THE POWER OF PYRAMID DECOMPOSITION IN NORMALIZ

WINFRIED BRUNS, BOGDAN ICHIM, AND CHRISTOF SÖGER

ABSTRACT. We describe the use of pyramid decomposition in Normaliz, a software tool for the computation of Hilbert bases and enumerative data of rational cones and affine monoids. Pyramid decomposition in connection with efficient parallelization and streamlined evaluation of simplicial cones has enabled Normaliz to process triangulations of size $\approx 5 \cdot 10^{11}$ that arise in the computation of Hilbert series related to combinatorial voting theory.

1. INTRODUCTION

Normaliz is a software tool for the computation of Hilbert bases and enumerative data of rational cones and affine monoids. In the 14 years of its existence it has found numerous applications; for example, see Bogart, Raymond and Thomas [3], Craw, Maclagan and Thomas [10], Kappl, Ratz and Staudt [17] or Sturmfels and Welker [22]. Normaliz is used in polymake [16] and Regina [9].

The mathematics of the very first version was described in [8], and the ideas leading to version 2.2 (2009) are contained in [6]. In this article we want to document the most recent development[1] that has extended the scope of Normaliz by several orders of magnitude.

From the very beginning Normaliz has used lexicographic triangulations. (It also contains a triangulation free Hilbert basis algorithm; see [6]). Lexicographic triangulations are essentially characterized by being incremental in the following sense. Suppose that the cone C is generated by vectors $x_1, \ldots, x_n \in \mathbb{R}^d$ and set $C_i = \mathbb{R}_+ x_1 + \cdots + \mathbb{R}_+ x_i$, $i = 0, \ldots, n$. Then the lexicographic triangulation Λ (for the ordered system $x_1, \ldots, x_n$) restricts to a triangulation of C_i for $i = 0, \ldots, n$. Lexicographic triangulations are easy to compute, and go very well with Fourier-Motzkin elimination that computes the support hyperplanes of C by successive extension from C_i to C_{i+1}, $i = 0, \ldots, n-1$. The triangulation Λ_i of C_i is extended to C_{i+1} by all simplicial cones $F + \mathbb{R}_+ x_{i+1}$ where $F \in \Lambda_i$ is visible from x_{i+1}.

As simple as the computation of the lexicographic triangulation is, the algorithm in the naive form just described has two related drawbacks: (i) one must store Λ_i and this becomes very difficult for sizes $\geq 10^8$; (ii) in order to find the facets F that are visible from x_{i+1} we must match the simplicial cones in Λ_i with the support hyperplanes of C_i that are visible from x_{i+1}. While (i) is a pure memory problem, (ii) quickly leads to impossible computation times.

Pyramid decomposition is the basic idea that has enabled Normaliz to compute dimension 24 triangulations of size $\approx 5 \cdot 10^{11}$ in acceptable time on standard multiprocessor

---

[1]Version 2.8 has been uploaded to `http://www.math.uos.de/normaliz` on June 29, 2012.

systems such as SUN xFire 4450 or Dell PowerEdge R910. Instead of going for the lex-
icographic triangulation directly, we first decompose C into the pyramids generated by
x_{i+1} and the facets of C_i that are visible from x_{i+1}, $i = 0, \ldots, n-1$. These pyramids (of
level 0) are then decomposed into pyramids of level 1 etc. While the level 0 decomposi-
tion need not be a polyhedral subdivision in the strict sense, pyramid decomposition stops
after finitely many iterations at the lexicographic triangulation (see Section 2).

Pure pyramid decomposition is extremely memory friendly, but its computation times
are even more forbidding than those of pure lexicographic triangulation since too many
Fourier-Motzkin eliminations become necessary, and almost all of them are inevitably
wasted. That Normaliz can nevertheless cope with extremely large triangulations relies
on a well balanced combination of both strategies that we outline in Section 3.

In Section 4 we describe the steps by which Normaliz evaluates the simplicial cones in
the triangulation for the computation of Hilbert bases, volumes and Hilbert series. The
evaluation almost always takes significantly more time than the triangulation. Therefore
it must be streamlined as much as possible. For the Hilbert series Normaliz uses a Stanley
decomposition [21]. That it can be found efficiently relies crucially on an idea of Köppe
and Verdoolaege [18].

We document the scope of Normaliz' computations in Section 5. Our main examples
come from combinatorial voting theory that we found in Schürmann's paper [20]. The de-
sire to master the Hilbert series computations asked for in [20] was an important stimulus
in the recent development of Normaliz.

For Hilbert basis computations pyramid decomposition has a further and sometimes
tremendous advantage: one can avoid the triangulation of those pyramids for which it is
a priori clear that they will not supply new candidates for the Hilbert basis. This observa-
tion, on which the contribution of the authors to [5] is based, triggered the use of pyramid
decomposition as a general principle. See Remark 8 for a brief discussion and Section 5.5
for data of computations.

It is an important aspect of pyramid decomposition that it is very parallelization friendly
since the pyramids can be treated independently of each other. Normaliz uses OpenMP
for shared memory systems. Needless to say that triangulations of the size mentioned
above can hardly be reached in serial computation.

2. LEXICOGRAPHIC TRIANGULATION AND PYRAMID DECOMPOSITION

Consider vectors $x_1, \ldots, x_n \in \mathbb{R}^d$. For Normaliz these must be integral vectors, but
integrality is irrelevant in this section. We want to compute a triangulation of the cone

$$C = \text{cone}(x_1, \ldots, x_n) = \mathbb{R}_+ x_1 + \cdots + \mathbb{R}_+ x_n$$

with rays through $x_1, \ldots, x_n$. Such a triangulation is a polyhedral subdivision of C into
simplicial subcones σ generated by linearly independent subsets of $\{x_1, \ldots, x_n\}$.

For a triangulation Σ of a cone C and a subcone C' we set

$$\Sigma|C' = \{\sigma \cap C' : \sigma \in \Sigma, \ \dim \sigma \cap C' = \dim C'\}.$$

In general $\Sigma|C'$ need not be a triangulation of C', but it is so if C' is a face of C.

The *lexicographic* (or *placing*) triangulation $\Lambda(x_1, \ldots, x_n)$ of $\text{cone}(x_1, \ldots, x_n)$ can be
defined recursively as follows: (i) the triangulation of the zero cone is the trivial one,

(ii) $\Lambda(x_1, \ldots, x_n)$ is given by

$$\Lambda(x_1, \ldots, x_n) = \Lambda(x_1, \ldots, x_{n-1}) \cup \{ \mathrm{cone}(\sigma, x_n) : \sigma \in \Lambda(x_1, \ldots, x_{n-1}) \text{ visible from } x_n \}$$

where σ is *visible* from x_n if $x_n \notin \mathrm{cone}(x_1, \ldots, x_{n-1})$ and the line segment $[x_n, y]$ for every point y of σ intersects $\mathrm{cone}(x_1, \ldots, x_{n-1})$ only in y. Note that a polyhedral complex is always closed under the passage to faces, and the definition above takes care of it. In the algorithms below a polyhedral subdivision can always be represented by its maximal faces which for convex full dimensional polyhedra are the full dimensional cones in the subdivision. For simplicial subdivisions of cones one uses of course that the face structure is completely determined by set theory: every subset E of the set of generators spans a conical face of dimension $|E|$.

We state some useful properties of lexicographic triangulations:

Proposition 1. *With the notation introduced, let $C_i = \mathrm{cone}(x_1, \ldots, x_i)$ and $\Lambda_i = \Lambda(x_1, \ldots, x_i)$ for $i = 1, \ldots, n$.*

(1) *Λ_n is the unique triangulation of C with rays through a subset of $\{x_1, \ldots, x_n\}$ that restricts to a triangulation of C_i for $i = 1, \ldots, n$ and $\Lambda | C_i$ has rays through a subset of $\{x_1, \ldots, x_i\}$.*

(2) *For every face F of C the restriction $\Lambda | F$ is the lexicographic triangulation $\Lambda(x_{i_1}, \ldots, x_{i_m})$ where $\{x_{i_1}, \ldots, x_{i_m}\} = F \cap \{x_1, \ldots, x_n\}$ and $i_1 < \cdots < i_m$.*

(3) *If $\dim C_i > \dim C_{i-1}$, then $\Lambda = \Lambda(x_1, \ldots, x_{i-2}, x_i, x_{i-1}, x_{i+1}, \ldots, x_n)$.*

(4) *$\Lambda = \Lambda(x_{i_1}, \ldots, x_{i_d}, x_{j_1}, \ldots, x_{j_{n-d}})$ where $(i_1, \ldots, i_d)$ is the lexicographic smallest index vector of a rank d subset of $\{x_1, \ldots, x_n\}$ and $j_1 < \cdots < j_{n-d}$ lists the complementary indices.*

Proof. (1) By construction it is clear that Λ_n satisfies the properties of which we claim that they determine Λ uniquely. On the other hand, the extension of Λ_{i-1} to a triangulation of C_i is uniquely determined if one does not introduce further rays: the triangulation of the part V of the boundary of C_{i-1} that is visible from x_i has to coincide with the restriction of Λ_{i-1} to V.

(2) One easily checks that $\Lambda | F$ satisfies the conditions in (1) that characterize $\Lambda(x_{i_1}, \ldots, x_{i_m})$.

(3) It is enough to check the claim for $i = n$. Then the only critical point for the conditions in (1) is whether $\Lambda(x_1, \ldots, x_{n-2}, x_n, x_{n-1})$ restricts to C_{n-1}. But this is the case since C_{n-1} is a facet of C if $\dim C > \dim C_{n-1}$.

(4) follows by repeated application of (3). $\qquad\square$

In the following we will assume that C is full dimensional: $\dim C = d = \dim \mathbb{R}^d$. Part (4) helps us to keep the data structure of lexicographic triangulations simple: right from the start we need only to work with the list of dimension d simplicial cones of Λ by searching $x_{i_1}, \ldots, x_{i_d}$ first, choosing $\mathrm{cone}(x_{i_1}, \ldots, x_{i_d})$ as the first d-dimensional simplicial cone and subsequently extending the list as prescribed by the definition of the lexicographic triangulation. In other words, we can assume that $x_1, \ldots, x_d$ are linearly independent, and henceforth we will do so.

In order to extend the triangulation we must of course know which facets of C_{i-1} are visible from x_i. Recall that a cone C of dimension d in $\mathbb{R}^d$ has a unique irredundant

representation as an intersection of linear halfspaces:

$$C = \bigcap_{H \in \mathscr{H}(C)} H^+,$$

where $\mathscr{H}(C)$ is a finite set of hyperplanes and the orientation of the closed half spaces H^- and H^+ is chosen in such a way that $C \subset H^+$ for $H \in \mathscr{H}(C)$. For $H \in \mathscr{H}(C_{i-1})$ the facet $H \cap C_{i-1}$ is visible from x_i if and only if x_i lies in the open halfspace $H^< = H^- \setminus H$. When we refer to support hyperplane in the following we always mean those that appear in the irredundant decomposition of C since only they are important in the algorithmic context.

Hyperplanes are represented by linear forms $\lambda \in (\mathbb{R}^d)^*$, and we always work with the basis $e_1^*, \ldots, e_d^*$ that is dual to the basis $e_1, \ldots, e_d$ of unit vectors. For rational hyperplanes the linear form λ can always be chosen in such a way that it has integral coprime coefficients and satisfies $\lambda(x) \geq 0$ for $x \in C$. This choice determines λ uniquely. (If one identifies $e_1^*, \ldots, e_d^*$ with $e_1, \ldots, e_d$ via the standard scalar product, then λ is nothing but the primitive integral inner (with respect to C) normal vector of H.) For later use we define the *(lattice) height* of $x \in \mathbb{R}^d$ over H by

$$\mathrm{ht}_H(x) = |\lambda(x)|.$$

If $F = C \cap H$ is the facet of C cut out by H, we set $\mathrm{ht}_F(x) = \mathrm{ht}_H(x)$.

We can now describe the computation of the triangulation $\Lambda(x_1, \ldots, x_n)$ in a more formal way in Table 1. For simplicity we will identify a simplicial cone σ with its generating set $\subset \{x_1, \ldots, x_n\}$. It should be clear from the context what is meant. For further use we introduce the notation

$$\mathscr{H}^*(C, x) = \{H \in \mathscr{H}(C), x \in H^*\} \qquad \text{where } * \in \{<, >, +, -\}.$$

Table 1 formalizes the computation of the lexicographic triangulation.

LEXTRIANGULATION$(x_1, \ldots, x_n)$
1   ADDSIMPLEX$(x_1, \ldots, x_d)$
2   **for** $i \leftarrow d + 1$ **to** n
3   **do**
4   EXTENDTRI(i)
5   FINDNEWHYP(i)

EXTENDTRI(i)
1   **parallel for** $H \in \mathscr{H}^<(C_{i-1}, x_i)$
2   **do**
3   **for** $\sigma \in \Lambda(x_1, \ldots, x_{i-1})$
4   **do**
5   **if** $|\sigma \cap H| = d - 1$
6   **then** ADDSIMPLEX$(x_i \cup (\sigma \cap H))$

TABLE 1. Incremental building of lexicographic triangulation

The function ADDSIMPLEX adds a simplicial cone to the (initially empty) list of simplicial cones that, upon completion, contains the lexicographic triangulation of C. The function FINDNEWHYP computes $\mathscr{H}(C_i)$ from $\mathscr{H}(C_{i-1})$ by Fourier-Motzkin elimination. (It does nothing if $x_i \in C_{i-1}$.) Its Normaliz implementation has been described in great detail in [6]; therefore we skip it here. The function EXTENDTRI does exactly what its name says: it extends the triangulation $\Lambda(x_1, \ldots, x_{i-1})$ of C_{i-1} to the triangulation $\Lambda(x_1, \ldots, x_i)$ of C_i (again doing nothing if $x_i \in C_{i-1}$).

Note that the set of hyperplanes over which the loop in EXTENDTRI runs is given by $\mathscr{H}^<(C_{i-1}, x_i)$.

One is tempted to improve EXTENDTRI by better bookkeeping and using extra information on triangulations of cones. We discuss our more or less fruitless attempts in the following remark.

Remark 2. (a) If one knows the restriction of $\Lambda(x_1, \ldots, x_{i-1})$ to the facets of C_{i-1}, then $\Lambda(x_1, \ldots, x_i)$ can be computed very fast. However, unless $i = n$, the facet triangulation must now be extended to the facets of C_i, and this step eats up the previous gain, as experiments have shown, at least for the relatively small triangulations to which EXTENDTRI is really applied after the pyramid decomposition described below.

(b) The test of the condition $|\sigma \cap H| = d - 1$ is satisfied if and only if $d - 1$ of the generators of σ lie in H. Its verification can be accelerated if one knows which facets of the d-dimensional cones in $\Lambda(x_1, \ldots, x_{i-1})$ are already shared by another simplicial cone in $\Lambda(x_1, \ldots, x_{i-1})$, and are therefore not available for the formation of a new simplicial cone. But the extra bookkeeping requires more time than is gained by its use.

(c) One refinement is used in our implementation, though its influence is almost unmeasurable. Each simplicial cone in $\Lambda(x_1, \ldots, x_{i-1})$ has been added with a certain generator x_j, $j < i$. (The first cone is considered to be added with each of its generators.) It is not hard to see that only those simplicial cones that have been added with a generator $x_j \in H$ can satisfy the condition $|\sigma \cap H| = d - 1$, and this information is used to reduce the number of pairs (H, σ) to be tested.

(d) If $|H \cap \{x_1, \ldots, x_{i-1}\}| = d - 1$, then $H \in \mathscr{H}^<(C_{i-1}, x_i)$ produces exactly one new simplicial cone of dimension d, namely $\mathrm{cone}(x_i, H \cap \{x_1, \ldots, x_{i-1}\})$, and therefore the loop over σ can be suppressed.

The product $|\mathscr{H}^<(C_{i-1}, x_i)| \cdot |\Sigma|$ determines the complexity of EXTENDTRI. Even though the loop over H is parallelized (as indicated by **parallel for**), the time spent in EXTENDTRI can be extremely large. (The "exterior" loops in FINDNEWHYP are parallelized as well.) The second limiting factor for EXTENDTRI is memory: it is already critical to store triangulations of size 10^8 and impossible for size $\geq 10^9$. Therefore the direct approach to lexicographic triangulations does not work for truly large cones.

Now we present a radically different way to lexicographic triangulations via iterated *pyramid decompositions*. The cones that appear in this type of decomposition are called *pyramids* since their cross-section polytopes are pyramids in the usual sense, namely of type $\mathrm{conv}(F, x)$ where F is a facet and x is a vertex not contained in F.

Definition 3. The *pyramid decomposition* $\Pi(x_1, \ldots, x_n)$ of $C = \mathrm{cone}(x_1, \ldots, x_n)$ is recursively defined as follows: it is the trivial decomposition for $n = 0$, and

$$\Pi(x_1, \ldots, x_n) = \Pi(x_1, \ldots, x_{n-1}) \cup \{\mathrm{cone}(F, x_n) :$$
$$F \text{ a face of } C(x_1, \ldots, x_{n-1}) \text{ visible from } x_n\}.$$

Note that the pyramid decomposition is not a polyhedral subdivision in the strong sense: the intersection of two faces F and F' need not be a common face of F and F' (but is always a face of F or F'). See Figure 1 for an example.
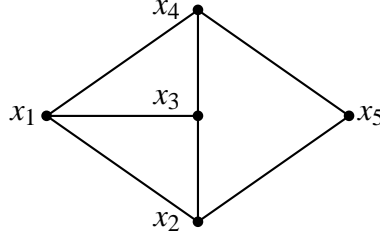
FIGURE 1. Cross-section of a pyramid decomposition

In order to iterate the pyramid decomposition we set $\Pi^0(x_1,\ldots,x_n) = \Pi(x_1,\ldots,x_n)$, and

$$\Pi^k(x_1,\ldots,x_n) = \bigcup_{P\in\Pi^{k-1}(x_1,\ldots,x_n)} \{\Pi(x_i : x_i \in P)\} \qquad \text{for } k > 0.$$

Note that this recursion cannot descend indefinitely, since the number of generators goes down in each recursion level. We denote the *total pyramid decomposition* by $\Pi^\infty(x_1,\ldots,x_n)$. More precisely:

Proposition 4. *One has $\Pi^\infty(x_1,\ldots,x_n) = \Pi^{n-d}(x_1,\ldots,x_n) = \Lambda(x_1,\ldots,x_n)$.*

Proof. In the case $n = d$, the pyramid decomposition is obviously the face lattice of C, and therefore coincides with the lexicographic triangulation. For $n > d$ the first full dimensional pyramid reached is the simplicial cone $\operatorname{cone}(x_1,\ldots,x_d)$. All the other pyramids have at most $n-1$ generators, and so we can use induction: For each $P \in \Pi(x_1,\ldots,x_n)$ the total pyramid decomposition of P is the lexicographic triangulation $\Lambda(x_i : x_i \in P)$. According to Proposition 1(2) these triangulations match along the common boundaries of the pyramids, and therefore constitute a triangulation of C. It evidently satisfies the conditions in Proposition 1(1).   □

This leads to a recursive computation of $\Lambda(x_1,\ldots,x_n)$ by the algorithms in Table 2. The

TOTALPYRDEC$(x_1,\ldots,x_n)$	PROCESSPYRSREC(i)
1 ADDSIMPLEX$(x_1,\ldots,x_d)$	1 **for** $H \in \mathscr{H}^<(C_{i-1},x_i)$
2 **for** $i \leftarrow d+1$ **to** n	2 **do** $key \leftarrow \{x_i\} \cup (\{x_1,\ldots,x_{i-1}\} \cap H)$
3 **do**	3 TOTALPYRDEC(key)
4 PROCESSPYRSREC(i)	

TABLE 2. Total pyramid decomposition

first realizes the building of $\Pi(x_1,\ldots,x_n)$ (represented by its full dimensional members) and the second takes care of the recursion that defines $\Pi^\infty(x_1,\ldots,x_n)$.

Pyramid decomposition has the virtue of requiring very little memory since one needs not store the triangulation in order to produce all the simplicial cones in it. However, there is a severe drawback: as above, one must compute the support hyperplanes in $\mathscr{H}(P)$ for all pyramids encountered. In a "pure" approach, one computes the support hyperplanes of the simplicial cones at the bottom of the pyramid decomposition; this is essentially the

inversion of the matrix of its generators (see equation (4.1)). Then one passes them back from a pyramid to its "mother", discarding those that fail to have all generators of the "mother" in its positive halfspace or have been found previously. These two conditions are easily tested. Suppose P is the pyramid to which TOTALPYRDEC is applied in PRO-CESSPYRSREC and $G \in \mathcal{H}(P)$. Then $G \in \mathcal{H}(x_1 \ldots, x_i) \setminus \mathcal{H}(x_1, \ldots, x_{i-1})$ if and only if the following two conditions are satisfied:

   (i) $x_j \in G^+$ for $j = 1, \ldots, i-1$;
   (ii) $x_j \in G^>$ for all $j = 1, \ldots, i-1$ such that $x_j \notin P$.

One should note that pyramids effectively reduce the dimension: the complexity of cone(F, x_n) is completely determined by the facet F, which has dimension $d - 1$.

While being extremely memory efficient, total pyramid decomposition is usually much slower than building the lexicographic triangulation directly. For one of our standard test examples ($4 \times 4 \times 3$ contingency tables, dimension 30 with 48 extreme rays; see [5]) the lexicographic triangulation with respect to the order of generators in the input file has $2,654,272$ full dimensional cones. In serial computation on an Intel i7 2600 PC, LEXTRI-ANGULATION computes it in approximately 2 minutes, whereas TOTALPYRDEC needs about 11 minutes. The current implementation, described in the next section, reduces the serial computation time to 13 seconds.

Remark 5. Pyramid decomposition is not only extremely useful for the computation of triangulations, but also helps in finding support hyperplanes. For them the critical complexity parameter is $|\mathcal{H}^<(C_{i-1}, x_i)| \cdot |\mathcal{H}^>(C_{i-1}, x_i)|$, and as in its use for triangulation, pyramid decomposition lets us replace a very large product of the sizes of two "global" lists by a sum of small "local" products–the price to be paid is the computational waste invested for the support hyperplanes of the pyramids that are forgotten later on. Nevertheless pyramid decomposition leads to a substantial reduction in computing time also for support hyperplanes, and Normaliz uses this effect. We illustrate this by computation times for the $5 \times 5 \times 3$ contingency tables (dimension 55 with 75 extreme rays; see [5]). The cone has $306,955$ support hyperplanes. On a Sun xFire 4450 we measured a serial computation time of $16,822$ seconds if only FINDNEWHYP is used. The current implementation reduces this to $4,334$ seconds.

3. THE CURRENT IMPLEMENTATION

Since version 2.7 (and partly since 2.5) Normaliz has combined lexicographic triangulation with pyramid decomposition. The support hyperplanes and the triangulation are extended from one generator to the next only until certain bounds are exceeded. From that point on, the algorithm BUILDCONE described in Table 3 switches to pyramid decomposition, and the same mixed strategy is then applied to the pyramids.

We now use two types of passage to pyramids, a recursive one via PROCESSPYRSREC and a nonrecursive one via PROCESSPYRS. The main reason for this split approach is that on the one hand recursion limits the effect of parallelization (as it does in Normaliz 2.7), and, on the other hand, the recursive approach nevertheless saves time in the computation of support hyperplanes for the top cone.

BUILDCONE$(x_1,\ldots,x_n;recursion)$ PROCESSPYRSREC(i)
 1   ADDSIMPLEX$(x_1,\ldots,x_d)$ 1   **for** $H \in \mathcal{H}^<(C_{i-1},x_i)$
 2   **for** $i \leftarrow d+1$ **to** n 2   **do** $key \leftarrow \{x_i\} \cup (\{x_1,\ldots,x_{i-1}\} \cap H)$
 3   **do**   3   BUILDCONE$(key,true)$
 4   **if** *MakePyramidsSupp* & *recursion*
 5   **then** PROCESSPYRSREC(i) PROCESSPYRS$(i,level)$
 6   **else** **if** *MakePyramidsTri*   1   **for** $H \in \mathcal{H}^<(C_{i-1},x_i)$
 7   **then** PROCESSPYRS$(i,level)$ 2   **do** $key \leftarrow \{x_i\} \cup (\{x_1,\ldots,x_{i-1}\} \cap H)$
 8   **else** EXTENDTRI(i) 3   STOREPYR$(key,level+1)$
 9   FINDNEWHYP(i)
10   **if** *TopCone*
11   **then** EVALUATEPYRS(0)

TABLE 3. Combining lexicographic triangulation and pyramid decomposition

The boolean *recursion* indicates whether the recursive passage to pyramids is allowed. For the top cone BUILDCONE is called with *recursion = true*. The boolean *MakePyramidsSupp* combines two conditions:

(1) while set to *false* initially, it remains *true* once the branch PROCESSPYRSREC has been taken the first time;
(2) it is set *true* if the complexity parameter $|\mathcal{H}^<(C_{i-1},x_i)| \cdot |\mathcal{H}^>(C_{i-1},x_i)|$ exceeds a threshold.

In the nonrecursive passage to pyramids we cut the umbilical cord between a pyramid and its mother and just store the pyramid for later evaluation. The nonrecursive call is controlled by the boolean *MakePyramidsTri* that combines three conditions:

(1) while set to *false* initially, it remains *true* once the branch PROCESSPYRS has been taken the first time;
(2) it is set *true* if the complexity parameter $|\mathcal{H}^<(C_{i-1},x_i)| \cdot |\Sigma|$ exceeds a threshold;
(3) it is set *true* if the memory protection threshold is exceeded;

The last point needs to be explained. BUILDCONE is not only called for the processing of the top cone C, but also for the parallelized processing of the stored pyramids. Since each of the "parallel" pyramids produces simplicial cones, the buffer in which the simplicial cones are collected for evaluation, may be severely overrun without condition (3), especially if $|\mathcal{H}^<(C_{i-1},x_i)|$ is small, and therefore condition (2) is reached only for large $|\Lambda(x_1,\ldots,x_{i-1})|$. The variable *level* indicates the generation of the pyramid; for the top cone it has value -1, and increases by 1 with each new generation.

At the end of BUILDCONE for the top cone C we start the evaluation of the stored pyramids as described in Table 4.

Remark 6. (a) For efficiency Normaliz completely avoids nested parallelization. Therefore the parallelization in FINDNEWHYP and EXTENDTRI is switched off when the parallelization in EVALUATEPYRS is active. On the other hand, these are active when the top cone or recursively built pyramids are being processed.

EVALUATEPYRS($level$)
1   **if** $PyramidList[level] = \emptyset$
2   **then return**
3   **parallel for** $P \in PyramidList[level]$
4   **do**
5   BUILDCONE($P, false$)
6   EVALUATEPYRS($level + 1$)

TABLE 4. Evaluation of pyramids

(b) Despite of considerable efforts we have not found a completely satisfactory solution in which pyramids could always be processed recursively and simultaneously in parallel. Because of (a) we can only parallelize the pyramids directly produced from the top cone in the recursive approach, and then parallelization may be limited by an insufficient number of hyperplanes in $\mathscr{H}^<(C_{i-1}, x_i)$ or, more often, by enormous differences in the sizes of the pyramids, so that one of them may be running solo for a long time—recognizing the size in advance has turned out difficult. Parallelization in FINDNEWHYP and EXTENDTRI is then the better solution.

Moreover, a large pyramid together with its children may produce a huge number of simplicial cones and overrun the evaluation buffer. Serial loops can be interrupted at any time, and therefore the memory problem cannot arise.

(c) As soon as BUILDCONE switches to pyramids, the triangulation $\Lambda(x_1, \ldots, x_{i-1})$ is no longer needed for further extension. Therefore it is shipped to the evaluation buffer. The buffer is emptied whenever it has exceeded its preset size and program flow allows its parallelized evaluation. (Because of (a) this is not always possible.)

(d) The strategy for the evaluation of pyramids is similar. If the buffer for $level + 1$ is exceeded, evaluation on that level will be started as soon as possible. Usually this results in a tree of evaluations over several levels.

We add a few minor details of the implementation.

Remark 7. (a) For nonrecursive pyramids the support hyperplanes arising from the last generator need not be computed since they are irrelevant for triangulation and pyramid decomposition.

(b) Simplicial facets of C_{i-1} produce exactly one simplicial pyramid in C_i. They are treated directly by ADDSIMPLEX.

(c) If the extreme rays of C have been singled out from the given generators $x_1, \ldots, x_n$ before BUILDCONE is called, then only the extreme rays are used in the pyramid decomposition and the lexicographic triangulation.

(d) If a grading is defined explicitly (see Section 4), then Normaliz orders the generators by degree and those of the same degree by input order before building the cone C. This is an attempt to cover as much ground as possible by using generators of small degree. On the whole, we have reached good results with this choice.

Remark 8. *(Partial triangulation)* The idea of pyramid decomposition was born when the authors observed that the computation of Hilbert bases usually does not need a full

triangulation of C. If a simplicial cone σ cannot contribute new candidates for the Hilbert basis of C, it need not be evaluated, and if a pyramid consists only of such simplicial cones, it need not be triangulated at all. This is the case if $\text{ht}_H(x_i) = 1$ in PROCESSPYRS.

The resulting strategy has sometimes striking results and was already described in [5]. It is mentioned here only for completeness. If a full triangulation is not required, then PROCESSPYRS discards all pyramids of height 1 from further processing. (However, their support hyperplanes must be computed if processed recursively.) If followed strictly, the recursion will not stop before the simplicial cones at the bottom of the pyramid decomposition. As for full triangulations, this is usually not optimal. Normaliz therefore switches to EXTENDTRI for pyramids of height ≥ 2 from a certain level on.

4. EVALUATION OF SIMPLICIAL CONES

The fast computation of triangulations via pyramid decomposition must be accompanied by an efficient evaluation of the simplicial cones in the triangulation, which is almost always the more time consuming step.

Let σ be a simplicial cone generated by the linearly independent vectors $v_1, \ldots, v_d$. The evaluation is based on the *generator matrix G_σ* whose *rows* are $v_1, \ldots, v_d$. Before we outline the evaluation procedure, let us substantiate the remark made in Section 2 that finding the support hyperplanes amounts to the inversion of G_σ. Let H_i be the support hyperplane of σ opposite to v_i, given by the linear form $\lambda_i = a_{1i}e_1^* + \cdots + a_{di}e_d^*$ with coprime integer coefficients a_j. Then

$$(4.1) \qquad \lambda_i(v_k) = \sum_{j=1}^d v_{kj}a_{ji} = \begin{cases} \text{ht}_{H_i}(v_i), & k = i, \\ 0, & k \neq i. \end{cases}$$

Thus the matrix (a_{ij}) is G_σ^{-1} up to scaling of its columns. Usually the inverse is computed only for the first simplicial cone in every pyramid since its support hyperplanes are really needed. But matrix inversion is rather expensive, and Normaliz goes to great pains to avoid it.

Normaliz computes sets of vectors, primarily Hilbert bases, but also measures, for example the volumes of rational polytopes. A polytope P arises from a cone C by cutting C with a hyperplane, and for Normaliz such hyperplanes are defined by gradings: a *grading* is a linear form $\deg : \mathbb{Z}^d \to \mathbb{Z}$ (extended naturally to $\mathbb{R}^d$) with the following properties: (i) $\deg(x) > 0$ for all $x \in C$, $x \neq 0$, and (ii) $\deg(\mathbb{Z}^d) = \mathbb{Z}$. The first condition guarantees that the intersection $P = C \cap A_1$ for the affine hyperplane

$$A_1 = \{x \in \mathbb{R}^d : \deg(x) = 1\}$$

is compact, and therefore a rational polytope. The second condition is harmless for integral linear forms since it can be achieved by extracting the greatest common divisor of the coefficients of deg with respect to the dual basis.

The grading deg can be specified explicitly by the user or chosen implicitly by Normaliz. The implicit choice makes only sense if there is a natural grading, namely one under which the extreme integral generators of C all have the same degree. (If it exists, it is of course uniquely determined.)

At present, Normaliz evaluates the simplicial cones σ in the triangulation of C for the computation of the following data:

(HB) the Hilbert basis of C,
(LP) the lattice points in the rational polytope $P = C \cap A_1$,
(Vol) the normalized volume $\mathrm{vol}(P)$ of the rational polytope P (also called the *multiplicity* of C),
(Ehr) the *Hilbert* or *Ehrhart function* $H(C, k) = |kP \cap \mathbb{Z}^d|$, $k \in \mathbb{Z}_+$.

Task (Vol) is the easiest, and Normaliz computes $\mathrm{vol}(P)$ by summing the volumes $\mathrm{vol}(\sigma \cap A_1)$ where σ runs over the simplicial cones in the triangulation. With the notation introduced above, one has

$$\mathrm{vol}(\sigma \cap A_1) = \frac{|\det(G_\sigma)|}{\deg(v_1) \cdots \deg(v_d)}.$$

For the justification of this formula note that the simplex $\sigma \cap A_1$ is spanned by the vectors $v_i / \deg(v_i)$, $i = 1, \ldots, d$, and that the vertex 0 of the d-simplex $\delta = \mathrm{conv}(0, \sigma \cap A_1)$ has (lattice) height 1 over the opposite facet $\sigma \cap A_1$ of δ so that $\mathrm{vol}(\sigma \cap A_1) = \mathrm{vol}(\delta)$.

The remaining tasks depend on the set E of lattice points in the semi-open parallelotope

$$\mathrm{par}(v_1, \ldots, v_d) = \{q_1 v_1 + \cdots + q_d v_d : 0 \le q_i < 1\}.$$

For the efficiency of the evaluation it is important to generate $E = \mathbb{Z}^d \cap \mathrm{par}(v_1, \ldots, v_d)$ as fast as possible. The basic observation is that E is a set of representatives of the group $\mathbb{Z}^d / U_\sigma$ where the subgroup U_σ is spanned by $v_1, \ldots, v_d$. Thus one finds E in two steps:

(Rep) find a representative of every residue class, and
(Mod) reduce its coefficients with respect to the $\mathbb{Q}$-basis $v_1, \ldots, v_d$ modulo 1.

The first idea for (Rep) that comes to mind (and used in the first version of Normaliz) is to decompose $\mathbb{Z}^d / U_\sigma$ into a direct sum of cyclic subgroups $\mathbb{Z}\overline{u}_i$, $i = 1, \ldots, d$ where $u_1, \ldots, u_d$ is a $\mathbb{Z}$-basis of $\mathbb{Z}^d$ and $^-$ denotes the residue class modulo U_σ. The elementary divisor theorem guarantees the existence of such a decomposition, and finding it amounts to a diagonalization of G_σ over $\mathbb{Z}$. But diagonalization is even more expensive than matrix inversion, and therefore it is very helpful that a filtration of $\mathbb{Z}^d / U_\sigma$ with cyclic quotients is sufficient. Such a filtration can be based on trigonalization:

Proposition 9. *With the notation introduced, let $e_1, \ldots, e_d$ denote the unit vectors in $\mathbb{Z}^d$ and let $X \in \mathrm{GL}(d, \mathbb{Z})$ such that XG_σ is an upper triangular matrix D with diagonal elements $a_1, \ldots, a_d \ge 1$. Then the vectors*

(4.2) $$b_1 e_1 + \cdots + b_d e_d, \qquad 0 \le b_i < a_i, \; i = 1, \ldots, d,$$

represent the residue classes in $\mathbb{Z}^d / U_\sigma$.

Proof. Note that the rows of XG_σ are a $\mathbb{Z}$-basis of U_σ. Since $|\mathbb{Z}^d / U_\sigma| = |\det G_\sigma| = a_1 \cdots a_d$ it is enough to show that the elements listed represent pairwise different residue classes. Let p be the largest index such that $a_p > 1$. Note that a_p is the order of the cyclic group $\mathbb{Z}\overline{e}_p$, and that we obtain a $\mathbb{Z}$-basis of $U'_\sigma = U_\sigma + \mathbb{Z}e_p$ if we replace the p-th row of XG_σ by e_p. If two vectors $b_1 e_1 + \cdots + b_p e_p$ and $b'_1 e_1 + \cdots + b'_p e_p$ in our list represent the same residue class modulo U_σ, then even more so modulo U'_σ. It follows that $b_i = b'_i$ for

$i = 1, \ldots, p-1$, and taking the difference of the two vectors, we conclude that $b_p = b'_p$ as well. □

The first linear algebra step that comes up is therefore the trigonalization

$$(4.3) \qquad\qquad\qquad X G_\sigma = D.$$

Let G_σ^{tr} be the transpose of G_σ. For (Mod) it is essentially enough to reduce those e_i modulo 1 that appear with a coefficient > 0 in (4.2), and thus we must solve the simultaneous linear systems

$$(4.4) \qquad\qquad\qquad G_\sigma^{\mathrm{tr}} x_i = e_i, \qquad a_i > 1,$$

where we consider x_i and e_i as column vectors. In a crude approach one would simply invert the matrix G_σ^{tr} (or G_σ), but in general the number of i such that $a_i > 1$ is small compared to d (especially if d is large), and it is much better to solve a linear system with the specific multiple right hand side given by (4.4). The linear algebra is of course done over $\mathbb{Z}$, using $a_1 \cdots a_d$ as a common denominator. Then Normaliz tries to produce the residue classes and to reduce them modulo 1 (or, over $\mathbb{Z}$, modulo $a_1 \cdots a_d$) as efficiently as possible.

For task (LP) one extracts the vectors of degree 1 from E, and the degree 1 vectors collected from all σ from the set of lattice points in $P = C \cap A_1$. For (HB) one first reduces the elements of $E \cup \{v_1, \ldots, v_d\}$ to a Hilbert basis of σ, collects these and then applies "global" reduction in C. This procedure has been described in [6], and nothing essential has been added meanwhile.

The most difficult and mathematically most interesting task is (Ehr). For its solution one uses the well-known fact that the *Hilbert* or *Ehrhart series*, the generating function

$$H_C(t) = \sum_{k=0}^{\infty} H(C,k) t^k,$$

is a rational function of t. For σ one has

$$H_\sigma(t) = \frac{h_0 + h_1 t + \cdots + h_s t^s}{(1 - t^{g_1}) \cdots (1 - t^{g_d})}, \qquad g_i = \deg v_i, \qquad h_j = |\{x \in E : \deg x = j\}|.$$

This follows immediately from the disjoint decomposition

$$(4.5) \qquad\qquad\qquad \mathbb{Z}^d \cap \sigma = \bigcup_{x \in E} x + M_\sigma$$

where M_σ is the (free) monoid generated by $v_1, \ldots, v_d$.

However, one cannot compute $H_C(t)$ by simply adding these functions since points in the intersections of the simplicial cones σ would be counted several times. Fortunately, the intricate inclusion-exclusion problem can be avoided since there exist *disjoint* decompositions of C by semi-open simplicial cones $\sigma \setminus S$ where S is a union of facets (and not just arbitrary faces!) of σ. The series $H_{\sigma \setminus S}(t)$ is as easy to compute as $H_\sigma(t)$ itself. Let $x \in E$, $x = \sum q_i v_i$. Then we define $\varepsilon(x)$ as the sum of all v_i such that (i) $q_i = 0$ and (ii) the facet opposite to v_i belongs to S. Then

$$(4.6) \qquad\qquad\qquad H_{\sigma \setminus S}(t) = \frac{\sum_{x \in E} t^{\deg \varepsilon(x) + \deg x}}{(1 - t^{g_1}) \cdots (1 - t^{g_d})}.$$

This follows from the fact that $(x + M_\sigma) \setminus S = \varepsilon(x) + x + M_\sigma$, and so we just sum over the disjoint decomposition of $\mathbb{Z}^d \cap (\sigma \setminus S)$ induced by (4.5). (Also see [6, Lemma 11].)

The existence of a disjoint decomposition of C into sets of type $\sigma \setminus S$ was shown by Stanley [21] using the existence of a line shelling of C proved by Bruggesser and Mani. Instead of finding a shelling order for the lexicographic triangulation (which is in principle possible), Normaliz 2.0–2.5 used a line shelling for the decomposition, as discussed in [6]. This approach works well for cones of moderate size, but has a major drawback: finding the sets S requires searching over the shelling order, and in particular the whole triangulation must be stored. Köppe and Verdoolaege [18] proved a much simpler principle for the disjoint decomposition (already implemented in Normaliz 2.7). As a consequence, each simplicial cone in the triangulation can be treated in complete independence from the others, and can therefore be forgotten once it has been evaluated (unless the user insists on seeing the triangulation):

Lemma 10. *Let O_C be a vector in the interior of C such that O_C is not contained in a support hyperplane of any simplicial σ in a triangulation of C. For σ choose S_σ as the union of the support hyperplanes $\mathscr{H}^<(\sigma, O_C)$. Then the semi-open simplicial cones $\sigma \setminus S_\sigma$ form a disjoint decomposition of C.*

See [18] for a proof. It is of course not possible to choose an *order vector* O_C that avoids all hyperplanes in advance, but this is not a real problem. Normaliz chooses O_C in the interior of the first simplicial cone, and works with a lexicographic infinitesimal perturbation O'_C. (This trick is known as "simulation of simplicity" in computational geometry; see [11]). If $O_C \in H^<$ (or $O_C \in H^>$), then $O'_C \in H^<$ (or $O'_C \in H^>$). In the critical case $O_C \in H$, we take the linear form λ representing H and look up its coordinates in the dual basis $e_1^*, \ldots, e_d^*$. If the first nonzero coordinate is negative, then $O'_C \in H^<$, and else $O'_C \in H^>$.

At first it seems that one must compute the support hyperplanes of σ in order to apply Lemma 10. However, it is much better to solve the system

$$(4.7) \qquad\qquad G_\sigma^{\mathrm{tr}} I^\sigma = O_C.$$

The solution I^σ is called the *indicator* of σ. One has $O_C \in H^<$ (or $O_C \in H^>$) if $I_i^\sigma < 0$ (or $I_i^\sigma > 0$) for the generator v_i opposite to H (λ vanishes on H). Let us call σ *generic* if all entries of I^σ are nonzero.

If $I_i^\sigma = 0$—this happens rarely, and extremely rarely for more than one index i—then we are forced to compute the linear form representing the support hyperplane opposite of v_i. In view of (4.1) this amounts to solving the systems

$$(4.8) \qquad\qquad G_\sigma x = e_i, \qquad I_i^\sigma = 0,$$

simultaneously for the lexicographic decision.

If σ is unimodular, in other words, if $|\det G_\sigma| = 1$, then the only system to be solved is (4.7), provided that σ is generic. Normaliz tries to take advantage of this fact by guessing whether σ is unimodular, testing two necessary conditions:

(PU1) Every σ (except the first) is inserted into the triangulation with a certain generator x_i. Let H be the facet of σ opposite to x_i. If $\mathrm{ht}_H(x_i) > 1$, then σ is nonunimodular. (The number $\mathrm{ht}_H(x_i)$ has been computed in the course of the triangulation.)

(PU2) If $\gcd(\deg v_1, \ldots, \deg v_d) > 1$, then σ is not unimodular.

If σ passes both tests, we call it *potentially unimodular*. (Data on the efficiency of this test will be given in Remark 12(e).

After these preparations we can describe the order in which Normaliz treats the trigonalization (4.3) and the linear systems (4.4), (4.7) and (4.8):

(L1) If σ is potentially unimodular, then (4.7) is solved first. It can now be decided whether σ is indeed unimodular.

(L2) If σ is not unimodular, then the trigonalization (4.3) is carried out next. In the potentially unimodular, but nongeneric case, the trigonalization is part of the solution of (4.8) (with multiple right hand side).

(L3) In the nonunimodular case, we now solve the system (4.4) (with multiple right hand side).

(L4) If σ is not potentially unimodular and not generic, it remains to solve the system (4.8) (with multiple right hand side).

As the reader may check, it is never necessary to perform all 4 steps. In the unimodular case, (L1) must be done, and additionally (L2) if σ is nongeneric. If σ is not even potentially unimodular, (L2) and (L3) must be done, and additionally (L4) if it is nongeneric. In the potentially unimodular, but nonunimodular case, (L1), (L2) and (L3) must be carried out.

Remark 11. (a) If one stores the transformation matrix X of (4.3) and its inverse (for example as a sequence of row exchanges and elementary transformations), then one can solve the remaining systems without further trigonalization. However, in general the bookkeeping needs more time than it saves as tests have shown.

(b) The simplicial cones stored in the evaluation buffer are processed in parallel, and parallelization works extremely well for them.

(c) Simplicial cones of height 1 need not be evaluated for (HB) and (LP); see Remark 8.

We conclude this section with a brief discussion of the computation and the representation of the Hilbert series by Normaliz. The reader can find the necessary background in [4, Chapter 6].

Adding the Hilbert series (4.6) is very simple if they all have the same denominator, for example in the case in which the generators of C (or at least the extreme integral generators) have degree 1. For efficiency, Normaliz first forms "denominator classes" in which the Hilbert series with the same denominator are accumulated. At the end, the class sums are added over a common denominator that is extended whenever necessary. This yields a "raw" form of the Hilbert series of type

$$(4.9) \qquad H_C(t) = \frac{R(t)}{(1 - t^{s_1}) \cdots (1 - t^{s_r})}, \qquad R(t) \in \mathbb{Z}[t],$$

whose denominator in general has $> d$ factors.

In order to find a presentation with d factors, Normaliz proceeds as follows. First it reduces the fraction to lowest terms by factoring the denominator of (4.9) into a product

of cyclotomic polynomials:

$$(4.10) \qquad H_C(t) = \frac{Z(t)}{\zeta_{z_1} \cdots \zeta_{z_w}}, \qquad Z(t) \in \mathbb{Z}[t], \quad \zeta_{z_j} \nmid Z(t),$$

which is of course the most economical way for representing $H_C(t)$ (as a single fraction). The orders and the multiplicities of the cyclotomic polynomials can easily be bounded since all denominators in (4.6) divide $(1 - t^\ell)^d$ where ℓ is the least common multiple of the degrees $\deg x_i$. So we can find a representation

$$(4.11) \qquad H_C(t) = \frac{F(t)}{(1 - t^{e_1}) \cdots (1 - t^{e_d})}, \qquad F(t) \in \mathbb{Z}[t],$$

in which e_d is the least common multiple of the orders of the cyclotomic polynomials that appear in (4.10), e_{d-1} is the least common multiple of the orders that have multiplicity ≥ 2 etc. Normaliz produces the presentation (4.11) whenever the degree of the numerator remains of reasonable size.

It is well-known that the Hilbert function itself is a quasipolynomial:

$$(4.12) \qquad H(C, k) = q_0(k) + q_1(k)k + \cdots + q_{d-1}(k)k^{d-1}, \qquad k \geq 0,$$

where the coefficients $q_j(k) \in \mathbb{Q}$ are periodic functions of k whose common period is the least common multiple of the orders of the cyclotomic polynomials in the denominator of (4.10). Normaliz computes the quasipolynomial, with the proviso that its period is not too large. It is not hard to see that the periods of the individual coefficients are related to the representation (4.11) in the following way: e_k is the common period of the coefficients $q_{d-1}, \ldots, q_{d-k}$. The leading coefficient q_{d-1} is actually constant (hence $e_1 = 1$), and related to the multiplicity by the equation

$$(4.13) \qquad q_{d-1} = \frac{\mathrm{vol}(P)}{(d-1)!}.$$

Since q_{d-1} and $\mathrm{vol}(P)$ are computed completely independently from each other, equation (4.13) can be regarded as a test of correctness for both numbers.

The choice (4.11) for $H_C(t)$ is motivated by the desire to find a standardized representation whose denominator conveys useful information. The reader should note that this form is not always the expected one. For example, for $C = \mathbb{R}_+^2$ with $\deg(e_1) = 2$ and $\deg(e_2) = 3$, the three representations (4.9)–(4.11) are

$$\frac{1}{(1 - t^2)(1 - t^3)} = \frac{1}{\zeta_1^2 \zeta_2 \zeta_3} = \frac{1 - t + t^2}{(1 - t)(1 - t^6)}.$$

Actually, it is unclear what the most natural standardized representation of the Hilbert series as a fraction of two polynomials should look like, unless the denominator is $(1 - t)^d$. Perhaps the most satisfactory representation should use a denominator $(1 - t^{p_1}) \cdots (1 - t^{p_d})$ in which the exponents p_i are the degrees of a homogeneous system of parameters (for the monoid algebra $K[\mathbb{Z}^d \cap C]$ over an infinite field K). At present Normaliz cannot find such a representation (except the one with the trivial denominator $(1 - t^\ell)^d$), but future versions may contain it.

5. Computational results

A driving force for the recent improvements in Normaliz that we have described in the previous sections was the desire to compute the volumes and Ehrhart series of certain polytopes related to combinatorial voting theory.

5.1. **Voting schemes and volumes of rational polytopes.**

We briefly sketch the connection between rational polytopes and combinatorial voting theory, referring the reader to [15], [20] or [23] for a more extensive treatment. We consider the three voting schemes discussed in [20].

Consider an election in which each of the k voters fixes a linear preference order of n candidates. In other words, voter i chooses a linear order $j_1 \succ_i \cdots \succ_i j_n$ of the candidates $1, \ldots, n$. Set $N = n!$. Counting the preference orders gives an N-tuple $(v_1, \ldots, v_N)$ in which v_p is the number of voters that have chosen the preference order p. Then $v_1 + \cdots + v_N = k$, and $(v_1, \ldots, v_N)$ can be considered as a lattice point in the positive orthant of $\mathbb{R}_+^N$, more precisely, as a lattice point in the simplex

$$\mathscr{U}_k^{(n)} = \mathbb{R}_+^N \cap A_k = k\big(\mathbb{R}_+^N \cap A_1\big) = k\mathscr{U}^{(n)}$$

where A_k is the hyperplane defined by $x_1 + \cdots + x_N = k$, and $\mathscr{U}^{(n)} = \mathscr{U}_1^{(n)}$ is the unit simplex of dimension $N - 1$ naturally embedded in N-space. All further discussion is based on the *Impartial Anonymous Culture* assumption that all lattice points in the simplex $\mathscr{U}_k^{(n)}$ have equal probability of being the outcome of the election.

We fix a specific outcome $v = (v_1, \ldots, v_N)$. Let us say that candidate j *beats* candidate j' with respect to v if

$$(5.1) \qquad |\{i : j \succ_i j' : i = 1, \ldots, k\}| > |\{i : j' \succ_i j : i = 1, \ldots, k\}|.$$

As the Marquis de Condorcet observed, the relation "beats" is nontransitive in general, and one must ask for the probability of Condorcet's paradoxon, namely an outcome without a Condorcet winner where candidate j is a *Condorcet winner* if j beats all other candidates j'. Let $C_k^{(n)}(j)$ denote the probability that candidate j is the Condorcet winner, and $C_k^{(n)}$ the probability that there is a Condorcet winner. By symmetry and by mutual exclusion $C_k^{(n)} = nC_k^{(n)}(1)$. Usually the number k of voters is very large, and therefore one is mainly interested in the limit

$$C^{(n)} = \lim_{k \to \infty} C_k^{(n)} = n \lim_{k \to \infty} C_k^{(n)}(1) = nC^{(n)}(1).$$

Let us fix candidate 1. It is not hard to see that the $n - 1$ inequalities (5.1) for $j = 1$ and $j' = 2, \ldots, n$ constitute homogeneous linear inequalities in the variables $v_1, \ldots, v_N$. Together with the inequalities $v_i \geq 0$ they define a semi-open subpolytope $\mathscr{C}_k^{(n)}$ of $\mathscr{U}_k^{(n)}$. Then

$$(5.2) \qquad C^{(n)}(1) = \lim_{k \to \infty} \frac{|\mathscr{C}_k^{(n)} \cap \mathbb{Z}^N|}{|\mathscr{U}_k^{(n)} \cap \mathbb{Z}^N|} = \frac{\operatorname{vol} \mathscr{C}_1^{(n)}}{\operatorname{vol} \mathscr{U}_1^{(n)}} = \operatorname{vol} \overline{\mathscr{C}}^{(n)}$$

where $\overline{}$ denotes closure and $\mathscr{C}^{(n)} = \mathscr{C}_1^{(n)}$. For the validity of (5.2) note that we work with the lattice normalized volume in which the unit simplex has volume 1.

In the case of two candidates Concordet's paradox cannot occur (if one excludes draws), and for 3 candidates the relevant volume is not hard to compute. The situation changes significantly for 4 candidates since $\mathscr{C}^{(4)}$ has dimension 23 and 234 vertices. As a subpolytope of $\mathscr{U}^{(4)}$, $\mathscr{C}^{(4)}$ is cut out by the inequalities $\lambda_i(v) > 0$, $i = 1, 2, 3$ whose coefficients are in the first 3 rows displayed in Table 5. For the assignment of indices the preference orders are listed lexicographically, starting with $1 \succ 2 \succ 3 \succ 4$ and ending with $4 \succ 3 \succ 2 \succ 1$.

$$\lambda_1\colon\ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ \ \ 1\ \ \ 1\ {-}1\ {-}1\ \ \ 1\ {-}1\ \ \ 1\ \ \ 1\ {-}1\ {-}1\ \ \ 1\ {-}1$$
$$\lambda_2\colon\ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ {-}1\ {-}1\ \ \ 1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ \ \ 1\ \ \ 1\ \ \ 1\ {-}1\ {-}1\ {-}1$$
$$\lambda_3\colon\ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ {-}1\ {-}1\ {-}1\ \ \ 1\ \ \ 1\ \ \ 1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1$$

$$\lambda_4\colon\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0$$
$$\lambda_5\colon\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0$$
$$\lambda_6\colon\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 1\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ \ \ 0\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1\ {-}1$$

TABLE 5. Inequalities for $\mathscr{C}^{(4)}$ and $\mathscr{P}^{(4)}$

The lexicographic triangulation used by Normaliz has (only) $1,473,107$ simplicial cones. Normaliz computes

$$\operatorname{vol}\mathscr{C}^{(4)} = \frac{1717}{8192}$$

in a few seconds. It follows that $C^{(4)} = 1717/2048 \approx 0.8384$. According to [20], this value was first determined by Gehrlein [14].

The simplest way out of the dilemma that there may not exist a Condorcet winner is *plurality voting*: candidate j is the *plurality winner* if j has more first places in the preference orders of the voters than any of the other $n-1$ candidates. The *Condorcet efficiency* of plurality voting (and similarly of other voting schemes) is the conditional probability that the Condorcet winner is elected, provided it exists, as $k \to \infty$. Therefore one must compute the probability of the event that candidate j is the Condorcet winner, but candidate $j' \neq j$ is the plurality winner. By symmetry, one can assume $j = 1$ and $j' = 2$. The semi-open polytope $\mathscr{P}_k^{(n)}$ whose lattice points represent this unexpected outcome is cut out from $\mathscr{C}_k^{(n)}$ by $n-1$ further inequalities saying that 2 has more first places than the other $n-1$ candidates. Thus one obtains

$$\frac{C^{(n)} - n(n-1)\operatorname{vol}\mathscr{P}^{(n)}}{C^{(n)}}$$

as the Condorcet efficiency of plurality voting where $\mathscr{P}^{(n)} = \mathscr{P}_1^{(n)}$.

The extra 3 inequalities $\lambda_i(v) > 0$, $i = 4, 5, 6$, given in the last 3 lines of Table 6 increase the complexity of the polytope $\mathscr{P}^{(4)}$ enormously. It has 3928 vertices, and the triangulation increases to $347,225,775,338$ simplicial cones. Nevertheless, Normaliz computes the volume and the Ehrhart series in acceptable time. We have obtained

$$\operatorname{vol}\mathscr{P}^{(4)} = \frac{369403718529016355068149157}{205426954327818240000000000000}$$

so that the Condorcet efficiency of plurality voting turns out to be

$$\frac{C^{(4)} - 12\,\mathrm{vol}\,\mathscr{P}^{(4)}}{C^{(4)}} = \frac{10658098255011916449318509}{14352135440302080000000000} \approx 0.7426$$

in perfect accordance with [20].

Additional support (and tests for Normaliz) can be added by a chain of volume computations as follows. We start from the unit simplex $P_0^+ = \mathscr{U}^{(4)}$ and set

$$P_i^+ = \{x \in \mathscr{U}^{(4)} : \lambda_1(x), \ldots, \lambda_i(x) \geq 0\} \quad \text{and} \quad P_i^- = P_{i-1}^+ \cap \{x : \lambda_i(x) \leq 0\}, \quad i = 1, \ldots, 6.$$

Then $P_6^+ = \mathscr{P}^{(4)}$, and we must have

$$\mathrm{vol}\,P_{i-1}^+ = \mathrm{vol}\,P_i^+ + \mathrm{vol}\,P_i^-, \qquad i = 1, \ldots, 6.$$

The volumes computed by Normaliz satisfy all these equations, as they should. We start from $i = 3$ since $P_3^+ = \mathscr{C}^{(4)}$:

$$\mathrm{vol}\,P_3^+ = \frac{1717}{8192},$$

$$\mathrm{vol}\,P_4^+ = \frac{418988423262545}{16231265527136256}, \qquad \mathrm{vol}\,P_4^- = \frac{2982999236660911}{16231265527136256},$$

$$\mathrm{vol}\,P_5^+ = \frac{16228863391807757733501803}{77035107872931840000000000}, \qquad \mathrm{vol}\,P_5^- = \frac{36567199778794370091947}{77035107872931840000000000},$$

$$\mathrm{vol}\,P_6^+ = \frac{3694037185290163550681491}{20542695432781824000000000}, \qquad \mathrm{vol}\,P_6^- = \frac{19009791575757152159699 51}{61628086298345472000000000}.$$

The largest triangulation of $463,613,250,401$ simplicial cones was produced by P_6^-.

The last problem discussed in [20] is *plurality voting versus plurality cutoff*. It works as follows. In the first round of the election the two top candidates in plurality voting are selected, and in the second round the preference orders are restricted to these two candidates. In order to model this situation by inequalities one must fix an outcome of the first round that involves all n candidates, say $1, \ldots, n$ in this order. This condition gives rise to $n - 1$ inequalities. Then the n-th inequality expresses that 2 is the winner of the second round, despite the fact that 1 was the winner of the first round. The volume of the corresponding polytope gives the probability of this event. By mutual exclusion and symmetry, we must multiply the volume by $n!$ in order to obtain the probability for the event that the winner of the first plurality round looses after cutoff.

As a subpolytope of $\mathscr{U}^{(4)}$, the polytope $\mathscr{Q}^{(4)}$ is defined by the inequalities in Table 6.

1	1	1	1	1	1	−1	−1	−1	−1	−1	−1	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	1	1	1	1	1	1	−1	−1	−1	−1	−1	−1	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	−1	−1	−1	−1	−1	−1
−1	−1	−1	−1	−1	−1	1	1	1	1	1	1	−1	1	−1	1	−1	1	−1	1	−1	1	−1	1

TABLE 6. Inequalities for $\mathscr{Q}^{(4)}$

It has 1872 vertices and the triangulation computed by Normaliz has $257,744,341,008$ simplicial cones. The volume is

$$\operatorname{vol} \mathscr{Q}^{(4)} = \frac{2988379676768359}{292162779488452608}.$$

The total probability of the failure of the winner of the first round is therefore $24 \cdot \operatorname{vol} \mathscr{Q}^{(4)}$ ≈ 0.2455, in accordance with the results of [20] for this model. Further support has been given by De Loera, Dutra, Köppe, Moreinis, Pinto and Wu in [12], where LattE Integrale [13] was used for the volume computation. (Additionally we have applied the same verification as for $\mathscr{P}^{(4)}$).

5.2. **Ehrhart series and quasipolynomials.** Normaliz has not only computed the volumes, but also the Ehrhart series and quasipolynomials for the closures of the semi-open polytopes $\mathscr{C}^{(4)}$, $\mathscr{P}^{(4)}$ and $\mathscr{Q}^{(4)}$. The Hilbert series of $\overline{\mathscr{C}}^{(4)}$ is the rational function with numerator

$$1 + t + 5t^1 + 133t^2 + 363t^3 + 4581t^4$$
$$+ 8655t^5 + 69821t^6 + 100915t^7 + 596834t^8 + 697232t^9$$
$$+ 3255226t^{10} + 3176870t^{11} + 12235441t^{12} + 10182887t^{13} + 33268048t^{14}$$
$$+ 23917200t^{15} + 67509138t^{16} + 42243510t^{17} + 104272000t^{18} + 56990048t^{19}$$
$$+ 123966919t^{20} + 59177761t^{21} + 113925878t^{22} + 47336170t^{23} + 80758791t^{24}$$
$$+ 28993857t^{25} + 43770180t^{26} + 13415068t^{27} + 17837843t^{28} + 4580485t^{29}$$
$$+ 5320122t^{30} + 1111974t^{31} + 1113216t^{32} + 180850t^{33} + 152891t^{34}$$
$$+ 17845t^{35} + 12346t^{36} + 890t^{37} + 481t^{38} + 15t^{39} + 6t^{40}.$$

and denominator

$$(1-t)(1-t^2)^{14}(1-t^4)^9$$

Numerator and denominator are coprime.

The quasipolynomial of $\overline{\mathscr{C}}^{(4)}$ has period 4. We give the numerators of its first two and its last two coefficients; the denominator is $d = 693959790182222163590774840000$:

$$q_0(0) = 693959790182222163590774840000/d = 1,$$
$$q_1(0) = 2089922514833674795902566400000/d,$$
$$q_{22}(0) = q_{22}(2) = 15982652919/d,$$
$$q_{23}(0) = q_{23}(1) = q_{23}(2) = q_{23}(3) = 56262656/d = \operatorname{vol} \mathscr{P}^{(4)}/23!,$$
$$q_0(1) = 203475031022335179700809216000/d,$$
$$q_1(1) = 709276434214253918714297164800/d,$$
$$q_{22}(1) = q_{22}(3) = 15528493056/d,$$
$$q_0(2) = 693308184929915219977568256000/d,$$
$$q_1(2) = 208924553117357562368549191680/d,$$
$$q_0(3) = 203475031022335179700809216000/d,$$

$$q_1(3) = 7092764342142539187142971648000/d.$$

The reader may have noticed that the coefficients of $q(1)$ and $q(3)$ coincide, as far as listed. In fact, these two polynomials are equal.

For the other two polytopes we only list the denominators for the representation (4.11) (with non-coprime numerators in both cases):

$$\overline{\mathscr{P}}^{(4)}: \quad (1-t)(1-t^2)^2(1-t^4)^5(1-t^{12})^4(1-t^{24})(1-t^{120})^{11},$$

$$\overline{\mathscr{Q}}^{(4)}: \quad (1-t)(1-t^2)^2(1-t^4)^5(1-t^{12})^{16}.$$

So $\overline{\mathscr{P}}^{(4)}$ has period 120 and $\overline{\mathscr{Q}}^{(4)}$ has period 12. On request the authors will provide full sets of data.

For the exact counting of the frequency of Condorcet's paradoxon, and similarly for the other events considered, one must compute the Ehrhart series of the semi-open polytopes. At present Normaliz cannot do this directly so that one depends on cumbersome inclusion/exclusion for the semi-open polytopes. We plan an extension for semi-open polytopes (or cones) the next version.

One should note that semi-open polytopes present an inherent difficulty since in general a disjoint decomposition into sets $\sigma \setminus S$ as discussed in Section 4 does not exist. (Recall that S is a union of facets, not just arbitrary faces.) Suppose that all extreme integral generators of the cone C have degree 1. Then all the Hilbert series $H_{\sigma \setminus S}$ have denominator $(1-t)^d$ and a numerator polynomial with nonnegative coefficients. If a disjoint decomposition into sets $\sigma \setminus S$ exists, the resulting Hilbert series must have a numerator polynomial with nonnegative coefficients as well. However, one can easily find semi-open polytopes for which this is not the case: remove two opposite edges from the unit square; then the Ehrhart series of the remaining semi-open polytope is $(2t^2 - t^3)/(1-t)^3$.

5.3. **Computation times.** Table 7 gives an indication of the computation times to be expected for the volumes and Ehrhart series of the rational polytopes discussed in the previous subsection. The times have been collected at different stages of the development on a SUN xFire 4450 with 20 threads (of the maximal number of 24). However, these have differed only little in this respect, except that the computation for $\mathscr{Q}^{(4)}$ has become about 20% faster now.

Polytope	computation	triangulation size	real time	parllelization
$\mathscr{C}^{(4)}$	Ehrhart series	1,473,107	00:00:30 h	serial
$\mathscr{P}^{(4)}$	Ehrhart series	347,225,775,338	292:50:22 h	1981%
$\mathscr{Q}^{(4)}$	Ehrhart series	257,744,341,008	175:11:26 h	1991%
P_5^+	volume	383,986,938,515	126:06:57 h	1953%
$\mathscr{Q}^{(4)}$	triangulation	271,164,705,162	18:50:47 h	1779%

TABLE 7. Computation times

Remark 12. (a) The size of the lexicographic triangulation depends on the order in which the extreme rays are processed. The polytopes in the table above are defined by their support hyperplanes, and therefore Normaliz first computes the extreme rays from them. The order used in the computations mentioned in Table 7 is not necessarily identical with the order produced by the most recent version. In it we have eliminated any of the unpredictable effects of parallelization in the function FINDNEWHYP (see Section 2). See also Remark 7(d).

That the order of the generators has some influence is shown by the two computations of $\mathcal{Q}^{(4)}$ in Table 7.

(b) The table shows that the times needed for (i) pure triangulation, (ii) volume computation and (iii) Ehrhart series are in approximate proportion $1 : 5 : 10$ for this class of cones.

(c) For the calculation of the Hilbert bases of the cones defined by the polytopes $\mathcal{C}^{(4)}$, $\mathcal{P}^{(4)}$ and $\mathcal{Q}^{(4)}$ one should use the dual mode of Normaliz. Then the Hilbert basis calculations are a matter of seconds. The Hilbert bases have the following numbers of elements: 242 for $\mathcal{C}^{(4)}$, 25192 for $\mathcal{P}^{(4)}$, 9621 for $\mathcal{Q}^{(4)}$.

(d) Normaliz needs relatively little memory. All the computations mentioned run stably with < 1 GB of RAM.

(e) From the Ehrhart series calculation of $\mathcal{Q}^{(4)}$ we have obtained the following statistics on the types of simplicial cones: $61,845,707,957$ are unimodular, $108,915,272,879$ are not unimodular, but satisfy condition (PU1), of which $62,602,898,779$ are potentially unimodular. This shows that condition (PU2) that was added at a later stage has a satisfactory effect. (The number of potentially unimodular, but nonunimodular simplicial cones is rather high in this class.) The average value of $|\det G_\sigma|$ is ≈ 10.

The number of nongeneric simplicial cones is $129,661,342$. The total number s of linear systems that had to be solved for the computation of the Ehrhart series is bounded by $516,245,872,838 \le s \le 516,375,534,180$.

The total number of pyramids was $80,510,681$. It depends on the number of parallel threads that are allowed.

5.4. **The exploitation of symmetry.** The elegant approach of Schürmann in [20] for the computation of the volumes of $\mathcal{C}^{(4)}$, $\mathcal{P}^{(4)}$ and $\mathcal{Q}^{(4)}$ uses the high degree of symmetries of these polytopes. If certain variables $v_{i_1}, \ldots, v_{i_u}$ occur in all of the linear forms given in Tables 5 and 6, then any permutation of them acts as a symmetry on the corresponding polytope, and the variables $v_{i_1}, \ldots, v_{i_u}$ can be replaced by their sum $v_{i_1} + \cdots + v_{i_u}$ in them. (The polytopes have further symmetries.) The substitution can be used for a projection into a space of much lower dimension, mapping the polytope P under consideration to a polytope Q (this requires that the grading affine hyperplane A_1 is mapped onto an affine hyperplane by the projection). Instead of counting the lattice points in kP one counts the lattice points in kQ weighted with their number of preimage lattice points in kP. This amounts to the consideration of a generalized Ehrhart function

$$k \mapsto \sum_{x \in kQ \cap \mathbb{Z}^d} f(x).$$

The theory of generalized Ehrhart functions has recently been developed in several papers; see [19], [1], [2]. An extension of Normaliz to the computation of generalized Ehrhart functions and their generating functions is envisaged.

In [20], only the leading term of the polynomial f is used. Integration with respect to Lesbesgue measure then yields the volume.

5.5. **Previous challenging computations by Normaliz.** We conclude by listing some more performance data of previous computations in Table 8. For the $5 \times 5 \times 3$ contingency tables see [5], and the computations of the statistical rank models were done for [22]. These computations, for which the triangulations are very large, but not as monstrous as those of $\mathscr{P}^{(4)}$ and $\mathscr{Q}^{(4)}$, are now doable in comfortable time. The Hilbert basis computations show the efficiency of partial triangulations (see Remark 8). The computations were done on our SUN xFire 4450 with 20 parallel threads at various stages of the development. The degree of parallelization varies somewhat because of the problems discussed in Remark 6(b). The computation times for triangulation, volume and Hilbert

	dim	|extr|	computation	|triangulation|	real time
$5 \times 5 \times 3$	43	75	Hilbert series	9,248,527,905	07:07:30 h
contingency tables			Hilbert basis	448,64	00:20:31 h
linear rank	16	720	Hilbert series	5,745,903,354	02:23:04 h
model for S_6			Hilbert basis	7,783,191	00:09:56 h
ascending rank	27	120	Hilbert series	20,853,141,970	07:28:11 h
model for S_5			Hilbert basis	0	00:00:01 h

TABLE 8. Previous challenging computations

series of the $5 \times 5 \times 3$ contingency tables are in approximate proportion $2 : 5 : 6$. This is not surprising since almost all simplicial cones are unimodular.

6. ACKNOWLEDGEMENT

REFERENCES

[1] V. Baldoni, N. Berline, J.A. De Loera, M. Köppe and M. Vergne, *How to integrate a polynomial over a simplex.* Math. Comp. **80** (2011), 297–325.
[2] V. Baldoni, N. Berline, J.A. De Loera, M. Köppe and M. Vergne, *Computation of the highest coefficients of weighted Ehrhart quasi-polynomials of rational polyhedra.* Found. Comp. Math., in press.

[3] T. Bogart, A. Raymond and R.R. Thomas, *Small Chvatal rank.* Mathematical Programming **124** (2010), 45–68.

[4] W. Bruns, J. Gubeladze, *Polytopes, rings and K-theory*, Springer, 2009.

[5] W. Bruns, R. Hemmecke, B. Ichim, M. Köppe, and C. Söger, *Challenging computations of Hilbert bases of cones associated with algebraic statistics.* Exp. Math. **20** (2011), 1–9.

[6] W. Bruns and B. Ichim, *Normaliz: algorithms for affine monoids and rational cones.* J. Algebra **324** (2010), 1098–1113.

[7] W. Bruns, B. Ichim and C. Söger, *Normaliz. Algorithms for rational cones and affine monoids.* Available from `http://www.math.uos.de/normaliz`.

[8] W. Bruns and R. Koch, *Computing the integral closure of an affine semigroup.* Univ. Iagel. Acta Math. **39** (2001), 59–70.

[9] B. A. Burton, *Regina: software for 3-manifold theory and normal surfaces,* available from `http://regina.sourceforge.net/`

[10] A. Craw, D. Maclagan and R.R. Thomas, *Moduli of McKay quiver representations II: Gröbner basis techniques.* J. Algebra **316** (2007), 514–535.

[11] H. Edelsbrunner, *algorithms in combinatorial geometry.* Springer 1987.

[12] J.A. De Loera, B. Dutra, M. Köppe, S. Moreinis, G. Pinto and J. Wu, *Software for exact integration of polynomials over polyhedra.* Preprint arXiv:1108.0117v3.

[13] J.A. DeLoera, M. Köppe et al., *LattE integrale.* Available at `http://www.math.ucdavis.edu/~latte/`

[14] W.V. Gehrlein, *Condorcet winners on four candidates with anonymous voters.* Economics Letters **71** (2001), 335–340.

[15] D. Lepelley, A. Louichi and H. Smaoui, *On Ehrhart polynomials and probability calculations in voting theory.* Social Choice and Welfare **30** (2008), 363–383.

[16] M. Joswig, B. Müller and A. Paffenholz, *Polymake and lattice polytopes.* In *DMTCS proc. AK,* C. Krattenthaler (ed.) et al., Proceedings of FPSAC 2009, pp. 491–502.

[17] R. Kappl, M. Ratz und C. Staudt, *The Hilbert basis method for D-flat directions and the superpotential.* Journal of High Energy Physics **10** (2011), 27, 1–11.

[18] M. Köppe and S. Verdoolaege, *Computing parametric rational generating functions with a Primal Barvinok algorithm.* Electr. J. Comb. **15** (2008), R16, 1–19.

[19] M. Schechter, *Integration over a polyhedron: an application of the Fourier- Motzkin elimination method.* Amer. Math. Monthly **105** (1998), 246–251.

[20] A. Schürmann, *Exploiting polyhedral symmetries in social choice.* Social Choice and Welfare, in press.

[21] R. P. Stanley, *Linear Diophantine equations and local cohomology.* Invent. math. **68**, 175–193 (1982).

[22] B. Sturmfels and V. Welker, *Commutative algebra of statistical ranking.* J. Algebra **361** (2012), 264–286.

[23] M.C. Wilson and G. Pritchard, *Probability calculations under the IAC hypothesis.* Math. Social Sci. **54** (2007), 244–256.

WINFRIED BRUNS, UNIVERSITÄT OSNABRÜCK, FB MATHEMATIK/INFORMATIK, 49069 OSNABRÜCK, GERMANY

E-mail address: `wbruns@uos.de`

BOGDAN ICHIM, INSTITUTE OF MATHEMATICS "SIMION STOILOW" OF THE ROMANIAN ACADEMY, C.P. 1-764, 010702 BUCHAREST, ROMANIA

E-mail address: `bogdan.ichim@imar.ro`

CHRISTOF SÖGER, UNIVERSITÄT OSNABRÜCK, FB MATHEMATIK/INFORMATIK, 49069 OSNABRÜCK, GERMANY

E-mail address: `csoeger@uos.de`