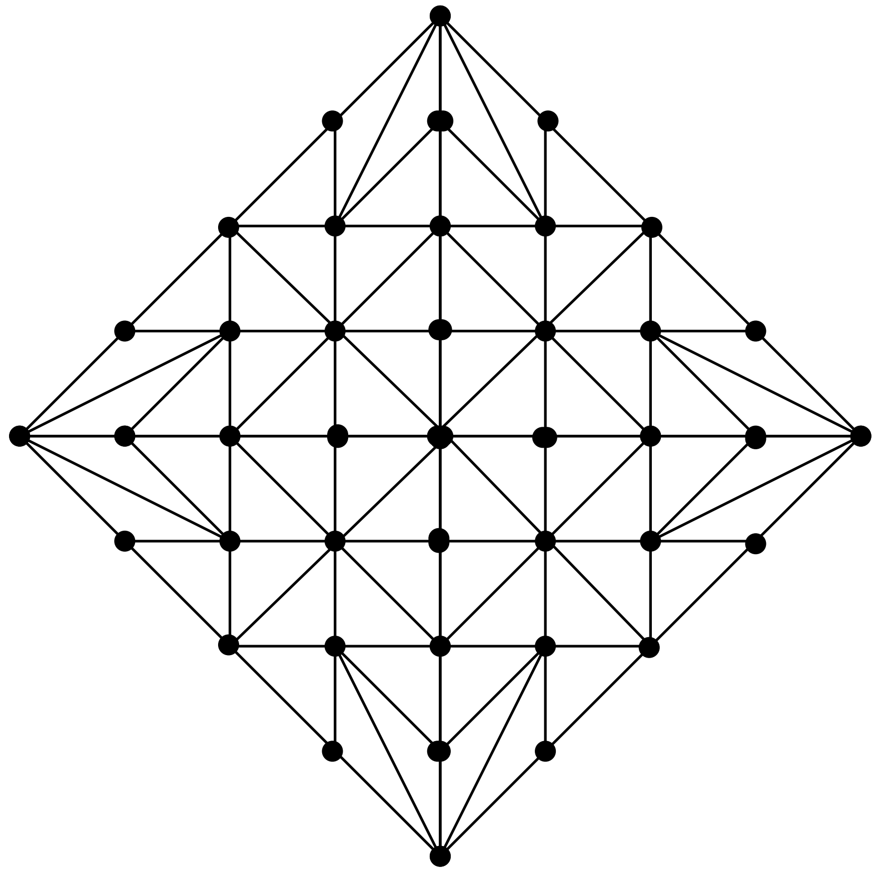

Triangulate the positive quadrant arbitrarily . . .

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

Viro's method:

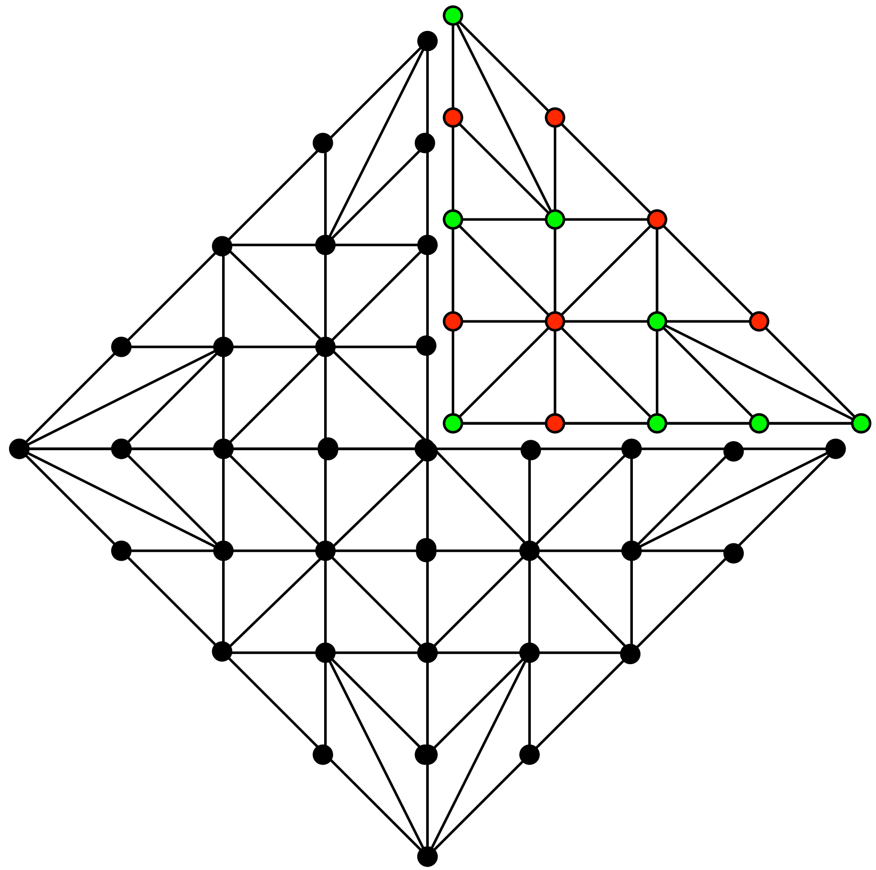

Choose arbitrary signs for the points in the first quadrant

Viro's method:

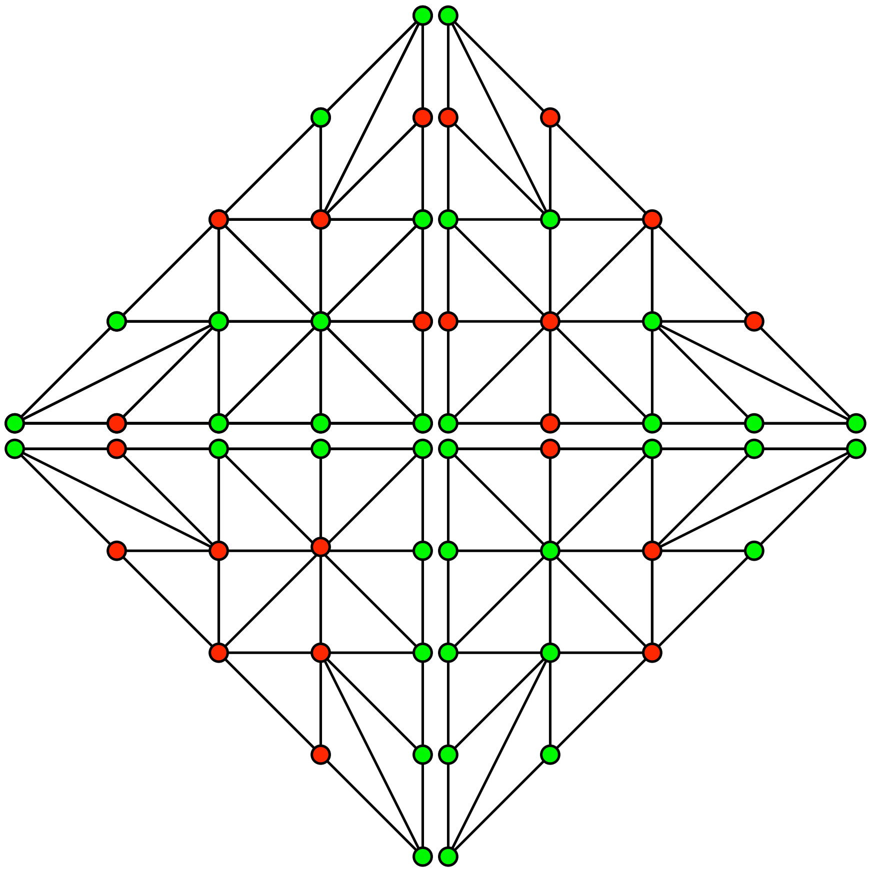

Choose arbitrary signs for the points in the first quadrant . . . and replicate them to the other three quadrants, taking parity of the corresponding coordinate into account.

Viro's method:

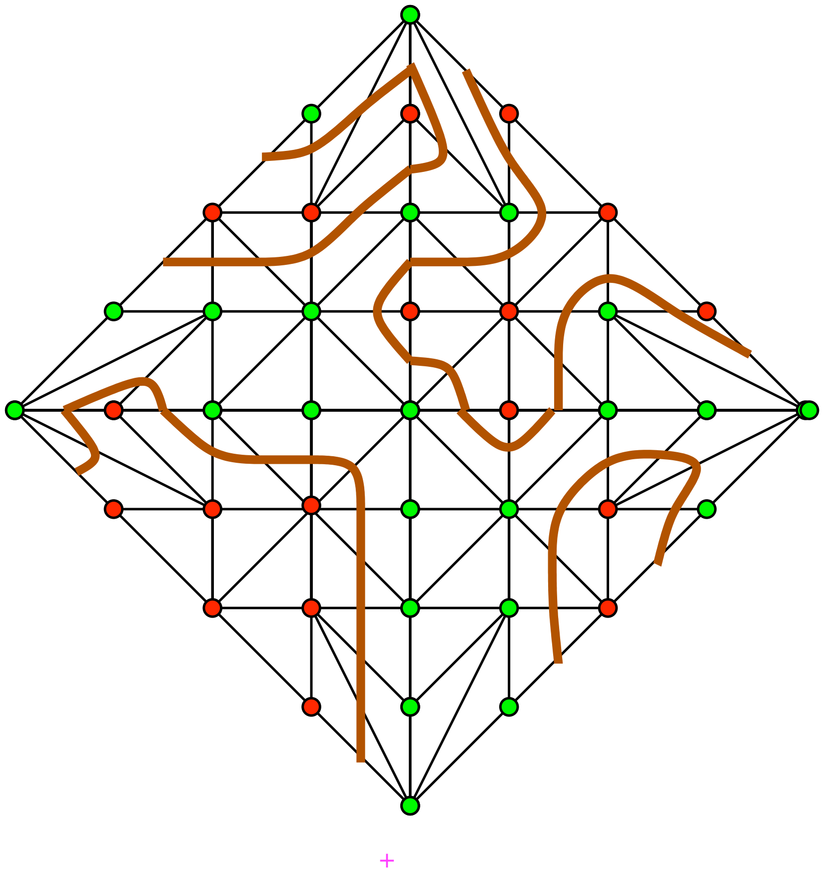

Finally draw your curve in such a way that it separates positive from negative points.

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). 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.

• Applied to a non-regular triangulation, the method may in principle produce curves not isotopic to algebraic curves of that degree (although not explicit example is known in the projective plane, there are examples of such curves in other two-dimensional toric varieties [Orevkov-Shustin, 2000]).

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.

• Applied to a non-regular triangulation, the method may in principle produce curves not isotopic to algebraic curves of that degree (although not explicit example is known in the projective plane, there are examples of such curves in other two-dimensional toric varieties [Orevkov-Shustin, 2000]).

• Still, the curves constructed with Viro's method (with non-regular triangulations) can be realized as **pseudo-holomorphic** curves in $\mathbb{CP}^2$ [Itenberg-Shustin, 2002].

3. Toric Hilbert schemes

A-graded ideals

Let $A = (a_1, \ldots, a_n) \in \mathbb{Z}^{n \times d}$ be an acyclic integer vector configuration. Let K be a field.

In the polynomial ring $K[x_1, \ldots, x_n]$ we consider the variable x_i to have (multi-)degree a_i and the monomial $\mathbf{x}^c := x_1^{c_1} \ldots x_n^{c_n}$ to have multi-degree Ac.

Example: $A = (1, \ldots, 1)$ defines the standard grading.

An ideal $I \subset K[x_1, \ldots, x_n]$ is said to be A-homogeneous if it can be generated by polynomials with all its monomials of the same multi-degree.

Example: The toric ideal I_A:

$$I_A = \langle \mathbf{x}^c - \mathbf{x}^d : \quad c, d \in \mathbb{Z}^n_{\geq 0}, \quad Ac = Ad \rangle$$

The multi-graded Hilbert function

Every A-homogeneous ideal I decomposes as

$$I = \bigoplus I_b,$$

where $b \in A(\mathbb{Z}^n_{\geq 0})$ ranges over all possible multidegrees.

The A-graded Hilbert function of I is the map

$$A(\mathbb{Z}^n_{\geq 0}) \to \mathbb{Z}_{\geq 0}$$

that sends each I_b to its linear dimension over K.

Remark: $\dim_K(I_b) \leq \#A^{-1}(b) = \#$ of monomials of degree b.

A-graded ideals and the toric Hilbert scheme

The toric ideal I_A has "codimension 1 in every degree":

$$\dim_K(I_b) = \#A^{-1}(b) - 1.$$

An A-homogeneous ideal is called A-graded if its Hilbert function equals that of the toric ideal I_A. The prototypical examples are all the initial ideals of I_A.

Remark: All A-graded ideals are binomial ideals.

The toric Hilbert scheme of A is the "space" of all possible A-graded ideals, together with certain scheme structure on it; each irreducible component is a (perhaps not normal) toric variety.

History:

[Arnol'd, 1989], [Korkina-Post-Roelofs, 1995]: study the case $d = 1$.

[Sturmfels 1995]: defines A-graded ideals in full-generality and shows their relation to triangulations of A. The relation is specially well-behaved for unimodular triangulations of A.

[Peeva-Stilman 2001]: introduce the scheme structure and coin the name "toric Hilbert scheme". They pose the connectivity question.

[Maclagan-Thomas, 2002]: define a graph of "monomial A-graded ideals" and show that the toric Hilbert scheme is connected if and only if the graph is. They use a result of Sturmfels to conclude:

Theorem: If A has unimodular triangulations that are not connected by flips, then the toric Hilbert scheme of A is not connected.

History:

[Arnol'd, 1989], [Korkina-Post-Roelofs, 1995]: study the case $d = 1$.

[Sturmfels 1995]: defines A-graded ideals in full-generality and shows their relation to triangulations of A. The relation is specially well-behaved for unimodular triangulations of A.

[Peeva-Stilman 2001]: introduce the scheme structure and coin the name "toric Hilbert scheme". They pose the connectivity question.

[Maclagan-Thomas, 2002]: define a graph of"'monomial A-graded ideals" and show that the toric Hilbert scheme is connected if and only if the graph is.

[S., 2002]: constructs a point set A $(\dim = 5, \#A = 50)$ that has unimodular triangulations not connected by flips. Hence, **its toric Hilbert scheme is not connected**.

The Sturmfels map

In 1991, Sturmfels had proved:

Theorem: The Gröbner fan of a toric ideal I_A refines the secondary fan of the corresponding vector configuration A. (Equivalently, the secondary polytope of A is a Minkowski summand of the state polytope of I_A).

In particular, there is a well-defined map

$$\text{initial ideals of } I_A \to \text{regular polyhedral subdivisions of } A$$

(the map sends monomial initial ideals to regular triangulations, and is surjective).

In 1995, he extended the map to

$$\Phi : A\text{-graded ideals} \to \text{polyhedral subdivisions of } A,$$

and the map sends monomial ideals to triangulations.

The map is now **not surjective** [Peeva 1995], but its image contains all the unimodular triangulations of A [Sturmfels 1995] (moreover, each unimodular triangulation T is the image of a unique monomial A-graded ideal, namely the Stanley-Reisner ring of T).

Maclagan and Thomas defined flips between A-graded monomial ideals, and showed that if I_1 and I_2 are related by a flip, then $\Phi(I_1)$ and $\Phi(I_2)$ either coincide or are related by a bistellar flip.

Open question: Is the toric Hilbert scheme connected when $d = 1$?

Observe that a vector configuration of dimension 1 gives a "point configuration of dimension 0". That is, n copies of a single point! This has n triangulations, one using each of the points; the secondary polytope is an $n-1$-simplex.

Even the case $n = 4$ is open!

4. Tropical polytopes

Tropical hypersurfaces

The tropical semiring (or min-plus algebra) is $(\mathbb{R}, \oplus, \odot)$ where

$$a \oplus b := \min(a, b) \qquad a \odot b := a + b.$$

A tropical polynomial is $$F(x_1, \dots, x_n) = \bigoplus_{i=1}^{r} c_i \; x_1^{\odot a_{i1}} \odot \cdots \odot x_n^{\odot a_{in}}.$$

(In usual arithmetics: $F(x_1, \dots, x_n) = \min\{c_i + \sum_{i=1}^n a_{ij} x_j : i = 1, \dots, r\}$.)

We define the "zero-set" (or tropical hypersurface) of F as the set of points for which the minimum is achieved twice. That is to say, the set of points were the function $F : \mathbb{R}^n \to \mathbb{R}$ is **not** linear. It is a polyhedral complex of codimension one, with cells in directions normal to faces of the simplex $\operatorname{conv}(O, e_1, \dots, e_n)$.

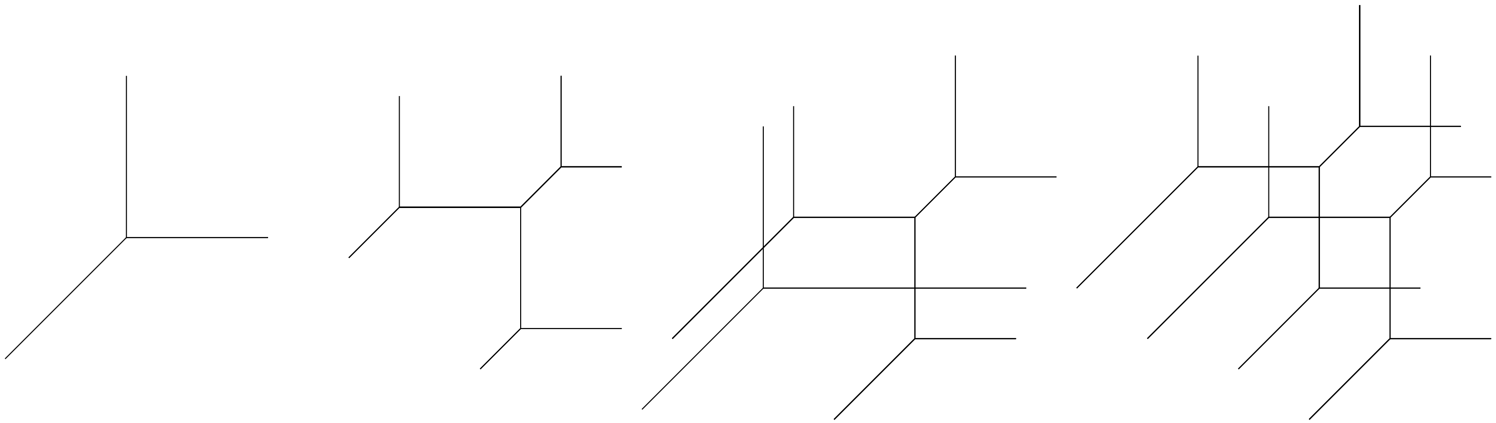

Tropical curves of degrees 1, 2, 3 and 4

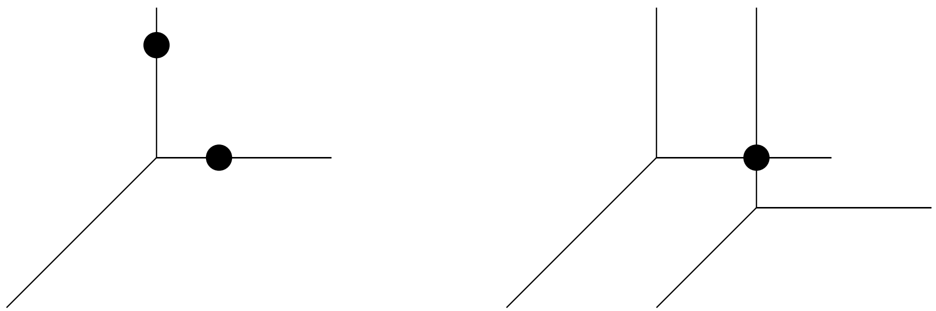

Trough every two points in $\mathbb{R}^2$ there is a unique tropical line, every two tropical lines meet in a single point

Tropical algebraic geometry

Let $K = \mathbb{C}\{\{t\}\}$ (field of Puiseux series over $\mathbb{C}$. The order of the series $c \cdot t^\alpha + \cdots$ is its minimum exponent, $\alpha \in \mathbb{Q}$. Now, we look at polynomials in $K[x_1, \ldots, x_n]$. To each such polynomial

$$f = c_1(t) \cdot x^{a_1} + \cdots c_r(t) \cdot x^{a_r}$$

one can associate a tropical polynomial

$$\mathrm{trop}(f) := \bigoplus_{i=1}^{r} \mathrm{order}(c_i(t)) \odot x^{\odot a_i}.$$

Similarly, for each point $c = (c_1, \ldots, c_n) \in (K^*)^n$, we define $\mathrm{order}(c) := (\mathrm{order}(c_1), \ldots, \mathrm{order}(c_n))$.

Tropical algebraic geometry

Theorem [Kapranov 2000, Sturmfels 2002] Let $I \subset K[x_1, \ldots, x_n]$ be an ideal, $V \subset (K^*)n$ its variety (intersected with $(K^*)n$) and G_I auniversal Gröbner basis of I. Then, the following subsets of $\mathbb{R}^n$ coincide:

1. The closure of the image of V under the order map $(K^*)n \to \mathbb{R}^n$.
2. The intersection of all tropical hypersurfaces defined by $\{\mathrm{trop}(f) : f \in G_I\}$.
3. The set of weight vectors $w \in \mathbb{R}^n$ for which the initial ideal $\mathrm{in}_w(I)$ contains no monomial.

Such a set is called the tropical variety of I. It is a polyhedral complex in $\mathbb{R}^n$.

Applications of tropical algebraic geometry

1. The “tropical Grassmannian of lines” $\mathbb{T}G(r, 2)$ equals the space of phylogenetic trees with r nodes (Billera-Holmes-Vogtmann, 2001).

2. Mikhalkin (2002) computes the number (dimension) of curves of a certain degree through a certain number of generic points in $\mathbb{CP}^2$ (or any other toric surface) “tropically” (counting integer lattice paths in the defining polygon). In particular, he computes the Gromov-Witten invariants of $\mathbb{CP}^2$ (or other non-singular toric surfaces). Shustin (2003) does a similar thing for real curves (Welshinger invariant).

3. Better understanding of A-discriminants (E.-M. Feichtner’s talk).

Tropical polyhedral geometry

Where do triangulations arise?

The graphs of the functions that define tropical hypersurfaces are polyhedral hypersurfaces polar to liftings of regular subdivisions.

Tropical polyhedral geometry

Where do triangulations arise?

Moreover, in 2003, Develin and Sturmfels initiated the study of tropical polytopes. That is, tropical convex hulls of finite point sets. They are bounded polyhedral complexes, with cells in the directions normal to faces of the standard simplex. They proved:

Theorem: There is a 1-to-1 correspondence between tropical polytopes with $n+1$ vertices in $\mathbb{R}^d$ and regular subdivisions of the product $\Delta^n \times \Delta^d$ of two simplices.

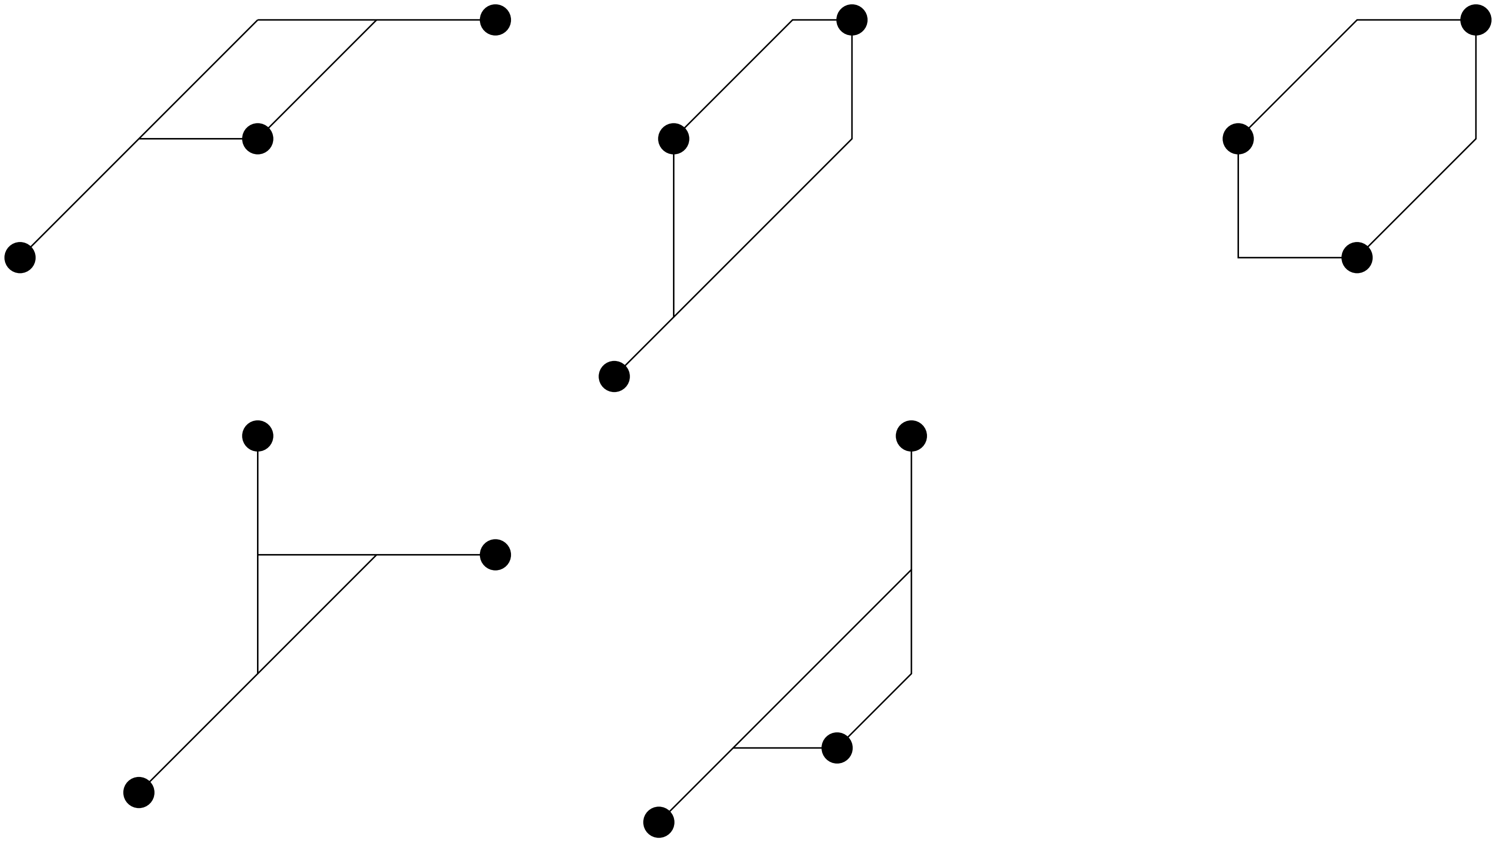

Some (generic) tropical triangles in the plane

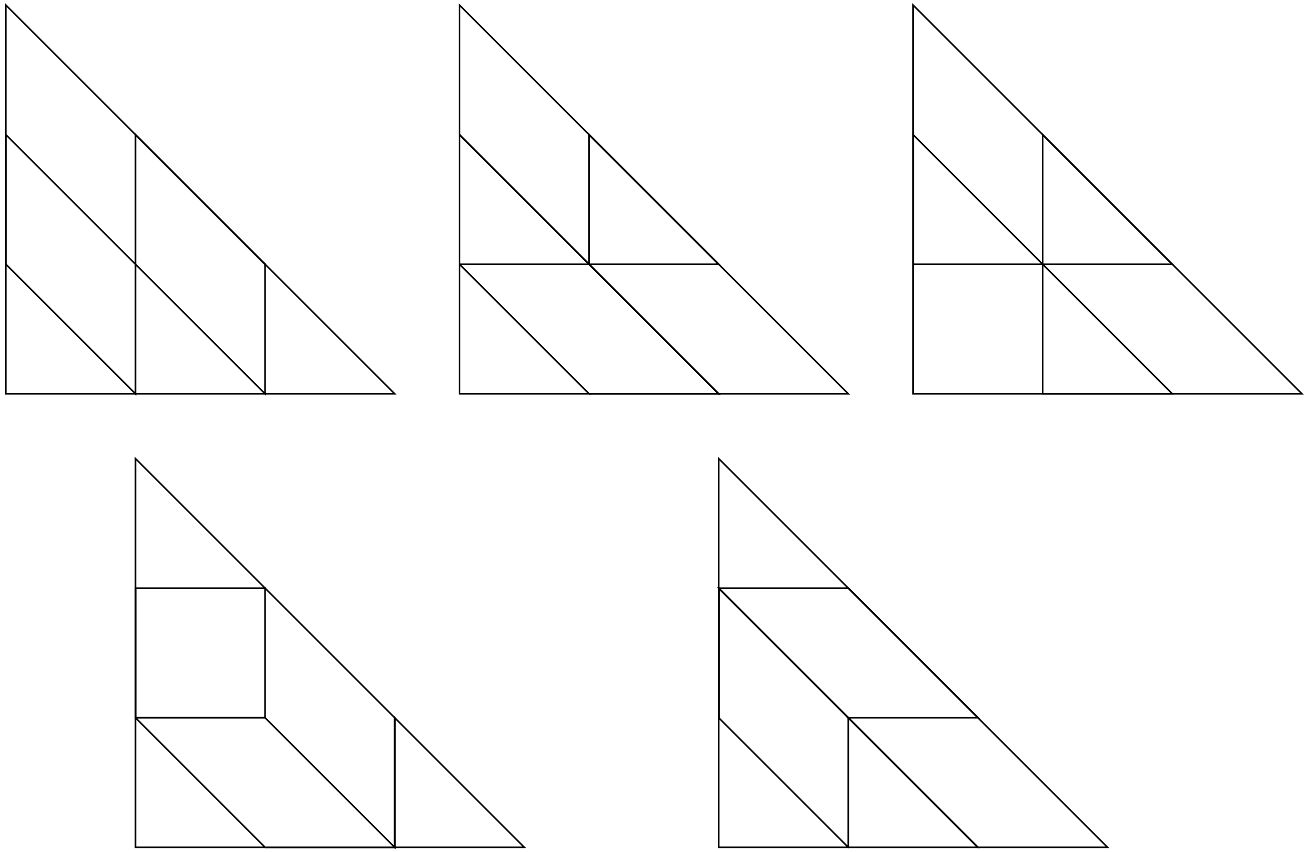

and the corresponding triangulations of $\Delta^2 \times \Delta^2$ (pictured, via the Cayley Trick, as regular mixed subdivisions of $3\Delta^2$)

An application of tropical geometry to the study of triangulations

By "construction", (the codimension 1 skeleton of) a tropical polytope of dimension d and with $n+1$ vertices lies in the union of $(d+1)(n+1)$ (usual) hyperplanes.

In particular, the number of combinatorial types of tropical polytopes is bounded above by that of hyperplane arrangements. Using known bounds [Goodman-Pollack, Alon]:

Theorem [Santos 2004] For a fixed d, the number of regular triangulations of $\Delta^d \times \Delta^n$ grows as $n^{\Theta(n)}$. In contrast, (if $d \geq 2$) the number of all triangulations grows as $n^{\Omega(n^2)}$.