Computing Triangulations Using Oriented Matroids

Julian Pfeifle[1] and Jörg Rambau[2]

[1] Dept. of Mathematics, MA 6-2, TU Berlin, 10623 Berlin, Germany
[2] Zuse-Institute Berlin, Takustr. 7, 14195 Berlin, Germany

Abstract. Oriented matroids are combinatorial structures that encode the combinatorics of point configurations. The set of all triangulations of a point configuration depends only on its oriented matroid. We survey the most important ingredients necessary to exploit oriented matroids as a data structure for computing all triangulations of a point configuration, and report on experience with an implementation of these concepts in the software package TOPCOM. Next, we briefly overview the construction and an application of the secondary polytope of a point configuration, and calculate some examples illustrating how our tools were integrated into the POLYMAKE framework.

1 Introduction

This paper surveys a selection of efficient combinatorial methods to compute triangulations of point configurations. We present results obtained for the first time by a software implementation (TOPCOM [20]) of these ideas. It turns out that a subset of all triangulations of a point configuration has a structure useful in different areas of mathematics, and we highlight one particular instance of such a connection. Finally, we calculate some examples by integrating TOPCOM into the POLYMAKE [7] framework. See also [29] for more related information on the combinatorics of triangulations.

Let us begin by motivating the use of triangulations and providing a precise definition.

1.1 Why triangulations?

Triangulations are widely used as a standard tool to decompose complicated objects into simple objects. A solution to a problem on a complicated object can sometimes be found by gluing solutions on the simple objects. Some examples are the following:

- Numerics: Finite Elements Method
- Algebraic Topology: computation of topological invariants
- Computer Graphics: Raytracing

Besides these applications, structures on whole sets of triangulations have interesting connections to seemingly distant disciplines, among them:

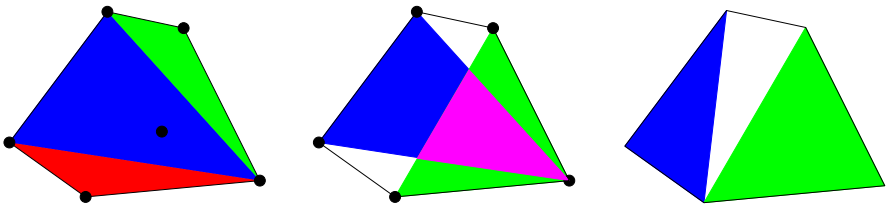

Fig. 1. A correct triangulation, an unwanted intersection (IP not met), and an incomplete triangulation (UP not met)

- Algebraic Geometry: Connection to Toric Varieties
- Algebra: Polynomial System Solving
- Homotopy Theory: Structure of Loop Spaces

Therefore, the study of *spaces of triangulations* has become a subject in its own right in the field of discrete geometry [18].

1.2 What exactly are triangulations?

For the rest of the paper, let $\mathcal{A}$ be a d-dimensional point configuration with n points. We assume that the points are labeled $1, 2, \ldots, n$, and denote the coordinate vector of the point i by a_i.

Definition 11 *A subset T of $(d+1)$-subsets of $\mathcal{A}$ is a* triangulation of $\mathcal{A}$ *if and only if*

$$\bigcup_{\sigma \in T} \operatorname{conv} \sigma = \operatorname{conv} \mathcal{A} \tag{UP}$$

$$\operatorname{conv} \sigma \cap \operatorname{conv} \sigma' = \operatorname{conv}(\sigma \cap \sigma') \quad \forall \sigma, \sigma' \in T. \tag{IP}$$

Condition UP makes sure that the *union* of all (convex hulls of) simplices in T covers (the convex hull) of $\mathcal{A}$. Condition IP takes care of unwanted intersections. Note that we do not require all points to be used in a triangulation. Figure 1 provides a sketch of the situation.

2 The oriented matroid of a point configuration

In a naive approach, checking for non-empty interior intersection of two simplices or for a complete covering are linear programming problems. Since we need numerically exact results for our purposes, exact arithmetics is a must. However, linear programming with exact arithmetics is computationally expensive.

Fortunately, solving linear programs exactly can be avoided by rigorously separating exact arithmetics and combinatorial computations in simplicial complexes. It turns out that all arithmetic calculations can be reduced to computing $\binom{\#\mathcal{A}}{d+1}$ determinants of $((d+1)\times(d+1))$-matrices, possibly in a preprocessing step. All subsequent computations concerning triangulations of $\mathcal{A}$ that exclusively use information about the relative orientation of points can then be performed without using the coordinates ever again.

More specifically, we show in this section how the conditions of Definition 11 can be checked purely combinatorially, provided we have the *oriented matroid* of $\mathcal{A}$ at hand. The resulting combinatorial characterization has been formulated and extensively used, e.g., in [19]. More rigorous proofs can also be found there. Related characterizations and applications thereof can be found in [12,13]. More general information about oriented matroids can be found in [3].

The combinatorial concepts presented in this paper also hold for *triangulations of oriented matroids*, purely combinatorial objects that are not necessarily derived from any (realizable) point configuration [27]. For the computational methods it does not make any difference whether the oriented matroid under consideration is given implicitly by the coordinates of a "real" point configuration or directly by a list of combinatorial data of a general, maybe non-realizable oriented matroid.

2.1 Geometric Problem I: Proper intersection of simplices

We would like to use combinatorial data of $\mathcal{A}$ to check (IP). It turns out that there is a finite set of minimal obstructions, the *circuits* of $\mathcal{A}$, that describe the intersections of simplices in $\mathcal{A}$ completely. How does this work?

Assume that the convex hulls of the d-simplices σ and σ' intersect improperly, i.e., σ and σ' violate (IP). Then, by Radon's theorem, we find $Z^+ \subseteq \sigma$ and $Z^- \subseteq \sigma'$ such

(i) $Z^+ \cap Z^- = \emptyset$,
(ii) the relative interiors of Z^+ and Z^- intersect,
(iii) Z^+ and Z^- are inclusion minimal with these properties.

The pair (Z^+, Z^-) is called an *intersection circuit* of σ and σ'. In particular, a subset of $\mathcal{A}$ with properties (i)–(iii) it is a *circuit* of $\mathcal{A}$.

Another interpretation of a circuit is that $\sum_{i\in Z^+}\lambda_i a_i = \sum_{i\in Z^-}\lambda_i a_i$ is an affine dependence with minimal support, for suitable $\lambda_i > 0$ with $\sum_{i\in Z^+}\lambda_i = \sum_{i\in Z^-}\lambda_i$, $i = 1,\ldots,n$. If we set $\Lambda := \sum_{i\in Z^+}\lambda_i$ then the (unique) intersection point in $\operatorname{conv} Z^+ \cap \operatorname{conv} Z^-$ is given by $\frac{1}{\Lambda}\sum_{i\in Z^+}\lambda_i a_i$.

This connection to affine dependences shows that we have exactly one circuit (modulo exchanging Z^+ and Z^-) for every affinely dependent set of points in $\mathcal{A}$. Since there are at most $\binom{n}{d+2}$ affinely dependent sets of points, there are at most that many circuits (modulo exchanging Z^+ and Z^-). Using the set of all circuits of $\mathcal{A}$, we can check (IP) easily as follows:

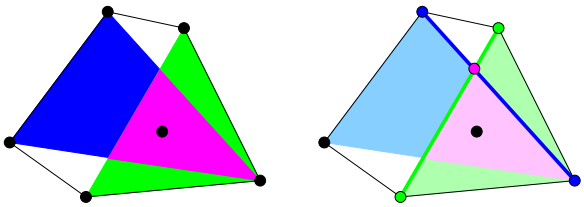

Fig. 2. An unwanted intersection can be detected by a circuit

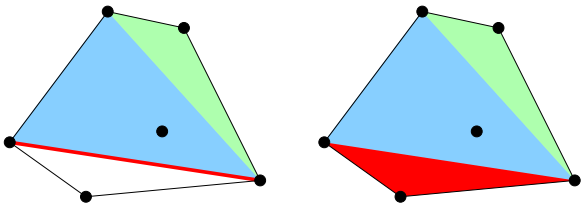

Fig. 3. An interior facet covered by only one simplex detects a (UP)-violation; another simplex containing it is required

> Two simplices σ and σ' in T violate (IP) if and only if there exists a circuit (Z^+, Z^-) of $\mathcal{A}$ with $Z^+ \subseteq \sigma$ and $Z^- \subseteq \sigma'$.

Circuits are, in other words, obstructions to (IP). See Figure 2 for a sketch of two simplices intersecting improperly and a corresponding intersection circuit.

2.2 Geometric Problem II: Proper covering by simplices

In order to check (UP) purely combinatorially, we will assume (IP') for all simplices in T. A non-empty set of simplices satisfying (IP'), but not necessarily covering $\operatorname{conv}\mathcal{A}$, can be seen as a *partial triangulation* of $\mathcal{A}$. How can we detect an uncovered area in $\operatorname{conv}\mathcal{A}$ purely combinatorially? To this end, we look at facets of d-simplices in T. Such a facet is *an interior facet of* T if it is not a facet of $\mathcal{A}$. (A facet of $\mathcal{A}$ is a $(d-1)$-dimensional subset of $\mathcal{A}$ that is the intersection of $\mathcal{A}$ with a supporting hyperplane. A supporting hyperplane is an affine hyperplane that does not separate A.)

The following is easy to see [19]: a partial triangulation T violates (UP) if and only if we find an interior facet of T that is not contained in any other simplex in T. Let us assume for the moment that we have a list of all facets of $\mathcal{A}$. Then we can simply go through the set of all facets of a partial triangulation T and count the number of simplices in T containing them. If there is an interior facet of T contained in only one simplex, then T violates (UP). This test is purely combinatorial.

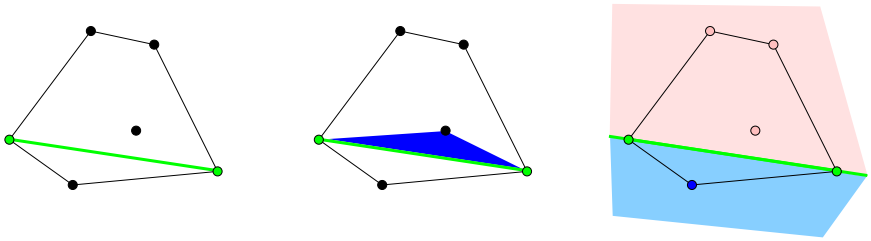

Fig. 4. A facet F in T, i.e., a facet of some simplex in T, is interior if the cocircuit of $\mathcal{A}$ spanned by F has both positive and negative elements

Now, how do we get the list of all facets? It turns out that the set of *cocircuits* of $\mathcal{A}$ contains all the necessary information. Consider the set of all affine hyperplanes spanned by $(d-1)$-dimensional subsets of $\mathcal{A}$, oriented arbitrarily. Each such hyperplane defines a *signature* on the points in $\mathcal{A}$: the signature of a point is *zero* if it lies *on* the hyperplane; it is *positive* if it lies strictly on the *positive side* of the hyperplane; it is *negative* if it lies strictly on the *negative side* of the hyperplane. Such a signature on $\mathcal{A}$ is called a *cocircuit* of $\mathcal{A}$. We denote such a cocircuit by (C^+, C^-), where C^+ contains all points with positive signature and C^- those with negative signature.

A subset of $\mathcal{A}$ is a facet of $\mathcal{A}$ if and only if it is the zero set of a cocircuit having no positive elements, or if it is the zero set of a cocircuit having no negative elements.

There are at most $\binom{n}{d}$ different hyperplanes spanned by subsets of a d-dimensional point configuration. Therefore, there are at most $\binom{n}{d}$ cocircuits of $\mathcal{A}$ modulo reversing signs. These give us all facets of $\mathcal{A}$ and thus all interior facets of T.

Summarizing this section, we can use the set of all cocircuits of $\mathcal{A}$ to check (UP) as follows:

> A partial triangulation T of $\mathcal{A}$ violates (UP) if there is an interior facet of T lying in only one simplex of T.

In other words: interior facets (determined by the cocircuits) incident to exactly one simplex are obstructions for (UP).

2.3 Triangulations depend only on the oriented matroid

The considerations in the previous sections justify the name 'combinatorial characterization of triangulations' in the following theorem (see, e.g., [19]):

Theorem 21 (Combinatorial Characterization of Triangulations) *A subset T of $(d+1)$-subsets of $\mathcal{A}$ is a* triangulation of $\mathcal{A}$ *if and only if*

(IP') For every pair of simplices $\sigma, \sigma' \in T$ there is no circuit (Z^+, Z^-) of $\mathcal{A}$ with $Z^+ \subseteq \sigma$ and $Z^- \subseteq \sigma'$.

(UP') *For every interior facet F of T there are at least two simplices in T containing F.*

The following fact allows a unified view on the previous two sections: The circuits of $\mathcal{A}$ determine, purely combinatorially, the cocircuits of $\mathcal{A}$, and vice versa. The *oriented matroid* of $\mathcal{A}$ is now defined by the set of circuits or by the set of cocircuits of $\mathcal{A}$, depending on what is more convenient in a particular situation.

Thus, the set of triangulations of $\mathcal{A}$ depends only on the oriented matroid of $\mathcal{A}$.

2.4 An interface from geometry to combinatorics: The chirotope

How do we compute the circuits and the cocircuits from the coordinates of the points in $\mathcal{A}$? There are actually several ways to do this. The most commonly used approach is the computation of a third equivalent combinatorial structure of $\mathcal{A}$: *the chirotope*.

The chirotope of $\mathcal{A}$ is the following alternating function on the set of $(d+1)$-subsets of $\mathcal{A}$:

$$\chi : \begin{cases} \binom{\mathcal{A}}{d+1} \to \{+,-,0\} \\ (i_1, i_2, \ldots, i_{d+1}) \mapsto \operatorname{sign}\big(\det(a_{i_1}, a_{i_2}, \ldots, a_{i_{d+1}})\big) \end{cases} \tag{1}$$

That means, in particular, that the chirotope assigns to each ordered basis of $\mathcal{A}$ its orientation. The orientation is usually normalized such that the affine standard basis of $\mathbb{R}^d$ has orientation $+$. The chirotope value on $(d+1)$-subsets of $\mathcal{A}$ that are not independent is zero. One of many ways to accelerate the computation of the chirotope is the *elimination tree method* of TOPCOM, indicated in Figure 5. Column normal forms for all k-subsets of columns are maintained as nodes in computation tree for all $k \leq d+1$; they can be reused in order to save elimination steps. The tree is traversed in depth-first-search; determinants are produced in the leaves in lexicographic order w.r.t. the indices of the point coordinates.

While in a naïve way of of computing determinants we need $O((d+1)^3)$ arithmetic operations per determinant (if we use an elimination algorithm), the elimination tree method just needs $O\big(\frac{n}{n-d}(d+1)^2\big)$ arithmetic operations per determinant [22].

There is another way of exploiting structure on the set of all determinants of submatrices of $\mathcal{A}$: the *exterior algebra method*. See [12] for a description; it is also used in the context of generating oriented matroids [4,6]. The basic idea is that there are certain relations between determinants coming from the distributive law in the exterior algebra. These relations correspond to the *Grassmann-Plücker-relations*.

The circuits of $\mathcal{A}$ can now be computed as follows: for any $(d+2)$-subset $\underline{Z} = \{z_1, z_2, \ldots, z_{d+2}\}$ with $z_1 < z_2 < \cdots < z_{d+2}$ of $\mathcal{A}$, we have the following

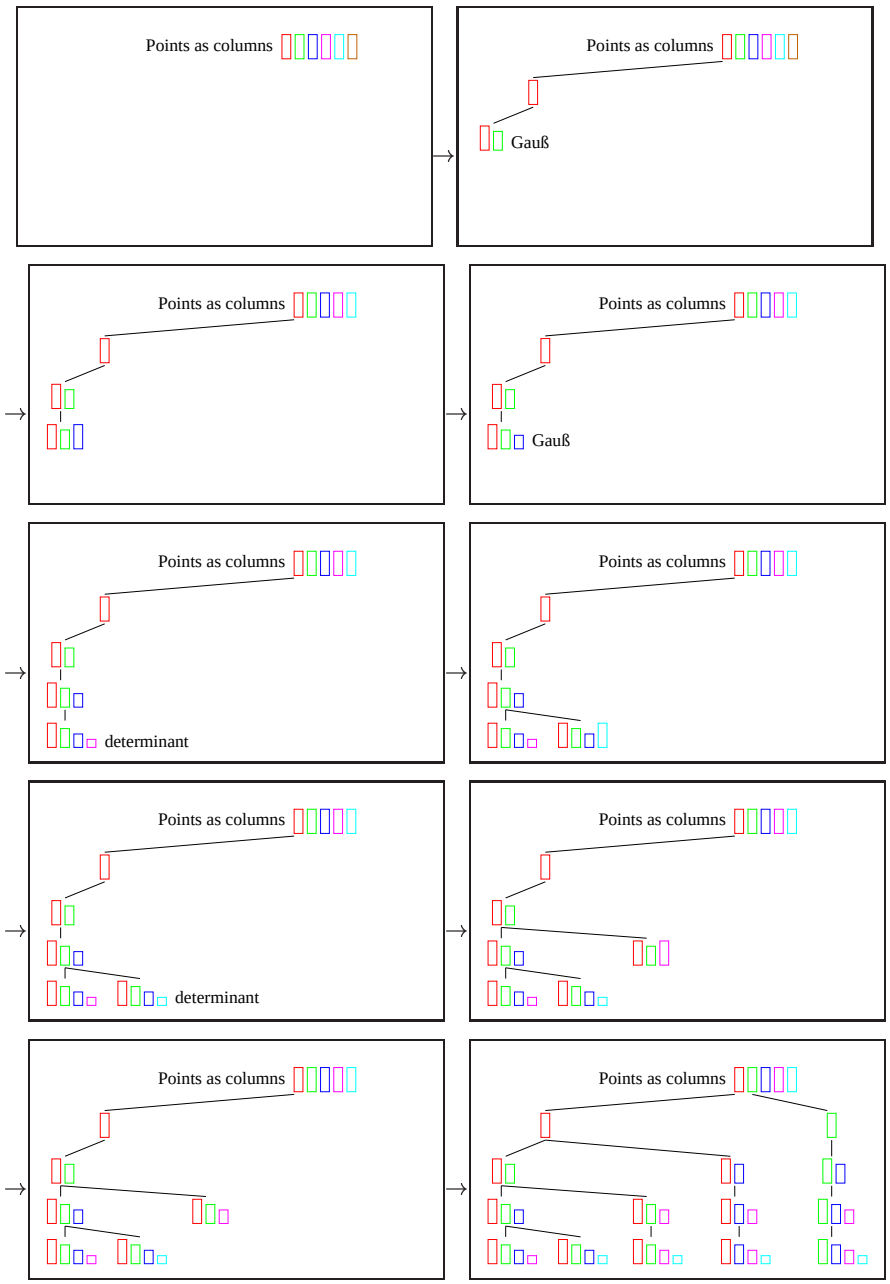

Fig. 5. Building the computation tree for the chirotope: the matrices in the nodes are all transformed in column normal form before new columns are added (indicated are regions of possible non-zeroes in the columns); the determinants can be computed easily at the leaves

circuit (Z^+, Z^-):

$$Z^+ = \left\{ z_k \in \underline{Z} \mid (-1)^k \chi(z_1, \ldots, z_{k-1}, z_{k+1}, \ldots, z_{d+2}) = + \right\} \tag{2}$$
$$Z^- = \left\{ z_k \in \underline{Z} \mid (-1)^k \chi(z_1, \ldots, z_{k-1}, z_{k+1}, \ldots, z_{d+2}) = - \right\} \tag{3}$$

The cocircuits can be computed as follows: for any d-subset C of $\mathcal{A}$, we have the following cocircuit:

$$C^+ = \left\{ i \in \mathcal{A} \setminus C \mid \chi(C, i) = + \right\} \tag{4}$$
$$C^- = \left\{ i \in \mathcal{A} \setminus C \mid \chi(C, i) = - \right\} \tag{5}$$

Note that computing circuits and cocircuits is possible without further access to the coordinates of the points in $\mathcal{A}$: the chirotope forms an interface from geometry to combinatorics.

3 Applications of the oriented matroid: how to find triangulations

So far, we have seen that the oriented matroid of $\mathcal{A}$ determines the set of triangulations of $\mathcal{A}$. Moreover, once the chirotope has been computed, we can check purely combinatorially whether or not a set of simplices is a triangulation of $\mathcal{A}$.

What we are still lacking is a method to generate triangulations; we certainly cannot check all possible sets of simplices in reasonable time. Therefore, we show in this section how the oriented matroid of $\mathcal{A}$ can be used to compute

- some triangulation (Application I),
- local changes in triangulations (Application II),
- many triangulations (Application III),
- all triangulations (Application IV).

3.1 Application I: The placing triangulation

Constructing a triangulation by placing [11] is an incremental method: we add the points in $\mathcal{A}$ one by one, starting from an affine basis. In each step, we cone the new point to all *visible facets* of what we have already. This way, we get a triangulation for each permutation of the points in $\mathcal{A}$. In special cases, some of the triangulations obtained from distinct permutations might coincide, in general they do not. We will see that all necessary ingredients can be computed via the oriented matroid.

For a more accurate definition of the placing triangulation, assume that we already have a triangulation T' of the subconfiguration $\mathcal{A}' := \mathcal{A} \setminus \{a\}$ for some $a \in \mathcal{A}$. There are two cases: $a \in \operatorname{conv} \mathcal{A}'$ or $a \notin \operatorname{conv} \mathcal{A}'$.

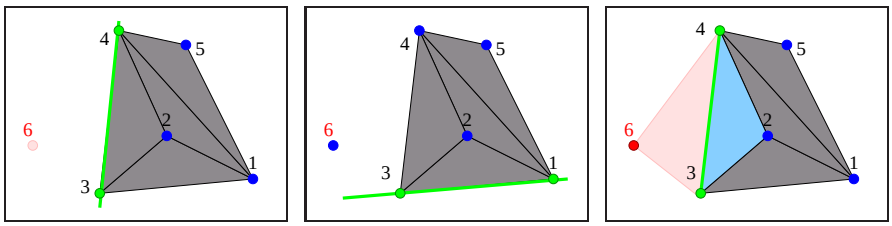

Fig. 6. Cocircuits of $\mathcal{A}$ detect visible facets: the cocircuit spanned by $\{3,4\}$ induces different signs on the new point 6 and the old points. Thus, $\{1,3\}$ is a visible facet. The cocircuit spanned by $\{1,3\}$ induces the same signs on the new and the old points: $\{1,3\}$ is not visible. Note that the addition of the cone over $\{3,4\}$ removes an obstruction for (UP) because $\{3,4\}$ is an interior facet of the new configuration including 6 and was included in only one simplex so far

In the first case we do nothing: T' is also a triangulation of $\mathcal{A}$. Indeed: (IP') holds because no new simplex has been added; for the same reason (UP') holds because all facets of $\mathcal{A}'$ are still facets of $\mathcal{A}$.

In the second case we define visible facets of T': a facet F, supported by the hyperplane H_F, of a d-simplex of T' is called a *facet of T' visible from a* if

- a is not on H_F.
- $\mathcal{A}$ and a are not on the same side of H_F.

In this case, we add all cones of visible facets with apex a to T':

$$T := T' \cup \{F \cup \{a\} \mid F \text{ is a facet of } T' \text{ visible from } a\}$$

Why does this give us a triangulation of $\mathcal{A}$?

First, the new simplices intersect properly with each other, since they form a cone over properly intersecting simplices. Second, the interior intersection between a new and an old simplex must be empty because none of the new simplices has an interior intersection with $\operatorname{conv}\mathcal{A}'$: it is separated from $\operatorname{conv}\mathcal{A}'$ by the supporting hyperplane of the corresponding facet of $\mathcal{A}$. Therefore, (IP) holds.

(UP) holds as well: The only facets in T' that become interior facets in $\operatorname{conv}\mathcal{A}$ are those visible from a, by construction.

Now, we want to solely use the oriented matroid data in the step from T' to T. The only thing we need is the set of visible facets, a subset of the boundary facets of T' (which are those facets of T' that are contained in exactly one simplex in T'). Visibility can now be checked for every boundary facet F in T' by looking at the cocircuit C_F spanned by F: F is visible from a if the following holds:

- The new point a must not be in C_F^0.

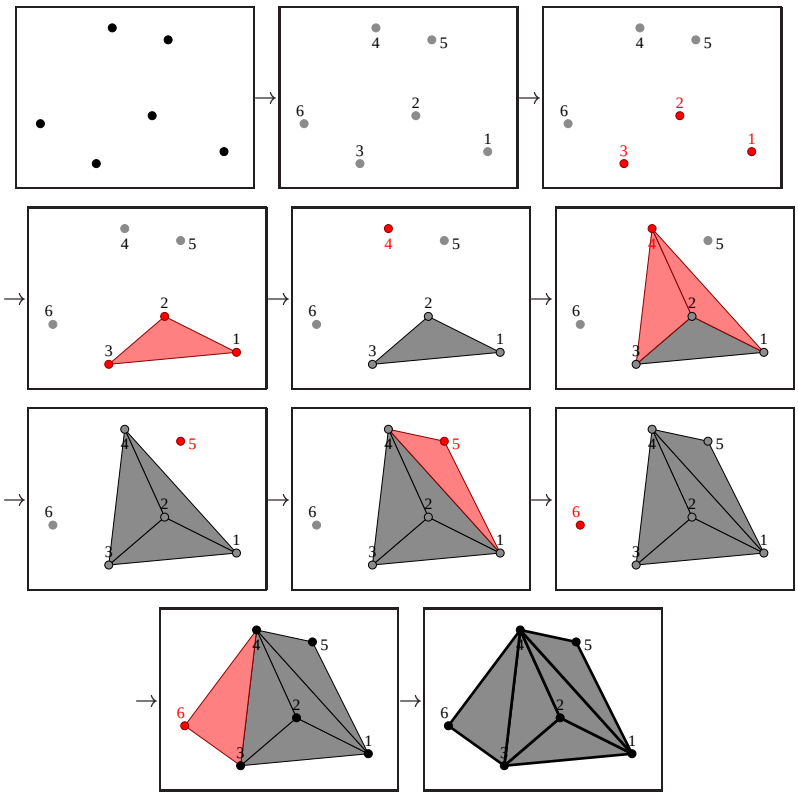

Fig. 7. Constructing a placing triangulation

- If one of the points in $\mathcal{A}'$ is in C_F^+ then a must be in C_F^-.

An example for these conditions is indicated in Figure 6.

Note that it suffices to find a non-zero signature of only one point a' in $\mathcal{A}'$; then all points in $\mathcal{A}'$ with non-zero signature automatically have the same signature as a'.

By induction, we can construct a triangulation of the whole point configuration by this placing operation. Furthermore, to do this we only need the cocircuits of the oriented matroid. Figure 7 shows the corresponding steps for our running example configuration.

3.2 Application II: Flips

Now that we know how to produce one triangulation we will show that flipping—the standard method to produce new triangulations from old ones—can also be done by using the oriented matroid.

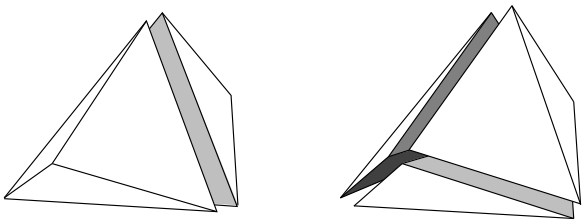

Fig. 8. The two possible triangulations of five points in general position in dimension three; Z^- contains, w.l.o.g., the three vertices of the interior cutting face in the left picture; Z^+ contains the two vertices of the interior cutting edge in the right picture

A flip in a two-dimensional triangulation is, e.g., an edge flip: replacing the diagonal in a convex quadrilateral. But in our context also the removal of an interior point lying in the convex hull of exactly three edges is a flip.

In general a flip is the following: Consider a circuit $Z := (Z^+, Z^-)$ in $\mathcal{A}$. Let $\underline{Z} := Z^+ \cup Z^-$. Then it is easy to see that there are exactly two triangulations $T^+(Z)$ and $T^-(Z)$ of $\underline{Z}$. The triangulations are:

$$T^+(Z) := \{\underline{Z} \setminus \{z_-\} | z_- \in Z^-\} \tag{6}$$

$$T^-(Z) := \{\underline{Z} \setminus \{z_+\} | z_+ \in Z^+\} \tag{7}$$

Figure 8 sketches T^+ and T^- in dimension three for a circuit on five points in general position. Note that T^+ and T^- differ in their number of simplices.

Assume that a triangulation of $\mathcal{A}$ contains one of $T^+(Z)$ and $T^-(Z)$ as a subcomplex, say $T^+(Z)$. If Z spans a full-dimensional subconfiguration, then flipping simply means replacing $T^+(Z)$ by $T^-(Z)$ in T.

If the circuit spans a lower-dimensional subconfiguration then things are more complicated: still, the circuit itself has exactly two triangulations $T^+(Z)$ and $T^-(Z)$. These are lower-dimensional triangulations. If one of them, say $T^+(Z)$, is a subcomplex in T we cannot simply replace $T^+(Z)$ by $T^-(Z)$ and still get a triangulation. If, however, we encounter that every maximal simplex in $T^+(Z)$ has the same link $L(Z)$ in T then we can replace $T^+(Z) * L(Z)$ by $T^-(Z) * L(Z)$ in T. (Here, "$*$" denotes the simplicial join.)

So, finding a flip in terms of oriented matroids amounts to the following:

1. Pick a circuit Z.
2. Check whether $T^+(Z)$ or $T^-(Z)$ is a full-dimensional subcomplex of T.
3. If $T^+(Z)$ is in T, then replacing $T^+(Z)$ by $T^-(Z)$ is a valid flip.
4. If not, check the links of all maximal simplices of $T^+(Z)$ or $T^-(Z)$.
5. If the links, w.l.o.g. the ones in $T^+(Z)$, are identical, say equal to $L(Z)$, then replacing $T^+(Z) * L(Z)$ by $T^-(Z) * L(Z)$ is a valid flip.

Figure 9 lists all flip types in dimension two.

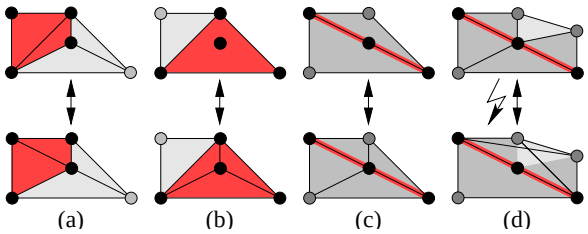

Fig. 9. All possible flip types in dimension two; the rightmost picture shows a situation where the links of all maximal simplices are not equal, whence the corresponding flip is impossible

Thus, the only information about the point configuration we need for flipping is its set of circuits. Given the notion of flipping, one can define the *flip graph* of $\mathcal{A}$:

- The nodes of the flip graph are all triangulations.
- Two nodes are connected by an edge if the corresponding triangulations can be flipped into each other by a single flip.

3.3 Application III: Computing a component of the flip-graph

So far, we have the following:

- We can compute one triangulation.
- We construct new triangulations from old ones by flipping.

Using standard enumeration techniques like breadth-first-search or depth-first-search on the flip graph we can now enumerate all triangulations in the component of the placing triangulation.

An interesting feature of the flip graph is that its equivariant version—i.e., the one consisting of combinatorial symmetry classes only—is compatible with a breadth-first-search. Thus, the combinatorial symmetries of the point configuration can be exploited in this framework.

Let us—in a short excursion—describe the ideas of exploiting symmetries. We assume that we already have all those permutations of $\mathcal{A}$ that maintain the combinatorial structure of $\mathcal{A}$. This means we know all permutations π with the following property: whenever (Z^+, Z^-) is a circuit in $\mathcal{A}$, then $(\pi(Z^+), \pi(Z^-))$ is also a circuit in $\mathcal{A}$.

In an ordinary breadth-first-search (BFS) in the flip graph we would like to maintain only two sets of nodes: the set of currently *open* (unprocessed) nodes C and the set of new open nodes N. A step in the BFS corresponds to the following: for all triangulations in C we flip all their unmarked flips, add the resulting triangulation to N, and mark all the "backward"-flips; finally, we set $C := N$ and $N := \emptyset$. In other words, we forget about closed nodes, nodes all of whose edges have been processed already.

In the equivariant BFS, we do not want to store the complete orbits of triangulation. We rather want to store just those representatives that we meet first in the enumeration process. For every new triangulation we check whether one of the stored representatives is in the orbit of the new triangulation.

However, then it is not at all clear that we can never re-enter a symmetry class whose representative we have forgotten already. The trick is that we additionally mark flips that are equivalent to "backward flips" via the same combinatorial symmetry that transforms the new triangulation into one of the known representatives. More accurately, we do the following: whenever we find via the flip f a triangulation T with $\pi(T) = T'$ for some stored representative $T' \in C$ and some combinatorial symmetry π, we mark $\pi(f^{-1})$ in T'; one can prove that this way we can forget about all nodes whose edges have been processed completely; we will never enter their symmetry class again. See the appendix for an example.

Moreover, whenever a flip in a triangulation node is marked, we also mark all equivalent flips corresponding to the automorphism group of the triangulation: they would hit the same symmetry classes as the originally marked flip.

3.4 Application IV: Computing all triangulations

For a long time it was an open problem whether or not the flip graph of any point configuration is connected. Recently, in [25] an example of a six-dimensional point configuration with a triangulation without flips was presented. The configuration has 324 points; the triangulations has over $25,000$ simplices [21]. Therefore, in general, we cannot compute all triangulations by flipping. Very recently, Santos has presented two brandnew constructions of point configurations with disconnected flip graph with 50 and 26 points, resp., in dimension five that can even be tweaked into convex position [26].

Here is another method of enumerating triangulations using the combinatorial characterization. We build up triangulations by adding one simplex at a time. Then we try to add new simplices maintaining (IP') for the complex built so far. As soon as (UP') is satisfied we have reached a triangulation.

Trying all possibilities by backtracking in this method yields an enumeration tree of all partial triangulations with triangulations sitting in some of the leaves. It can be seen as an enumeration of independent sets in the graph of forbidden intersections between simplices [30,10]. Since there are far more partial triangulations than triangulations, this method is not as fast as the flipping method, whenever the number of simplices in a triangulation is large (say at least ten).

4 Implementing the ideas: TOPCOM

All of the above (and some more) has been implemented in the package TOPCOM (*Triangulations of Point Configurations and Oriented Matroids*). TOPCOM is software under the GNU public license [20] and contains collection of clients for computations in point configurations, oriented matroids, and triangulations. Given a point configuration, TOPCOM can compute, for example,

- its chirotope,
- its circuits,
- its cocircuits,
- its facets,
- a placing triangulation,
- flips in a triangulation,
- the flip graph component of a triangulation,
- all triangulations.

This is done fairly quickly by using data structures custom-made for exploiting the combinatorial algorithms presented in this exposition.

Some numbers computed with TOPCOM that were unknown before can be found in Table 1. The number of triangulations of lattice points in dimension two can by now be computed much more efficiently by a special two-dimensional method [1]; for $(3 \times n)$- and $(4 \times n)$-grids a recursion formula for the number of triangulations was found by Ziegler.

Checking the result of Santos requires sophisticated use of a *lazy* chirotope: preprocessing is prohibitive for 324 points in dimension six.

Table 1. Some figures computed by TOPCOM for the first time (1.8GHz Pentium IV, 1GByte RAM, SuseLinux 7.3); cyclic polytopes have connected flip graphs [19] so have two-dimensional point sets; fine triangulations use all the vertices

what	*configuration*	*description*	*result*	*CPU/s*
# triangulations	$C(12,5)$	cyclic polytope	5,049,932	2,104
# triangulations	$C(13,7)$	cyclic polytope	6,429,428	4,055
# fine triangulations	(4×5)-lattice	two-dim. lattice, 20 points	2,822,648	501
# flip graph component	$\Delta_3 \times \Delta_3$	product of tetrahedra	4,533,408	499
# flip graph component	C^4	four-dimensional cube	92,487,256	15,624
check (IP') & (UP')	Santos triang.	six-dim. construction	okay	70
# flips	Santos triangulation	six-dim. construction	0	30

The latest addition in TOPCOM allows to check *regularity* or *coherence* of a triangulation with the help of an LP solver with exact arithmetics, like, e.g., `cdd` [5]. This concept is—in contrast to all the other ones presented in this exposition—not combinatorial: whether or not a triangulation is regular depends on the specific coordinates of the points.

Why this concept is important is the subject of the remaining sections of this paper.

5 Exploring further structures

We have seen in Section 3.2 that flipping is a natural way to locally modify a given triangulation. In this section, we will build on this concept and present one of the most striking and beautiful constructions of the theory of polyhedral subdivisions, which shows that a certain subclass of triangulations of a point configuration carries quite strong structural properties. Namely, we will outline the construction of the *secondary polytope* of a point configuration, briefly sketch one situation where it can be useful, and indicate how to calculate some interesting examples by integrating TOPCOM and the POLYMAKE [7] framework.

5.1 The convex hull of triangulations: Secondary polytopes

Let us define a *regular* or *coherent* triangulation of a point configuration $\mathcal{A}$ in $\mathbb{R}^d$ as one that arises by projecting the "lower" facets (with respect to some fixed direction) of a $(d+1)$-dimensional polytope $\widetilde{P}$ to $\mathbb{R}^d$, in such a way that the "lower" vertices of $\widetilde{P}$ project exactly to the points in $\mathcal{A}$. Another way of stating this condition is to ask for a convex lifting function from $\mathbb{R}^d$ to $\mathbb{R}^{d+1}$ that is linear on the simplices of the triangulation. More generally, we will also consider *regular subdivisions* of $\mathcal{A}$, that is to say collections of affinely independent subsets (*cells*) of $\mathcal{A}$ of cardinality at least $d+1$ that satisfy (UP), (IP), and (the obvious adaptation of) the regularity property. See Figure 10.

Regular triangulations correspond in a one-to-one fashion to the vertices of a convex polytope $\Sigma(\mathcal{A})$ that only depends on the point configuration, the so-called *secondary polytope of* $\mathcal{A}$. Moreover, this correspondence is not just bijective, but structural: Two regular triangulations T and T' are connected by a flip if and only if the vertices v_T and $v_{T'}$ lie on an edge of the convex hull of $\Sigma(\mathcal{A})$. It turns out that this correspondence extends to the whole face lattice of the secondary polytope, such that to each face F of $\Sigma(\mathcal{A})$ there corresponds some regular subdivision $\sigma(F)$ of $\mathcal{A}$. Furthermore, if $F \subset G$ are two faces of $\Sigma(\mathcal{A})$, then $\sigma(F)$ is a *refinement* of $\sigma(G)$, which means that any cell of $\sigma(G)$ is the union of cells of $\sigma(F)$.

As an example, let us construct the secondary polytope of the point configuration $\mathcal{A}$ formed by the vertices of a prism P over a triangle. The homogeneous coordinates of P are given by the columns of the following matrix.

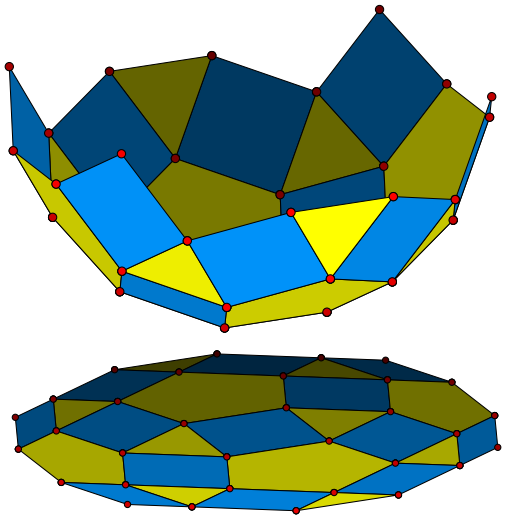

Fig. 10. Convex lifting function and regular subdivision [23]

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix}$$

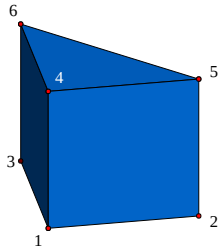

Any triangulation of P must contain one of the tetrahedra formed by the base $\{1,2,3\}$ and one vertex i in the set $\{4,5,6\}$, where the point labels correspond to the column indices of A. This leaves two choices for the apex of the tetrahedron with base $\{4,5,6\}$, and each one determines the last tetrahedron of the triangulation uniquely. We see that there are six distinct triangulations of P in total, namely,

{{1,2,3,4},{2,3,4,5},{3,4,5,6}}, {{1,2,3,4},{2,3,4,6},{2,4,5,6}},
{{1,2,3,5},{1,3,4,5},{3,4,5,6}}, {{1,2,3,5},{1,3,5,6},{1,4,5,6}},
{{1,2,3,6},{1,2,4,6},{2,4,5,6}}, {{1,2,3,6},{1,2,5,6},{1,4,5,6}}.

It turns out that all these triangulations are regular, and therefore we know that they all correspond to vertices of $\Sigma(\mathcal{A})$.

One way to construct the secondary polytope is to start by calculating a basis for the (right) kernel of A, i.e., a matrix B with $AB = 0$. Since A has full rank, its kernel has dimension 2, and one possible basis is given by the

columns of the following matrix B.

$$B = \begin{pmatrix} 1 & 1 \\ -1 & 0 \\ 0 & -1 \\ -1 & -1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}$$

By interpreting the *rows* of B as six points $b_1, b_2, \ldots, b_6$ in $\mathbb{R}^2$, we arrive at the *Gale transform* $\mathcal{A}^*$ of $\mathcal{A}$. In general, if $\mathcal{A}$ consists of n points in d-space (and $\mathcal{A}$ does not lie in any lower-dimensional subspace), then $\mathcal{A}^*$ is made up of n points in $(n-d-1)$-space. Now consider the set $\mathcal{C}(\mathcal{A})$ of all full-dimensional positive cones spanned by the points in $\mathcal{A}^*$ with apex in 0, together with the set R of all their facets. The *chamber complex* $\widetilde{\mathcal{C}}(\mathcal{A})$ of $\mathcal{C}(\mathcal{A})$ is the union of all full-dimensional polyhedral cones whose facets are facets of cones in $\mathcal{C}(\mathcal{A})$, but whose relative interior is not crossed by any member of R. In our two-dimensional example, the set R consists of the six rays

$$R = \Big\{ {}_{\mathbb{R}_{\geq 0}}\langle b_i \rangle : 1 \leq i \leq 6 \Big\},$$

so $\widetilde{\mathcal{C}}(\mathcal{A})$ is given by the following list of cones. See Figure 11 (left).

$$\widetilde{\mathcal{C}}(\mathcal{A}) = \Big\{ {}_{\mathbb{R}_{\geq 0}}\langle b_1, b_6 \rangle,\ {}_{\mathbb{R}_{\geq 0}}\langle b_6, b_2 \rangle,\ {}_{\mathbb{R}_{\geq 0}}\langle b_2, b_4 \rangle,\ {}_{\mathbb{R}_{\geq 0}}\langle b_4, b_3 \rangle,\ {}_{\mathbb{R}_{\geq 0}}\langle b_3, b_5 \rangle,\ {}_{\mathbb{R}_{\geq 0}}\langle b_5, b_1 \rangle \Big\}$$

We now consider each cone $\sigma \in \widetilde{\mathcal{C}}(\mathcal{A})$ in turn, and write down the generators of all cones in $\mathcal{C}(\mathcal{A})$ that contain σ. For instance, $\sigma = {}_{\mathbb{R}_{\geq 0}}\langle b_1, b_6 \rangle$ lies in the cones ${}_{\mathbb{R}_{\geq 0}}\langle b_5, b_6 \rangle$, ${}_{\mathbb{R}_{\geq 0}}\langle b_1, b_6 \rangle$, and ${}_{\mathbb{R}_{\geq 0}}\langle b_1, b_2 \rangle$ of $\mathcal{C}(\mathcal{A})$, and the *complements* $\{1,2,3,4\}$, $\{2,3,4,5\}$, and $\{3,4,5,6\}$ of these index sets correspond precisely to a triangulation of P! Since there are six maximal cones in $\widetilde{\mathcal{C}}(\mathcal{A})$, we expect each one of them to correspond to one of the six regular triangulations of P.

In fact this is true, and even more: The set $\widetilde{\mathcal{C}}(\mathcal{A})$ is a *complete polyhedral fan*, which means that the cones in $\widetilde{\mathcal{C}}(\mathcal{A})$ intersect precisely in common faces, and together span all of $\mathbb{R}^{n-d-1}$. This fan is called the *secondary fan* of $\mathcal{A}$. It has the additional property that it is the *normal fan* of a certain polytope in $\mathbb{R}^{n-d-1}$, which says that the vectors contained in a fixed cone of $\widetilde{\mathcal{C}}(\mathcal{A})$ are just the normal vectors of hyperplanes supporting exactly one face of this polytope. It now comes as no surprise that this polytope is the one defined to be the *secondary polytope* $\Sigma(\mathcal{A})$ of $\mathcal{A}$. Of course, this construction only determines $\Sigma(\mathcal{A})$ up to *normal equivalence*, i.e., any polytope with the same normal fan is also a secondary polytope of $\mathcal{A}$. In any case, passing from one maximal cone of $\widetilde{\mathcal{C}}(\mathcal{A})$ to an adjacent one corresponds to going from one

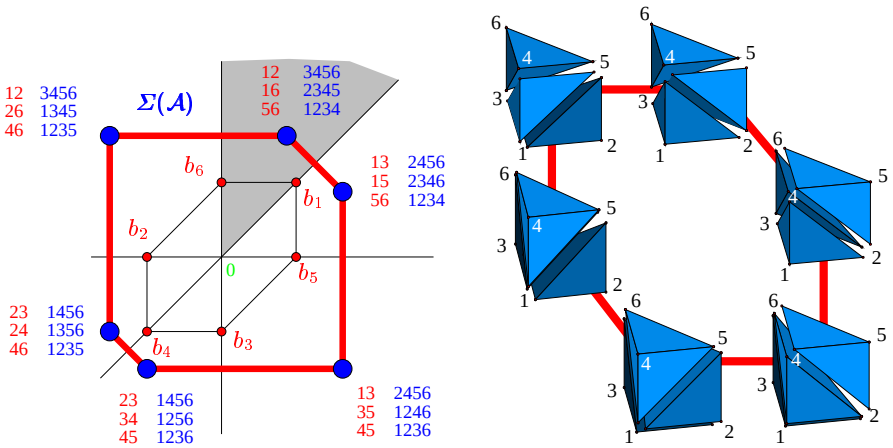

Fig. 11. The hexagon as the secondary polytope of the prism P. *Left:* One maximal cone of the secondary fan is highlighted. Pairs of digits inside such a cone σ index vertices b_i in whose positive span σ lies, and the complementary 4-tuples label the simplex of the triangulation of $\mathcal{A}$ that σ corresponds to. *Right:* Triangulations corresponding to vertices of $\Sigma(P)$. Edges of $\Sigma(P)$ representing flips between triangulations.

vertex of the secondary polytope to the next, and therefore to a flip between these triangulations. This is illustrated in Figure 11 (right).

We summarize our discussion in the following theorem.

Theorem 51 (Gel'fand, Kapranov, and Zelevinsky [8])

1. *The dimension of the secondary polytope $\Sigma(\mathcal{A})$ of a configuration of n points in $\mathbb{R}^d$ is $n-d-1$.*
2. *The faces of $\Sigma(\mathcal{A})$ correspond to the regular subdivisions of $\mathcal{A}$.*
3. *If $F \subset G$ are faces of $\Sigma(\mathcal{A})$, then the subdivision of $\mathcal{A}$ corresponding to F refines the subdivision corresponding to G. In particular, the vertices of $\Sigma(\mathcal{A})$ encode the regular triangulations of $\mathcal{A}$.*

5.2 Hypergeometric Differential Equations and Secondary Polytopes

In this section, we briefly present the connection between secondary polytopes and (initial ideals of) a certain class of systems of partial differential equations.

To the matrix A from the preceding section we can associate the following ideal in the (commutative) polynomial ring of differential operators $k[\partial] =$

$k[\partial_1, \partial_2, \ldots, \partial_n]$ with $n = 6$:

$$\begin{aligned} I_A &= \langle \partial^u - \partial^v : Au = Av,\ u, v \in \mathbb{N}^6 \rangle \\ &= \langle \partial_1\partial_5 - \partial_2\partial_4,\ \partial_1\partial_6 - \partial_3\partial_4,\ \partial_3\partial_5 - \partial_2\partial_6 \rangle, \end{aligned}$$

which corresponds to the system of differential equations

$$\begin{aligned} \frac{\partial^2}{\partial x_1\, \partial x_5} f(x_1, x_2, \ldots, x_6) &= \frac{\partial^2}{\partial x_2\, \partial x_4} f(x_1, x_2, \ldots, x_6), \\ \frac{\partial^2}{\partial x_1\, \partial x_6} f(x_1, x_2, \ldots, x_6) &= \frac{\partial^2}{\partial x_3\, \partial x_4} f(x_1, x_2, \ldots, x_6), \\ \frac{\partial^2}{\partial x_3\, \partial x_5} f(x_1, x_2, \ldots, x_6) &= \frac{\partial^2}{\partial x_2\, \partial x_6} f(x_1, x_2, \ldots, x_6) \end{aligned} \tag{8}$$

for a (formal) power series f in six variables. Notice how the differential operators that generate $I_{\mathcal{A}}$ correspond to elements of the kernel of A. The general theory developed in [24] tells us that a first step in constructing a series solution of (8) is to calculate the initial ideals $\mathrm{in}_\omega(I_A)$ for all term orders $\prec_\omega$ on $k[\partial]$ induced by weight vectors $\omega \in \mathbb{Z}^n$. The positive hull of the weight vectors that select a given initial ideal of $I_{\mathcal{A}}$ is a polyhedral cone in $\mathbb{R}^n$, and it is readily seen that the set of all such cones forms a polyhedral fan, the *Gröbner fan* of I_A. It is then also clear that the weight vectors in the *maximal* cones of the Gröbner fan select *monomial* initial ideals, while those in lower-dimensional cones lead to initial ideals whose generators have more than one term.

Just as for the secondary fan, there exists an equivalence class of polytopes whose normal fan coincides with the Gröbner fan. Any representative from this class is called a *state polytope* [28, Chapter 2] of $\mathcal{A}$. By the preceding paragraph, the vertices of the state polytope exactly correspond to the monomial initial ideals of $I_{\mathcal{A}}$.

In general [28, Prop. 8.15], the Gröbner fan refines the secondary fan; an equivalent way of putting this is to say that the secondary polytope is a *Minkowski summand* of the state polytope. However, for a certain subclass of point configurations it is known that the Gröbner fan coincides with the secondary fan, and that therefore also the state polytope and the secondary polytope are the same (up to normal equivalence). These are the *unimodular* point configurations: those configurations *all* of whose triangulations are entirely made up of simplices of unit volume (appropriately normalized for the dimension of the ambient space).

Therefore, for differential ideals coming from unimodular point configurations $\mathcal{A}$, we can calculate the Gröbner fan via geometrical techniques. We only need to enumerate all triangulations T of $\mathcal{A}$, and for each of them construct

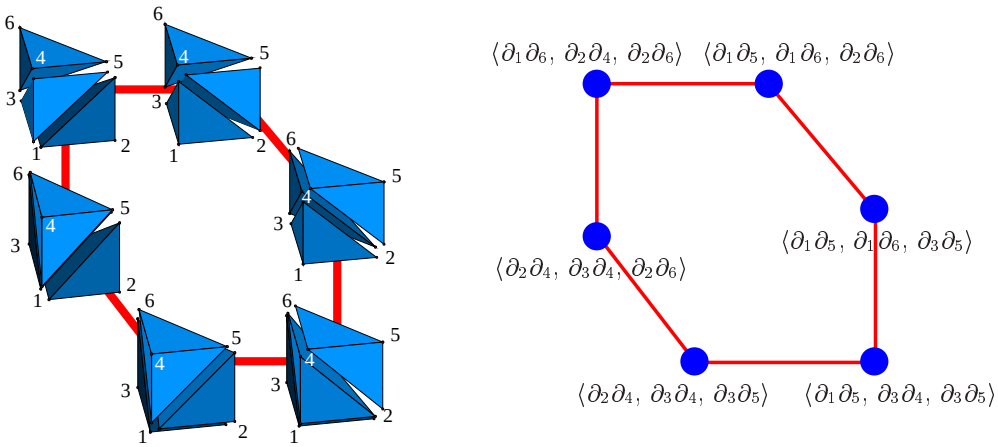

Fig. 12. How to construct geometrically the six initial ideals of the unimodular ideal $I_{\mathcal{A}} = \langle \partial_1\partial_5 - \partial_2\partial_4,\ \partial_1\partial_6 - \partial_3\partial_4,\ \partial_3\partial_5 - \partial_2\partial_6 \rangle$: The generators of each initial ideal are precisely the minimal non-faces of the corresponding regular triangulation of $\mathcal{A}$.

the following ideal, called the *Stanley-Reisner ideal* of T:

$$\left\langle \prod_{j\in J} \partial_j : J \text{ does not index a face of } T \right\rangle \;=\; \bigcap_{\sigma\in T} \langle \partial_j : j \notin \sigma \rangle \;\subset\; k[\partial].$$

By unraveling definitions, this is exactly the initial ideal of I_A selected by any weight vector in the cone of the secondary fan which is dual to the vertex v_T of $\Sigma(\mathcal{A})$. See Figure 12. Note that this initial ideal is square-free by construction.

If the point configuration is not unimodular, i.e., if it admits some triangulation with at least one simplex of non-unit volume, then as we saw above the Gröbner fan is a proper refinement of the secondary fan. The Stanley-Reisner ideal of such a triangulation T is then only the *radical* of the initial ideals selected by weight vectors in those Gröbner cones that refine the cone corresponding to T. It therefore does not properly reflect the algebraic structure of I_A anymore. However, calculating all regular triangulations of $\mathcal{A}$ at least gives a lower bound for the number of monomial initial ideals of I_A.

We now proceed to present a way of actually computing the secondary polytope of a point configuration.

5.3 The GKZ vectors

The original construction of the secondary polytope—presented by Gelfand and Zelevinsky in 1989—remained somewhat mysterious; as we will see, it

gives rise to a straightforward recipe for calculating secondary polytopes, but it is not at all so straightforward to understand what is happening geometrically. In 1992, BILLERA and STURMFELS [2] finally presented secondary polytopes as the *fiber polytopes* of the projection of the $(n-1)$-dimensional simplex to a configuration $\mathcal{A}$ of n points. We will not develop this theory here, but instead refer the interested reader to [2], where also the formulation in terms of Gale transforms was first given, and especially to Chapter 9 of [31].

The GKZ construction proceeds as follows. We associate an n-dimensional vector v_T to any given triangulation T of $\mathcal{A}$, in such a way that the i-th coordinate of v_T is the sum of the volumes of all simplices in T incident to the point i.

$$(v_T)_i \;=\; \sum_{\sigma\colon \sigma\in T,\, i\in\sigma} \operatorname{vol}\operatorname{conv}\sigma$$

This gives us one n-dimensional vector for each triangulation of $\mathcal{A}$. The *secondary polytope* $\Sigma(\mathcal{A}) \subset \mathbb{R}^n$ of $\mathcal{A}$ is then defined as the convex hull of all such vectors obtained by considering every possible triangulation of $\mathcal{A}$.

$$\Sigma(\mathcal{A}) \;=\; \operatorname{conv}\{v_T : T \text{ triangulation of } \mathcal{A}\}$$

It turns out that the secondary polytope defined in this way is not full-dimensional, but resides in an $(n-d-1)$-dimensional subspace. However, the fact that this polytope coincides, up to scaling and normal equivalence, with the secondary polytope as defined earlier definitely comes as a surprise!

5.4 How to find the face lattice of the secondary

To actually calculate the face lattice of the secondary polytope of a point configuration in $\mathbb{R}^d$, we combine the methods presented in Sections 3.1, 3.3, and 5.3. First, we calculate a placing triangulation of $\mathcal{A}$, which is known to be regular. Now we could proceed as in Section 3.4 by completing partial triangulations, but in fact the faster method described in Section 3.3 of generating the connected component of the flip graph that contains the placing triangulation is sufficient for our purposes. Since flips correspond to edges of the secondary polytope, and the 1-skeleton of any convex polytope of dimension at least 2 is connected, we know that this component contains at least all regular triangulations of $\mathcal{A}$—possibly along with some non-regular ones.

Next, we embed the nodes of the flip graph in $\mathbb{R}^n$ via their GKZ coordinates, project the resulting point configuration to $\mathbb{R}^{n-d-1}$, and calculate (the vertex-facet incidence matrix of) the convex hull of the result.

We have achieved an embedding not only of the vertices corresponding to regular triangulations, but of the entire flip graph. This allows to investigate, for example, the *flip distance* of a non-regular triangulation to the nearest regular one. Of course, our "fast" procedure following Section 3.3 misses all connected components of the flip graph that do not contain any regular triangulation.

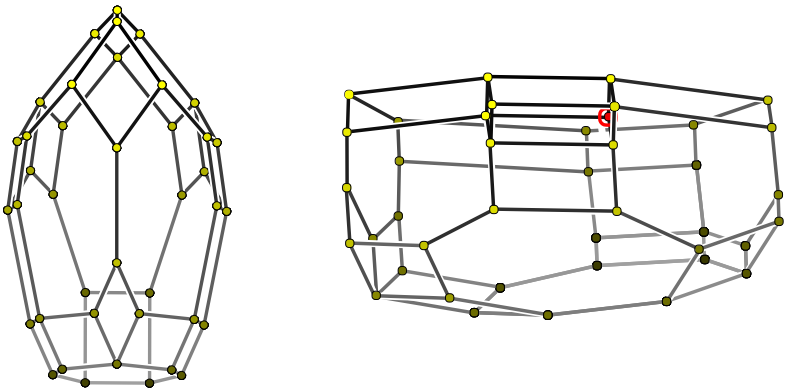

Fig. 13. *Left:* Secondary polytope of the cyclic 4-polytope with 8 vertices. All 40 triangulations are regular. *Right:* Secondary of a different neighborly 4-polytope with 8 vertices [9, Ch. 7], which has one nonregular triangulation (circled; in the interior of the convex hull) among 41

6 Implementing the ideas: Software integration with polymake

We have integrated the tools available in TOPCOM and POLYMAKE by writing clients that interchange between the respective data representations and implement the procedure described in the previous section using the standard POLYMAKE rule base.

To calculate the secondary polytope of a configuration $\mathcal{A}$ of n points in $\mathbb{R}^d$, the user executes the command `secondary point-conf` , where `point-conf` is the name of a POLYMAKE file containing the homogeneous coordinates of $\mathcal{A}$. The client `secondary` converts this data to TOPCOM format, requests a list of all triangulations of $\mathcal{A}$, and for each one calculates the coordinates of its GKZ embedding in $\mathbb{R}^n$. Next, the client asks for these points to be projected to $\mathbb{R}^{n-d-1}$, and then to calculate the convex hull of these projected points. These requests are all answered by the POLYMAKE server, which in turn calls the appropriate clients for each task as specified in the rule base. Finally, the `secondary` client outputs the flip graph of $\mathcal{A}$ both with its embedding and as an abstract (`.gml`-)graph, and marks points corresponding to non-regular triangulations. If dimension permits, this embedded flip graph can then be visualized, e.g. with *Javaview* [17].

In Figures 13 and 14, we present the results of four such calls to `secondary`. All running times, excluding the computation of the convex hull, remained well under one minute on a Sun Blade. The bottleneck is calculating the convex hull: The longest such computation with `cdd` [5] for the secondary of the cyclic 8-polytope with 12 vertices realized on the Carathéodory curve took 2 minutes.

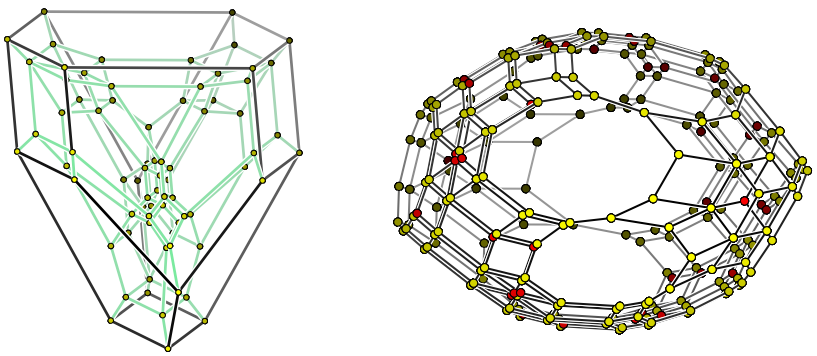

Fig. 14. *Left:* Schlegel diagram of the secondary polytope of the 3-cube [16]. All 74 triangulations are regular. *Right:* Secondary of the cyclic 8-polytope with 12 vertices, realized on the Carathéodory curve [15]. There are 42 nonregular triangulations among 244 in all

With a view towards future developments, we remark that the computation of the entire vertex-facet incidence matrix of the secondary polytope seems wasteful if all one is interested in is the information which edges of the embedded flip graph actually lie on the convex hull of the secondary polytope.

Moreover, while TOPCOM is fine-tuned to exploit symmetries of a point configuration, at this moment there is no convex hull code available that could do likewise. However, implementing an algorithm that simultaneously inserts all points of an orbit under a given symmetry group may well prove to be a non-trivial task.

In the future, perhaps it will be possible to exploit the fact that TOPCOM not only provides a list of points corresponding to triangulations, but in fact a connected component of the flip-graph that includes the entire 1-skeleton of the polytope.

There have been other approaches to computing the edge graph of the secondary polyope by reverse search, e.g. in [14]. Unfortunately, we do not have access to any code based on these algorithms. If memory limitations dominate then reverse search is certainly the method of choice.

7 Conclusion

Oriented matroids are a suitable interface between calculations in coordinates and computations in combinatorial geometry. In this exposition, we have presented the computation of triangulations of a point configuration using oriented matroids. After the computation of the chirotope, all other required operations can be performed in a purely combinatorial way. The package TOPCOM implements this concept. Regular triangulations provide

a beautiful connection to algebraic structures. Their handling, however, requires the integration of additional software; we have shown examples that POLYMAKE is a suitable tool for this.

Acknowledgements

We would like to thank Michael Joswig for many stimulating discussions and his unfailing support.

References

1. Oswin Aichholzer, *The path of a triangulation*, Proceedings of the 15th Annual ACM Symposium on Computational Geometry (Miami Beach, Florida, USA), 1999, pp. 14–23.
2. Louis J. Billera and Bernd Sturmfels, *Fiber polytopes*, Annals of Mathematics **135** (1992), 527–549.
3. Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Günter M. Ziegler, *Oriented matroids*, Encyclopedia of Mathematics, vol. 46, Cambridge University Press, Cambridge, 1993.
4. Jürgen Bokowski and A. Guedes de Oliveira, *On the generation of oriented matroids*, Discrete & Computational Geometry **24** (2000), 197–208, Branko Grünbaum Birthday Issue.
5. Komei Fukuda, *cdd—an implementation of the double description method with LP solver*, `http://www.ifor.math.ethz.ch/~fukuda/cdd_home/cdd.html`.
6. Komei Fukuda and Lukas Finschi, *Generation of oriented matroids—a graph theoretical approach*, Discrete & Computational Geometry **27** (2002), 117–136.
7. Ewgenij Gawrilow and Michael Joswig, *Polymake—a versatile tool for the algorithmic treatment of polytopes and polyhedra*, 2001.
8. Izrail M. Gelfand, Mikhail M. Kapranov, and Andrei V. Zelevinsky, *Discriminants, resultants, and multidimensional determinants*, Mathematics: Theory & Applications, Birkhäuser, Boston, 1994.
9. Branko Grünbaum, *Convex Polytopes*, Interscience, London, 1967.
10. E. Lawler, J. Lenstra, and A. H. G. Rinnooy Kan, *Generating all maximal independent sets: NP-hardness and polynomial-time algorithms*, SIAM Journal on Computing **9** (1980), 558–565.
11. Carl Lee, *Triangulations of polytopes*, CRC Handbook of Discrete and Computational Geometry (Jacob E. Goodman and Joseph O'Rourke, eds.), CRC Press LLC, Boca Raton, 1997, pp. 271–190.
12. Jesús A. de Loera, *Triangulations of polytopes and computational algebra*, Ph.D. thesis, Cornell University, 1995.
13. Jesús A. de Loera, Serkan Hoşten, Francisco Santos, and Bernd Sturmfels, *The polytope of all triangulations of a point configuration*, Documenta Mathematika **1** (1996), 103–119.
14. Tomonari Masada, Hiroshi Imai, and Keiko Imai, *Enumeration of regular triangulations*, Proceedings of the 12th Annual Symposium on Computational Geometry (Philadelphia, PA, USA), ACM Press, 1996, pp. 224–233.

15. Julian Pfeifle, *Secondary polytope of $c_4(8)$ on the carathéodory curve*, Electronic Geometry Model No. 2000.09.032, 2000, `http://www.eg-models.de/`.
16. Julian Pfeifle, *Secondary polytope of the 3-cube*, Electronic Geometry Model No. 2000.09.031, 2000, `http://www.eg-models.de/`.
17. Konrad et al. Polthier, *Javaview*, `http://www.javaview.de`, 2001.
18. Jörg Rambau, *Projections of polytopes and polyhedral subdivisions*, Berichte aus der Mathematik, Shaker, Aachen, 1996, Dissertation, TU Berlin.
19. Jörg Rambau, *Triangulations of cyclic polytopes and higher Bruhat orders*, Mathematika **44** (1997), 162–194.
20. Jörg Rambau, *TOPCOM—Triangulations of Point Configurations and Oriented Matroids*, Software under the Gnu Public Licence, `http://www.zib.de/rambau/TOPCOM.html`, 1999.
21. Jörg Rambau, *Point configuration and a triangulation without flips as constructed by Santos*, Electronic Geometry Model No. 2000.08.005, 2000, `http://www.eg-models.de/`.
22. Jörg Rambau, *Topcom: Triangulations of point configurations and oriented matroids*, ZIB-Report 02-17, 2002, Proceedings of the International Congress of Mathematical Software, to appear.
23. Ulrich Reitebuch, *Rhombicosidodecahedron*, Electronic Geometry Model No. 2000.09.013, 2000, `http://www.eg-models.de/`.
24. Mutsumi Saito, Bernd Sturmfels, and Nobuki Takayama, *Gröbner Deformations of Hypergeometric Differential Equations*, Algorithms and Computation in Mathematics, vol. 6, Springer, 2000.
25. Francisco Santos, *A point configuration whose space of triangulations is disconnected*, Journal of the American Mathematical Society **13** (2000), 611–637.
26. Francisco Santos, *Non-connected toric Hilbert schemes*, Preprint CO/0204044 v1, arXiv:math, April 2002.
27. Francisco Santos, *Triangulations of oriented matroids*, Memoirs of the American Mathematical Society **156** (2002).
28. Bernd Sturmfels, *Gröbner Bases and Convex Polytopes*, University Lecture Series, vol. 8, AMS, 1996.
29. Fumihiko Takeuchi, *Combinatorics of triangulations*, Ph.D. thesis, Graduate School of Information Science, University of Tokyo, 2001.
30. Fumihiko Takeuchi and Hiroshi Imai, *Enumerating triangulations for products of two simplices and for arbitrary configurations of points*, Computing and Combinatorics, 1997, pp. 470–481.
31. Günter M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, vol. 152, Springer, New York, 1995, *Updates, Corrections, and more* available at `http://www.math.tu-berlin.de/~ziegler`.

A Equivariant BFS: An example run for the six-gon

Here, we show how the equivariant BFS behaves on the graph of triangulations of the six-gon. We invite the interested reader to look for the symmetries that induce the marking operations on flips. Only the marking in picture 12 is critical and prevents the algorithm from returning to an old symmetry class. The other marking operations hit already marked edges.

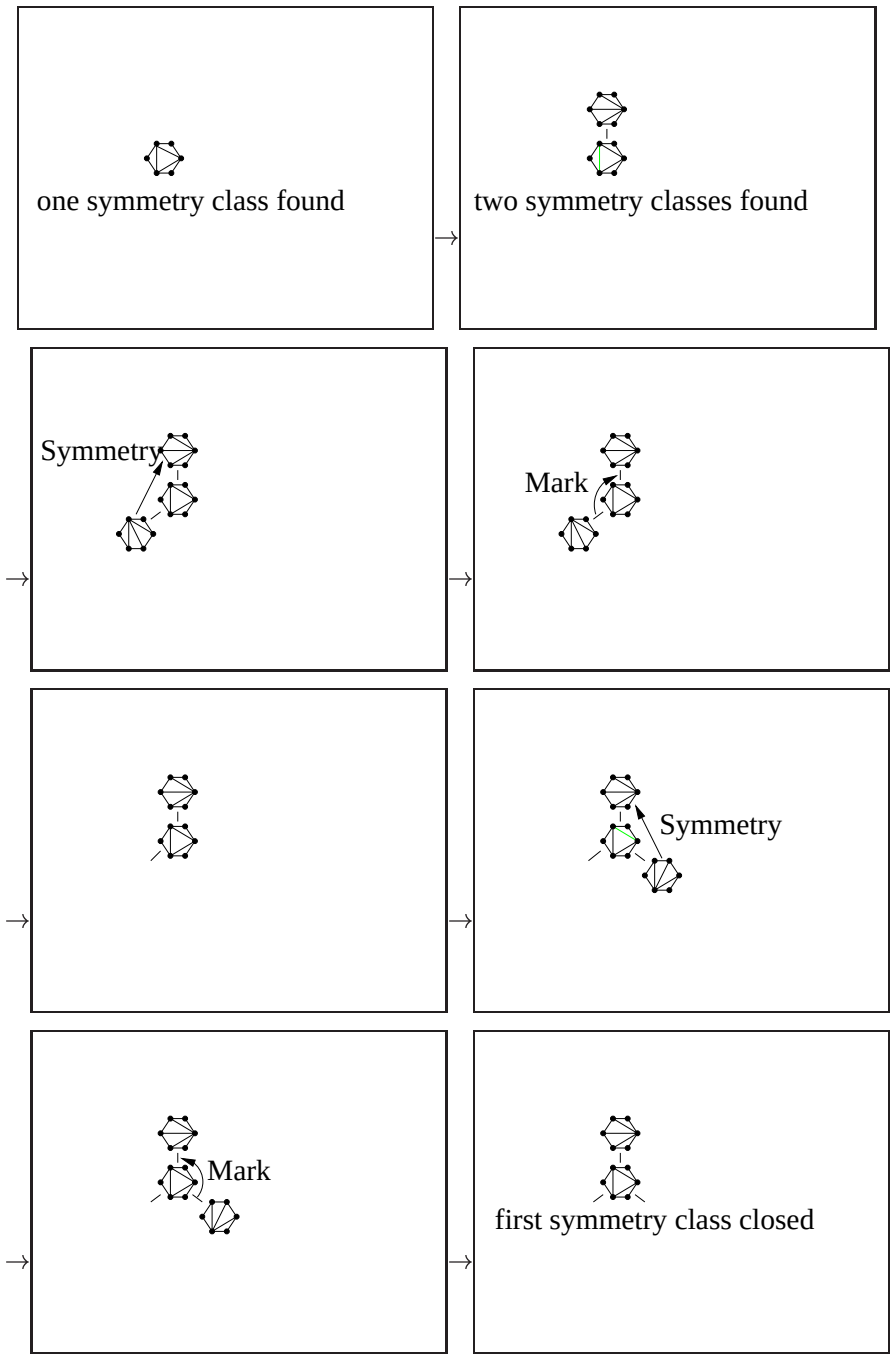
one symmetry class found
two symmetry classes found
Symmetry
Mark
Symmetry
Mark
first symmetry class closed

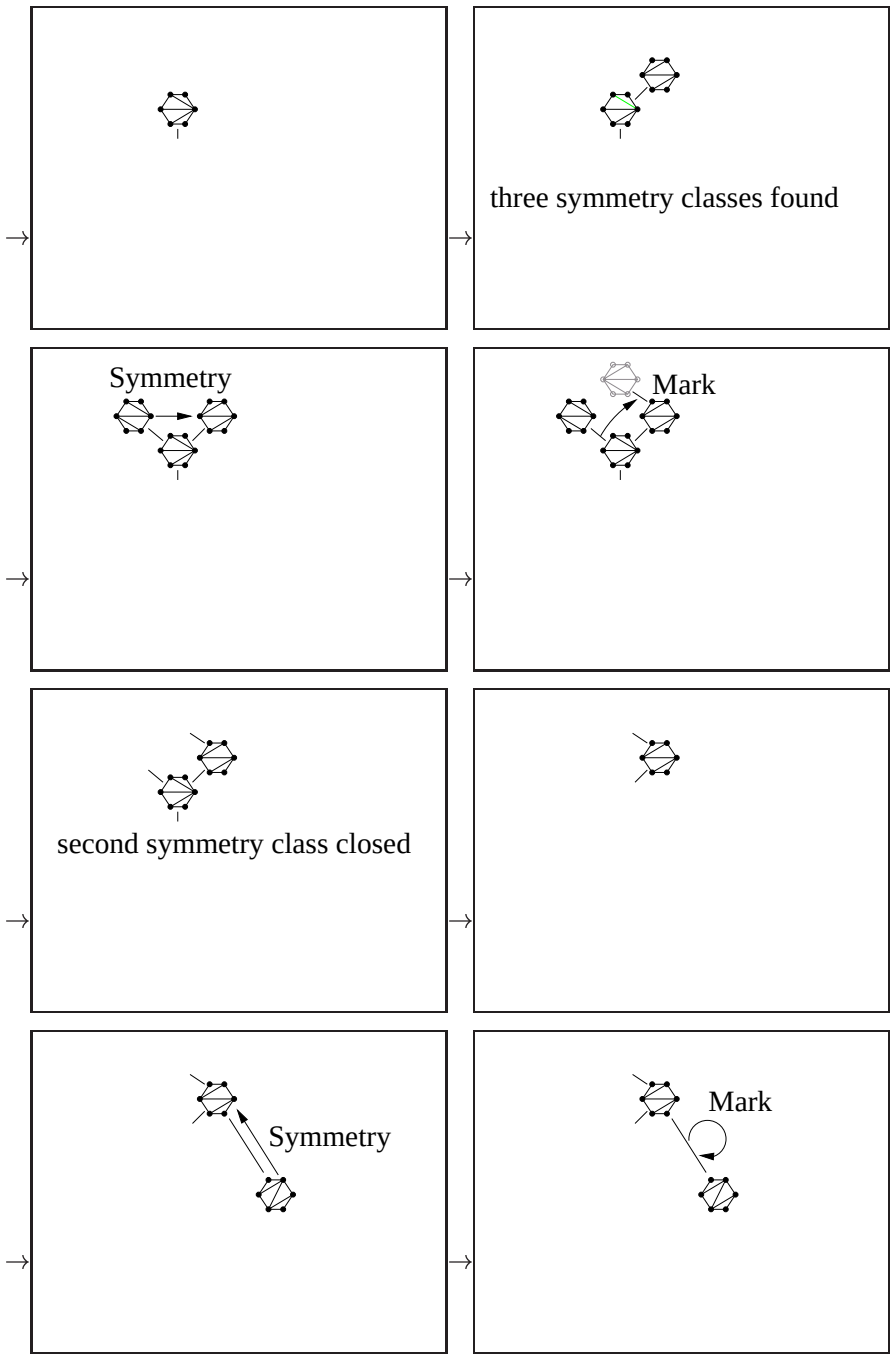
three symmetry classes found
Symmetry
Mark
second symmetry class closed
Symmetry
Mark

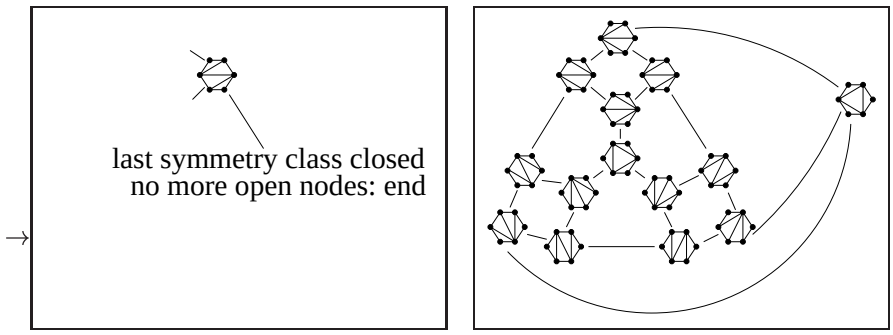
last symmetry class closed
no more open nodes: end