

Parallel Enumeration of Triangulations

Charles Jordan*

Graduate School of Information Science and Technology
Hokkaido University
N-14 W-9, 060-0814 Sapporo, Japan
skip@ist.hokudai.ac.jp

Michael Joswig[†] Lars Kastner[‡]

Institut für Mathematik
TU Berlin
Str. des 17. Juni 136, 10623 Berlin, Germany
{kastner,joswig}@math.tu-berlin.de

Submitted: Sep 14, 2017; Accepted: May 22, 2018; Published: Jul 13, 2018

© The authors. Released under the CC BY-ND license (International 4.0).

Abstract

We report on the implementation of an algorithm for computing the set of all regular triangulations of finitely many points in Euclidean space. This algorithm, which we call *down-flip reverse search*, can be restricted, e.g., to computing full triangulations only; this case is particularly relevant for tropical geometry. Most importantly, down-flip reverse search allows for massive parallelization, i.e., it scales well even for many cores. Our implementation allows to compute the triangulations of much larger point sets than before.

Mathematics Subject Classifications: 52B55 (68U05)

1 Introduction

Triangulations are ubiquitous in combinatorics, optimization, algebra and other parts of mathematics. For an overview about the range of applications we recommend the first

*Research by C. Jordan supported in part by JSPS Kakenhi Grants 16H02785 and 18K18027.

[†]Research by M. Joswig is carried out in the framework of Matheon supported by Einstein Foundation Berlin. Further partial support by Deutsche Forschungsgemeinschaft (SFB-TRR 109: “Discretization in Geometry and Dynamics” and SFB-TRR 195: “Symbolic Tools in Mathematics and their Application”) is gratefully acknowledged.

[‡]Research by L. Kastner is supported by Deutsche Forschungsgemeinschaft (SFB-TRR 195: “Symbolic Tools in Mathematics and their Application”).

chapter of the monograph [7] by De Loera, Rambau and Santos, and the rest of that book is recommended for the foundations. The software which defines the state of the art is Rambau's TOPCOM [25]; see also [22]. An implementation of a different method is part of *Gfan* by Jensen [16]. While this comes with some advantages of its own, TOPCOM is superior in most cases. In this article we describe an algorithm using different ideas but which still shares some of TOPCOM's features. Most importantly we report on new computational results obtained with our implementation MPTOPCOM; some of these are out of the reach of the other software systems.

In order to explain the method it is best to start by looking into TOPCOM's algorithm; see also [7, §8.3]. As a first step, for a fixed point configuration, P , TOPCOM creates one triangulation as a *seed*. Then TOPCOM generates all triangulations of P which can be obtained by local transformations, known as *flips*. The flips induce a graph structure on the set of triangulations of P , and TOPCOM makes a breadth first search starting at the seed. In practical applications it is often important to restrict the attention to those triangulations of P which are *regular*, i.e., they are induced by a convex lifting function. Regularity can be checked via solving a linear programming feasibility problem, and this is supported by TOPCOM, e.g., via calling `cddlib` [10]. Any two regular triangulations of P are connected by a sequence of flips. Checking for regularity on the way is costly but is also beneficial since the regular triangulations of P have a particularly nice structure. The reason is that they correspond to the vertices of a convex polytope, the *secondary polytope* of P , which was introduced by Gel'fand, Kapranov and Zelevinsky [12], and the edges of the secondary polytope correspond to flips (but there are flips which do not arise from edges). The coordinates of a regular triangulation Δ , seen as a vertex of the secondary polytope, form the *GKZ-vector* of Δ . Already in 2002 Imai et al. [13] proposed an algorithm which exploits the GKZ-vectors to tailor a reverse search scheme [2] for enumerating the regular triangulations of P , and we call their method *down-flip reverse search*. Surprisingly, the down-flip reverse search algorithm apparently received little attention, and some of its properties have been discovered independently by Pournin and Liebling [24]. A key challenge in practice is that relevant point sets P exhibit a great deal of symmetry. Usually this symmetry is given in terms of generators of some finite group G acting on P by permutations. The real task now is to compute the (regular) triangulations of P up to symmetry, i.e., exactly one representative from each G -orbit. This reduction is strictly necessary since the total number of triangulations is often too large. Imai et al. [13] also describe down-flip reverse search up to symmetry.

One of the main advantages of the reverse search algorithm is that it is output sensitive. In particular, if one has bounds on the number of neighbors of a node in the search graph, one can derive an effective upper bound on the memory consumption of reverse search, see [2, Theorem 3.2]. The price one pays is recomputation of intermediate results, whenever a node is analyzed, but not visited. This obstacle can be partially overcome by using caches. While TOPCOM's limit often is dictated by the amount of memory available, MPTOPCOM's limit is given by the time one is willing to let the program run. This is where parallelization comes in.

While the general algorithmic idea is known, a practical implementation of down-flip

reverse search requires one to overcome several challenges. That is the focus of this paper. We will show that our implementation `MPTOPCOM` of down-flip reverse search, which can enumerate hundreds of millions of triangulations of a given point set, is superior to other methods in practice. Two aspects are most important. First, for large input, the implementation must be parallelized in order to be able to benefit from modern hardware. This is where the reverse search scheme shows its full strength; `mts` [3] is a competitive parallel implementation of the abstract reverse search method based on `MPI` [30]. Second, by far the most frequent subtask is to compute a canonical representative for each G -orbit of a given triangulation. This allows to distinguish two orbits by comparing their representatives. A minor drawback of parallelized reverse search is that those representatives are recomputed by several workers, even for the same orbit and the same triangulation. This is because we want to avoid *any* communication between the nodes in the computation tree since this is what makes parallelized reverse search so successful. Our main new algorithmic contribution is a procedure for computing these canonical representatives efficiently. This is a variation of the Schreier–Sims algorithm adapted to our special situation; cf. [29, §4.2]. Despite our highly optimized setup, depending on the input and the desired output, the computation of the canonical representatives still may take more than 90% of the combined total time. The benefit of our approach, however, is tremendous. It turns out that, even if we restrict our parallel implementation to a single core, we are able to beat `TOPCOM` by a factor of five or more on medium size input. More importantly, on large input `MPTOPCOM` scales rather well even for more than a hundred cores. This means that we are now able to compute triangulations of point sets an order of magnitude larger than before.

This paper is organized as follows. In Section 2 we collect the basic facts concerning the reverse search method of Avis and Fukuda [2]. This is a powerful general scheme for organizing a large enumeration via a rooted tree. The paradigmatical example is a dual convex hull computation. Here the input is a system of linear inequalities, and the output are the vertices and the rays of the feasible region, which is a convex polyhedron. A generic linear objective function induces an orientation on the edges of that polyhedron, and any pivoting strategy for the simplex method yields a tree, in which each edge is directed toward the global optimum. The basic idea of down-flip reverse search is to mimic this behavior on the secondary polytope of a point configuration. We close this section by describing the principles underlying the parallelization employed by `mts` [3]. Section 3 deals with the basic notions concerning triangulations, and we formulate the down-flip reverse search algorithm. Then, in Section 4 we briefly explain the general approach of Imai et al. [13] to combine enumeration up to symmetry with down-flip reverse search. The core of our paper is Section 5. *The* major challenge in `MPTOPCOM`'s implementation is to efficiently find canonical representatives for each orbit of triangulations. To this end we introduce the concept of *switch tables* and *evaluations* to facilitate that computation (cf. Definition 9). This is developed in a way which specifically addresses the enumeration of triangulations, and yet we believe that it may also be useful in other circumstances. Our main theoretical contribution is a procedure for computing canonical representatives via switch tables (cf. Algorithm 4) and its analysis. In Section 7 we report on experimental results. First we consider a few standard examples, such as the 16 vertices of the 4-dimensional regular cube.

More interesting are our new results. This concerns, e.g., products of simplices; these computations have been employed by Schröter for deriving results on coarsest matroid subdivision of hypersimplices [28]. Further, for the first time we are able to enumerate all regular and full triangulations of the 3-simplex with dilation factor of three. This is a configuration of 20 points in \mathbb{R}^3 , and it turns out that it has precisely 21 125 102 regular and full triangulations, up to the symmetry induced by the symmetric group of degree four, acting on the vertices of the tetrahedron (cf. Theorem 19). This outcome is relevant, e.g., for tropical geometry, as this leads to the classification of the smooth tropical cubic surfaces in the tropical 3-torus; see [18, §4.5]. A full account of the consequences from that one computation is beyond the scope of this paper.

MPTOPCOM is available as open source software which can be downloaded from <https://www.polymake.org/mptopcom>.

We are grateful to Jörg Rambau, Francisco Santos, Benjamin Schröter and an anonymous reviewer for valuable comments and suggestions and to Benjamin Lorenz for assistance with the implementation.

2 Budgeted reverse search and its parallelization

Reverse search [2] is a technique for enumerating large sets of objects. Essentially, one explores a graph $\Gamma = (V, E)$ where V is the set of objects to be enumerated and the edges are given by an adjacency oracle. For every node v , the number $\delta(v)$ denotes the outdegree of v in Γ . The adjacency oracle $\text{Adj}(v, j)$ returns the j th neighbor of $v \in V$ for $j \in [\delta]$ or *null* if no such neighbor exists. To apply reverse search to a problem, one provides the adjacency oracle, a local search function $\pi(v)$ and an element $v^* \in V$. It is required that $\pi(v)$ returns a tuple (u, j) such that $\text{Adj}(u, j) = v$ and that repeated application of π to any $v \in V$ results in a path from v to v^* . The local search function therefore generates a spanning tree of Γ with root v^* . We sometimes omit the second component of $\pi(v)$ when it is not needed, e.g., in line 6 of Algorithm 1.

Reverse search (cf. Algorithm 1) traverses this spanning tree in depth-first fashion starting from v^* , and at each node outputs the object corresponding to that node. One can use the local search function to backtrack. This allows to implement reverse search in a way which does not require any additional memory after an initial setup. However, it is also possible to cache auxiliary information needed by the application, for example to avoid recomputation when backtracking. This way the memory usage can be controlled even for computations with extremely large output. For more details on reverse search see [2].

It was recognized from the beginning that reverse search can be easily parallelized. The enumeration process can be restarted from any node given only a description of that node, and no node is reachable via multiple paths. This means that one does not have to store previously-visited nodes, and communication between processes is minimal. Many reverse search trees are highly unbalanced and so the underlying problem is to explore an unbalanced tree in parallel.

The first parallel reverse search implementation was the generic reverse search layer

Algorithm 1 Reverse Search

```
1: procedure RS( $v^*$ , Adj,  $\pi$ )
2:    $v \leftarrow v^*$ ,  $j \leftarrow 0$ ,  $depth \leftarrow 0$ 
3:   repeat
4:     while  $j < \delta(v)$  do
5:        $j \leftarrow j + 1$ 
6:       if  $\pi(\text{Adj}(v, j)) = v$  then
7:          $v \leftarrow \text{Adj}(v, j)$ 
8:          $j \leftarrow 0$ 
9:          $depth \leftarrow depth + 1$ 
10:        Output  $v$ 
11:       end if
12:     end while
13:     if  $depth > 0$  then
14:        $(v, j) \leftarrow \pi(v)$ 
15:        $depth \leftarrow depth - 1$ 
16:     end if
17:   until  $depth = 0$  and  $j = \delta(v)$ 
18: end procedure
```

in ZRAM [6] which was applied in various areas, including the first parallel program for vertex/facet enumeration [6] and a program for certain quadratic maximization problems [9]. Other parallel reverse search applications include the computation of Minkowski sums [31]. Furthermore, Jensen computed exact mixed volumes [14] and homotopy continuation [15] via parallel reverse search implemented in **Gfan** [16].

Recently budgeted reverse search [4] was introduced as a simple scheme for load-balancing in parallel reverse search. There, one adds an additional parameter (the *budget*) and each of the parallel processes explores its assigned subtree subject to the budget. Once the budget is exhausted, the process backtracks to the root of the subtree while reporting unexplored nodes along the backtrack path. These unexplored nodes are scheduled for later exploration. This is a particularly simple approach to parallel tree exploration that in practice can scale beyond 1000 cores [4]. If properties of the trees generated by an application are known, it can be possible to prove certain performance guarantees [1].

Budgeted reverse search inherits a number of other features from reverse search. One can checkpoint and restart the overall process by waiting for all processes to exhaust their budget and writing the descriptions of unexplored nodes to a file. The budget can be tuned dynamically based on the number of unexplored subtrees available. While the implementation of [4] is specific to vertex/facet enumeration, a generic implementation of parallel budgeted reverse search (**mts**) is also available along with a tutorial [3].

Implementing a budgeted reverse search application with **mts** allows a clean separation of the parallelization layer and application code. This permits independent development of the application and the parallelization layer. The current **mts** implementation uses

MPI [30] and dedicates a process as *master* and another to handling output. This overhead is insignificant when many processes are available, but limits parallel efficiency when only few processes are used. We refer to Section 7 below for further details on how MPTOPCOM employs `mts` for enumerating triangulations.

3 Triangulations of point configurations

Let $P \subset \mathbb{R}^d$ be a finite set of n points that affinely spans the entire space. A (*polyhedral*) *subdivision* Σ of P is a polytopal complex which covers the convex hull $\text{conv } P$, such that the vertices of each cell form a subset of the given points P ; cf. [7, §2.3.1]. The subdivision Σ is *regular* if it is induced by a height function $h : P \rightarrow \mathbb{R}$ in the sense that the lower convex hull of

$$\text{conv}\{(p, h(p)) \mid p \in P\} \subset \mathbb{R}^{d+1}$$

projects to Σ by omitting the last coordinate. The subdivision Σ is a *triangulation* if all its cells are simplices. The set of all subdivisions of P is partially ordered by refinement, and the triangulations are precisely the finest subdivisions. Our goal in this section is to discuss and present algorithms for enumerating all regular triangulations of P .

For a given subdivision Σ the set of all height functions which induce Σ on P is a relatively open polyhedral cone, the *secondary cone* of Σ . The secondary cone of Σ is non-empty if and only if Σ is regular. The set of all secondary cones forms a polyhedral fan, the *secondary fan* $\Sigma\text{-fan}(P)$. The relatively open secondary cones partition the space \mathbb{R}^n of height functions, i.e., the secondary fan is *complete*. Each non-empty secondary cone contains a $(d + 1)$ -dimensional linear subspace, the space of linealities of the secondary fan. Fixing the heights on an affine basis in P amounts to intersecting the secondary fan of P in a way such that each cone in the resulting fan is pointed. We call that $(n - d - 1)$ -dimensional pointed polyhedral fan the *pointed secondary fan* of P . It is unique up to linear transformations.

As a key fact the pointed secondary fan is *polytopal*, i.e., it is the normal fan of a convex polytope of dimension $n - d - 1$. We quickly review the construction. For a triangulation Δ of P the *GKZ-vector* is

$$\text{gkz}_\Delta = (\text{gkz}_\Delta(p) \mid p \in P) ,$$

where $\text{gkz}_\Delta(p)$ is the sum of the (Euclidean) volumes of those simplices in Δ which contain p as a vertex. The convex hull

$$\Sigma\text{-poly}(P) = \text{conv}\{\text{gkz}_\Delta \mid \Delta \text{ triangulation of } P\}$$

of all GKZ-vectors is the *secondary polytope* of P ; its normal fan is the (pointed) secondary fan of P ; cf. [7, §5.2.2]. The vertices of $\Sigma\text{-poly}(P)$ are precisely the GKZ-vectors of the regular triangulations. As we assumed the point configuration P is spanning, the dimension of its secondary polytope equals $n - d - 1$; cf. [7, §5.1.3].

We will now sketch an algorithm by Imai et al. [13] to enumerate (regular) triangulations of P based on applying reverse search to the vertex–edge graph of $\Sigma\text{-poly}(P)$ (or a suitable

supergraph). See [7, §5.3.2] for an account of the geometric facts. We begin by defining a linear objective function on the secondary polytope $\Sigma\text{-poly}(P)$. To this end choose a positive real number M . Then the vector

$$\lambda = (M^n, M^{n-1}, \dots, M)$$

of length n defines a linear form on \mathbb{R}^n . For all sufficiently large $M \gg 0$ that linear form is injective on the finite set $\{\text{gkz}_\Delta \mid \Delta \text{ triangulation of } P\}$. Moreover, if $M \gg 0$, then comparing any two GKZ-vectors with respect to λ amounts to checking the lexicographic ordering. In particular, there is no need to determine any valid choices for M . Like the GKZ-vectors the following relies on the choice of a fixed ordering of the points in P .

Definition 1. For any two triangulations, Δ and Δ' , of P we let $\Delta' > \Delta$ if

$$\lambda(\text{gkz}_{\Delta'}) > \lambda(\text{gkz}_\Delta) ,$$

or if $\lambda(\text{gkz}_{\Delta'}) = \lambda(\text{gkz}_\Delta)$ and Δ' is lexicographically larger than Δ .

This defines a total ordering on the set of all triangulations of P . Since λ induces a total ordering on the vertices of $\Sigma\text{-poly}(P)$ the lexicographic ordering is not required as a tie-breaker if restricted to regular triangulations. However, it is important for nonregular triangulations; see Example 8 below.

A *flip* f is a local modification of a triangulation, Δ , of P which yields another triangulation, Δ' ; we write $f = [\Delta \rightsquigarrow \Delta']$. Each edge of the secondary polytope comes from a flip, but the converse does not hold [7, §5.3.1]. We call f an *up-flip* if $\Delta' > \Delta$. Otherwise $\Delta' < \Delta$, and f is a *down-flip*. The *flip graph* Φ of P is the graph whose nodes are the triangulations of P and whose edges are given by the flips. The flip graph Φ is directed with respect to up-flips. Note that Φ is not necessarily connected [7, §7.3 and §7.4]. However, the subgraph Φ_{reg} induced on the subset of regular triangulations is always connected.

Remark 2. It is known that if $f = [\Delta \rightsquigarrow \Delta']$ is a flip and Δ is regular, then Δ' does not need to be regular. That is, the connected component Φ_{reg}^c of the regular triangulations may be strictly larger than Φ_{reg} . We call Φ_{reg}^c the *regular component* of the flip graph, and we call a triangulation *sub-regular* if it is reachable from a regular triangulation via down-flips. Note that sub-regularity depends on the choice of ordering of triangulations. Moreover, even if Δ and Δ' both are regular then f does not necessarily correspond to an edge of $\Sigma\text{-poly}(P)$. In general the vertex–edge graph of $\Sigma\text{-poly}(P)$ may be a proper subgraph of Φ_{reg} as it may have fewer edges.

The reverse search algorithm (cf. Algorithm 1) works on a graph which is given implicitly by means of an oracle. It is only after the termination of the procedure that the entire graph (or rather a spanning tree) is known. Therefore, we introduce the following notation. Our adjacency oracle $\Phi_{\text{reg}}(\Delta, j)$ for the regular flip graph Φ_{reg} at a given regular triangulation Δ of P returns a regular triangulation which can be obtained from Δ by some down-flip, and it is the j th one in the total ordering from Definition 1. If there are

fewer than j regular triangulations accessible via down-flips, then the oracle returns *null*. We call our predecessor function π , and it returns the maximal triangulation which can be obtained via an up-flip, if it exists. There is a unique regular triangulation Δ^* whose GKZ-vector is maximal among all triangulations, regular or not. We set $\pi(\Delta^*) = \textit{null}$, and we call Δ^* the *optimal* triangulation of P . Recall that our ordering of the triangulations and thus this notion of optimality relies on the choice of an ordering of the points in P .

Algorithm 2 computes all regular triangulations of P . While it can easily be modified to, e.g., obtain all not necessarily regular triangulations which can be obtained from Δ^* by down-flips, the computation of the regular triangulations is probably the most interesting use case.

Algorithm 2 Down-flip reverse search

```

1: procedure DFRS( $P$ )
2:    $\Delta \leftarrow$  some regular triangulation of  $P$ 
3:   while  $\pi(\Delta) \neq \textit{null}$  do
4:      $\Delta \leftarrow \pi(\Delta)$ 
5:   end while
6:   RS( $\Delta, \Phi_{\text{reg}}, \pi$ )
7: end procedure

```

To analyze Algorithm 2 we introduce three parameters. First, δ_{\max} denotes the maximum degree of a regular triangulation in the flip graph. This depends on the point configuration P and the ordering of its points, as this defines the optimal triangulation Δ^* . The following bounds hold:

$$n - d - 1 \leq \delta_{\max} \leq \binom{n}{d+2}. \quad (1)$$

The lower bound is the dimension of the secondary polytope (cf. [7, Corollary 5.3.2]), and that binomial coefficient to the right is an upper bound for the number of circuits of P .

Remark 3. The upper bound in (1) is extremely coarse, and any improvement would be very interesting. In practice, the degree seems to be close to the lower bound, which means that a typical secondary polytope is somewhat close to being simple.

The second parameter, s_{\max} , is the maximal number of facets a triangulation of P may have. It is trivially bounded by

$$s_{\max} \leq \binom{n}{d+1}. \quad (2)$$

Yet, if P is a lattice polytope, we obtain a much better bound from the observation that each lattice simplex must have normalized volume at least one. In this case, this entails

$$s_{\max} \leq d! \cdot \text{vol}(\text{conv } P) ,$$

where vol is the Euclidean volume. For instance, this is tight for all cubes $[0, 1]^d$. Finally, the third parameter, N , is the number of regular triangulations of P . The trivial upper bound

$$N \leq 2^{s_{\max}} \leq 2^{\binom{n}{d+1}} \quad (3)$$

is due to the encoding of a triangulation as the set of its maximal cells. Note that the upper bound does not benefit from counting regular triangulations only. That is, the same bound also holds for the number of all triangulations of P .

The running time of Algorithm 2 is dominated by the combined costs for checking the regularity via linear optimization. Hence it is natural to measure the runtime complexity in the number of LPs to be solved. In MPTOPCOM linear programs are solved via TOPCOM's interface to `cddlib` [10]. This leads to the following straightforward analysis, which generalizes [2, Theorem 3.6] to arbitrary dimensions. For each fixed ordering of the point set P the beneath-and-beyond method provides the corresponding placing triangulation, which is known to be regular; cf. [7, Lemma 4.3.5]. In MPTOPCOM this serves as the initial triangulation required in Step 2. The following is essentially [13, Theorem 13].

Theorem 4. *Given some initial regular triangulation, the down-flip reverse search algorithm computes all regular triangulations of P . Its running time is bounded from above by the time required to solve $O(\delta_{\max} \cdot N)$ linear programs in dimension n with at most $\binom{n}{d}$ linear constraints. The space requirement is bounded from above by $O(d \cdot s_{\max})$ plus the space required for solving the linear programs.*

Here and below we use a unit cost model for representing indices corresponding to points in P . This is adequate since, if n does not fit into a machine size `Int`, there will be far too many triangulations to allow for any reasonable enumeration.

Proof. The correctness follows from (1) the fact that from each regular triangulation we can reach the optimal triangulation Δ^* via up-flips (cf. [7, Theorem 5.3.2]), and (2) the correctness of the reverse search method (cf. Algorithm 1).

For the asymptotic running time we can neglect the time for up-flipping from the initial triangulation to Δ^* since we will revisit all triangulations on that path during the reverse search. Each flip from a regular triangulation to another triangulation, regular or not, is considered at most once. The total number of such flips is bounded by $\delta_{\max} \cdot N$. For each flip we need to decide if the resulting triangulation is regular or not. Given a triangulation Δ its regularity can be determined by solving a linear program in dimension n , the number of points. Each cell of codimension one gives rise to one linear constraint, and there are not more than $\binom{n}{d}$ of these.

The space required comes from storing one triangulation as the set of its maximal cells. Each such maximal cell in turn is a subset of the vertices of cardinality $d+1$, encoded as a list of integers. As pointed out above each of these integers is assumed to be small, whence it is accounted for by a constant space requirement in our analysis. \square

By employing (1), (2) and (3) the above complexities can be translated into (horrendous) bounds in terms of the input parameters d and n . Note that, by [7, Corollary 5.3.11], the

number of up-flips to the optimal triangulation does not exceed

$$\min \left\{ (d+2) \binom{n}{\lfloor \frac{d}{2} + 1 \rfloor}, \binom{n}{d+2} \right\} .$$

Example 5. Figure 1 shows the secondary polytope of the “mother of all examples” (MOAE) from [7, Example 5.5.7]; see also [21]. This is a configuration of six points in \mathbb{R}^2 with three vertices of the convex hull and three points in the interior of the outer triangle. There are 18 triangulations, 16 of which are regular. The 18 triangulations are grouped into five orbits, and the coloring of the vertices in Figure 1 shows those orbits. The two nonregular triangulations (are sub-regular and) share the same GKZ-vector, and these occur as the central blue point in a hexagon with yellow vertices. The optimal triangulation corresponds to the one black vertex (at the top and toward the back). The thick edges form the reverse search tree of the graph $\Phi_{\text{reg}}^c = \Phi$; the optimal triangulation is the root.

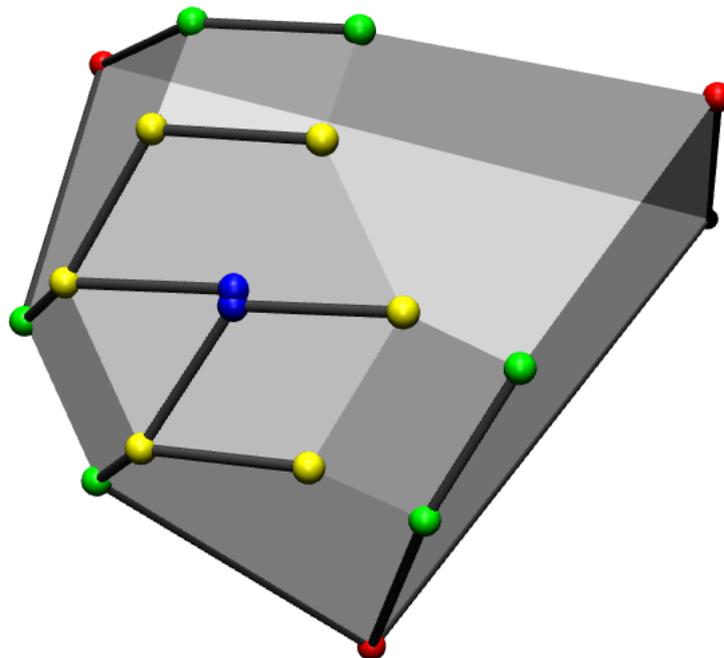


Figure 1: Reverse search tree of MOAE from Example 5. The two blue points correspond to the two nonregular triangulations. Since they share the same GKZ-vector they actually coincide; here we chose to draw them slightly apart in order to reveal the tree.

Remark 6. A triangulation of P is *full* if it uses all the points in P . It was proved in [13, §5] that the subgraph of Φ_{reg} induced on the full triangulations is connected; see also [24] and [7, Corollary 5.3.14]. As a consequence, the down-flip reverse search algorithm can be applied to enumerate the full triangulations only.

4 Triangulations up to symmetry

In many applications for which the set of all (regular) triangulation is sought after the point set exhibits a great deal of symmetry. Typical examples are the set of vertices of a high-dimensional cube (cf. Section 7.1) or the set of lattice points in a dilated simplex (cf. Section 7.3). It is natural to exploit this symmetry, and this is a standard feature of TOPCOM and MPTOPCOM. An obvious drawback of enumerating via reverse search is that the group of, say, affine automorphisms of the point set does not operate on the reverse-search tree. That is, applying reverse search up to symmetry requires some extra considerations; see Bremner, Dutour Sikirić and Schürmann [5, §7] for a brief discussion in the context of convex hull computations. The simple idea is to apply the reverse search scheme to the graph whose nodes correspond to the orbits of the (possibly regular) triangulations and whose edges are induced by the edges in Φ , i.e., by flips.

Following Imai et al. [13] we suggest to adapt the approach via GKZ-vectors from Section 3 to the symmetric setting. Let $P \subset \mathbb{R}^d$ be a finite point set, and let

$$G \leq \mathrm{SL}_d(\mathbb{R}) \rtimes \mathbb{R}^d$$

be a finite group of affine unimodular automorphisms which acts on the set P by permutations. In particular, this action is faithful, and it preserves the volume. Clearly, there is an induced action of G on the set of all triangulations of P , which leaves the set of regular triangulations invariant.

Lemma 7. *Let g be an element of G , and let Δ be a triangulation. Then we have*

$$\mathrm{gkz}_{g \cdot \Delta}(g(p)) = \mathrm{gkz}_{\Delta}(p)$$

for all points $p \in P$.

Proof. This is an immediate consequence of G preserving the volume. \square

The main challenge in implementing reverse search for enumerating triangulations is to find a fast implementation for computing the canonical representatives. To this end we will exhibit a procedure based on the Schreier–Sims algorithm (cf. [29, §4.2]) that works especially well if the orbit size is nearly as large as the group size. In most relevant cases G contains a few thousand elements, while the set of triangulations on which G acts often has millions of orbits. Moreover, despite the fact that the point set is symmetric, few triangulations exhibit much symmetry. This entails that most orbits are about the size of the entire group. See Section 7 for more details and precise numbers.

In order to employ reverse search for traversing (a connected component of) the flip graph, the total ordering on the set of triangulations introduced in Definition 1 is crucial. This leads us to represent the G -orbit of a triangulation Δ of P by

$$\rho(\Delta) := \max(G \cdot \Delta) . \tag{4}$$

That is, the *canonical representative* $\rho(\Delta)$ of an orbit $G \cdot \Delta$ is characterized as follows: (i) its GKZ-vector is lexicographically maximal among all GKZ-vectors of the triangulations

in $G \cdot \Delta$, and (ii) among all triangulations in $G \cdot \Delta$ which satisfy (i) it is lexicographically maximal, considered as a characteristic vector of maximal simplices.

In particular, ρ maps triangulations to triangulations in the same orbit such that $\rho(\Delta) = \rho(\Delta')$ if and only if $G \cdot \Delta = G \cdot \Delta'$, and it follows that $\rho(\rho(\Delta)) = \rho(\Delta)$. Recall that the above definition still depends on the ordering of the points in P as well as on the ordering of the maximal simplices. If the triangulation is regular, then the entire orbit $G \cdot \Delta$ consists of regular triangulations and thus their GKZ-vectors are pairwise distinct.

The following example shows that the GKZ-vectors of nonregular triangulations may behave in an unexpected way, and this shows that condition (ii) is necessary.

Example 8. Let I^4 be the set of 16 vertices of the 4-dimensional 0/1-cube. Here and below we use $I = [0, 1]$ to denote the unit interval on the real line. We may read each vertex as a bitstring of length four, and in this way we obtain a natural encoding in terms of the hexadecimal digits 0 through F. The full group G of affine unimodular automorphisms is the wreath product of a cyclic group of order two (corresponding to a reflection which sends x_i to $1 - x_i$) by the symmetric group $\text{Sym}(4)$ (acting on the four neighbors of a fixed vertex); the total size of G equals 384. Consider the triangulation Δ whose maximal simplices read

01278 0157D 0178D 02478 0457D 0478D 1237A 1278A 137AB
 178AB 178BD 189BD 2467C 2478C 267AC 278AC 478CD 67ACE
 78ABC 78BCD 7ABCE 7BCDF 7BCEF

There is a flip, $f = [\Delta \rightsquigarrow \Delta']$, to another triangulation, Δ' , which replaces the underlined simplices with

2467A 478AC 467AC 2478A .

The triangulations Δ and Δ' lie in different G -orbits (both of which have the maximal length 384), and their respective GKZ-vectors are

$$(6, 10, 8, 2, 6, 2, 3, 23, 14, 1, 9, 10, 11, 10, 3, 2)$$

and

$$(6, 10, 6, 2, 8, 2, 3, 23, 14, 1, 11, 10, 9, 10, 3, 2) .$$

Yet the lexicographically maximal GKZ-vector in both orbits is the same, and it reads

$$(23, 3, 2, 8, 2, 6, 10, 6, 2, 3, 10, 9, 10, 11, 1, 14) .$$

It follows that neither Δ nor Δ' are regular (but they turn out to be sub-regular). The same holds for the canonical representatives $\rho(\Delta)$ and $\rho(\Delta')$.

This example illustrates that, in general, the GKZ-vectors cannot distinguish between G -orbits. Moreover, it also shows that GKZ-vectors alone do not suffice to establish a total ordering on the flip graph. For a tie-breaker we need, e.g., the lexicographic ordering as in Definition 1.

5 Canonical representatives via switch tables

The purpose of this section is to explain why the canonical representatives of triangulations are chosen as in (4) and how this can be exploited. To this end we start out with a more abstract setting. We assume that the finite group G acts on the set $[m] := \{0, \dots, m-1\}$. The following concept is our main tool. The underlying idea is a variation of the classical Schreier–Sims algorithm; cf. [29, 4.2].

Definition 9. An m -switch table for G is a function

$$\text{st} : [m] \times [m] \rightarrow G$$

$$(i, j) \mapsto \begin{cases} g \in G, \text{ with } g(k) = k, \text{ for } k < i, \text{ and } g(j) = i & \text{if it exists} \\ \text{id} & \text{otherwise} \end{cases} .$$

Note that we have $j > i$ if $\text{st}(i, j) \neq \text{id}$. We denote by $\mu(\text{st})$ the *depth* of the switch table, which is defined as

$$\mu(\text{st}) := \max \{i + 1 \mid \text{there is an index } j \text{ with } \text{st}(i, j) \neq \text{id}\} .$$

That is, μ is the index of the first row of st which only contains the identity. In general, a switch table is by no means unique as there may be many candidates in G for $\text{st}(i, j)$.

The i th row of a switch table tells us which elements $\geq i$ can be moved to position i while leaving the first i elements of $[m]$ unchanged. That is, the switch $\text{st}(i, j)$ lies in the stabilizer

$$G_{[i]} := \{g \in G \mid g(j) = j \text{ for all } j < i\} .$$

Formally, we have $G_{[0]} = G$. Moreover, the stabilizer $G_{[m]}$ is the kernel of the action of G on the set $[m]$, and this is a normal subgroup. In particular, that action is faithful if and only if $G_{[m]}$ is the trivial group. Let

$$\sigma_i := 1 + \# \{ \text{st}(i, j) \neq \text{id} \mid j > i \}$$

be one plus the number of nontrivial entries in the i th row of the switch table.

Proposition 10. *The i th row of the switch table forms a (left) transversal of the subgroup $G_{[i+1]}$ in $G_{[i]}$. In particular, the number σ_i is the index of $G_{[i+1]}$ in $G_{[i]}$.*

Proof. Consider the case $i = 0$, where we need to show that the 0th row of the switch table gives a transversal of $H := G_{[1]}$ in $G = G_{[0]}$. The group G acts transitively on the left cosets of H by multiplication on the left. Let g be an element in the complement $G \setminus H$. If this does not exist there is nothing to show. Since g does not stabilize $i = 0$ there is an index $j > 0$ such that $g \cdot j = 0$. By definition of the switch table there must also be a switch $s = s(0, j)$ with $s \cdot j = 0$. We obtain that $s^{-1}g$ lies in H and thus $sH = gH$, which proves the claim for $i = 0$.

Erasing the 0th row and the 0th column of the switch table yields a switch table for $H = G_{[1]}$. This shows that, inductively, the above argument also resolves the general case. \square

Corollary 11. For each element $g \in G$ there is a unique sequence $\text{st}(0, j_0), \text{st}(1, j_1), \dots, \text{st}(m-1, j_{m-1})$ of switches such that the product

$$\text{st}(0, j_0) \text{st}(1, j_1) \cdots \text{st}(m-1, j_{m-1}) g^{-1}$$

lies in the kernel $G_{[m]}$.

Proof. From Proposition 10 we know that there is a unique switch $\text{st}(0, j_0)$ such that $\text{st}(0, j_0)^{-1} g \in G_{[0]}$. Inductively, we obtain $\text{st}(i+1, j_{i+1})$ by requiring

$$\text{st}(i+1, j_{i+1})^{-1} \text{st}(i, j_i)^{-1} \cdots \text{st}(0, j_0)^{-1} g \in G_{[i+1]} .$$

The uniqueness follows as we have $\text{st}(i, j) = \text{id}$ whenever $j \leq i$. □

Let $\sigma := \sigma_0 \sigma_1 \cdots \sigma_{m-1}$ be the product of the σ_i .

Corollary 12. We have $\sigma \leq m!$, and equality holds if and only if the quotient $G/G_{[m]}$ is the symmetric group $\text{Sym}(m)$. Moreover, σ agrees with the order of G if the action of G on $[m]$ is faithful.

Proof. We have $\sigma_i \leq m-i$, and this shows the upper bound. By Corollary 11 the number σ is the index of the kernel $G_{[m]}$ in G . The quotient $G/G_{[m]}$ is isomorphic to a subgroup of $\text{Sym}(m)$, and this characterizes the equality case. □

A switch table is a special notation for a system of strong generators for G (cf. [29, §4]), tailored for finding the optimal element in an orbit; cf. Proposition 15 below. Strong generators are used, e.g., in `permlib` [27]. However, the concepts differ as can be seen from the example below.

Example 13. Let $G = \text{Sym}(4)$ be the symmetric group of degree four acting naturally on the set $\{0, 1, 2, 3\}$. Using cycle notation a switch table st is given by

$$\begin{aligned} \text{st}(0, 1) &= (0\ 3\ 2\ 1), & \text{st}(0, 2) &= (0\ 2)(1\ 3), & \text{st}(0, 3) &= (0\ 1\ 2\ 3), \\ \text{st}(1, 2) &= (1\ 3\ 2), & \text{st}(1, 3) &= (1\ 2\ 3), & \text{st}(2, 3) &= (2\ 3), \end{aligned}$$

and all other entries of the switch table are the identity element. The depth of this switch table is 3. We have $\sigma_0 = 4$, $\sigma_1 = 3$, $\sigma_2 = 2$ and $\sigma = 24$, which is the order of $\text{Sym}(4)$.

Let G_i^{st} be the subgroup of G , called the *i th switch group*, which is generated by the i th row of the switch table st . The switch group G_0^{st} is generated by the 4-cycle $(0\ 1\ 2\ 3)$, while G_1^{st} is generated by the 3-cycle $(1\ 2\ 3)$, and G_2^{st} is generated by the transposition $(2\ 3)$. Consequently, G_0^{st} is a proper subgroup of $G_{[0]} = G$. Similarly, G_1^{st} is a proper subgroup of $G_{[1]}$, which is the symmetric group of degree three acting on $\{1, 2, 3\}$, whereas G_2^{st} agrees with $G_{[2]}$. Observe that G_2^{st} is not a subgroup of G_1^{st} , which in turn is not a subgroup of G_0^{st} .

Now we consider a second action of G on some other set Ω and a map $v : \Omega \rightarrow \mathbb{R}^m$. Our first action of G on $[m]$ induces a (linear) action on the vector space \mathbb{R}^m by permuting the coordinates. If the compatibility condition

$$v(g \cdot \omega) = g \cdot v(\omega)$$

is met we call v an m -*evaluation map* of the G -action on the set Ω . Due to the compatibility we obtain an action of G on the set $v(\Omega) \subset \mathbb{R}^m$. Now we define the *canonical representative* $\rho(v(\omega))$ for that action as the lexicographically maximal vector in the orbit $G \cdot v(\omega)$. Below we will discuss the relationship of this definition with the canonical representative of a triangulation from (4). The following example is natural.

Example 14. If G acts on $[m]$ then this induces a second action on the set Ω of all subsets of $[m]$. Sending $\omega \in \Omega$ to its characteristic vector of length m yields an evaluation map, say χ . The canonical representative of an orbit is the lexicographically smallest set in that orbit.

This constitutes the second part of our approach for MPTOPCOM. The two-step approach can be summarized as choosing v to be $v := (\text{gkz}, \chi)$. The canonical representative then has lexicographically largest gkz-vector, and is the largest element among those with the same gkz-vector. If one is looking at regular triangulations exclusively, one may as well forget about the second part and choose $v := \text{gkz}$.

Our main method for computing canonical representatives is the procedure CANONICAL in Algorithm 4. It essentially relies on the function GOODSWITCHES in Algorithm 3 which determines all switches that may lexicographically improve a given vector $z \in \mathbb{R}^m$ in the orbit of $G_{[i]}$. The idea is to employ a depth first search.

Algorithm 3 Find switches which may yield a larger vector in $G_{[i]} \cdot z$

```

1: procedure GOODSWITCHES( $z, i, \text{st}$ )
2:    $Y = (y_0, y_1, \dots, y_{\ell-1}) \leftarrow$  sort descending  $\{z_j \mid z_j > z_i\}$ 
3:   for  $k = 0, 1, \dots, \ell - 1$  do
4:      $J \leftarrow \{j \in [m] \mid z_j = y_k\}$ 
5:      $S \leftarrow \{\text{st}(i, j) \mid j \in J, \text{st}(i, j) \neq \text{id}\}$ 
6:     if  $S \neq \emptyset$  then
7:       return  $S$ 
8:     end if
9:   end for
10:   $J \leftarrow \{j \in [m] \mid z_j = z_i\}$ 
11:  return  $\{\text{st}(i, j) \mid j \in J\}$ 
12: end procedure

```

Proposition 15. Let $\omega \in \Omega$ be an arbitrary element. For all $i \in [m]$, Algorithm 4 computes an element $\omega' \in G_{[i]} \cdot \omega$ such that $v(\omega')$ is lexicographically maximal among all elements in the orbit $G_{[i]} \cdot \omega$. In particular, for $i = 0$, the evaluation $v(\omega')$ is the canonical representative $\rho(v(\omega))$.

Algorithm 4 Canonical representative of an orbit

```
1: procedure CANONICAL( $\omega, v, i, \text{st}$ )
2:   if  $i > \mu(\text{st})$  then
3:     return  $\omega$ 
4:   end if
5:    $S \leftarrow \text{GOODSWITCHES}(v(\omega), i, \text{st})$ 
6:    $\omega' \leftarrow \omega$ 
7:   for  $s \in S$  do
8:      $\omega'' \leftarrow \text{CANONICAL}(s \cdot \omega, v, i + 1, \text{st})$ 
9:     if  $v(\omega'') > v(\omega')$  then
10:       $\omega' \leftarrow \omega''$ 
11:    end if
12:  end for
13:  return  $\omega'$ 
14: end procedure
```

Proof. The correctness of Algorithm 4 essentially follows from Corollary 11, which provides a transversal of $G_{[m]}$ as a subgroup of G in terms of switches. As that transversal arises from an inductive construction Algorithm 4 works recursively.

It remains to discuss the case where $v(\omega')_0 = v(\omega)_0$. Then we need to consider all possible switches which may or may not change the entry with index 0 but keep the v -value. This may include the identity. \square

Example 16. Again we consider the symmetric group of degree four, G , acting on the set $\{0, 1, 2, 3\}$ and the 4-switch table st from Example 13. As a second action of G we take the induced action on the powerset $\Omega = 2^{\{0,1,2,3\}}$. Then the map v which sends $\omega \in \Omega$ to its characteristic vector is a 4-evaluation; see Example 14. We want to compute $\rho(\{1, 2\}, 0)$ via Algorithm 4, i.e., we are taking the two-element set $\omega = \{1, 2\}$ as our input. Hence we have

$$v(\omega) = (0, 1, 1, 0) .$$

The set S originates from Step 7 in Algorithm 3, and it comprises the two permutations $(0\ 3\ 2\ 1)$ and $(0\ 2)(1\ 3)$.

In the first round of the for-loop we have $\text{st}(0, 1) = (0\ 3\ 2\ 1)$, and thus $\text{st}(0, 1) \cdot \omega = \{0, 1\}$. Recursively we compute $\rho(\{0, 1\}, 1) = \{0, 1\}$. For the second switch $\text{st}(0, 2) = (0\ 2)(1\ 3)$ with $\text{st}(0, 2) \cdot \omega = \{0, 3\}$ we get $\rho(\{0, 3\}, 1) = \{0, 1\}$. Hence, both branches return $\{0, 1\}$ and the algorithm produces $\{0, 1\}$ as its output.

For $\omega \in \Omega$ denote by $\psi(\omega)$ the maximum number of identical entries of the vector $v(\omega)$. Since G acts on the set $v(\Omega)$ by coordinate permutations, the number $\psi(\omega)$ is an invariant of the orbit $G \cdot \omega$. We let ψ be the maximal $\psi(\omega)$, taken over all elements in Ω .

Define ϕ_i as the minimum of ψ and σ_i ; this depends both on the switch table and on the evaluation function. Furthermore, we set $\phi := \phi_0 \phi_1 \cdots \phi_{\mu(\text{st})-1}$. By construction we have $\phi \leq \sigma \leq \#G$ and $\phi \leq m!$. The benefit of our somewhat elaborate setup comes

from the fact that the number ϕ may be much smaller than the order of the group G ; see Example 18 below.

Corollary 17. *The worst-case time complexity of Algorithm 4 is of order $O(\max\{\phi \cdot m, m \log m\})$ and the worst-case space complexity is of order $O(m^3)$.*

Proof. Time complexity: In the procedure `GOODSWITCHES`, called with the vector z of length m and the index i , the set E constructed in Step 2 has size at most $m - 1$. Hence sorting E is of order $O(m \log m)$. The size of set J constructed in Step 5 or Step 11 is bounded by the number of identical entries $\psi(\omega)$ in the vector $v(\omega)$. Thus, the set S constructed in Step 6 or Step 12 is bounded by the size of J as well. But the set S also contains only non-trivial entries from the i th row of the switch table, in case of Step 12 a single copy of the identity. Hence the set of switches returned is of size at most ϕ_i .

The height of the recursion tree of `CANONICAL` is at most $m - 1$, and hence the number of leaves is bounded by $\phi_0 \phi_1 \cdots \phi_{m-1} = \phi$. Consequently, the total number of nodes and also the total number of edges is of order $O(\phi)$. For each edge we lexicographically compare two vectors of length m . All other costs are dominated by the total complexity of these comparisons.

Space complexity: The number of entries in the switch table is given by the expression

$$\sum_{i=0}^{\mu(\text{st})-1} \phi_i \leq \sum_{i=0}^{\mu(\text{st})-1} m - i \leq m^2 . \quad (5)$$

Algorithm 3 returns at most one row of the switch table, and the recursion depth is bounded by the depth of the switch table. Each element of the switch table is a permutation of length m . Thus the total space requirement amounts to $O(m^3)$. \square

Note that $O(m^3)$ is a coarse estimate. In practice the depth of the switch table is often small, and then (5) provides better bounds.

Finally, we can explain our choice for the canonical representative of a triangulation as in (4), which rests on Definition 1. Let us recall our setup. The point set $P \subset \mathbb{R}^d$, of cardinality n , is affinely spanning. It is equipped with a group G of affine unimodular automorphisms. This action induces an action on the set Ω of all triangulations of P . We can encode a triangulation as its set of d -simplices, which are the maximal cells. The latter are encoded as those $(d+1)$ -subsets of P which are affinely spanning. If m is the number of all d -simplices then encoding a triangulation Δ as its characteristic vector among the set of all d -simplices yields an m -evaluation map of the action of G on Ω ; see Example 14. For computing the canonical representative of a triangulation $\Delta \in \Omega$ in a brute-force approach all elements of G are applied and the lexicographically largest one is picked. Of course, the elements of G can be precomputed once in the initialization. Since most orbits are expected to be about the size of G this is often superior to the more traditional approach of trying generators of G until no new triangulations are found, since it requires fewer comparisons of (characteristic vectors of) triangulations.

Yet evaluating at GKZ-vectors leads to a significant improvement. The map which sends a triangulation Δ to its GKZ-vector is an n -evaluation, and n is always much smaller than m . More importantly most entries of a typical GKZ-vector are distinct and hence the parameter ψ which enters the complexity analysis in Corollary 17 is very small. This explains criterion (i) in Definition 1. The tie-breaker criterion (ii) is only necessary for dealing with nonregular triangulations. The following example occurs in a computation which should be considered small by current standards. Larger input results in larger gains.

Example 18. Let us again look at the set $P = I^4$ of vertices of the 4-dimensional regular cube. The group G is the full group of affine (unimodular) automorphisms of order 384. Mapping to GKZ-vectors yields a 16-evaluation. Any switch table has depth four with

$$\sigma_0 = 16, \sigma_1 = 4, \sigma_2 = 3, \sigma_3 = 2,$$

Their product $\sigma = \sigma_0\sigma_1\sigma_2\sigma_3$ equals 384, which is the order of the group G .

To assess the complexity of computing the canonical representatives, the shape of the GKZ-vectors is the key. Here is the full distribution of values of $\psi(\cdot)$ for the 247 451 orbits of sub-regular triangulations of I^4 .

1	2	3	4	5	6	7	8	9	10	11	12
38 673	134 773	58 835	11 699	2 985	364	107	11	2			2

The average value of ψ (taken over all orbits) is approximately 3.22. This results in an average complexity parameter of $\phi \approx 66.8$, which is much smaller than 384, the order of the group.

6 Implementation details

Our software MPTOPCOM builds on and uses existing code from `mts` [3], `polymake` [11] and `TOPCOM` [25]. The `mts` setup dedicates one process to the master and a second one for dealing with the output; the remaining processes are reserved for the workers.

Let us give a brief overview of the interplay of the different software systems: While parallelization is handled by `mts`, the single worker processes are in the domain of `polymake` and `TOPCOM`. Triangulations and groups are handled by `TOPCOM`'s highly optimized code. Since `TOPCOM` provides vectors and matrices already, a first implementation used these for gkz-computation. However, replacing `TOPCOM`'s vectors and matrices by the corresponding objects from `polymake` vastly improved performance. At this point, all vectors, matrices, and sets are handled by `polymake`. Checking regularity is handled by `TOPCOM`'s internal interface to `cddlib` [10].

Parallelizing with more threads will scale almost linearly at first, but depending on the size of the example, the curve of time consumption over number of threads will flatten sooner or later, as can be seen in the Figures 2, 3, and 4. In order to determine whether a down-flip is valid, the reverse search algorithm descends to the target node and checks

whether the predecessor of the target is the original node. This happens essentially for every down-flip into a triangulation. Thus, for a given triangulation the predecessor is computed multiple times, and hence, its neighbors are computed multiple times. If one considers regular triangulations exclusively, then regularity is checked multiple times as well. In the symmetric case, the canonical representative is computed several times. We attack this problem by maintaining three caches. The key type in every cache is a triangulation Δ , so we just list the values.

- (i) *Flip cache*: Contains the list of all flips of Δ .
- (ii) *Orbit cache*: The canonical representative from the orbit of Δ .
- (iii) *Regularity cache*: A boolean whether Δ is regular or not.

Each worker is equipped with three such caches; they follow the least-recently-used paradigm and can store a fixed number of keys, subject to change by the user. Caching is combined with hashing such that previously computed data is instantly available, if it is still cached. There are a number of further caches with dynamic size, e.g., a cache containing the volumes of all maximal simplices. These caches are filled at startup before the reverse search begins and never changed again. Sharing these caches among the workers as hinted at before is possible, but not realized yet, and would probably damage the flexibility of MPTOPCOM.

Parallelizing does not change the number of down-flips going into a triangulation. But since the caches are not shared among the workers, every worker populates its own caches, leading to the non-linearity in scaling. In theory, working without caches would scale linearly, but is infeasible for larger examples, even starting with I^4 .

We use several encodings of the same triangulation simultaneously. For instance, enumerating all maximal simplices spanned by our point configuration first allows to store a triangulation as a set of machine-size integers. Those integers are the indices pointing into the array of maximal simplices.

The implementation is flexible enough to deal with several scenarios. In our fastest setup we assume that the coordinates of the point configuration are machine size integers. This means that we may use `int` for the entries of the GKZ-vectors. The condition on the integrality is natural for the applications for algebraic and tropical geometry we have in mind.

As pointed out several times the total ordering on the triangulations from Definition 1 depends on an ordering of the points. In practice it is often beneficial to pick a random ordering. This may reduce the height of the search tree for the canonical representative.

We experimented with several compilers on various kinds of hardware. By and large we found `clang`, version 3.8.0, to be about 10% faster than various versions of `gcc` compilers. Therefore our timings below employ `clang`. For the hardware we tried the following:

- A cluster with four nodes, each of which is equipped with $2 \times$ Intel Xeon E5-2630 v2 Hexa-Core (2600–3100MHz, 5201.45 bogomips) and 64GB RAM per node. On this machine we used 40 threads. The operating system is openSUSE 42.2 with kernel version 4.4.79.

- A desktop machine with an AMD Ryzen 7 1700 CPU. This has 8 cores/16 threads (3000 MHz, 5967.87 bogomips) and 32GB RAM. The operating system is openSUSE 42.1 with kernel version 4.12.1.

Recall that two processes are reserved for the master and the output processes. That is, n processes means $n - 2$ workers.

7 Experimental results

We tried our implementation on a number of point configurations which occur naturally in geometric combinatorics and related areas. Our notation is as follows. The point set Δ_d comprises the $d + 1$ vertices $0, e_1, e_2, \dots, e_d$ of the standard d -dimensional simplex, where e_i is the i th standard basis vector of \mathbb{R}^d . We abbreviate $I = [0, 1]$, and so I^d is the d -dimensional unit cube. From these several interesting point configurations can be formed, e.g., by taking products. The point configurations of type I^d or $\Delta_p \times \Delta_q$ or $I^p \times \Delta_q$ are in convex position, i.e., they form the vertices of polytope. For such point configurations Table 3 shows the number of orbits of regular and sub-regular triangulations; in all cases the group G is the full group of affine symmetries.

To the best of our knowledge the results for $\Delta_2 \times \Delta_6$, $\Delta_3 \times \Delta_4$ and $3 \cdot \Delta_3$ are new.

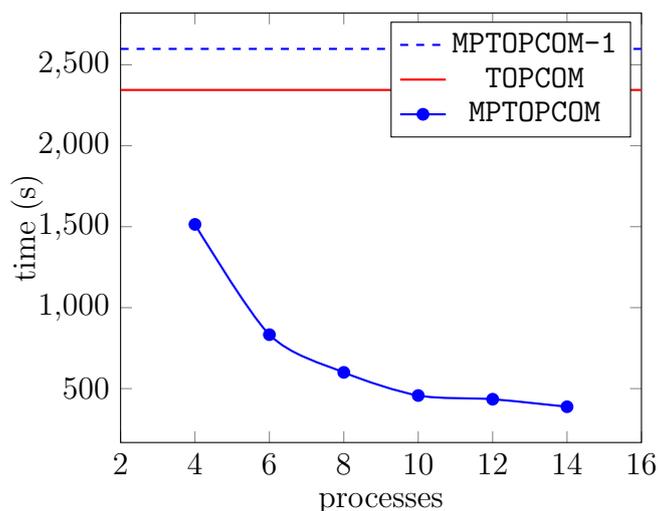


Figure 2: Timings for enumerating the sub-regular triangulations of I^4 taken on AMD Ryzen 7 1700 with 32GB RAM, depending on the number of processes. The timings for the single-threaded version of MPTOPCOM (marked with “-1”) and TOPCOM are added for reference.

7.1 The 4-cube

A standard test case is the four-dimensional cube I^4 ; see also Example 18. The group of affine unimodular automorphisms has order 384, and there are 247 451 orbits of triangulations in the flip-component of the regular triangulations. By today’s standard this is

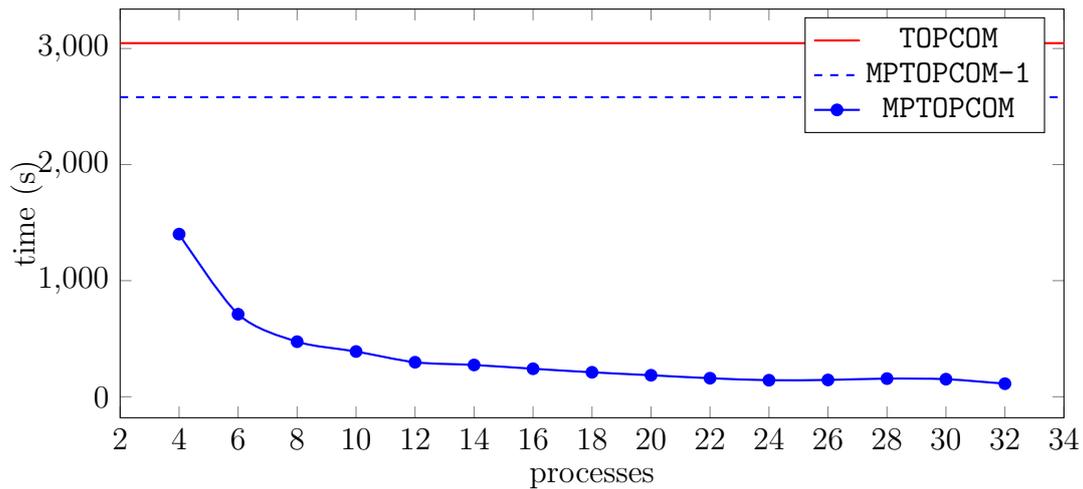


Figure 3: Timings for enumerating the sub-regular triangulations of I^4 taken on the Intel Xeon E5-2630 v2 cluster with 64GB RAM per node, depending on the number of processes.

to be considered small input. Out of these, actually 235 277 are regular [7, Thm. 6.3.12]. Pournin proved that the flip graph of I^4 is connected [23]. Since MPTOPCOM also counts 247 451 orbits it follows that each triangulation of I^4 is sub-regular. The total number of triangulations of I^4 adds up to 92 487 256.

On an AMD Ryzen 7 1700 TOPCOM takes 2345 seconds to enumerate the regular component, while MPTOPCOM requires 2598 seconds (single-threaded) and 457 seconds (with 10 processes), respectively. Figure 2 shows how MPTOPCOM scales with the number of processes. It shows that, on a standard desktop computer, already four processes, i.e., two workers, suffice for MPTOPCOM to be substantially faster than TOPCOM. This should be compared with Figure 3 which shows very similar behavior on our cluster. The only exception is that, on that hardware, even a single-threaded version of MPTOPCOM beats TOPCOM. The single-threaded MPTOPCOM-1 is the pure down-flip reverse search algorithm built without the overhead of `mts` and MPI.

Next we try to give an idea about which percentage of the total running time is spent on which subtask. Table 1 shows the values for MPTOPCOM-1 during the computation of all sub-regular triangulations of I^4 . For comparison the relative timings for TOPCOM are given in Table 2. All these numbers were determined with Valgrind’s tool `callgrind` [19]. In both cases the bulk of the time is spent on finding and processing the flips. For MPTOPCOM the major subtask is to determine the canonical representatives, while TOPCOM will explicitly compute the full orbits of each triangulation that it visits. For both programs the cost for finding the initial triangulation (e.g., in Algorithm 2 before Step 6) is negligible. The picture changes entirely if one restricts the algorithms to enumerate regular triangulations only. Then MPTOPCOM spends 95% of its time on solving linear programs, while that mark for TOPCOM reaches 98%. The remaining time is used in a similar fashion to that seen in Tables 1 and 2. For larger examples the overall pattern stays the same, but the most costly subtasks tend to take up even higher proportions of the total running times.

7.2 Products of two simplices

Another interesting class of point configurations are the products $\Delta_p \times \Delta_q$ of two simplices [7, §6.2]. The natural group action is by the product $\text{Sym}(p+1) \times \text{Sym}(q+1)$ of symmetric groups. In tropical geometry, e.g., their regular subdivisions control the combinatorial types of tropical polytopes [18, §5.2]. The special case where one of the factors is one-dimensional, i.e., when the product of simplices is a prism, is fully understood [7, §6.2.1]. Therefore, in our experiments we restrict our attention to cases with $2 \leq p \leq q$. There is a formula for the number of all triangulations of $\Delta_2 \times \Delta_q$ [7, 9.2.5], but this does not immediately yield the number of (semi-)regular triangulations or the number of orbits.

Figure 4 shows the speed for computing the triangulations in the regular component of $\Delta_2 \times \Delta_6$ with MPTOPCOM on the cluster, depending on the number of processes. This computation is medium size, i.e., a bit larger than the previous, and so it pays to use more processes. The timings are as follows: one hour and 22 minutes with MPTOPCOM (10 processes), five hours with the single-threaded version MPTOPCOM-1 and eight days and 16 hours with TOPCOM. A computation of this kind essentially marks the end of TOPCOM's range.

Our new results for $\Delta_2 \times \Delta_6$ and $\Delta_3 \times \Delta_4$ helped Schröter [28] to obtain new results on coarsest subdivisions of hypersimplices. An attempt to handle $\Delta_3 \times \Delta_5$ is currently under way (running on more than a hundred cores for some weeks). So far it has found more than 900 million orbits of sub-regular triangulations.

7.3 Dilated simplices

A third class of point configurations is denoted as $k \cdot \Delta_d$. These are the $\binom{n+d}{n} = \binom{n+d}{d}$ lattice points in the simplex Δ_d which is dilated by the factor k . For any polynomial in $d+1$ indeterminates which is homogeneous of degree k the monomials correspond to points in the point configuration $k \cdot \Delta_d$. In particular, the vertices of the Newton polytope form a subconfiguration. It follows that the tropical hypersurfaces in the tropical d -torus $\mathbb{T}^{d+1}/\mathbb{R}\mathbf{1}$ of homogeneous degree k are dual to regular subdivisions of $k \cdot \Delta_d$; see [18, §3.1]. The regular unimodular triangulations of $k \cdot \Delta_d$, which are necessarily full, correspond to those tropical hypersurfaces which are smooth. For the first time, we computed the full

Table 1: Percentages of total running time spent by MPTOPCOM-1 for computing all (sub-regular) triangulations of the 4-cube. The value for flip processing further refined in second column.

95%	Process flips
68%	Compute canonical representative
20%	Partition into up- and down-flips
7%	all remaining
3%	Check whether flip is edge of reverse search tree
2%	all remaining

triangulations of $3 \cdot \Delta_3$, and these classify the smooth tropical cubics in 3-space [18, §4.5].

Theorem 19. *There are exactly 21 125 102 orbits of regular and full triangulations of $3 \cdot \Delta_3$ with respect to the natural action of the symmetric group of degree four. Out of these, 14 373 645 are unimodular.*

This is the largest experiment that we completed so far. The computation took about four days on the Intel Xeon E5-2630 v2 cluster with 40 threads.

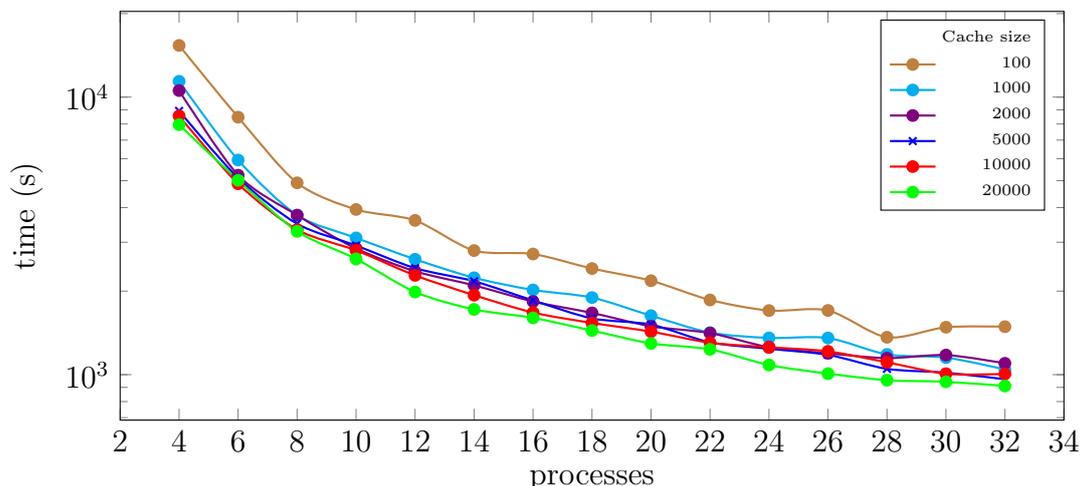


Figure 4: Timings for enumerating the triangulations with MPTOPCOM of $\Delta_2 \times \Delta_6$ in the regular component taken on the Intel Xeon E5-2630 v2 cluster with 64GB RAM per node, depending on the number of processes and cache sizes. Each data point is the result of averaging over ten runs with the same parameters.

7.4 Using more memory for caching

On the one hand, as a key benefit, the memory consumption of the down-flip reverse search algorithm allows for an excellent a priori estimate which is also quite low; cf. Theorem 4.

Table 2: Percentages of total running time spent by TOPCOM for computing all triangulations of the 4-cube. The values for the most expensive subtasks are refined in the second and third columns.

99%		Process flips
	92%	Check whether class was already found
		82%
		Enumerate orbits
		10%
		all remaining
	7%	all remaining
1%		all remaining

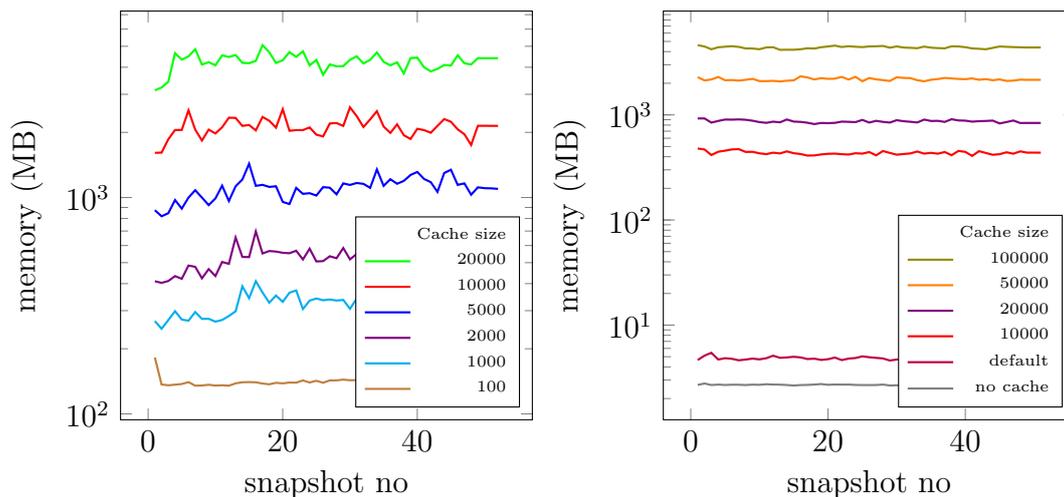


Figure 5: Memory usage determined with `Valgrind`. We took 52 `massif` snapshots for $\Delta_2 \times \Delta_6$ (left) and $3 \cdot \Delta_3$ (right).

On the other hand this approach results in a considerable amount of duplication. To avoid at least some of this it is natural to employ caching, as explained in Section 6. Here we want to report on some experiments concerning the impact of caching on the overall running time.

First we investigate the medium-size example $\Delta_2 \times \Delta_6$ from Section 7.2. It has 533 242 sub-regular triangulations up to symmetry. Figure 4 shows how the running-time depends on the number of workers and on the cache sizes of the three main caches introduced in Section 6. Their size can be varied individually, but for this plot we gave all of them the same size. As a default `MPTOPCOM` stores 2000 triangulations in each of these caches. In this case increasing the cache size from 100 to 2000 reduces the running time by about one third, independent of the number of threads. Increasing the cache to 20 000 only results in a further reduction by another 5%.

Figure 5 shows how the total amount of memory consumed depends on the cache sizes. The measurements have been taken by `Valgrind`'s tool `massif` which records memory snapshots in fixed time intervals [19]. The left hand side corresponds to the computation in Figure 4 for $\Delta_2 \times \Delta_6$. By and large the overall memory consumption depends linearly on the cache sizes. Note that the sizes of the triangulations as well as the sizes of the various objects for the cache values vary. So some fluctuations should be expected. Indeed, this is visible for $\Delta_2 \times \Delta_6$, which is not very large. In the much larger example $3 \cdot \Delta_3$ from Section 7.3 we see a similar behavior, but the fluctuations are nearly gone.

8 Concluding remarks

Table 3 also contains some empty rows, where we do not know the respective number of triangulations. Most of these will be out of reach for the current implementations, including `MPTOPCOM`. The reason for listing these nonetheless is to give a feel for the orders

Table 3: Summary of enumerations.

Points P	$n = \#P$	$d = \dim P$	$n - d$	$\#G$	#triangulation orbits (full) regular sub-regular	
I^3	8	3	5	48	6	6
I^4	16	4	12	384	235 277	247 451
I^5	32	5	27	3 840		
$\Delta_2 \times \Delta_2$	9	4	5	36	5	5
$\Delta_2 \times \Delta_3$	12	5	7	144	35	35
$\Delta_2 \times \Delta_4$	15	6	9	720	530	530
$\Delta_2 \times \Delta_5$	18	7	11	4 320	13 621	13 629
$\Delta_2 \times \Delta_6$	21	8	13	30 240	531 862	533 242
$\Delta_3 \times \Delta_3$	16	6	10	576	7 869	7 955
$\Delta_3 \times \Delta_4$	20	7	13	2 880	7 051 957	7 402 421
$\Delta_3 \times \Delta_5$	24	8	16	17 280		$> 9 \cdot 10^8$
$2\Delta_3$	10	3	7	24	15	59
$3\Delta_3$	20	3	17	24	21 125 102	925 148 763
$4\Delta_3$	35	3	32	24		

of magnitude involved. One main complexity parameter for enumerating triangulations is the difference $n - d$ of the number of points and the dimension. This is also one plus the dimension of the secondary fan, modulo linealities. Our experiments suggest that, as a very rough estimate, the range for TOPCOM seems to be limited by $n - d \approx 13$. This bar is raised substantially by MPTOPCOM to cover point configurations with $n - d = 17$ such as $3 \cdot \Delta_3$. It is an interesting question if MPTOPCOM can, e.g., deal with $I^3 \times \Delta_2$ where $n - d = 19$. That particular point configuration played a role in work of Orden and Santos [20] on efficient triangulations of cubes; see also [7, §6.3.3].

The empty rows of Table 3 show some cases which seem to be rather difficult challenges, with the current techniques. This includes the five-dimensional cube I^5 or the dilated simplex $4 \cdot \Delta_3$. Proving results about their triangulations might require clever strategies for random probing.

Another direction which looks promising is to investigate the triangulations of the cyclic polytopes. This is related to the higher Stasheff–Tamari orders which were introduced by Kaparanov and Voevodsky [17] and studied, e.g., by Edelman and Reiner [8]; see Rambau and Reiner [26] for a survey.

References

- [1] David Avis and Luc Devroye, *An analysis of budgeted parallel search on conditional Galton-Watson trees*, 2017, Preprint [arXiv:1703.10731](https://arxiv.org/abs/1703.10731).

- [2] David Avis and Komei Fukuda, *Reverse search for enumeration.*, Discrete Appl. Math. **65** (1996), no. 1-3, 21–46 (English).
- [3] David Avis and Charles Jordan, *A parallel framework for reverse search using mts*, 2016, Preprint [arXiv:1610.07735](https://arxiv.org/abs/1610.07735).
- [4] ———, *mplrs: A scalable parallel vertex/facet enumeration code*, Mathematical Programming Computation **10** (2018), no. 2, 267–302.
- [5] David Bremner, Mathieu Dutour Sikirić, and Achill Schürmann, *Polyhedral representation conversion up to symmetries*, Polyhedral computation, CRM Proc. Lecture Notes, vol. 48, Amer. Math. Soc., Providence, RI, 2009, pp. 45–71. MR 2503772
- [6] Adrian Brügger, Ambros Marzetta, Komei Fukuda, and Jurg Nievergelt, *The parallel search bench ZRAM and its applications*, Annals of Operations Research **90** (1999), 45–63.
- [7] Jesús A. De Loera, Jörg Rambau, and Francisco Santos, *Triangulations*, Algorithms and Computation in Mathematics, vol. 25, Springer-Verlag, Berlin, 2010, Structures for algorithms and applications. MR 2743368 (2011j:52037)
- [8] Paul H. Edelman and Victor Reiner, *The higher Stasheff-Tamari posets.*, Mathematika **43** (1996), no. 1, 127–154 (English).
- [9] J.-A. Ferrez, K. Fukuda, and Th. M. Liebling, *Solving the fixed rank convex quadratic maximization in binary variables by a parallel zonotope construction algorithm*, European J. Oper. Res. **166** (2005), no. 1, 35–50. MR 2128976
- [10] Komei Fukuda, *cddlib, version 0.94h*, http://www.inf.ethz.ch/personal/fukudak/cdd_home/, 2015.
- [11] Ewgenij Gawrilow and Michael Joswig, *polymake: a framework for analyzing convex polytopes*, Polytopes—combinatorics and computation (Oberwolfach, 1997), DMV Sem., vol. 29, Birkhäuser, Basel, 2000, pp. 43–73. MR MR1785292 (2001f:52033)
- [12] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky, *Discriminants, resultants and multidimensional determinants*, Modern Birkhäuser Classics, Birkhäuser Boston Inc., Boston, MA, 2008, Reprint of the 1994 edition. MR MR2394437 (2009a:14065)
- [13] Hiroshi Imai, Tomonari Masada, Fumihiko Takeuchi, and Keiko Imai, *Enumerating triangulations in general dimensions*, Internat. J. Comput. Geom. Appl. **12** (2002), no. 6, 455–480. MR 1945594
- [14] Anders N. Jensen, *An implementation of exact mixed volume computation*, Mathematical Software – ICMS 2016: 5th International Conference, Berlin, Germany, July 11-14, 2016, Proceedings (Gert-Martin Greuel, Thorsten Koch, Peter Paule, and Andrew Sommese, eds.), Springer International Publishing, Cham, 2016, pp. 198–205.
- [15] ———, *Tropical homotopy continuation*, 2016, Preprint [arXiv:1601.02818](https://arxiv.org/abs/1601.02818).
- [16] ———, *Gfan, a software system for Gröbner fans and tropical varieties, version 0.6*, Available at <http://home.imf.au.dk/jensen/software/gfan/gfan.html>, 2017.
- [17] Mikhail M. Kapranov and Vladimir A. Voevodsky, *Combinatorial-geometric aspects of polycategory theory: pasting schemes and higher Bruhat orders (list of results)*,

- Cahiers Topologie Géom. Différentielle Catég. **32** (1991), no. 1, 11–27, International Category Theory Meeting (Bangor, 1989 and Cambridge, 1990). MR 1130400
- [18] Diane Maclagan and Bernd Sturmfels, *Introduction to tropical geometry*, Graduate Studies in Mathematics, vol. 161, American Mathematical Society, Providence, RI, 2015. MR 3287221
- [19] Nicholas Nethercote and Julian Seward, *Valgrind: A framework for heavyweight dynamic binary instrumentation*, Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation (New York, NY, USA), PLDI '07, ACM, 2007, pp. 89–100.
- [20] David Orden and Francisco Santos, *Asymptotically efficient triangulations of the d -cube*, Discrete Comput. Geom. **30** (2003), no. 4, 509–528. MR 2013970
- [21] Julian Pfeifle, *Secondary polytope of the “mother of all examples”*, Electronic Geometry Models (2000), No. 2000.09.033.
- [22] Julian Pfeifle and Jörg Rambau, *Computing triangulations using oriented matroids*, Algebra, geometry, and software systems, Springer, Berlin, 2003, pp. 49–75. MR 2011753 (2004i:68233)
- [23] Lionel Pournin, *The flip-graph of the 4-dimensional cube is connected*, Discrete Comput. Geom. **49** (2013), no. 3, 511–530. MR 3038527
- [24] Lionel Pournin and Thomas M. Liebling, *Constrained paths in the flip-graph of regular triangulations*, Comput. Geom. **37** (2007), no. 2, 134–140. MR 2310598
- [25] Jörg Rambau, *TOPCOM: triangulations of point configurations and oriented matroids*, Mathematical software (Beijing, 2002), World Sci. Publ., River Edge, NJ, 2002, pp. 330–340. MR 1932619
- [26] Jörg Rambau and Victor Reiner, *A survey of the higher Stasheff-Tamari orders.*, Associahedra, Tamari lattices and related structures. Tamari memorial Festschrift, Basel: Birkhäuser, 2012, pp. 351–390 (English).
- [27] Thomas Rehn and Achill Schürmann, *C++ tools for exploiting polyhedral symmetries.*, Mathematical software – ICMS 2010. Third international congress on mathematical software, Kobe, Japan, September 13–17, 2010. Proceedings, Berlin: Springer, 2010, pp. 295–298 (English).
- [28] Benjamin Schröter, *Multi-splits and tropical linear spaces from nested matroids*, 2017, Preprint [arXiv:1707.02814](https://arxiv.org/abs/1707.02814).
- [29] Ákos Seress, *Permutation group algorithms*, Cambridge Tracts in Mathematics, vol. 152, Cambridge University Press, Cambridge, 2003. MR 1970241
- [30] Marc Snir, Steve Otto, Steven Huss-Lederman, David Walker, and Jack Dongarra, *MPI– the complete reference, vol 1: The MPI core*, 2nd ed., MIT Press, 1998.
- [31] Christophe Weibel, *Implementation and parallelization of a reverse-search algorithm for Minkowski sums*, 2010 Proceedings of the Twelfth Workshop on Algorithm Engineering and Experiments (ALENEX), 2010, pp. 34–42.