

THE TROPICAL DOUBLE DESCRIPTION METHOD

XAVIER ALLAMIGEON¹ AND STÉPHANE GAUBERT² AND ÉRIC GOUBAULT³

¹ Direction du Budget, 4ème sous-direction, Bureau des transports, Paris, France

² INRIA Saclay and CMAP, Ecole Polytechnique, France

³ CEA, LIST MeASI – Gif-sur-Yvette, France

E-mail address: `firstname.lastname@{polytechnique.org,inria.fr,cea.fr}`

ABSTRACT. We develop a tropical analogue of the classical double description method allowing one to compute an internal representation (in terms of vertices) of a polyhedron defined externally (by inequalities). The heart of the tropical algorithm is a characterization of the extreme points of a polyhedron in terms of a system of constraints which define it. We show that checking the extremality of a point reduces to checking whether there is only one minimal strongly connected component in an hypergraph. The latter problem can be solved in almost linear time, which allows us to eliminate quickly redundant generators. We report extensive tests (including benchmarks from an application to static analysis) showing that the method outperforms experimentally the previous ones by orders of magnitude. The present tools also lead to worst case bounds which improve the ones provided by previous methods.

Introduction

Tropical polyhedra are the analogues of convex polyhedra in tropical algebra. The latter deals with structures like the max-plus semiring \mathbb{R}_{\max} (also called *max-plus algebra*), which is the set $\mathbb{R} \cup \{-\infty\}$, equipped with the addition $x \oplus y := \max(x, y)$ and the multiplication $x \otimes y := x + y$.

The study of the analogues of convex sets in tropical or max-plus algebra is an active research topic, and has been treated under various guises. It arose in the work of Zimmerman [Zim77], following a way opened by Vorobyev [Vor67], motivated by optimization theory. Max-plus cones were studied by Cuninghame-Green [CG79]. Their theory was independently developed by Litvinov, Maslov and Shpiz [LMS01] (see also [MS92]) with

1998 ACM Subject Classification: F.2.2.Geometrical problems and computations, G.2.2 Hypergraphs; Algorithms, Verification.

Key words and phrases: convexity in tropical algebra, algorithmics and combinatorics of tropical polyhedra, computational geometry, discrete event systems, static analysis.

This work was performed when the first author was with EADS Innovation Works, SE/IA – Suresnes, France and CEA, LIST MeASI – Gif-sur-Yvette, France.

This work was partially supported by the Arpege programme of the French National Agency of Research (ANR), project “ASOPT”, number ANR-08-SEGI-005 and by the Digiteo project DIM08 “PASO” number 3389.



motivations from variations calculus and asymptotic analysis, and by Cohen, Gaubert, and Quadrat [CGQ04] who initiated a “geometric approach” of discrete event systems [CGQ99], further developed in [Kat07, DLGKL09]. Other motivations arise from abstract convexity, see the book by Singer [Sin97], and also the work of Briec and Horvath [BH04]. The field has attracted recently more attention after the work of Develin and Sturmfels [DS04], who pointed out connections with tropical geometry, leading to several works by Joswig, Yu, and the same authors [Jos05, DY07, JSY07, Jos09].

A tropical polyhedron can be represented in two different ways, either internally, in terms of extreme points and rays, or externally, in terms of linear inequalities (see Sect. 1 for details). As in the classical case, passing from the external description of a polyhedron to its internal description is a fundamental computational issue. This is the object of the present paper.

Butkovič and Hegedus [BH84] gave an algorithm to compute the generators of a tropical polyhedral cone described by linear inequalities. Gaubert gave a similar one and derived the equivalence between the internal and external representations [Gau92, Ch. III] (see [GK09] for a recent discussion). Both algorithms rely on a successive elimination of inequalities, but have the inconvenience of squaring at each step the number of candidate generators, unless an elimination technique is used, as in the Maxplus toolbox of SCILAB [CGMQ]. Joswig developed a different approach, implemented in POLYMAKE [GJ], in which a tropical polytope is represented as a polyhedral complex [DS04, Jos09].

The present work grew out from two applications: to discrete event systems [Kat07, DLGKL09], and to software verification by static analysis [AGG08]. In these applications, passing from the external to the internal representation is a central difficulty. A further motivation originates from mean payoff games [AGG09b]. These motivations are reviewed in Section 2.

Contributions. We develop a new algorithm which computes the extreme elements of tropical polyhedra. It is based on a successive elimination of inequalities, and a result (Th. 4.1) allowing one, given a polyhedron \mathcal{P} and a tropical halfspace \mathcal{H} , to construct a list of candidates for the generators of $\mathcal{P} \cap \mathcal{H}$. The key ingredient is a combinatorial characterization of the extreme generators of a polyhedron defined externally (Th. 3.5 and 3.7): we reduce the verification of the extremality of a candidate to the existence of a strongly connected component reachable from any other in a directed hypergraph. We include a complexity analysis and experimental results (Sect. 4), showing that the new algorithm outperforms the earlier ones, allowing us to solve instances which were previously by far inaccessible. Our result also leads to worst case bounds improving the ones of previously known algorithms.

1. Definitions: tropical polyhedra and polyhedral cones

The neutral elements for the addition \oplus and multiplication \otimes , *i.e.*, the zero and the unit, will be denoted by $\mathbf{0} := -\infty$ and $\mathbf{1} := 0$, respectively. The tropical analogues of the operations on vectors and matrices are defined naturally. The elements of \mathbb{R}_{\max}^d , the d th fold Cartesian product of \mathbb{R}_{\max} , will be thought of as vectors, and denoted by bold symbols, like $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_d)$.

A *tropical halfspace* is a set of the vectors $\mathbf{x} = (\mathbf{x}_i) \in \mathbb{R}_{\max}^d$ verifying an inequality constraint of the form

$$\max_{1 \leq i \leq d} a_i + \mathbf{x}_i \leq \max_{1 \leq i \leq d} b_i + \mathbf{x}_i, \quad a_i, b_i \in \mathbb{R}_{\max}.$$

A *tropical polyhedral cone* is defined as the intersection of n halfspaces. It can be equivalently written as the set of the solutions of a system of inequality constraints $A\mathbf{x} \leq B\mathbf{x}$. Here, $A = (a_{ij})$ and $B = (b_{ij})$ are $n \times d$ matrices with entries in \mathbb{R}_{\max} , concatenation denotes the matrix product (with the laws of \mathbb{R}_{\max}), and \leq denotes the standard partial ordering of vectors. For sake of readability, tropical polyhedral cones will be simply referred to as *polyhedral cones* or *cones*.

Tropical polyhedral cones are known to be generated by their extreme rays [GK06, GK07, BSS07]. Recall that a *ray* is the set of scalar multiples of a non-zero vector \mathbf{u} . It is *extreme* in a cone \mathcal{C} if $\mathbf{u} \in \mathcal{C}$ and if $\mathbf{u} = \mathbf{v} \oplus \mathbf{w}$ with $\mathbf{v}, \mathbf{w} \in \mathcal{C}$ implies that $\mathbf{u} = \mathbf{v}$ or $\mathbf{u} = \mathbf{w}$. A finite set $G = (\mathbf{g}^i)_{i \in I}$ of vectors is said to *generate* a polyhedral cone \mathcal{C} if each \mathbf{g}^i belongs to \mathcal{C} , and if every vector \mathbf{x} of \mathcal{C} can be written as a *tropical linear combination* $\bigoplus_i \lambda_i \mathbf{g}^i$ of the vectors of G (with $\lambda_i \in \mathbb{R}_{\max}$). Note that in tropical linear combinations, the requirement that λ_i be nonnegative is omitted. Indeed, $\mathbb{0} = -\infty \leq \lambda$ holds for all scalar $\lambda \in \mathbb{R}_{\max}$.

The tropical analogue of the Minkowski theorem [GK07, BSS07] shows in particular that every generating set of a cone that is minimal for inclusion is obtained by selecting precisely one (non-zero) element in each extreme ray.

A tropical polyhedron of \mathbb{R}_{\max}^d is the affine analogue of a tropical polyhedral cone. It is defined by a system of inequalities of the form $A\mathbf{x} \oplus \mathbf{c} \leq B\mathbf{x} \oplus \mathbf{d}$. It can be also expressed as the set of the tropical affine combinations of its generators. The latter are of the form $\bigoplus_{i \in I} \lambda_i \mathbf{v}^i \oplus \bigoplus_{j \in J} \mu_j \mathbf{r}^j$, where the $(\mathbf{v}^i)_{i \in I}$ are the extreme points, the $(\mathbf{r}^j)_{j \in J}$ a set formed by one element of each extreme ray, and $\bigoplus_i \lambda_i = \mathbf{1}$. It is known [CGQ04, GK07] that every tropical polyhedron of \mathbb{R}_{\max}^d can be represented by a tropical polyhedral cone of \mathbb{R}_{\max}^{d+1} thanks to an analogue of the homogenization method used in the classical case (see [Zie98, Sect. 1.5]). Then, the extreme rays of the cone are in one-to-one correspondence with the extreme generators of the polyhedron. That is why, in the present paper, we will only state the main results for cones, leaving to the reader the derivation of the affine analogues, along the lines of [GK07].

In the sequel, we will illustrate our results on the polyhedral cone \mathcal{C} given in Fig. 1, defined by the system in the right side. The left side is a representation of \mathcal{C} in barycentric coordinates: each element $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ is represented as a barycenter with weights $(e^{\mathbf{x}_1}, e^{\mathbf{x}_2}, e^{\mathbf{x}_3})$ of the three vertices of the outermost triangle. Then two elements of a same ray are represented by the same point. The cone \mathcal{C} is depicted in solid gray (the black border is included), and is generated by the extreme elements $\mathbf{g}^0 = (0, 0, 0)$, $\mathbf{g}^1 = (-2, 1, 0)$, $\mathbf{g}^2 = (2, 2, 0)$, and $\mathbf{g}^3 = (0, 0, 0)$.

2. Motivations from static analysis, discrete event systems, and mean pay-off games

Tropical polyhedra have been recently involved in static analysis by abstract interpretation [AGG08]. It has been shown that they allow to automatically compute complex invariants involving the operators \min and \max which hold over the variables of a program. Such invariants are disjunctive, while most existing techniques in abstract interpretation are only able to express conjunctions of affine constraints, see in particular [CC77, CH78, Min01].

For instance, tropical polyhedra can handle notorious problems in verification of memory manipulations. Consider the well-known memory string manipulating function `memcpy` in C. A call to `memcpy(dst, src, n)` copies exactly the first n characters of the string buffer



Figure 1: A tropical polyhedral cone in \mathbb{R}_{\max}^3 

Figure 2: memcpy invariant

`src` to `dst`. In program verification, precise invariants over the length of the strings are needed to ensure the absence of string buffer overflows. Recall that the length of a string is defined as the position of the first null character in the string. To precisely analyze the function `memcpy`, two cases have to be distinguished:

- (i) either `n` is strictly smaller than the source length `len_src`, so that only non-null characters are copied into `dst`, hence `len_dst` \geq `n`,
- (ii) or `n` \geq `len_src` and the null terminal character of `src` will be copied into `dst`, thus `len_dst` = `len_src`.

Thanks to tropical polyhedra, the invariant $\min(\text{len_src}, n) = \min(\text{len_dst}, n)$, or equivalently $\max(-\text{len_src}, -n) = \max(-\text{len_dst}, -n)$, can be automatically inferred. It is the *exact* encoding of the disjunction of the cases (i) and (ii). The invariant is represented by the non-convex set of \mathbb{R}^3 depicted in Figure 2. In the application to static analysis, the performance of the algorithm computing the extreme elements of tropical polyhedra plays a crucial role in the scalability of the analyzer (see [AGG08] for further details).

A second motivation arises from the “geometric approach” of max-plus linear discrete event systems [CGQ99], in which the computation of feedbacks ensuring that the state of the system meets a prescribed constraint (for instance that certain waiting times remain bounded) reduces [Kat07] to computing the greatest fixed point of an order preserving map on the set of tropical polyhedra. Similar computations arise when solving dual observability problems [DLGKL09]. Again, the effective handling of these polyhedra turns out to be the bottleneck.

A third motivation arises from the study of mean payoff combinatorial games. In particular, it is shown in [AGG09b] that checking whether a given initial state of a mean payoff game is winning is equivalent to finding a vector in an associated tropical polyhedral cone (with a prescribed finite coordinate). This polyhedron consists of the super-fixed points of the dynamic programming operator (potentials), which certify that the game is winning.

3. Characterizing extremality from inequality constraints

3.1. Preliminaries on extremality

The following lemma, which is a variation on the proof of Th. 3.1 of [GK07] and on Th. 14 of [BSS07], shows that extremality can be expressed as a minimality property:

Proposition 3.1. *Given a polyhedral cone $\mathcal{C} \subset \mathbb{R}_{\max}^d$, \mathbf{g} is extreme if and only if there exists $1 \leq t \leq d$ such that \mathbf{g} is a minimal element of the set $\{\mathbf{x} \in \mathcal{C} \mid \mathbf{x}_t = \mathbf{g}_t\}$, i.e. $\mathbf{g} \in \mathcal{C}$ and for each $\mathbf{x} \in \mathcal{C}$, $\mathbf{x} \leq \mathbf{g}$ and $\mathbf{x}_t = \mathbf{g}_t$ implies $\mathbf{x} = \mathbf{g}$. In that case, \mathbf{g} is said to be extreme of type t .*

In Fig. 3, the light gray area represents the set of the elements $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ of \mathbb{R}_{\max}^3 such that $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) \leq \mathbf{g}^2$ implies $\mathbf{x}_1 < \mathbf{g}_1^2$. It clearly contains the whole cone except \mathbf{g}^2 , which shows that \mathbf{g}^2 is extreme of type 1.

A *tropical segment* is the set of the tropical linear combinations of two points. Using the fact that a tropical segment joining two points of a polyhedral cone \mathcal{C} yields a continuous path included in \mathcal{C} , one can check that \mathbf{g} is extreme of type t in \mathcal{C} if and only if there is a neighborhood N of \mathbf{g} such that \mathbf{g} is minimal in $\{\mathbf{x} \in \mathcal{C} \cap N \mid \mathbf{x}_t = \mathbf{g}_t\}$. Thus, extremality is a local property.

Finally, the extremality of an element \mathbf{g} in a cone \mathcal{C} can be equivalently established by considering the vector formed by its non-0 coordinates. Formally, let $\text{supp}(\mathbf{x}) := \{i \mid \mathbf{x}_i \neq 0\}$ for any $\mathbf{x} \in \mathbb{R}_{\max}^d$. Then \mathbf{g} is extreme in \mathcal{C} if and only if it is extreme in $\{\mathbf{x} \in \mathcal{C} \mid \text{supp}(\mathbf{x}) \subset \text{supp}(\mathbf{g})\}$. This allows to assume that $\text{supp}(\mathbf{g}) = \{1, \dots, d\}$ without loss of generality.

3.2. Expressing extremality using the tangent cone

For now, the polyhedral cone \mathcal{C} is supposed to be defined by a system $A\mathbf{x} \leq B\mathbf{x}$ of n inequalities.

Consider an element \mathbf{g} of the cone \mathcal{C} , which we assume, from the previous discussion, to satisfy $\text{supp}(\mathbf{g}) = \{1, \dots, d\}$. In this context, the *tangent cone* of \mathcal{C} at \mathbf{g} is defined as the tropical polyhedral cone $\mathcal{T}(\mathbf{g}, \mathcal{C})$ of \mathbb{R}_{\max}^d given by the system of inequalities

$$\max_{i \in \arg \max(A_k \mathbf{g})} \mathbf{x}_i \leq \max_{j \in \arg \max(B_k \mathbf{g})} \mathbf{x}_j \quad \text{for all } k \text{ such that } A_k \mathbf{g} = B_k \mathbf{g}, \quad (3.1)$$

where for each row vector $\mathbf{c} \in \mathbb{R}_{\max}^{1 \times d}$, $\arg \max(\mathbf{c}\mathbf{g})$ is defined as the argument of the maximum $\mathbf{c}\mathbf{g} = \max_{1 \leq i \leq d} (\mathbf{c}_i + \mathbf{g}_i)$, and where A_k and B_k denote the k th rows of A and B , respectively.

The tangent cone $\mathcal{T}(\mathbf{g}, \mathcal{C})$ provides a local description of the cone \mathcal{C} around \mathbf{g} :

Proposition 3.2. *There exists a neighborhood N of \mathbf{g} such that for all $\mathbf{x} \in N$, \mathbf{x} belongs to \mathcal{C} if and only if it is an element of $\mathbf{g} + \mathcal{T}(\mathbf{g}, \mathcal{C})$.*

As an illustration, Fig. 4 depicts the set $\mathbf{g}^2 + \mathcal{T}(\mathbf{g}^2, \mathcal{C})$ (in semi-transparent light gray) when \mathcal{C} is the cone given in Fig. 1. Both clearly coincide in the neighborhood of \mathbf{g}^2 .

Since extremality is a local property, it can be equivalently characterized in terms of the tangent cone. Let $\mathbf{1}$ be the element of \mathbb{R}_{\max}^d whose all coordinates are equal to 1.

Proposition 3.3. *The element \mathbf{g} is extreme in \mathcal{C} iff the vector $\mathbf{1}$ is extreme in $\mathcal{T}(\mathbf{g}, \mathcal{C})$.*





Figure 3: Extremality of \mathbf{g}^2 Figure 4: The set $\mathbf{g}^2 + \mathcal{T}(\mathbf{g}^2, \mathcal{C})$ Figure 5: The $\{0, 1\}$ -elements of $\mathcal{T}(\mathbf{g}^2, \mathcal{C})$ 

Figure 6: A directed hypergraph

The problem is now reduced to the characterization of the extremality of the vector $\mathbf{1}$ in a $\{0, 1\}$ -cone, *i.e.* a polyhedral cone defined by a system of the form $C\mathbf{x} \leq D\mathbf{x}$ where $C, D \in \{0, 1\}^{n \times d}$. The following proposition states that only $\{0, 1\}$ -vectors, *i.e.* elements of the tropical regular cube $\{0, 1\}^d$, have to be considered:

Proposition 3.4. *Let $\mathcal{D} \subset \mathbb{R}_{\max}^d$ be a $\{0, 1\}$ -cone. Then $\mathbf{1}$ is extreme of type t if and only if it is the unique element \mathbf{x} of $\mathcal{D} \cap \{0, 1\}^d$ satisfying $\mathbf{x}_t = \mathbf{1}$.*

The following criterion of extremality is a direct consequence of Prop. 3.3 and 3.4:

Theorem 3.5. *Let $\mathcal{C} \subset \mathbb{R}_{\max}^d$ be a polyhedral cone. Then $\mathbf{g} \in \mathcal{C}$ is extreme of type t if and only if the vector $\mathbf{1}$ is the unique $\{0, 1\}$ -element of the tangent cone $\mathcal{T}(\mathbf{g}, \mathcal{C})$ whose t -th coordinate is $\mathbf{1}$.*

Figure 5 shows that in our running example, the $\{0, 1\}$ -elements of $\mathcal{T}(\mathbf{g}^2, \mathcal{C})$ distinct from $\mathbf{1}$ (in squares) all satisfy $\mathbf{x}_1 = 0$. Naturally, testing, by exploration, whether the set of 2^{d-1} $\{0, 1\}$ -elements \mathbf{x} verifying $\mathbf{x}_t = \mathbf{1}$ belonging to $\mathcal{T}(\mathbf{g}, \mathcal{C})$ consists only of $\mathbf{1}$ does not have an acceptable complexity. Instead, the approach of the next section will rely on the equivalent formulation of the criterion of Th. 3.5:

$$\forall l \in \{1, \dots, d\}, [\forall \mathbf{x} \in \mathcal{T}(\mathbf{g}, \mathcal{C}) \cap \{0, 1\}^d, \mathbf{x}_l = 0 \implies \mathbf{x}_t = 0]. \quad (3.2)$$

3.3. Characterizing extremality with directed hypergraphs

A *directed hypergraph* is a couple (N, E) such that each element of E is of the form (T, H) with $T, H \subset N$.

The elements of N and E are respectively called *nodes* and *hyperedges*. Given a hyperedge $e = (T, H) \in E$, the sets T and H represent the *tail* and the *head* of e respectively, and are also denoted by $T(e)$ and $H(e)$. Figure 6 depicts an example of hypergraph whose nodes are u, v, w, x, y, t , and of hyperedges $e_1 = (\{u\}, \{v\})$, $e_2 = (\{v\}, \{w\})$, $e_3 = (\{w\}, \{u\})$, $e_4 = (\{v, w\}, \{x, y\})$, and $e_5 = (\{w, y\}, \{t\})$.

Reachability is extended from digraphs to directed hypergraphs by the following recursive definition: given $u, v \in N$, then v is *reachable from u in \mathcal{H}* , which is denoted $u \rightsquigarrow_{\mathcal{H}} v$, if one of the two conditions holds: $u = v$, or there exists $e \in E$ such that $v \in H(e)$ and all the elements of $T(e)$ are reachable from u . In our example, t is reachable from u .

The size $\text{size}(\mathcal{H})$ of a hypergraph $\mathcal{H} = (N, E)$ is defined as $|N| + \sum_{e \in E} (|T(e)| + |H(e)|)$. In the rest of the paper, directed hypergraphs will be simply referred to as hypergraphs.

We associate to the tangent cone $\mathcal{T}(\mathbf{g}, \mathcal{C})$ the hypergraph $\mathcal{H}(\mathbf{g}, \mathcal{C}) = (N, E)$ defined by:

$$N = \{1, \dots, d\} \quad E = \{(\arg \max(B_k \mathbf{g}), \arg \max(A_k \mathbf{g})) \mid A_k \mathbf{g} = B_k \mathbf{g}, 1 \leq k \leq n\}.$$

The extremality criterion of Eq. (3.2) suggests to evaluate, given an element of $\mathcal{T}(\mathbf{g}, \mathcal{C}) \cap \{0, 1\}^d$, the effect of setting its l -th coordinate to the other coordinates. Suppose that it has been discovered that $\mathbf{x}_l = 0$ implies $\mathbf{x}_{j_1} = \dots = \mathbf{x}_{j_n} = 0$. For any hyperedge e of $\mathcal{H}(\mathbf{g}, \mathcal{C})$ such that $T(e) \subset \{l, j_1, \dots, j_n\}$, \mathbf{x} satisfies: $\max_{i \in H(e)} \mathbf{x}_i \leq \max_{j \in T(e)} \mathbf{x}_j = 0$, so that $\mathbf{x}_i = 0$ for all $i \in H(e)$. Thus, the propagation of the value 0 from the l -th coordinate to other coordinates mimicks the inductive definition of the reachability relation from the node l in $\mathcal{H}(\mathbf{g}, \mathcal{C})$:

Proposition 3.6. *For all $l \in \{1, \dots, d\}$, the statement given between brackets in Eq. (3.2) holds if and only if t is reachable from l in the hypergraph $\mathcal{H}(\mathbf{g}, \mathcal{C})$.*

Hence, the extremality criterion can be restated thanks to some considerations on the strongly connected components of $\mathcal{H}(\mathbf{g}, \mathcal{C})$. The *strongly connected components* (SCCs for short) of a hypergraph \mathcal{H} are the equivalence classes of the equivalence relation $\equiv_{\mathcal{H}}$, defined by $u \equiv_{\mathcal{H}} v$ if $u \rightsquigarrow_{\mathcal{H}} v$ and $v \rightsquigarrow_{\mathcal{H}} u$. They form a partition of the set of nodes of \mathcal{H} . They can be partially ordered by the relation $\preceq_{\mathcal{H}}$, defined by $C_1 \preceq_{\mathcal{H}} C_2$ if C_1 and C_2 admit a representative u and v respectively such that $v \rightsquigarrow_{\mathcal{H}} u$ (beware of the order of v and u in $v \rightsquigarrow_{\mathcal{H}} u$). Then Prop. 3.6 and Th. 3.5 imply the following statement:

Theorem 3.7. *Let $\mathcal{C} \subset \mathbb{R}_{\max}^d$ be a polyhedral cone, and $\mathbf{g} \in \mathcal{C}$. Then \mathbf{g} is extreme if and only if the set of the SCCs of the hypergraph $\mathcal{H}(\mathbf{g}, \mathcal{C})$, partially ordered by $\preceq_{\mathcal{H}(\mathbf{g}, \mathcal{C})}$, admits a least element.*

This theorem is reminiscent of a classical result, showing that a point of a polyhedron defined by inequalities is extreme if and only if the family of gradients of active inequalities at this point is of full rank. Here, the hypergraph encodes precisely the subdifferentials (set of generalized gradients) of the active inequalities but a major difference is that the rank condition must be replaced by the above minimality condition, which is essentially stronger. Indeed, using this theorem, it is shown in [AGK09] that an important class of tropical polyhedra has fewer extreme rays than its classical analogue.

An algorithm due to Gallo *et al.* [GLPN93] shows that one can compute the set of nodes that are reachable from a given node in linear time in an hypergraph. The following result shows that one can in fact compute the minimal SCCs with almost the same complexity. The algorithm is included in the extended version of the present paper [AGG09c]. Although it shows some analogy with the classical Tarjan algorithm, the hypergraph case differs




Figure 7: Intersecting a cone with a halfspace




Figure 8: Intersecting 10 affine hyperplanes in dimension 3

critically from the graph case in that one cannot compute all the SCCs using the same technique.

Theorem 3.8. *The set of minimal SCCs of a hypergraph $\mathcal{H} = (N, E)$ can be computed in time $O(\text{size}(\mathcal{H}) \times \alpha(|N|))$, where α denotes the inverse of the Ackermann function.*

4. The tropical double description method

Our algorithm is based on a successive elimination of inequalities. Given a polyhedral cone \mathcal{C} defined by a system of n constraints, the algorithm computes by induction on k ($0 \leq k \leq n$) a generating set G_k of the intermediate cone defined by the first k constraints. Then G_n forms a generating set of the cone \mathcal{C} . Passing from the set G_k to the set G_{k+1} relies on a result which, given a polyhedral cone \mathcal{K} and a tropical halfspace $\mathcal{H} = \{\mathbf{x} \mid \mathbf{a}\mathbf{x} \leq \mathbf{b}\mathbf{x}\}$, allows to build a generating set G' of $\mathcal{K} \cap \mathcal{H}$ from a generating set G of \mathcal{K} :

Theorem 4.1. *Let \mathcal{K} be a polyhedral cone generated by a set $G \subset \mathbb{R}_{\max}^d$, and $\mathcal{H} = \{\mathbf{x} \mid \mathbf{a}\mathbf{x} \leq \mathbf{b}\mathbf{x}\}$ a tropical halfspace ($\mathbf{a}, \mathbf{b} \in \mathbb{R}_{\max}^{1 \times d}$). Then the polyhedral cone $\mathcal{K} \cap \mathcal{H}$ is generated by the set $\{\mathbf{g} \in G \mid \mathbf{a}\mathbf{g} \leq \mathbf{b}\mathbf{g}\} \cup \{(\mathbf{a}\mathbf{h})\mathbf{g} \oplus (\mathbf{b}\mathbf{g})\mathbf{h} \mid \mathbf{g}, \mathbf{h} \in G, \mathbf{a}\mathbf{g} \leq \mathbf{b}\mathbf{g}, \text{ and } \mathbf{a}\mathbf{h} > \mathbf{b}\mathbf{h}\}$.*

For instance, consider the cone defined in Fig. 1 and the constraint $x_2 \leq x_3 + 2.5$ (depicted in semi-transparent gray in Fig. 7). The three generators \mathbf{g}^1 , \mathbf{g}^2 , and \mathbf{g}^3 satisfy the constraint, while \mathbf{g}^0 does not. Their combinations are the elements $\mathbf{h}^{1,0}$, $\mathbf{h}^{2,0}$, and $\mathbf{h}^{3,0}$ respectively. The resulting algorithm is given in Figure 9. As in the classical case, this inductive approach produces redundant generators, hence, the heart of the algorithm is the extremality test in Line 10. We use here the hypergraph characterization (Theorems 3.7 and 3.8).

Complexity analysis. The complexity of the elementary step of COMPUTEEXTREME, *i.e.* the computation of the elements provided by Th. 4.1 and the elimination of non-extreme ones (Lines 7 to 13), can be precisely characterized to $O(nd\alpha(d)|G|^2)$, where G is the generating set of the last intermediate cone. By comparison, for classical polyhedra, the same step in the refined double description method by Fukuda and Prodon [FP96] takes a time $O(n|G|^3)$. Note that $|G|$ can take values much larger than d .


```

1: procedure COMPUTEEXTREME( $A, B, n$ )  $\triangleright A, B \in \mathbb{R}_{\max}^{n \times d}$ 
2:   if  $n = 0$  then  $\triangleright$  Base case
3:     return the tropical canonical basis  $(\epsilon_i)_{1 \leq i \leq d}$ 
4:   else  $\triangleright$  Inductive case
5:     split  $A\mathbf{x} \leq B\mathbf{x}$  into  $C\mathbf{x} \leq D\mathbf{x}$  and  $\mathbf{a}\mathbf{x} \leq \mathbf{b}\mathbf{x}$ , with  $C, D \in \mathbb{R}_{\max}^{(n-1) \times d}$  and  $\mathbf{a}, \mathbf{b} \in \mathbb{R}_{\max}^{1 \times d}$ 
6:      $G := \text{COMPUTEEXTREME}(C, D, n-1)$ 
7:      $G^{\leq} := \{\mathbf{g}^i \in G \mid \mathbf{a}\mathbf{g}^i \leq \mathbf{b}\mathbf{g}^i\}$ ,  $G^{>} := \{\mathbf{g}^j \in G \mid \mathbf{a}\mathbf{g}^j > \mathbf{b}\mathbf{g}^j\}$ ,  $H := G^{\leq}$ 
8:     for all  $\mathbf{g}^i \in G^{\leq}$  and  $\mathbf{g}^j \in G^{>}$  do
9:        $\mathbf{h} := (\mathbf{a}\mathbf{g}^j)\mathbf{g}^i \oplus (\mathbf{b}\mathbf{g}^i)\mathbf{g}^j$ 
10:      if  $\mathbf{h}$  is extreme in  $\{\mathbf{x} \mid A\mathbf{x} \leq B\mathbf{x}\}$  then
11:        append  $\kappa\mathbf{h}$  to  $H$ , where  $\kappa$  is the opposite of the first non-0 coefficient of  $\mathbf{h}$ 
12:      end
13:    done
14:  end
15:  return  $H$ 
16: end

```

Figure 9: Our main algorithm computing the extreme rays of tropical cones

The overall complexity of the algorithm COMPUTEEXTREME depends on the size of the sets returned in the intermediate steps. In classical geometry, the upper bound theorem of McMullen [McM70] shows that the maximal number of extreme points of a convex polytope in \mathbb{R}^d defined by n inequality constraints is equal to

$$U(n, d) := \binom{n - \lfloor (d+1)/2 \rfloor}{n-d} + \binom{n - \lfloor (d+2)/2 \rfloor}{n-d}.$$

The polars of the *cyclic polytopes* (see [Zie98]) are known to reach this bound. Allamigeon, Gaubert, and Katz [AGK09] showed that a similar bound is valid in the tropical setting.

Theorem 4.2 ([AGK09]). *The number of extreme rays of a tropical cone in $(\mathbb{R} \cup \{-\infty\})^d$ defined as the intersection of n tropical half-spaces cannot exceed $U(n+d, d-1) = O((n+d)^{\lfloor (d-1)/2 \rfloor})$.*

The bound is asymptotically tight for a fixed n , as d tends to infinity, being approached by a tropical generalization of the (polar of) the cyclic polytope [AGK09]. The bound is believed not to be tight for a fixed d , as n tends to infinity. Finding the optimal bound is an open problem. By combining Theorem 4.2, Theorem 3.8, and Theorem 3.7, we readily get the following complexity result, showing that the execution time is smaller in the tropical case than in the classical case, even with the refinements of [FP96].

Proposition 4.3. *The hypergraph implementation of the tropical double description method returns the set of extreme rays of a polyhedral cone defined by n inequalities in dimension d in time $O(n^2 d \alpha(d) G_{\max}^2)$, where G_{\max} is the maximal number of extreme rays of a cone defined by a subsystem of inequalities taken from $A\mathbf{x} \leq B\mathbf{x}$. In particular, the time can be bounded by $O(n^2 d \alpha(d) (n+d)^{d-1})$.*

Alternative approaches. The existing approaches discussed in the introduction have a structure which is similar to COMPUTEEXTREME. However, their implementation in the Maxplus toolbox of SCILAB [CGMQ] and in our previous work [AGG08] relies on a much less efficient elimination of redundant generators. Its principle is the following: an element \mathbf{h} is extreme in the cone generated by a given set H if and only if \mathbf{h} can not be expressed as the tropical linear combination of the elements of H which are not proportional to it. This property can

Table 1: Benchmarks on a single core of a 3 GHz Intel Xeon with 3 Gb RAM

	d	n	# final	# inter.	T (s)	T' (s)	T/T'
rnd100	12	15	32	59	0.24	6.72	0.035
rnd100	15	10	555	292	2.87	321.78	$8.9 \cdot 10^{-3}$
rnd100	15	18	152	211	6.26	899.21	$7.0 \cdot 10^{-3}$
rnd30	17	10	1484	627	15.2	4667.9	$3.3 \cdot 10^{-3}$
rnd10	20	8	5153	1273	49.8	50941.9	$9.7 \cdot 10^{-4}$
rnd10	25	5	3999	808	9.9	12177.0	$8.1 \cdot 10^{-4}$
rnd10	25	10	32699	6670	3015.7	—	—
cyclic	10	20	3296	887	25.8	4957.1	$5.2 \cdot 10^{-3}$
cyclic	15	7	2640	740	8.1	1672.2	$5.2 \cdot 10^{-3}$
cyclic	17	8	4895	1589	44.8	25861.1	$1.7 \cdot 10^{-3}$
cyclic	20	8	28028	5101	690	45 days	$1.8 \cdot 10^{-4}$
cyclic	25	5	25025	1983	62.6	8 days	$9.1 \cdot 10^{-5}$
cyclic	30	5	61880	3804	261	—	—
cyclic	35	5	155040	7695	1232.6	—	—

	# var	# lines	T (s)	T' (s)	T/T'
oddeven8	17	118	7.6	152.1	0.050
oddeven9	19	214	128.0	22101.2	$5.8 \cdot 10^{-3}$
oddeven10	21	240	1049.0	—	—

be checked in $O(d \times |H|)$ time using residuation (see [BSS07] for algorithmic details). In the context of our algorithm, the worst case complexity of the redundancy test is therefore $O(d|G|^2)$, where G is the set of the extreme rays of the last intermediary cone. This is much worse than our method in $O(nd\alpha(d))$ based on directed hypergraphs, since the cardinality of the set G may be exponential in d (see Theorem 4.2). This is also confirmed by our experiments (see below).

We next sketch a different method relying on arrangement of tropical hyperplanes (arrangements of classical hyperplanes yield naive bounds). Indeed, Theorem 3.7 implies that every extreme ray belongs to the intersection of $d - 1$ tropical hyperplanes, obtained by saturating $d - 1$ inequalities among the $n + d$ taken from $Ax \leq Bx$ and $x_i \geq -\infty$, for $i \in [d]$. The max-plus Cramer theorem (see [AGG09a] and the references therein) implies that for generic values of the matrices A, B , every choice of $d - 1$ saturated inequalities yields at most one candidate to be an extreme ray, which can be computed in $O(d^3)$ time. This yields a list of $O((n + d)^{d-1})$ candidates, from which the extreme rays can be extracted by using the present hypergraph characterization (Theorems 3.7 and 3.8), leading to a $O((nd\alpha(d) + d^3)(n + d)^{d-1})$ execution time, which is better than the one of Proposition 4.3 by a factor $n/\alpha(d)$ when $n \gg d$. However, the resulting algorithm is of little practical use, since the worst case execution time is essentially always achieved, whereas the double description method takes advantage of the fact that G_{\max} is in general much smaller than the upper bound of Theorem 4.2 (which is probably not optimal in the case $n \gg d$).

A third approach, along the lines of [DS04, Jos09], would consist in representing tropical polyhedra by polyhedral complexes in the usual sense. However, an inconvenient of polyhedral complexes is that their number of vertices (called “pseudo-vertices” to avoid ambiguities) is exponential in the number of extreme rays [DS04]. Hence, the representations used here are more concise. This is illustrated in Figure 8 (generated using POLYMAKE), which shows an intersection of 10 signed tropical hyperplanes, corresponding to the “natural” pattern studied in [AGK09]. There are only 24 extreme rays, but 1215 pseudo-vertices.

Experiments. Allamigeon has implemented Algorithm COMPUTEEXTREME in OCaml, as part of the “Tropical polyhedral library” (TPLib), <http://penjili.org/tplib.html>. Table 1 reports some experiments for different classes of tropical cones: samples formed by several cones chosen randomly (referred to as rnd x where x is the size of the sample), and signed cyclic cones which are known to have a very large number of extreme elements [AGK09].

The successive columns respectively report the dimension d , the number of constraints n , the size of the final set of extreme rays, the mean size of the intermediary sets, and the execution time T (for samples of “random” cones, we give average results).

The result provided by COMPUTEEXTREME does not depend on the order of the inequalities in the initial system. This order may impact the size of the intermediary sets and subsequently the execution time. In our experiments, inequalities are dynamically ordered during the execution: at each step of the induction, the inequality $\mathbf{ax} \leq \mathbf{bx}$ is chosen so as to minimize the number of combinations $(\mathbf{ag}^j)\mathbf{g}^i \oplus (\mathbf{bg}^i)\mathbf{g}^j$. This strategy reports better results than without ordering.

We compare our algorithm with a variant using the alternative extremality criterion which is discussed in Sect. 4 and used in the other existing implementations [CGMQ, AGG08]. Its execution time T' is given in the seventh column. The ratio T/T' shows that our algorithm brings a huge breakthrough in terms of execution time. When the number of extreme rays is of order of 10^4 , the second algorithm needs several days to terminate. Therefore, the comparison could not be made in practice for some cases.

Table 1 also reports some benchmarks from applications to static analysis. The experiments *oddeven*i** correspond to the static analysis of the odd-even sorting algorithm of i elements. It is a sort of worst case for our analysis. The number of variables and lines in each program is given in the first columns. The analyzer automatically shows that the sorting program returns an array in which the last (resp. first) element is the maximum (minimum) of the array given as input. It clearly benefits from the improvements of COMPUTEEXTREME, as shown by the ratio with the execution time T' of the previous implementation of the static analyzer [AGG08].

References

- [AGG08] X. Allamigeon, S. Gaubert, and É. Goubault. Inferring min and max invariants using max-plus polyhedra. In *SAS'08*, volume 5079 of *LNCS*, pages 189–204. Springer, Valencia, Spain, 2008.
- [AGG09a] M. Akian, S. Gaubert, and A. Guterman. Linear independence over tropical semirings and beyond. In G.L. Litvinov and S.N. Sergeev, editors, *Proc. of the International Conference on Tropical and Idempotent Mathematics*, volume 495 of *Contemp. Math.*, pages 1–38. AMS, 2009.
- [AGG09b] M. Akian, S. Gaubert, and A. Guterman. Tropical polyhedra are equivalent to mean payoff games. arXiv:0912.2462, 2009.
- [AGG09c] X. Allamigeon, S. Gaubert, and E. Goubault. Computing the extreme points of tropical polyhedra. Eprint arXiv:math/0904.3436v2, 2009.
- [AGK09] X. Allamigeon, S. Gaubert, and R. D. Katz. The number of extreme points of tropical polyhedra. Eprint arXiv:math/0906.3492, accepted for publication in *JCTA*, 2009.
- [BH84] P. Butkovič and G. Hegedüs. An elimination method for finding all solutions of the system of linear equations over an extremal algebra. *Ekonomicko-matematicky Obzor*, 20, 1984.
- [BH04] W. Bricc and C. Horvath. \mathbb{B} -convexity. *Optimization*, 53:103–127, 2004.
- [BSS07] P. Butkovič, H. Schneider, and S. Sergeev. Generators, extremals and bases of max cones. *Linear Algebra Appl.*, 421(2-3):394–406, 2007.
- [CC77] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Conference Record of the Fourth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 238–252, Los Angeles, California, 1977. ACM Press, New York, NY.
- [CG79] R. A. Cuninghame-Green. *Minimax algebra*, volume 166 of *Lecture Notes in Economics and Mathematical Systems*. Springer, 1979.
- [CGMQ] G. Cohen, S. Gaubert, M. McGettrick, and J.-P. Quadrat. Maxplus toolbox of SCILAB. Available at <http://minimal.inria.fr/gaubert/maxplustoolbox/>; now integrated in SCICOSLAB. <http://www.scicoslab.org>.

- [CGQ99] G. Cohen, S. Gaubert, and J.P. Quadrat. Max-plus algebra and system theory: where we are and where to go now. *Annual Reviews in Control*, 23:207–219, 1999.
- [CGQ04] G. Cohen, S. Gaubert, and J. P. Quadrat. Duality and separation theorem in idempotent semimodules. *Linear Algebra and Appl.*, 379:395–422, 2004.
- [CH78] P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Conference Record of the Fifth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 84–97, Tucson, Arizona, 1978. ACM Press.
- [DLGKL09] M. Di Loreto, S. Gaubert, R. D. Katz, and J.-J. Loiseau. Duality between invariant spaces for max-plus linear discrete event systems. Eprint arXiv:0901.2915., 2009.
- [DS04] M. Develin and B. Sturmfels. Tropical convexity. *Doc. Math.*, 9:1–27 (electronic), 2004.
- [DY07] M. Develin and J. Yu. Tropical polytopes and cellular resolutions. *Experimental Mathematics*, 16(3):277–292, 2007.
- [FP96] K. Fukuda and A. Prodon. Double description method revisited. In *Selected papers from the 8th Franco-Japanese and 4th Franco-Chinese Conference on Combinatorics and Computer Science*, pages 91–111, London, UK, 1996. Springer.
- [Gau92] S. Gaubert. *Théorie des systèmes linéaires dans les dioïdes*. Thèse, École des Mines de Paris, July 1992.
- [GJ] E. Gawrilow and M. Joswig. POLYMAKE. <http://www.math.tu-berlin.de/polymake/>.
- [GK06] S. Gaubert and R. Katz. Max-plus convex geometry. In R. A. Schmidt, editor, *ReMiCS/AKA 2006*, volume 4136 of *Lecture Notes in Comput. Sci.*, pages 192–206. Springer, 2006.
- [GK07] S. Gaubert and R. Katz. The Minkowski theorem for max-plus convex sets. *Linear Algebra and Appl.*, 421:356–369, 2007.
- [GK09] S. Gaubert and R. Katz. Minimal half-spaces and external representation of tropical polyhedra. Eprint arXiv:math/0908.1586, 2009.
- [GLPN93] G. Gallo, G. Longo, S. Pallottino, and S. Nguyen. Directed hypergraphs and applications. *Discrete Appl. Math.*, 42(2-3):177–201, 1993.
- [Jos05] M. Joswig. Tropical halfspaces. In *Combinatorial and computational geometry*, volume 52 of *Math. Sci. Res. Inst. Publ.*, pages 409–431. Cambridge Univ. Press, Cambridge, 2005.
- [Jos09] M. Joswig. Tropical convex hull computations. In G.L. Litvinov and S.N. Sergeev, editors, *Proc. of the International Conference on Tropical and Idempotent Mathematics*, volume 495 of *Contemp. Math.* AMS, 2009.
- [JSY07] M. Joswig, B. Sturmfels, and J. Yu. Affine buildings and tropical convexity. *Albanian J. Math.*, 1(4):187–211, 2007.
- [Kat07] R. D. Katz. Max-plus (A, B) -invariant spaces and control of timed discrete event systems. *IEEE Trans. Aut. Control*, 52(2):229–241, 2007.
- [LMS01] G.L. Litvinov, V.P. Maslov, and G.B. Shpiz. Idempotent functional analysis: an algebraic approach. *Math. Notes*, 69(5):696–729, 2001.
- [McM70] P. McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.
- [Min01] A. Miné. The octagon abstract domain. In *AST 2001 in WCRE 2001*, IEEE, pages 310–319. IEEE CS Press, October 2001. <http://www.di.ens.fr/~mine/publi/article-mine-ast01.pdf>.
- [MS92] V. Maslov and S. Samborskii, editors. *Idempotent analysis*, volume 13 of *Adv. in Sov. Math.* AMS, RI, 1992.
- [Sin97] I. Singer. *Abstract convex analysis*. Wiley, 1997.
- [Vor67] N.N. Vorobyev. Extremal algebra of positive matrices. *Elektron. Informationsverarbeitung und Kybernetik*, 3:39–71, 1967. in Russian.
- [Zie98] G. M. Ziegler. *Lectures on Polytopes*. Springer, 1998.
- [Zim77] K. Zimmermann. A general separation theorem in extremal algebras. *Ekonom.-Mat. Obzor*, 13(2):179–201, 1977.