



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Journal of Combinatorial Theory, Series A

www.elsevier.com/locate/jcta


Minimal external representations of tropical polyhedra


 Xavier Allamigeon^a, Ricardo D. Katz^b
^a INRIA and CMAP, École Polytechnique, 91128 Palaiseau Cedex, France

^b CONICET, Instituto de Matemática “Beppo Levi”, Universidad Nacional de Rosario, Avenida Pellegrini 250, 2000 Rosario, Argentina

ARTICLE INFO

Article history:

Received 1 June 2012

Available online 11 February 2013

Keywords:

Tropical convexity

Max-plus convexity

Polyhedra

Polytopes

Supporting half-spaces

External representations

Cell complexes

ABSTRACT

Tropical polyhedra are known to be representable externally, as intersections of finitely many tropical half-spaces. However, unlike in the classical case, the extreme rays of their polar cones provide external representations containing in general superfluous half-spaces. In this paper, we prove that any tropical polyhedral cone in \mathbb{R}^n (also known as “tropical polytope” in the literature) admits an essentially unique minimal external representation. The result is obtained by establishing a (partial) anti-exchange property of half-spaces. Moreover, we show that the apices of the half-spaces appearing in such non-redundant external representations are vertices of the cell complex associated with the polyhedral cone. We also establish a necessary condition for a vertex of this cell complex to be the apex of a non-redundant half-space. It is shown that this condition is sufficient for a dense class of polyhedral cones having “generic extremities”.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

Tropical convex geometry consists in the study of the analogues of convex sets in tropical algebra. In this paper, we consider the max-plus semiring \mathbb{R}_{\max} instantiation of tropical algebra, dealing with the set $\mathbb{R} \cup \{-\infty\}$ equipped with $x \oplus y := \max\{x, y\}$ as addition and $x \otimes y := x + y$ as multiplication. Thus, in the max-plus semiring, $-\infty$ is the neutral element for addition and 0 is the neutral element for multiplication. The n -fold product space \mathbb{R}_{\max}^n carries the structure of a semimodule over \mathbb{R}_{\max}

E-mail addresses: xavier.allamigeon@inria.fr (X. Allamigeon), rkatz@fceia.unr.edu.ar (R.D. Katz).




Fig. 1. The real polyhedral cone generated by $v^1 = (0, 1, 3)$, $v^2 = (0, 4, 1)$ and $v^3 = (0, 9, 4)$ (gray), together with a tropical half-space with apex $(0, 6, 1)$ (green). (For interpretation of the references to color, the reader is referred to the web version of this article.)

when equipped with the tropical scalar multiplication $(\lambda, x) \mapsto \lambda x := (\lambda + x_1, \dots, \lambda + x_n)$ and the component-wise tropical addition $(x, y) \mapsto x \oplus y := (\max\{x_1, y_1\}, \dots, \max\{x_n, y_n\})$.

Since any number is “non-negative” in the max-plus semiring (i.e., greater than or equal to the neutral element for addition $-\infty$), following the analogy with ordinary convexity, a subset $\mathcal{C} \subset \mathbb{R}_{\max}^n$ is said to be a *tropical convex set* if

$$\lambda x \oplus \mu y \in \mathcal{C} \quad \text{for all } x, y \in \mathcal{C} \text{ and } \lambda, \mu \in \mathbb{R}_{\max} \text{ such that } \lambda \oplus \mu = 0. \tag{1}$$

Similarly, \mathcal{C} is called a *tropical (convex) cone* when (1) is satisfied without the condition $\lambda \oplus \mu = 0$.

A tropical cone \mathcal{C} is said to be *polyhedral* when it can be generated by finitely many vectors, meaning that there exists $\{v^1, \dots, v^p\} \subset \mathbb{R}_{\max}^n$ such that

$$\mathcal{C} = \{v \in \mathbb{R}_{\max}^n \mid v = \lambda_1 v^1 \oplus \dots \oplus \lambda_p v^p \text{ for some } \lambda_1, \dots, \lambda_p \in \mathbb{R}_{\max}\}. \tag{2}$$

In this paper, we mainly deal with the situation in which all the generators v^i belong to \mathbb{R}^n . In this case, we say that \mathcal{C} is a *real polyhedral cone*.

It is worth mentioning that our terminology differs from the one introduced by Develin and Sturmfels in [14], where tropical cones are called *tropical convex sets*, and real polyhedral cones are referred to as *tropical polytopes*.

Tropical cones are closed under tropical scalar multiplication. In consequence, it turns out to be convenient to identify a real polyhedral cone with its image in the *real projective space*

$$\mathbb{P}^{n-1} := \mathbb{R}^n / (1, \dots, 1)\mathbb{R}.$$

We adopt this approach here and, for visualization purposes, represent a vector $(x_1, \dots, x_n) \in \mathbb{R}^n$ by the point $(x_2 - x_1, \dots, x_n - x_1)$ of \mathbb{R}^{n-1} . For example, the real polyhedral cone generated by $v^1 = (0, 1, 3)$, $v^2 = (0, 4, 1)$ and $v^3 = (0, 9, 4)$ is depicted in Fig. 1. This tropical cone is given by the bounded gray region together with the line segments joining the points v^1 and v^3 to it.

As in classical convexity, any tropical polyhedral cone admits an “external” representation, as the intersection of finitely many tropical half-spaces [20]. A *tropical half-space* can be defined as a set of the form:

$$\mathcal{H} = \left\{x \in \mathbb{R}_{\max}^n \mid \max_{i \in I} \{x_i - \alpha_i\} \geq \max_{j \in J} \{x_j - \alpha_j\}\right\},$$

where I and J are non-empty disjoint subsets of $[n] := \{1, \dots, n\}$ and $\alpha_h \in \mathbb{R}$ for $h \in I \cup J$. We say that \mathcal{H} is *non-degenerate* when $I \cup J = [n]$. In this case, the vector $(\alpha_1, \dots, \alpha_n)$ is an element of \mathbb{R}^n , and

is called the *apex* of \mathcal{H} . Note that apices are defined up to a scalar multiple. Consequently, we think of them as elements of \mathbb{P}^{n-1} . For instance, the tropical half-space $\{x \in \mathbb{R}_{\max}^3 \mid x_3 - 1 \geq \max\{x_1, x_2 - 6\}\}$ is non-degenerate, and its apex is the point $(0, 6, 1)$. It corresponds to the region located above the green dashed half-lines in Fig. 1. We warn the reader that, unless explicitly specified, every tropical half-space considered in the sequel will be non-degenerate.

Tropical convex sets, which were introduced by K. Zimmermann [30] when studying discrete optimization problems, have been the topic of many works coming from different fields. Tropical cones have been studied in idempotent analysis. It came from an observation of Maslov implying the solutions of a Hamilton–Jacobi equation associated with a deterministic optimal control problem belong to structures similar to convex cones, called semimodules or idempotent linear spaces [25,12]. Besides, the invariant spaces that appear in the study of some discrete event systems are naturally equipped with structures of tropical cones, see [10]. This motivated the study of tropical cones or semimodules by Cohen, Gaubert and Quadrat [11,12], following the algebraic approach to discrete event systems initiated by Cohen, Dubois, Quadrat and Viot [9]. Another interest in the tropical analogues of convex sets comes from abstract convex analysis [28], see for instance [13,26]. With the same motivation, the notion of \mathbb{B} -convexity (which is another name for tropical convexity) was introduced and studied by Briec, Horvath and Rubinov [6,7]. The theory of tropical convexity has recently been developed in relation to tropical geometry. Real (tropical) polyhedral cones were considered by Develin and Sturmfels [14]. They developed a combinatorial approach, thinking of these tropical cones as polyhedral complexes in the usual sense. This was at the origin of several works (see for example [23,8,15]) by the same authors and by Joswig, Block and Yu, to quote but a few.

In classical convex geometry, any full-dimensional polyhedron admits a unique minimal external representation, which is provided by the facet-defining half-spaces. No analogues of facets nor faces are currently known for tropical polyhedra, see the work of Develin and Yu [15]. Minimal (inclusion-wise) tropical half-spaces containing a given tropical polyhedral cone have been studied by Joswig [23], Block and Yu [8], and Gaubert and Katz [20]. They can be seen as the tropical counterparts of supporting half-spaces. In [20], the authors show the surprising result that there can exist an infinite number of minimal half-spaces, as soon as $n \geq 4$. In a joint work with Allamigeon [4], they have provided a characterization of the extreme vectors of the polars of tropical polyhedral cones. They pointed out that these vectors do not provide in general minimal representations by half-spaces. They have also introduced a method to eliminate redundant half-spaces via a reduction to a problem in game theory (solving mean payoff games). In particular, it has been observed that the greedy elimination of superfluous half-spaces produces different non-redundant representations, depending on the order in which the half-spaces are considered. As far as we know, the question whether there exists some structure behind the different non-redundant representations of a tropical polyhedral cone has remained open.

The purpose of this work is to study the *minimal external representations*, also called *non-redundant external representations*, of a real polyhedral cone $\mathcal{C} \subset \mathbb{P}^{n-1}$. Without loss of generality, we restrict our attention to external representations composed of tropical half-spaces with apices located in \mathcal{C} . Indeed, as shown in Section 2.3, we can always replace any half-space containing \mathcal{C} by a smaller half-space (inclusion-wise), whose apex belongs to \mathcal{C} . In particular, the apices of minimal half-spaces are elements of \mathcal{C} .

Under this assumption on the apices, we show that a real polyhedral cone has an essentially unique non-redundant external representation. More precisely, we prove the following theorem.

Theorem 1. *For each real polyhedral cone $\mathcal{C} \subset \mathbb{P}^{n-1}$ there exist a subset \mathcal{A} of \mathcal{C} , and for each $a \in \mathcal{A}$, a collection \mathcal{C}_a of disjoint sets of tropical half-spaces with apex a , satisfying the following property:*

A finite set of tropical half-spaces with apices in \mathcal{C} is a non-redundant external representation of \mathcal{C} if, and only if, it is composed of precisely one half-space in each set of the collection \mathcal{C}_a for each $a \in \mathcal{A}$.

As a consequence of this result, any non-redundant external representation (composed of finitely many tropical half-spaces with apices in the cone) precisely involves the same set \mathcal{A} of apices. They are referred to as *non-redundant apices*. Moreover, the multiplicity of each non-redundant apex a (i.e., the number of half-spaces with this apex) is identical in any non-redundant external representation. It

is equal to the cardinality of the collection C_a . Theorem 34 below specifies that the sets in the collection C_a are given by some strongly connected components of a certain directed graph (see Section 4.3 for details). Consequently, two non-redundant external representations only differ on the choice of the representative of each strongly connected component.

This paper is organized as follows. The next section is devoted to recalling basic notions and results concerning tropical convexity. Moreover, given an external representation of a real polyhedral cone, we present a method to replace its half-spaces by half-spaces with apices in the cone (Proposition 7). Note that this method handles arbitrary half-spaces, including degenerate ones.

In Section 3 we establish a combinatorial criterion to determine whether a half-space \mathcal{H} is redundant in a given external representation of a tropical polyhedral cone \mathcal{C} . It is expressed as a certain reachability problem in a directed hypergraph, and plays a fundamental role in the subsequent results. It applies to the case where the half-space \mathcal{H} is non-degenerate, and its apex belongs to \mathcal{C} . In contrast, the half-spaces in the external representation of \mathcal{C} can be arbitrary.

The main result of this paper, Theorem 1 above, is proved in Section 4. Thus, in this section we consider only non-degenerate half-spaces containing a fixed real polyhedral cone \mathcal{C} , and whose apices belong to \mathcal{C} . The proof consists of several steps. Firstly, we show an anti-exchange result which applies to half-spaces with distinct apices (Theorem 21). Secondly, we prove that the set of apices arising in non-redundant external representations is always equal to a certain set \mathcal{A} (Theorem 29). Finally, we fix an apex $a \in \mathcal{A}$, and study which half-spaces with apex a appear in non-redundant representations. This leads to the characterization of the collections C_a (Theorem 34).

Section 5 studies the relationship between non-redundant apices and vertices of the cell complex associated with the cone. Theorem 43 establishes that all the non-redundant apices belong to a particular subset of vertices. We then provide a sufficient condition for a vertex in this subset to be a non-redundant apex of the cone (Theorem 45). Finally, we show (Theorem 51) that this sufficient condition is always satisfied when the cone has “generic extremities”, meaning that each of its extreme vectors belongs to a closed ball of positive radius (for the tropical projective Hilbert metric) contained in the cone.

2. Preliminaries

2.1. Basic notions in tropical convexity

Henceforth, we will use concatenation xy to denote tropical multiplication $x \otimes y$ of two scalars $x, y \in \mathbb{R}_{\max}$. When x, y are vectors of \mathbb{R}_{\max}^n , xy represents the tropical inner product of x and y , i.e.

$$xy := \bigoplus_{i \in [n]} x_i y_i.$$

To emphasize the semiring structure of \mathbb{R}_{\max} , we denote by $\mathbb{0}$ the neutral element for addition, i.e. $\mathbb{0} := -\infty$, and by $\mathbb{1}$ the neutral element for multiplication, i.e. $\mathbb{1} := 0$. The i th (tropical) unit vector will be denoted by e^i , i.e. $e^i \in \mathbb{R}_{\max}^n$ is the vector defined by

$$e^i_j := \mathbb{1} \quad \text{and} \quad e^i_j := \mathbb{0} \quad \text{for } j \neq i.$$

The multiplicative inverse of a non-zero (in the tropical sense) scalar $\lambda \in \mathbb{R}_{\max}$, i.e. $-\lambda$, will be represented by λ^- . When $x \in \mathbb{P}^{n-1}$, we denote by x^- the vector whose coordinates are x_i^- . Given $I \subset [n]$, the vector x_I^- is defined by

$$(x_I^-)_i := x_i^- \quad \text{if } i \in I \quad \text{and} \quad (x_I^-)_i := \mathbb{0} \quad \text{otherwise.}$$

The identification of a real polyhedral cone with its image in the real projective space \mathbb{P}^{n-1} can be generalized to any tropical cone $\mathcal{C} \subset \mathbb{R}_{\max}^n$ provided that we consider the tropical projective space

$$\mathbb{P}_{\max}^{n-1} := (\mathbb{R}_{\max}^n \setminus \{(-\infty, \dots, -\infty)\}) / (1, \dots, 1)\mathbb{R}.$$




Fig. 2. A tropical hyperplane with the corresponding sectors (left) and a tropical half-space (right), both of them with apex $a = (0, 4, 3)$.

We define the tropical projective Hilbert metric over \mathbb{P}^{n-1} by

$$d_H(x, y) := \max_{i \in [n]}(x_i - y_i) - \min_{i \in [n]}(x_i - y_i).$$

It can be extended to \mathbb{P}_{\max}^{n-1} by setting

$$d_H(x, y) := \begin{cases} \max_{y_i \neq 0}(x_i - y_i) - \min_{y_i \neq 0}(x_i - y_i) & \text{when } \{i \mid x_i \neq 0\} = \{i \mid y_i \neq 0\}, \\ +\infty & \text{otherwise.} \end{cases}$$

The sets \mathbb{P}^{n-1} and \mathbb{P}_{\max}^{n-1} are both endowed with the topology induced by the metric d_H . In the sequel, closed balls for this metric will be called (closed) Hilbert balls.

2.1.1. Extreme vectors of tropical cones

A (non-zero) vector x of a tropical cone $\mathcal{C} \subset \mathbb{P}_{\max}^{n-1}$ is said to be extreme in \mathcal{C} if for all $y, z \in \mathcal{C}$,

$$x = y \oplus z \text{ implies either } x = y \text{ or } x = z.$$

The tropical version of Minkowski theorem [17,18] in the case of cones shows that a tropical polyhedral cone \mathcal{C} is generated by a set $V \subset \mathbb{P}_{\max}^{n-1}$ if, and only if, V contains the extreme vectors of \mathcal{C} . Thus, a tropical polyhedral cone \mathcal{C} has a unique minimal generating set (as a subset of \mathbb{P}_{\max}^{n-1}).

2.1.2. Tropical half-spaces and hyperplanes

With the notation introduced above, observe that any non-degenerate half-space can be written in the form

$$\mathcal{H} = \{x \in \mathbb{P}_{\max}^{n-1} \mid a_i^- x \geq a_{[n] \setminus I}^- x\}, \tag{3}$$

where $a \in \mathbb{P}^{n-1}$ and I is a non-empty proper subset of $[n]$. In what follows, we shall also shortly denote such half-space by $\mathcal{H}(a, I)$.

Non-degenerate half-spaces are related to the notion of tropical hyperplanes. The (max-plus) tropical hyperplane with apex a is defined as the set of vectors $x \in \mathbb{P}_{\max}^{n-1}$ such that the maximum

$$a^- x = \max\{x_1 - a_1, \dots, x_n - a_n\}$$

is attained at least twice. The complement of such hyperplane is the disjoint union of n regions, the topological closure of which

$$\mathcal{S}(a, i) := \{x \in \mathbb{P}_{\max}^{n-1} \mid a_i^- x_i \geq a_j^- x_j \text{ for all } j \in [n]\},$$

are special tropical half-spaces called (closed) sectors, see the left-hand side of Fig. 2. Note that the half-space \mathcal{H} in (3) coincides with the union, for $i \in I$, of the sectors $\mathcal{S}(a, i)$ supported by the

hyperplane with apex a . This is illustrated in the right-hand side of Fig. 2. Besides, observe that the apex a and the set of sectors I are both uniquely determined by \mathcal{H} , and so they will be denoted by $\text{apex}(\mathcal{H})$ and $\text{sect}(\mathcal{H})$ respectively. We refer the reader to [23] for more information on hyperplanes and half-spaces, but we warn that the results of [23] are in the setting of the (real) min-plus semiring $(\mathbb{R}, \min, +)$, which is however equivalent to the setting considered here.

2.1.3. Cell decomposition

We now recall basic definitions and properties concerning the natural cell decomposition of \mathbb{P}^{n-1} induced by a finite set of vectors $\{v^r\}_{r \in [p]} \subset \mathbb{P}^{n-1}$. For a complete presentation in the equivalent setting of the (real) min-plus semiring $(\mathbb{R}, \min, +)$, we refer the reader to [14].

Given $x \in \mathbb{P}^{n-1}$, the *type* of x relative to $\{v^r\}_{r \in [p]}$ is the n -tuple $\text{type}(x) = (S_1(x), \dots, S_n(x))$ of subsets of $[p]$ defined as follows:

$$S_j(x) := \{r \in [p] \mid x_j^- v_j^r \geq x_i^- v_i^r \text{ for all } i \in [n]\},$$

for $j \in [n]$. An n -tuple (S_1, \dots, S_n) of subsets of $[p]$ is said to be a *type* if it arises in this way.

With each n -tuple $S = (S_1, \dots, S_n)$ of subsets of $[p]$, it can be associated the set X_S of all the vectors whose type contains S , i.e.

$$X_S := \{x \in \mathbb{P}^{n-1} \mid S_j \subset S_j(x), \text{ for all } j \in [n]\}.$$

Lemma 10 of [14] shows that these sets are given by

$$X_S = \{x \in \mathbb{P}^{n-1} \mid x_j v_i^r \leq x_i v_j^r, \text{ for all } i, j \in [n] \text{ and } r \in S_j\},$$

and so they are both closed convex polyhedra (in the usual sense) and tropical polyhedral cones. The *natural cell decomposition of \mathbb{P}^{n-1} induced by $\{v^r\}_{r \in [p]}$* is defined as the collection of convex polyhedra X_S , where S ranges over all the possible types.

A simple geometric construction of the natural cell decomposition of \mathbb{P}^{n-1} induced by $\{v^r\}_{r \in [p]}$ can be obtained if we consider the min-plus hyperplanes whose apices are these vectors. Recall that given $a \in \mathbb{P}^{n-1}$, the *min-plus hyperplane* with apex a is the set of vectors x such that the minimum

$$\min\{x_1 - a_1, \dots, x_n - a_n\}$$

is attained at least twice. By Proposition 16 of [14], the cell decomposition induced by $\{v^r\}_{r \in [p]}$ is the common refinement of the fans defined by the p min-plus hyperplanes whose apices are the vectors v^r , for $r \in [p]$.

Given a cell X_S , if we define the undirected graph G_S with set of nodes $[n]$ and an arc connecting nodes i and j if, and only if, $S_i \cap S_j \neq \emptyset$, then by Proposition 17 of [14] the dimension of X_S (in the projective space) is one less than the number of connected components of G_S . A zero-dimensional cell is called a *vertex* of the natural cell decomposition.

When \mathcal{C} is the tropical cone generated by $\{v^r\}_{r \in [p]}$, the natural cell decomposition of \mathbb{P}^{n-1} induced by $\{v^r\}_{r \in [p]}$ has in particular the property that \mathcal{C} is the union of its bounded cells, see [14] for details. Corollary 12 of [14] also shows that a cell X_S is bounded if, and only if, $S_j \neq \emptyset$ for all $j \in [n]$. It follows that $x \in \mathcal{C}$ if, and only if, $S_j(x) \neq \emptyset$ for all $j \in [n]$.

Example 2. The natural cell decomposition of \mathbb{P}^2 induced by $v^1 = (0, 1, 3)$, $v^2 = (0, 4, 1)$ and $v^3 = (0, 9, 4)$ is illustrated in Fig. 3. As explained above, it can be obtained by drawing three min-plus hyperplanes (dotted lines in Fig. 3) whose apices are the vectors v^1 , v^2 and v^3 . This cell decomposition consists of six zero-dimensional cells (vertices), fifteen one-dimensional cells (nine unbounded and six bounded) and ten two-dimensional cells (nine unbounded and only one bounded).

Fig. 3 also provides the type (relative to v^1 , v^2 and v^3) of any vector in the relative interior of each bounded cell. For instance, the type of $x = (0, 8, 3)$ is $(\{1, 2\}, \{3\}, \{1, 3\})$, and so this vector is a vertex (the undirected graph G_S , where $S = \text{type}(x)$, is connected). The line segment joining x with v^3 is the cell X_S for $S = (\{1, 2\}, \{3\}, \{3\})$, and the only bounded two-dimensional cell is X_S for $S = (\{2\}, \{3\}, \{1\})$.

Comparing Figs. 1 and 3, it can be seen that the tropical cone generated by v^1 , v^2 and v^3 is precisely the union of the bounded cells in the natural cell decomposition of \mathbb{P}^2 induced by these vectors.




Fig. 3. The natural cell decomposition of \mathbb{P}^2 induced by $v^1 = (0, 1, 3)$, $v^2 = (0, 4, 1)$ and $v^3 = (0, 9, 4)$.

2.2. Tropical polar cones

As in classical convex analysis, the polar \mathcal{C}° of a tropical cone $\mathcal{C} \subset \mathbb{P}_{\max}^{n-1}$ can be defined [19] to represent the set of all (tropical) linear inequalities satisfied by the vectors of \mathcal{C} :

$$\mathcal{C}^\circ := \{(u, u') \in \mathbb{P}_{\max}^{2n-1} \mid ux \geq u'x, \forall x \in \mathcal{C}\}. \tag{4}$$

However, note that tropical linear forms must be considered on both sides of the inequality due to the absence of a “minus sign”. This means that the polar of \mathcal{C} is a tropical cone of \mathbb{P}_{\max}^{2n-1} .

As a consequence of the separation theorem for tropical cones of [30,29,13], a tropical polyhedral cone \mathcal{C} is characterized by its polar cone, i.e.

$$\mathcal{C} = \{x \in \mathbb{P}_{\max}^{n-1} \mid ux \geq u'x, \forall (u, u') \in \mathcal{C}^\circ\}.$$

Moreover, when \mathcal{C} is polyhedral, it can be shown that \mathcal{C}° is also polyhedral, implying \mathcal{C} is the intersection of the (finite) set of tropical half-spaces associated with the extreme vectors of \mathcal{C}° .

An equivalent notion to the polar is that of the j th polar, see [4]. For $j \in [n]$, the j th polar \mathcal{C}_j° of \mathcal{C} is defined as the tropical cone

$$\mathcal{C}_j^\circ := \left\{ u \in \mathbb{P}_{\max}^{n-1} \mid \bigoplus_{i \in [n] \setminus \{j\}} u_i x_i \geq u_j x_j, \forall x \in \mathcal{C} \right\}, \tag{5}$$

which lies in \mathbb{P}_{\max}^{n-1} . As in the case of the polar, a tropical polyhedral cone \mathcal{C} is given by the (finite) intersection of the tropical half-spaces associated with the extreme vectors of \mathcal{C}_j° , for $j \in [n]$. Indeed, the set of extreme vectors of \mathcal{C}° precisely consists of the vectors (e^i, e^i) ($i \in [n]$) and the extreme vectors of the j th polars of \mathcal{C} ($j \in [n]$), see [4].

The extreme vectors of the polars of tropical polyhedral cones have been characterized in different ways, see [20, Theorem 5] or [4, Theorem 3]. We shall need the following variant of Theorem 3 of [4], which is more adapted to our setting.

Theorem 3. Let \mathcal{C} be a real polyhedral cone generated by the set $\{v^r\}_{r \in [p]} \subset \mathbb{P}^{n-1}$, and let $u \in \mathcal{C}_j^\circ$ be such that $u_j \neq 0$.

Then, u is extreme in \mathcal{C}_j° if, and only if, for each $i \neq j$ either $u_i = 0$ or there exists $r \in [p]$ such that $u_i v_i^r = u_j v_j^r > \bigoplus_{k \in [n] \setminus \{i, j\}} u_k v_k^r$.

Observe that the case $u_j = 0$ is not considered in Theorem 3. This is due to the fact that \mathcal{C}_j° contains the unit vectors e^i , for $i \neq j$, and so they are the only extreme vectors u of \mathcal{C}_j° satisfying $u_j = 0$. These extreme vectors of \mathcal{C}_j° will be called *trivial*, because they represent tautological inequalities $x_i \geq 0$, and so they play no role in the external representation of \mathcal{C} .

2.3. Saturation and minimal half-spaces

Let \mathcal{C} be the real polyhedral cone generated by the set $\{v^r\}_{r \in [p]} \subset \mathbb{P}^{n-1}$. A half-space is said to be *minimal* with respect to \mathcal{C} if it is minimal for inclusion among the set of half-spaces containing \mathcal{C} . Gaubert and Katz have proved in [20] that any minimal half-space with respect to \mathcal{C} is non-degenerate, and its apex can be characterized in terms of the natural cell decomposition of \mathbb{P}^{n-1} induced by the generating set $\{v^r\}_{r \in [p]}$.

Theorem 4. (See [20, Theorem 4].) *The half-space $\mathcal{H}(a, I)$ is minimal with respect to the real polyhedral cone \mathcal{C} if, and only if, the following conditions are satisfied:*

- (C1) $\bigcup_{i \in I} S_i(a) = [p]$,
- (C2) for each $j \in [n] \setminus I$ there exists $i \in I$ such that $S_i(a) \cap S_j(a) \neq \emptyset$,
- (C3) for each $i \in I$ there exists $j \in [n] \setminus I$ such that $S_i(a) \cap S_j(a) \not\subset \bigcup_{k \in I \setminus \{i\}} S_k(a)$,

where $(S_1(a), \dots, S_n(a)) = \text{type}(a)$ is the type of a relative to the generating set $\{v^r\}_{r \in [p]}$.

The apices of minimal half-spaces consequently form certain cells of the natural cell decomposition of \mathbb{P}^{n-1} induced by the generators of \mathcal{C} . It was shown in [20] that these cells need not be zero-dimensional, so the number of apices of minimal half-spaces can be infinite. Since Conditions (C2) and (C3) above imply $S_h(a) \neq \emptyset$ for all $h \in [n]$, we readily obtain the following corollary:

Corollary 5. *If \mathcal{H} is a minimal half-space with respect to the real polyhedral cone \mathcal{C} , then its apex belongs to \mathcal{C} .*

Remark 6. The three conditions of Theorem 4 do not depend on the choice of the generating set of \mathcal{C} . For instance, Condition (C1) amounts to $\mathcal{C} \subset \mathcal{H}(a, I)$. Similarly, assuming $\mathcal{C} \subset \mathcal{H}(a, I)$, Condition (C2) is equivalent to the fact that, for each $j \in [n] \setminus I$, there exists $x \in \mathcal{C}$ such that $a_j^- x_j = \bigoplus_{k \in I} a_k^- x_k$. Observe that the latter is trivially satisfied when $a \in \mathcal{C}$.

Given a (possibly degenerate) half-space \mathcal{H} containing \mathcal{C} , there always exists a minimal half-space \mathcal{H}' such that $\mathcal{C} \subset \mathcal{H}' \subset \mathcal{H}$, see [20, Theorem 3]. Using Corollary 5 and the fact that \mathcal{C} is a finite intersection of tropical half-spaces by (the conic form) of the tropical Minkowski–Weyl theorem [20], we conclude that \mathcal{C} is a finite intersection of half-spaces with apices in \mathcal{C} , and these half-spaces can be assumed to be minimal.

Since Theorem 3 of [20] is not constructive, in this section we explain a simple method, referred to as *saturation*, to compute a half-space \mathcal{H}' satisfying $\text{apex}(\mathcal{H}') \in \mathcal{C}$ and $\mathcal{C} \subset \mathcal{H}' \subset \mathcal{H}$. Suppose that $\mathcal{H} = \{x \in \mathbb{P}_{\max}^{n-1} \mid \bigoplus_{i \in I} \alpha_i^- x_i \geq \bigoplus_{j \in J} \alpha_j^- x_j\}$, where I and J are disjoint non-empty subsets of $[n]$ and $\alpha_h \in \mathbb{R}$ for all $h \in I \cup J$. Consider the half-space $\mathcal{H}(b, I')$ whose apex $b = (\beta_1, \dots, \beta_n) \in \mathbb{P}^{n-1}$ and sectors I' are defined as follows:

$$\beta_i := \bigoplus_{r \in [p]} \lambda_r v_i^r \quad \text{for all } i \in [n], \quad I' := \{i \in I \mid \alpha_i = \beta_i\},$$

with $\lambda_r \in \mathbb{R}$ being defined by $\lambda_r := (\bigoplus_{h \in I \cup J} \alpha_h^- v_h^r)^-$. Then, the following proposition holds:




Fig. 4. Saturation of a half-space. (For interpretation of the references to color, the reader is referred to the web version of this article.)

Proposition 7. The half-space $\mathcal{H}(b, I')$ satisfies the following properties:

- (i) its apex b belongs to \mathcal{C} ;
- (ii) $\mathcal{C} \subset \mathcal{H}(b, I') \subset \mathcal{H}$.

Proof. The first property readily follows from $b = \bigoplus_{r \in [p]} \lambda_r v^r$.

On the other hand, since $\alpha_k^- v_k^r \leq \bigoplus_{h \in I \cup J} \alpha_h^- v_h^r = \lambda_r^-$ for all $k \in I \cup J$ and $r \in [p]$, it follows that $\beta_k = \bigoplus_{r \in [p]} \lambda_r v_k^r \leq \alpha_k$ for all $k \in I \cup J$. Moreover, note that $\alpha_k = \beta_k$ if, and only if, there exists $r \in [p]$ such that $\alpha_k^- v_k^r = \bigoplus_{h \in I \cup J} \alpha_h^- v_h^r$.

Consider now any $s \in [p]$. Firstly, observe that there exists $i \in I$ such that $\alpha_i^- v_i^s = \bigoplus_{h \in I \cup J} \alpha_h^- v_h^s$, because $v^s \in \mathcal{C} \subset \mathcal{H}$. Since in that case we have $\alpha_i = \beta_i$ by the discussion above, it follows that $i \in I'$ and so

$$\bigoplus_{h \in I \cup J} \alpha_h^- v_h^s = \bigoplus_{i \in I'} \alpha_i^- v_i^s.$$

Now note that for any $j \in [n] \setminus I'$ we have

$$\beta_j^- v_j^s = \left(\bigoplus_{r \in [p]} \lambda_r v_j^r \right)^- v_j^s \leq \lambda_s^- = \bigoplus_{h \in I \cup J} \alpha_h^- v_h^s = \bigoplus_{i \in I'} \alpha_i^- v_i^s = \bigoplus_{i \in I'} \beta_i^- v_i^s,$$

and thus $v^s \in \mathcal{H}(b, I')$. Since this holds for any $s \in [p]$, we conclude that $\mathcal{C} \subset \mathcal{H}(b, I')$.

Finally, if we assume $\bigoplus_{i \in I'} \beta_i^- x_i \geq \bigoplus_{j \in [n] \setminus I'} \beta_j^- x_j$, it follows that

$$\bigoplus_{i \in I} \alpha_i^- x_i \geq \bigoplus_{i \in I'} \alpha_i^- x_i = \bigoplus_{i \in I'} \beta_i^- x_i \geq \bigoplus_{j \in [n] \setminus I'} \beta_j^- x_j \geq \bigoplus_{j \in J} \beta_j^- x_j \geq \bigoplus_{j \in J} \alpha_j^- x_j,$$

because $I' \subset I$, $J \subset [n] \setminus I'$ and $\beta_h \leq \alpha_h$ for all $h \in I \cup J$, where the equality holds for $h \in I'$. Then, we conclude that $\mathcal{H}(b, I') \subset \mathcal{H}$. \square

Example 8. Consider the cone of Fig. 1, and the half-space $\{x \in \mathbb{P}_{\max}^2 \mid x_1 \oplus x_3 \geq (-8)x_2\}$ with apex $a = (0, 8, 0)$, depicted in orange in Fig. 4. It can be verified that $\lambda_1 = -3$, $\lambda_2 = -1$, and $\lambda_3 = -4$, thus $\beta_1 = (-3)v_1^1 \oplus (-1)v_1^2 \oplus (-4)v_1^3 = -1$. Similarly, $\beta_2 = 5$ and $\beta_3 = 0$, so that $b = (-1, 5, 0)$, and $I' = \{3\}$. The half-space $\mathcal{H}(b, I')$ is represented in green in Fig. 4.

Remark 9. Note that when $I \cup J = [n]$, we have $\lambda_r = \max\{\lambda \in \mathbb{R}_{\max} \mid \lambda v^r \leq a\}$, where $a := (\alpha_1, \dots, \alpha_n)$ is the apex of \mathcal{H} (here \leq refers to the component-wise comparison over vectors of \mathbb{R}^n). Then, in that case, the apex b can be seen as the projection of the vector a onto the cone \mathcal{C} . This projection is known to minimize the tropical projective Hilbert metric, i.e. for all $x \in \mathcal{C}$, $d_H(a, b) \leq d_H(a, x)$, see [11,12] for details.

In general, the half-space $\mathcal{H}(b, I')$ is not minimal with respect to \mathcal{C} . However, we next show that $\mathcal{H}(b, I')$ is minimal in the important special case where \mathcal{H} is the half-space associated with a non-trivial extreme vector of the j th polar of \mathcal{C} ($j \in [n]$).

Proposition 10. Let $\mathcal{H} = \{x \in \mathbb{P}_{\max}^{n-1} \mid \bigoplus_{i \in [n] \setminus \{j\}} u_i x_i \geq u_j x_j\}$ be the half-space associated with a non-trivial extreme vector u of the j th polar of \mathcal{C} . The half-space obtained by saturation of \mathcal{H} is minimal respect to \mathcal{C} , and is of the form $\mathcal{H}(b, I)$ with $b \in \mathbb{P}^{n-1}$ satisfying $u = b^-_I \oplus b^-_j e_j$.

Proof. Let $\mathcal{H}(b, I')$ be the half-space obtained by saturation of \mathcal{H} . Note that using the notation of Proposition 7, we have $\mathcal{H} = \{x \in \mathbb{P}_{\max}^{n-1} \mid \bigoplus_{i \in I} \alpha^-_i x_i \geq \bigoplus_{j \in J} \alpha^-_j x_j\}$ where $J = \{j\}$, $I = \{i \in [n] \mid u_i \neq 0, i \neq j\}$ and $\alpha_h = u_h^-$ for $h \in I \cup J = I \cup \{j\}$.

By Theorem 3, for each $i \in I$ there exists $r \in [p]$ such that

$$u_i v_i^r = u_j v_j^r > \bigoplus_{k \in I \setminus \{i\}} u_k v_k^r. \tag{6}$$

As we have seen in the proof of Proposition 7, this implies $\beta_h = \alpha_h = u_h^-$ for all $h \in I \cup \{j\}$, and so in particular $I' = I$. Besides, by (6), we obtain that

$$b_i^- v_i^r = b_j^- v_j^r > \bigoplus_{k \in I \setminus \{i\}} b_k^- v_k^r.$$

It follows that both $b_i^- v_i^r$ and $b_j^- v_j^r$ are maximal among the $b_h^- v_h^r$ for $h \in I \cup \{j\}$, and even among the $b_h^- v_h^r$ for $h \in [n]$, since $b_i^- v_i^r \geq b_{[n] \setminus I'}^- v_i^r$ and $I = I'$. Thus $r \in S_i(b) \cap S_j(b)$. However, $r \notin \bigcup_{k \in I \setminus \{i\}} S_k(b)$, and so Condition (C3) holds for $\mathcal{H}(b, I')$.

Moreover, by Remark 6, Conditions (C1) and (C2) are satisfied as $\mathcal{C} \subset \mathcal{H}(b, I')$ and $b \in \mathcal{C}$. Therefore, $\mathcal{H}(b, I')$ is a minimal half-space with respect to \mathcal{C} . \square

3. A combinatorial criterion to determine whether a half-space is redundant

Let Γ be a set of (possibly degenerate) half-spaces. A half-space \mathcal{H} is said to be *redundant* with respect to Γ if \mathcal{H} is implied by the half-spaces in Γ , meaning that their intersection $\bigcap_{\mathcal{H}' \in \Gamma} \mathcal{H}'$ is contained in \mathcal{H} .

In this section, we show that the redundancy of a non-degenerate half-space \mathcal{H} with respect to Γ is a local property when the apex of \mathcal{H} is assumed to belong to all the half-spaces in Γ . As a consequence, under the same assumption, we show that the redundancy of a half-space in a finite set of half-spaces is equivalent to a reachability problem in directed hypergraphs.

Proposition–Definition 11. Let Γ be a set of (possibly degenerate) half-spaces, and \mathcal{H} a non-degenerate half-space whose apex belongs to each half-space in Γ . Then, \mathcal{H} is redundant with respect to Γ if, and only if, there exists a neighborhood \mathcal{N} of apex(\mathcal{H}) such that $(\bigcap_{\mathcal{H}' \in \Gamma} \mathcal{H}') \cap \mathcal{N} \subset \mathcal{H}$.

In the latter case, \mathcal{H} is said to be *locally redundant* with respect to Γ .

Proof. The “only if” part is obvious.

To prove the “if” part, let $a := \text{apex}(\mathcal{H})$, $I := \text{sect}(\mathcal{H})$ and $\mathcal{D} := \bigcap_{\mathcal{H}' \in \Gamma} \mathcal{H}'$. Assume there exists a neighborhood \mathcal{N} of a such that $\mathcal{D} \cap \mathcal{N} \subset \mathcal{H}$, but \mathcal{H} is non-redundant in Γ , i.e. $\mathcal{D} \not\subset \mathcal{H}$. Then, pick any $x \in \mathcal{D} \setminus \mathcal{H}$ and let $j \in [n] \setminus I$ be such that $a_j^- x_j > a_i^- x_i$ for all $i \in I$. Define λ as the maximal

scalar such that $\lambda x_i \leq a_i$ for all $i \in [n]$. Let us denote by R the (non-empty) set of the coordinates r such that $\lambda x_r = a_r$. Note that for any $i \in I$,

$$\lambda x_i < \lambda a_j^- x_j a_i \leq a_i,$$

and so $R \cap I = \emptyset$.

Now, define $y := a \oplus \mu x$, where $\mu > \lambda$. Due to the definition of R , if we take μ close enough to λ , we have

$$y_r > a_r \iff r \in R.$$

Then, since $R \cap I = \emptyset$, it follows that $y_i = a_i$ for all $i \in I$. As a consequence, $y \notin \mathcal{H}(a, I)$ while $y \in \mathcal{D}$ (because y is a tropical linear combination of $a, x \in \mathcal{D}$ and \mathcal{D} is a tropical cone). However, this contradicts the fact that $\mathcal{D} \cap \mathcal{N} \subset \mathcal{H}(a, I)$, because $y \in \mathcal{N}$ for μ close enough to λ . \square

To exploit the local characterization of Proposition 11, we use the notion of tangent cone [1,2]. Given a vector $z \in \mathbb{P}^{n-1}$ of a tropical polyhedral cone $\mathcal{D} \subset \mathbb{P}_{\max}^{n-1}$, the tangent cone of \mathcal{D} at z provides a description of \mathcal{D} in a neighborhood of z . We say that a tropical half-space $\{x \in \mathbb{P}_{\max}^{n-1} \mid \bigoplus_{i \in I} \alpha_i^- x_i \geq \bigoplus_{j \in J} \alpha_j^- x_j\}$ is active at z if the following equality holds:

$$\bigoplus_{i \in I} \alpha_i^- z_i = \bigoplus_{j \in J} \alpha_j^- z_j.$$

Definition 12. Let $\mathcal{D} = \bigcap_{\mathcal{H} \in \Gamma} \mathcal{H} \subset \mathbb{P}_{\max}^{n-1}$, where Γ is a finite set of (possibly degenerate) half-spaces, and let $z \in \mathcal{D} \cap \mathbb{P}^{n-1}$. With each half-space $\mathcal{H} = \{x \in \mathbb{P}_{\max}^{n-1} \mid \bigoplus_{i \in I} \alpha_i^- x_i \geq \bigoplus_{j \in J} \alpha_j^- x_j\}$ in Γ active at z , we associate the inequality

$$\max_{i \in M} y_i \geq \max_{j \in N} y_j, \tag{7}$$

where M and N are respectively the argument of the maxima $\bigoplus_{i \in I} \alpha_i^- z_i$ and $\bigoplus_{j \in J} \alpha_j^- z_j$.

Then, the tangent cone $\mathcal{T}(\mathcal{D}, z)$ of \mathcal{D} at z is given by the set of vectors $y \in \mathbb{P}^{n-1}$ satisfying all the inequalities of the form (7) associated with the (active) half-spaces in Γ .

The term *tangent cone* refers to the usual terminology used in optimization and convex analysis. In particular, the term *cone* refers here to the property that for all $y \in \mathcal{T}(\mathcal{D}, z)$ and $\lambda > 0$, the vector $\lambda \times y$ belongs to the set $\mathcal{T}(\mathcal{D}, z)$.

Proposition 13. (See [2].) Let $z \in \mathcal{D} \cap \mathbb{P}^{n-1}$, where \mathcal{D} is a tropical polyhedral cone. There exists a neighborhood \mathcal{N} of z such that for all $x \in \mathcal{N}$, $x \in \mathcal{D}$ if, and only if, $x \in z + \mathcal{T}(\mathcal{D}, z)$.

We now introduce an equivalent encoding of tangent cones in terms of directed hypergraphs. Recall that directed hypergraphs are generalizations of directed graphs, in which the tail and the head of arcs may consist of several nodes. More precisely, a directed hypergraph on the node set $[n] = \{1, \dots, n\}$ consists of a set of hyperarcs, each of which is of the form (T, H) , where $T, H \subset [n]$.

Reachability can be naturally extended to directed hypergraphs as follows. Given a directed hypergraph \mathcal{G} on the node set $[n]$, a node $j \in [n]$ is reachable from a set of nodes $I \subset [n]$ if one of the following two conditions holds:

- (i) j belongs to I ,
- (ii) or there is a hyperarc (T, H) in \mathcal{G} such that $j \in H$, and every $t \in T$ is reachable from I .

By extension, given two sets of nodes $I, J \subset [n]$, J is reachable from I if each node in J is reachable from I . Equivalently, J is reachable from I if there exists a hyperpath from I to J , i.e. a sequence $(T_1, H_1), \dots, (T_q, H_q)$ of hyperarcs of \mathcal{G} such that:




Fig. 5. A directed hypergraph.

$$T_i \subset \bigcup_{0 \leq j \leq i-1} H_j \quad \text{for all } i \in [q + 1],$$

with the convention $H_0 = I$ and $T_{q+1} = J$.

Remark 14. Given $I \subset [n]$, the set of the subsets of $[n]$ reachable from I admits a greatest element $R \subset [n]$, composed of all the nodes $j \in [n]$ reachable from I .

A directed hypergraph consequently provides a concise representation, in terms of a set of hyperarcs, of a possibly large set of relations between subsets of $[n]$. This representation also allows to efficiently determine the relation between two subsets. Indeed, the reachability from I to J can be determined in linear time in the size $\sum_{(T,H) \in \mathcal{G}} (|T| + |H|)$ of the hypergraph, see for instance [21].

Example 15. We provide an example of directed hypergraph on the node set $\{1, \dots, 7\}$ in Fig. 5. It consists of the hyperarcs $(\{1\}, \{2\})$, $(\{2, 3\}, \{4, 5\})$, and $(\{5, 6\}, \{7\})$. Each hyperarc is represented as a bundle of arrows decorated by a solid disk sector. For instance, nodes 4 and 5 are both reachable from the set $\{1, 3\}$, through a hyperpath formed by the first hyperarc (which leads to node 2) and the second one. Similarly, the greatest set reachable from $\{1, 3, 6\}$ is the whole set of nodes $[7]$.

In our setting, directed hypergraphs are used to represent inequalities of the form (7).

Definition 16. Let Γ be a finite set of (possibly degenerate) half-spaces, and $z \in \mathbb{P}^{n-1}$ such that $z \in \mathcal{H}$ for all \mathcal{H} in Γ . With each half-space $\mathcal{H} = \{x \in \mathbb{P}^{n-1} \mid \bigoplus_{i \in I} \alpha_i^- x_i \geq \bigoplus_{j \in J} \alpha_j^- x_j\}$ in Γ active at z , we associate the hyperarc (M, N) , where

$$M := \arg \max \left(\bigoplus_{i \in I} \alpha_i^- z_i \right) \quad \text{and} \quad N := \arg \max \left(\bigoplus_{j \in J} \alpha_j^- z_j \right).$$

The *tangent directed hypergraph at z induced by Γ* , denoted by $\mathcal{G}(\Gamma, z)$, is the directed hypergraph on the node set $[n]$ whose hyperarcs are the ones associated with the active half-spaces in Γ .

Observe that by definition, $\mathcal{G}(\Gamma, z)$ depends on the set of half-spaces Γ . However, the following proposition shows that the reachability relations in $\mathcal{G}(\Gamma, z)$ only depend on the tropical cone $\mathcal{D} = \bigcap_{\mathcal{H} \in \Gamma} \mathcal{H}$.

Proposition 17. Let $\mathcal{D} \subset \mathbb{P}^{n-1}$ be a tropical cone, and $z \in \mathcal{D} \cap \mathbb{P}^{n-1}$. Assume $\mathcal{D} = \bigcap_{\mathcal{H} \in \Gamma} \mathcal{H}$, where Γ is a finite set of (possibly degenerate) half-spaces. Then, for any $I, J \subset [n]$, the following statements are equivalent:

- (i) J is reachable from I in the directed hypergraph $\mathcal{G}(\Gamma, z)$,
- (ii) the inequality $\max_{i \in I} y_i \geq \max_{j \in J} y_j$ is valid for $\mathcal{T}(\mathcal{D}, z)$, meaning that it is satisfied for any $y \in \mathcal{T}(\mathcal{D}, z)$.

Proof. Assume J is reachable from I in $\mathcal{G}(\Gamma, z)$. By definition, there exists a (possibly empty) hyperpath $(T_1, H_1), \dots, (T_q, H_q)$ from I to J in $\mathcal{G}(\Gamma, z)$, meaning that $T_i \subset \bigcup_{0 \leq l \leq i-1} H_l$ for $i \in [q + 1]$, where $H_0 = I$ and $T_{q+1} = J$. By definition, each hyperarc (T_k, H_k) corresponds to an inequality

$$\max_{i \in T_k} y_i \geq \max_{j \in H_k} y_j$$

which is valid for $\mathcal{T}(\mathcal{D}, z)$. This allows us to prove by induction on k that

$$\max_{i \in I} y_i \geq \max_{j \in H_k} y_j$$

is a valid inequality for $\mathcal{T}(\mathcal{D}, z)$ for $k = 1, \dots, q$. Since $J \subset H_0 \cup \dots \cup H_q$, we conclude that $\max_{i \in I} y_i \geq \max_{j \in J} y_j$ is also valid for $\mathcal{T}(\mathcal{D}, z)$.

Now assume that for all $y \in \mathcal{T}(\mathcal{D}, z)$, the inequality $\max_{i \in I} y_i \geq \max_{j \in J} y_j$ holds. Let R be the biggest subset of $[n]$ reachable from I in $\mathcal{G}(\Gamma, z)$.

Given $\epsilon > 0$, define the vector $y' \in \mathbb{P}^{n-1}$ by $y'_i = 0$ if $i \in R$, and $y'_i = \epsilon$ otherwise. Consider any active half-space $\mathcal{H} = \{x \in \mathbb{P}^{n-1}_{\max} \mid \bigoplus_{i \in I'} \alpha_i^- x_i \geq \bigoplus_{j \in J'} \alpha_j^- x_j\}$ in Γ , and let M, N be as in Definition 16. We claim that y' satisfies the inequality

$$\max_{i \in M} y'_i \geq \max_{i \in N} y'_i$$

associated with \mathcal{H} . If $M \not\subset R$, then it is obviously satisfied. If $M \subset R$, then the set M , and subsequently the set N , are both reachable from I in $\mathcal{G}(\Gamma, z)$. Thus, $M \cup N \subset R$ and

$$\max_{i \in M} y'_i = 0 = \max_{i \in N} y'_i.$$

As this holds for any active half-space \mathcal{H} in Γ , we conclude that y' belongs to the tangent cone $\mathcal{T}(\mathcal{D}, z)$. Since the inequality $\max_{i \in I} y_i \geq \max_{j \in J} y_j$ is valid for $\mathcal{T}(\mathcal{D}, z)$ and $I \subset R$, we have

$$0 = \max_{i \in I} y'_i \geq \max_{j \in J} y'_j,$$

implying $y'_j = 0$ for all $j \in J$. This means that $J \subset R$, and so J is reachable from I in $\mathcal{G}(\Gamma, z)$. \square

We are going to use the reduction to local redundancy to characterize redundancy by means of the tangent hypergraph.

Proposition 18. *Let Γ be a finite set of (possibly degenerate) half-spaces and $\mathcal{H}(a, I)$ a half-space whose apex belongs to every half-space in Γ . Then, $\mathcal{H}(a, I)$ is redundant with respect to Γ if, and only if, $[n]$ is reachable from I in the tangent directed hypergraph $\mathcal{G}(\Gamma, a)$.*

Proof. Let $\mathcal{D} = \bigcap_{\mathcal{H}' \in \Gamma} \mathcal{H}'$. By Proposition 11, $\mathcal{H}(a, I)$ is redundant with respect to Γ if, and only if, there exists a neighborhood \mathcal{N} of a such that $\mathcal{D} \cap \mathcal{N} \subset \mathcal{H}(a, I)$. By Proposition 13, this is equivalent to the fact that

$$\mathcal{T}(\mathcal{D}, a) \cap \mathcal{N}' \subset \left\{ y \in \mathbb{P}^{n-1} \mid \max_{i \in I} y_i \geq \max_{j \in [n] \setminus I} y_j \right\}, \tag{8}$$

for some neighborhood \mathcal{N}' of the vector $(0, \dots, 0)$. Besides, we claim that (8) holds if, and only if,

$$\mathcal{T}(\mathcal{D}, a) \subset \left\{ y \in \mathbb{P}^{n-1} \mid \max_{i \in I} y_i \geq \max_{j \in [n] \setminus I} y_j \right\}.$$

To see this, assume (8) holds, and let $y \in \mathcal{T}(\mathcal{D}, a)$. Then, $\lambda \times y \in \mathcal{T}(\mathcal{D}, a)$ for all $\lambda > 0$, and if λ is sufficiently small, $\lambda \times y \in \mathcal{N}'$. It follows that $\lambda \times y$, and consequently y , satisfies the inequality $\max_{i \in I} y_i \geq \max_{j \in [n] \setminus I} y_j$, proving the claim.

Finally, using the first part of the proof and Proposition 17, we conclude that $\mathcal{H}(a, I)$ is redundant with respect to Γ if, and only if, the set $[n] \setminus I$, or equivalently $[n]$, is reachable from I in $\mathcal{G}(\Gamma, a)$. \square




Fig. 6. Determining the redundancy of the half-space $\mathcal{H}(v^1, \{2\})$ with respect to the half-spaces in orange. (For interpretation of the references to color, the reader is referred to the web version of this article.)

Example 19. The cone introduced in Fig. 1 can be expressed as the intersection of the collection Γ of half-spaces given by the following inequalities:

$$\begin{aligned}
 x_2 &\geq 1 + x_1, \\
 \max(-4 + x_2, \underline{-3 + x_3}) &\geq x_1, \\
 -1 + x_3 &\geq \max(x_1, -6 + x_2), \\
 x_1 &\geq -4 + x_3, \\
 \max(x_1, -8 + x_2) &\geq \underline{-3 + x_3}.
 \end{aligned} \tag{9}$$

These half-spaces are depicted in orange in Fig. 6. We illustrate Proposition 18 by establishing that the half-space $\mathcal{H}(v^1, \{2\})$ (in blue in Fig. 6) is redundant with respect to Γ . Only the first two and last half-spaces of the list are active at z . For each of the corresponding inequalities $\bigoplus_{i \in I} \alpha_i^- x_i \geq \bigoplus_{j \in J} \alpha_j^- x_j$, the terms attaining the maxima $\bigoplus_{i \in I} \alpha_i^- x_i$ and $\bigoplus_{j \in J} \alpha_j^- x_j$ are underlined. The directed hypergraph $\mathcal{G}(\Gamma, v^1)$ consequently consists of the hyperarcs $(\{2\}, \{1\})$, $(\{3\}, \{1\})$, and $(\{1\}, \{3\})$.¹ Node 1 is reachable from $\{2\}$ through the first hyperarc, and then node 3 is accessible through the last one. We conclude that the set $\{1, 2, 3\}$ is indeed reachable from $\{2\}$.

The interest of the criterion of Proposition 18 is not only theoretical, but also algorithmic, since it provides a polynomial-time method to eliminate superfluous half-spaces, assuming their apices belong to the other half-spaces:

Corollary 20. Given a finite set Γ of (possibly degenerate) half-spaces, and a non-degenerate half-space \mathcal{H} such that $\text{apex}(\mathcal{H}) \in \mathcal{H}^l$ for all \mathcal{H}^l in Γ , the redundancy of \mathcal{H} with respect to Γ can be determined in time $O(n|\Gamma|)$.

This result has to be compared with a criterion previously established in [4], and expressed in terms of strategies for mean payoff games. Although the latter criterion applies to any half-space (without any assumption on the apex), it is not known whether it can be evaluated in polynomial time (the corresponding decision problem belongs to the complexity class $\text{NP} \cap \text{coNP}$).

¹ Tangent directed hypergraphs can be computed with the library TPLib [5] (version 1.2 or later).

4. Non-redundant external representation of real polyhedral cones

Throughout this section, $\mathcal{C} \subset \mathbb{P}^{n-1}$ denotes a real polyhedral cone. Thanks to Proposition 7, we now focus on external representations of \mathcal{C} composed of (non-degenerate) half-spaces whose apices belong to \mathcal{C} .

We denote by Σ the set of half-spaces containing \mathcal{C} and with apices in \mathcal{C} , i.e.

$$\Sigma := \{ \mathcal{H} \mid \mathcal{C} \subset \mathcal{H}, \text{apex}(\mathcal{H}) \in \mathcal{C} \}.$$

To study the redundancy of a half-space in a set of half-spaces, it is convenient to introduce the function $\tau : 2^\Sigma \rightarrow 2^\Sigma$ defined by

$$\tau(\Gamma) := \left\{ \mathcal{H} \in \Sigma \mid \bigcap_{\mathcal{H}' \in \Gamma} \mathcal{H}' \subset \mathcal{H} \right\}. \tag{10}$$

This function is a closure operator, meaning that for any $\Gamma, \Lambda \in 2^\Sigma$ the following properties hold:

- (i) $\tau(\emptyset) = \emptyset$,
- (ii) $\Gamma \subset \tau(\Gamma)$,
- (iii) $\Gamma \subset \Lambda$ implies $\tau(\Gamma) \subset \tau(\Lambda)$,
- (iv) $\tau(\tau(\Gamma)) = \tau(\Gamma)$.

With this notation, a half-space $\mathcal{H} \in \Sigma$ is redundant with respect to a set $\Gamma \subset \Sigma$ if, and only if, $\mathcal{H} \in \tau(\Gamma)$ or, equivalently, $\tau(\Gamma) = \tau(\Gamma \cup \{\mathcal{H}\})$. A finite set $\Gamma \subset \Sigma$ will be called a *non-redundant external representation* of \mathcal{C} if

$$\begin{aligned} \mathcal{C} &= \bigcap_{\mathcal{H} \in \Gamma} \mathcal{H} \quad (\text{or equivalently, } \tau(\Gamma) = \Sigma), \quad \text{and} \\ \mathcal{H} &\notin \tau(\Gamma \setminus \{\mathcal{H}\}) \quad \text{for each half-space } \mathcal{H} \in \Gamma. \end{aligned}$$

This section is organized as follows. In Section 4.1, we prove a key result establishing that half-spaces with distinct apices satisfy an anti-exchange property. Section 4.2 deals with non-redundant apices, and Section 4.3 with non-redundant half-spaces with the same apex. These sections bring all the results to establish Theorem 1 in Section 4.4. Section 4.5 is devoted to the particular case of non-redundant external representations of pure cones.

4.1. The partial anti-exchange property

We want to show the following partial anti-exchange property:

Theorem 21. *Let $\Gamma \subset \Sigma$ be a finite set of half-spaces and $\mathcal{H}, \mathcal{H}' \in \Sigma$ with distinct apices. If $\mathcal{H}' \notin \tau(\Gamma)$ and $\mathcal{H}' \in \tau(\{\mathcal{H}\} \cup \Gamma)$, then $\mathcal{H} \notin \tau(\{\mathcal{H}'\} \cup \Gamma)$.*

To prove this theorem, we shall use the following lemma:

Lemma 22. *Let $\Gamma \subset \Sigma$ be a finite set of half-spaces and $\mathcal{H}(a, I) \in \tau(\Gamma)$. Then, for each non-empty subset P of $[n] \setminus I$ there exists a half-space $\mathcal{H}(b, J)$ in Γ such that*

$$b_J^- a = b_{[n] \setminus J}^- a, \quad \arg \max(b_J^- a) \cap P = \emptyset, \quad \text{and} \quad \arg \max(b_{[n] \setminus J}^- a) \cap P \neq \emptyset.$$

Proof. Since $\mathcal{H}(a, I) \in \tau(\Gamma)$ and $a \in \mathcal{C} \subset \mathcal{H}'$ for all $\mathcal{H}' \in \Gamma$, we know by Proposition 18 that any subset of $[n]$ is reachable from I in the tangent directed hypergraph $\mathcal{G}(\Gamma, a)$. In particular, P is reachable from I , thus the hypergraph $\mathcal{G}(\Gamma, a)$ must contain a hyperarc (T, H) such that $H \cap P \neq \emptyset$ and $T \subset [n] \setminus P$ (given a hyperpath $(T_1, H_1), \dots, (T_q, H_q)$ from I to P , it suffices to set $(T, H) = (T_k, H_k)$,

where $k \geq 1$ is the greatest integer such that $P \cap (\bigcup_{l=0}^{k-1} H_l) = \emptyset$, recalling that $H_0 = I$. By definition, this hyperarc is associated with a half-space $\mathcal{H}(b, J)$ in Γ active at a , meaning that

$$b^-_J a = b^-_{[n] \setminus J} a, \quad T = \arg \max(b^-_J a), \quad \text{and} \quad H = \arg \max(b^-_{[n] \setminus J} a).$$

This provides the expected result. \square

Proof of Theorem 21. Let $a := \text{apex}(\mathcal{H})$, $I := \text{sect}(\mathcal{H})$, $b := \text{apex}(\mathcal{H}')$, and $J := \text{sect}(\mathcal{H}')$.

Since $\text{apex}(\mathcal{H}') \in \mathcal{C}$, $\mathcal{H}' \notin \tau(\Gamma)$, and $\mathcal{H}' \in \tau(\Gamma \cup \{\mathcal{H}\})$, by Proposition 18 the set $[n]$ is reachable from J in the hypergraph $\mathcal{G}(\Gamma \cup \{\mathcal{H}\}, b)$, while it is not in the hypergraph $\mathcal{G}(\Gamma, b)$. Consequently, the two hypergraphs are not equal, which proves that the half-space \mathcal{H} necessarily provides a hyperarc in the hypergraph $\mathcal{G}(\Gamma \cup \{\mathcal{H}\}, b)$, i.e. \mathcal{H} is active at b . More precisely, the hypergraph $\mathcal{G}(\Gamma \cup \{\mathcal{H}\}, b)$ is obtained from $\mathcal{G}(\Gamma, b)$ by adding the hyperarc (M, N) , where $M = \arg \max(a^-_J b)$ and $N = \arg \max(a^-_{[n] \setminus J} b)$.

Let R be the biggest subset of $[n]$ reachable from J in $\mathcal{G}(\Gamma, b)$. From the previous discussion, we have $R \subsetneq [n]$. Let P be the complement of R in $[n]$ (note that in particular $P \subset [n] \setminus J$ because $J \subset R$). As $[n]$ is reachable from J in $\mathcal{G}(\Gamma \cup \{\mathcal{H}\}, b)$, we necessarily have $N \cap P \neq \emptyset$ and $M \subset R$ (otherwise, the set P would not be reachable from J in $\mathcal{G}(\Gamma \cup \{\mathcal{H}\}, b)$). Hence,

$$a^-_J b = a^-_{[n] \setminus J} b, \quad \arg \max(a^-_J b) \subset R, \quad \text{and} \quad \arg \max(a^-_{[n] \setminus J} b) \cap P \neq \emptyset. \tag{11}$$

Let $P' := \arg \max(a^-_{[n] \setminus J} b) \cap P$. As $a^-_{[n] \setminus J} b$ is equal to $a^-_J b$, it is also equal to $a^- b$, and so we have $P' \subset \arg \max(a^- b) \cap ([n] \setminus I)$.

We shall prove that $\mathcal{H} \notin \tau(\{\mathcal{H}'\} \cup \Gamma)$ by contradiction, so suppose that $\mathcal{H} \in \tau(\{\mathcal{H}'\} \cup \Gamma)$. Then, as $P' \subset [n] \setminus I$, by Lemma 22 we know that there exists a half-space \mathcal{H}'' in $\{\mathcal{H}'\} \cup \Gamma$, with apex c and sectors K , such that

$$c^-_K a = c^-_{[n] \setminus K} a, \quad \arg \max(c^-_K a) \cap P' = \emptyset, \quad \text{and} \quad \arg \max(c^-_{[n] \setminus K} a) \cap P' \neq \emptyset. \tag{12}$$

Consider an arbitrary element $i \in \arg \max(c^-_{[n] \setminus K} a) \cap P'$. Since $c^-_{[n] \setminus K} a = c^- a$, we have $i \in \arg \max(c^- a) \cap ([n] \setminus K)$.

Suppose that the half-space \mathcal{H}'' coincides with \mathcal{H}' , and so in particular $c = b$. Then, since $i \in \arg \max(c^- a) = \arg \max(b^- a)$, $a^-_i b_i$ is the minimum of $a^-_h b_h$ for $h \in [n]$. But as $i \in P' \subset \arg \max(a^- b)$, $a^-_i b_i$ is also the maximum of $a^-_h b_h$ for $h \in [n]$. This is impossible unless a and b are identical (as elements of \mathbb{P}^{n-1}). As a consequence, the half-space \mathcal{H}'' necessarily belongs to Γ .

Now, since $i \in \arg \max(c^- a)$ and $i \in P' \subset \arg \max(a^- b)$, we have

$$c^-_i b_i = c^-_i a_i a^-_i b_i \geq c^-_h a_h a^-_h b_h = c^-_h b_h \quad \text{for any } h \in [n],$$

and thus $i \in \arg \max(c^- b)$. Then, as $i \in [n] \setminus K$, we conclude that

$$i \in \arg \max(c^-_{[n] \setminus K} b) \quad \text{and} \quad c^-_K b = c^-_{[n] \setminus K} b,$$

because $c^-_K b \geq c^-_{[n] \setminus K} b$ according to the fact that $b \in \mathcal{C} \subset \mathcal{H}''$.

Observe that for all $j \in K$,

$$(c^-_K a)(a^- b) \geq (c^-_j a_j)(a^-_j b_j) = c^-_j b_j.$$

Moreover, the bound $(c^-_K a)(a^- b)$ is the maximum of $c^-_j b_j$ for $j \in K$. Indeed,

$$\bigoplus_{j \in K} c^-_j b_j = c^-_K b = c^-_{[n] \setminus K} b = c^-_i b_i = (c^-_i a_i)(a^-_i b_i),$$

and since $i \in \arg \max(c^-_{[n] \setminus K} a) \cap \arg \max(a^- b)$, we have $c^-_i a_i = c^-_{[n] \setminus K} a = c^-_K a$ and $a^-_i b_i = a^- b$. It follows that

$$\arg \max(c^-_K b) = \arg \max(c^-_K a) \cap \arg \max(a^- b).$$

We are now going to show that $\arg \max(c_K^- b) \subset R$. Given $h \in \arg \max(c_K^- b)$, we either have $h \in \arg \max(a_{[n] \setminus I}^- b)$ or $h \in \arg \max(a_I^- b)$. In the latter case, by (11) we have $h \in \arg \max(a_I^- b) \subset R$. Assume now that $h \in \arg \max(a_{[n] \setminus I}^- b)$. Then, since

$$\arg \max(c_K^- b) \cap \arg \max(a_{[n] \setminus I}^- b) \cap P \subset \arg \max(c_K^- a) \cap P' = \emptyset,$$

it follows that $h \notin P$, i.e. $h \in R$. As a consequence, $\arg \max(c_K^- b) \subset R$.

Finally, since $\mathcal{H}'' \in \Gamma$, and

$$c_K^- b = c_{[n] \setminus K}^- b, \quad \arg \max(c_K^- b) \subset R \quad \text{and} \quad i \in \arg \max(c_{[n] \setminus K}^- b),$$

we conclude that node i is reachable from J in the hypergraph $\mathcal{G}(\Gamma, b)$, i.e. $i \in R$. This contradicts the fact that $i \in P$, and completes the proof of the theorem. \square

We shall need the following corollary of the partial anti-exchange property.

Corollary 23. *Let $\Gamma_1, \Gamma_2 \subset \Sigma$ be two finite sets of half-spaces, and $\mathcal{H} \in \Sigma$ be such that $\text{apex}(\mathcal{H}) \neq \text{apex}(\mathcal{H}')$ for all $\mathcal{H}' \in \Gamma_2$. If $\tau(\{\mathcal{H}\} \cup \Gamma_1) = \tau(\Gamma_2)$, then $\mathcal{H} \in \tau(\Gamma_1)$.*

Proof. Let \mathcal{H}' be any half-space in Γ_2 and define $\Gamma'_2 := \Gamma_2 \setminus \{\mathcal{H}'\}$. Note that:

$$\mathcal{H}' \in \tau(\Gamma_2) = \tau(\{\mathcal{H}\} \cup \Gamma_1) \subset \tau(\{\mathcal{H}\} \cup \Gamma_1 \cup \Gamma'_2),$$

and

$$\mathcal{H} \in \tau(\{\mathcal{H}\} \cup \Gamma_1) = \tau(\Gamma_2) \subset \tau(\{\mathcal{H}'\} \cup \Gamma_1 \cup \Gamma'_2).$$

Since \mathcal{H} and \mathcal{H}' have distinct apices, we conclude by Theorem 21 that $\mathcal{H} \in \tau(\Gamma_1 \cup \Gamma'_2)$. If the set Γ'_2 is non-empty, we can repeat the same argument by choosing a new half-space \mathcal{H}'' in Γ'_2 . Since Γ_2 is a finite set, this completes the proof. \square

4.2. Apices of non-redundant external representations

Note that the boundary $\partial \mathcal{C}$ of \mathcal{C} is precisely the set of apices of half-spaces in Σ :

Lemma 24. *We have $\partial \mathcal{C} = \{\text{apex}(\mathcal{H}) \mid \mathcal{H} \in \Sigma\}$.*

Proof. Since no neighborhood of $\text{apex}(\mathcal{H})$ is contained in \mathcal{H} for any half-space \mathcal{H} , it readily follows that the apex of any half-space in Σ does not belong to the interior of \mathcal{C} .

Conversely, consider $a \in \partial \mathcal{C}$, and assume \mathcal{C} is not contained in any half-space with apex a . Then, for each $i \in [n]$ there exists $x^i \in \mathcal{C}$ such that $a_i^- x_i^i > a_{[n] \setminus \{i\}}^- x^i$. Let $\epsilon \in \mathbb{R}$ be such that $a_i^- x_i^i > \epsilon^- \geq a_{[n] \setminus \{i\}}^- x^i$, and define $y^i := a \oplus \epsilon x^i$. Thus, y^i satisfies $y_j^i > a_j$ and $y_j^i = a_j$ for all $j \neq i$. Now consider the cone \mathcal{N} generated by the vectors y^i for $i \in [n]$. This cone forms a neighborhood of a (it contains the Hilbert ball of center a and radius $\rho = \min_{i \in [n]} a_i^- y_j^i > 0$). Besides, $\mathcal{N} \subset \mathcal{C}$ since $y^i \in \mathcal{C}$ for all $i \in [n]$. Hence, a is in the interior of \mathcal{C} , which is a contradiction. \square

For each $a \in \partial \mathcal{C}$, we denote by Σ_a the set of half-spaces with apex a which contain \mathcal{C} , i.e.

$$\Sigma_a := \{\mathcal{H}(a, I) \mid \mathcal{C} \subset \mathcal{H}(a, I)\}.$$

Obviously, $\Sigma = \bigcup_{a \in \partial \mathcal{C}} \Sigma_a$. Now, define the function $\tau' : 2^{\partial \mathcal{C}} \rightarrow 2^{\partial \mathcal{C}}$ by

$$\tau'(X) := \left\{ a \in \partial \mathcal{C} \mid \Sigma_a \subset \tau \left(\bigcup_{b \in X} \Sigma_b \right) \right\}.$$

Then, as in the case of τ , we have:

Proposition 25. *The function τ' is a closure operator on $\partial\mathcal{C}$.*

Proof. First, $\tau'(\emptyset) = \emptyset$, as a consequence of the fact that $\tau(\emptyset) = \emptyset$. Similarly, for any $X \in 2^{\partial\mathcal{C}}$, we have $X \subset \tau'(X)$ because $\Sigma_a \subset \tau(\bigcup_{b \in X} \Sigma_b)$ for $a \in X$.

Besides, for $X, Y \in 2^{\partial\mathcal{C}}$, we have

$$\begin{aligned} X \subset Y &\implies \bigcup_{a \in X} \Sigma_a \subset \bigcup_{a \in Y} \Sigma_a \implies \tau\left(\bigcup_{a \in X} \Sigma_a\right) \subset \tau\left(\bigcup_{a \in Y} \Sigma_a\right) \\ &\implies \tau'(X) \subset \tau'(Y). \end{aligned}$$

Finally, let us show that $\tau'(\tau'(X)) = \tau'(X)$ for all $X \in 2^{\partial\mathcal{C}}$. If we define

$$\mathcal{D} := \bigcap_{\mathcal{H} \in \Sigma_a, a \in X} \mathcal{H} \quad \text{and} \quad \mathcal{D}' := \bigcap_{\mathcal{H} \in \Sigma_a, a \in \tau'(X)} \mathcal{H},$$

then $\mathcal{D}' \subset \mathcal{D}$ because $X \subset \tau'(X)$. Moreover, for any $\mathcal{H} \in \bigcup_{a \in \tau'(X)} \Sigma_a$, we have $\mathcal{D} \subset \mathcal{H}$ and thus $\mathcal{D} \subset \mathcal{D}'$. Therefore, we conclude that $\mathcal{D} = \mathcal{D}'$. Note that $a \in \tau'(X)$ if, and only if, $\mathcal{D} \subset \mathcal{H}$ for all $\mathcal{H} \in \Sigma_a$. Similarly, $a \in \tau'(\tau'(X))$ is equivalent to $\mathcal{D}' \subset \mathcal{H}$ for all $\mathcal{H} \in \Sigma_a$. This implies $\tau'(\tau'(X)) = \tau'(X)$. \square

Unlike τ , the closure operator τ' satisfies the anti-exchange property.

Proposition 26. *Let X be a finite subset of $\partial\mathcal{C}$, and $a, b \in \partial\mathcal{C}$ two distinct elements of \mathbb{P}^{n-1} . Then, $b \notin \tau'(X)$ and $b \in \tau'(X \cup \{a\})$ imply $a \notin \tau'(X \cup \{b\})$.*

Proof. Suppose that $a \in \tau'(X \cup \{b\})$. Then, if we define $\Gamma := \bigcup_{c \in X} \Sigma_c$, we have

$$\tau(\Gamma \cup \Sigma_a) = \tau(\Gamma \cup \Sigma_a \cup \Sigma_b) = \tau(\Gamma \cup \Sigma_b).$$

Note that b is distinct from the apices of half-spaces in $\Gamma = \bigcup_{c \in X} \Sigma_c$ because $b \notin \tau'(X)$. Therefore, by successive applications of Corollary 23 to the half-spaces in Σ_b , we conclude that these half-spaces belong to $\tau(\Gamma)$. This implies $b \in \tau'(X)$, which is a contradiction. \square

As a consequence of the previous two propositions, we obtain:

Corollary 27. *The pair $(\partial\mathcal{C}, \tau')$ is a convex geometry.*

Recall that $X \subset \partial\mathcal{C}$ is said to be a *spanning set* of $\partial\mathcal{C}$ if $\tau'(X) = \partial\mathcal{C}$. When the ground set G of a convex geometry (G, τ') is finite, it is known that G has a unique minimal spanning set, see for example [24]. This minimal spanning set is composed of the *extreme elements* of G , which are the elements $a \in G$ such that $a \notin \tau'(G \setminus \{a\})$. Even if in our case the ground set $\partial\mathcal{C}$ is infinite, we next show that it also admits a unique minimal finite spanning set.

Corollary 28. *There exists a unique minimal finite subset \mathcal{A} of $\partial\mathcal{C}$ satisfying $\tau'(\mathcal{A}) = \partial\mathcal{C}$.*

Proof. In the first place, observe that there exists a finite spanning set X of $\partial\mathcal{C}$. Indeed, as \mathcal{C} is a real polyhedral cone, there exists a finite set of half-spaces Γ , whose apices belong to \mathcal{C} , such that $\mathcal{C} = \bigcap_{\mathcal{H} \in \Gamma} \mathcal{H}$, see Section 2.3. Then, we have $\tau'(\{\text{apex}(\mathcal{H}) \mid \mathcal{H} \in \Gamma\}) = \partial\mathcal{C}$.

Assume now that X and Y are two distinct minimal finite spanning sets of $\partial\mathcal{C}$, and let $a \in X \setminus Y$. Let $\Gamma_1 := \bigcup_{b \in X \setminus \{a\}} \Sigma_b$ and $\Gamma_2 := \bigcup_{b \in Y} \Sigma_b$. Then, since $\tau'(X) = \tau'(Y) = \partial\mathcal{C}$, we have

$$\tau(\Sigma_a \cup \Gamma_1) = \tau(\Gamma_2) = \Sigma.$$

Now, as $a \notin Y$, we can repeatedly apply Corollary 23 to the half-spaces in Σ_a to conclude that $\Sigma_a \subset \tau(\Gamma_1)$, and so

$$\tau(\Gamma_1) = \tau(\Sigma_a \cup \Gamma_1) = \Sigma.$$

Therefore, $\tau'(X \setminus \{a\}) = \partial\mathcal{C}$ contradicting the fact that X is a minimal spanning set of $\partial\mathcal{C}$. \square

We can now establish the main theorem of this subsection, which shows that the set \mathcal{A} precisely characterizes the apices of the half-spaces in any finite non-redundant external representation of the cone \mathcal{C} . As indicated in the introduction, such apices will be referred to as *non-redundant apices*.

Theorem 29. *Let Γ be any non-redundant external representation of \mathcal{C} (composed of finitely many half-spaces with apices in \mathcal{C}). Then, $\mathcal{A} = \{\text{apex}(\mathcal{H}) \mid \mathcal{H} \in \Gamma\}$.*

Proof. Since $\bigcap_{\mathcal{H} \in \Gamma} \mathcal{H} = \mathcal{C}$, we have $\tau'(\{\text{apex}(\mathcal{H}) \mid \mathcal{H} \in \Gamma\}) = \partial\mathcal{C}$. So, by Corollary 28,

$$\mathcal{A} \subset \{\text{apex}(\mathcal{H}) \mid \mathcal{H} \in \Gamma\}.$$

Now suppose that for some $\mathcal{H}' \in \Gamma$, $\text{apex}(\mathcal{H}') \notin \mathcal{A}$. Since $\tau(\Gamma) = \Sigma = \tau(\bigcup_{\mathcal{A} \in \mathcal{A}} \Sigma_a)$, by Corollary 23 it follows that $\mathcal{H}' \in \tau(\Gamma \setminus \{\mathcal{H}'\})$. This contradicts the fact that Γ is a non-redundant external representation of \mathcal{C} . \square

4.3. Non-redundant half-spaces with the same apex

We now study those half-spaces which have the same apex $a \in \partial\mathcal{C}$ in non-redundant external representations of \mathcal{C} . With this aim, assume \mathcal{C} is given by the intersection of half-spaces $\{\mathcal{H}(a, I_l)\}_{l \in [q]} \subset \Sigma_a$ with apex a , and a tropical cone

$$\mathcal{D} := \bigcap_{\mathcal{H}' \in \Lambda} \mathcal{H}',$$

where $\Lambda \subset \Sigma \setminus \Sigma_a$ is a finite set of half-spaces whose apices are distinct from a . We want to characterize the minimal subsets L of $[q]$ satisfying:

$$\tau(\Lambda \cup \{\mathcal{H}(a, I_l)\}_{l \in L}) = \tau(\Lambda \cup \{\mathcal{H}(a, I_l)\}_{l \in [q]}) = \Sigma. \tag{13}$$

Observe that a is a non-redundant apex if, and only if, such minimal subsets are non-empty. In principle, these subsets depend on the half-spaces composing the set Λ . However, we next show that indeed this is not the case.

Proposition 30. *Let Λ_1 and Λ_2 be two finite sets of half-spaces in Σ , whose apices are distinct from a , such that*

$$\tau(\Lambda_1 \cup \{\mathcal{H}(a, I_l)\}_{l \in [q]}) = \tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{l \in [q]}). \tag{14}$$

Then, L is a minimal subset of $[q]$ satisfying

$$\tau(\Lambda_1 \cup \{\mathcal{H}(a, I_l)\}_{l \in L}) = \tau(\Lambda_1 \cup \{\mathcal{H}(a, I_l)\}_{l \in [q]}) \tag{15}$$

if, and only if, L is a minimal subset of $[p]$ satisfying

$$\tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{l \in L}) = \tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{l \in [q]}). \tag{16}$$

Proof. Observe that to prove the proposition, it is enough to show that a subset L of $[q]$ satisfies (15) only if it satisfies (16).

By the contrary, suppose that (15) is satisfied by some $L \subset [q]$ but (16) is not. In that case, we can always define a subset L' of $[q]$ such that $L \subsetneq L'$, (16) is satisfied with L' instead of L , but

$$\tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L' \setminus \{r\}}) \subsetneq \tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in [q]})$$

for some $r \in L' \setminus L$. Then, by (15) and the fact that $L \subset L' \setminus \{r\}$, we obtain

$$\mathcal{H}(a, I_r) \in \tau(\Lambda_1 \cup \Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L' \setminus \{r\}}).$$

Moreover, given $\mathcal{H}' \in \Lambda_1$, we have

$$\mathcal{H}' \in \tau((\Lambda_1 \setminus \{\mathcal{H}'\}) \cup \Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L'})$$

by using (14) and the fact that (16) is satisfied with L' instead of L . As $\mathcal{H}(a, I_r)$ and \mathcal{H}' have distinct apices, by Theorem 21, it follows that

$$\mathcal{H}(a, I_r) \in \tau((\Lambda_1 \setminus \{\mathcal{H}'\}) \cup \Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L' \setminus \{r\}}).$$

Repeating this argument, we conclude that $\mathcal{H}(a, I_r) \in \tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L' \setminus \{r\}})$. However, this is a contradiction, because it would imply

$$\tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L' \setminus \{r\}}) = \tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L'})$$

and so (16) would be satisfied with $L' \setminus \{r\}$ instead of L . \square

We now introduce a directed graph $\mathcal{R}(\mathcal{D}, a)$ defined as follows:

- (i) its nodes are the elements of the set

$$E_a := \{[n]\} \cup \{I \subset [n] \mid \mathcal{H}(a, I) \in \Sigma_a\} = \{[n]\} \cup \{I \subset [n] \mid \mathcal{C} \subset \mathcal{H}(a, I)\},$$

- (ii) there is an arc from I to J if, and only if, J is reachable from I in the tangent directed hypergraph $\mathcal{G}(\Lambda, a)$ at a induced by Λ .

Note that, by Proposition 17, the graph $\mathcal{R}(\mathcal{D}, a)$ does not depend on the choice of the set Λ of half-spaces representing the cone \mathcal{D} .

The following proposition shows that the redundancy of half-spaces with apex a can be characterized using $\mathcal{R}(\mathcal{D}, a)$.

Proposition 31. *Let $\{\mathcal{H}(a, I_l)\}_{I_l \in L}$ be a non-empty subset of Σ_a , and $\mathcal{H}(a, J) \in \Sigma_a$. Then,*

$$\mathcal{H}(a, J) \in \tau(\Lambda \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L}) \tag{17}$$

if, and only if, for some $r \in L$, I_r is reachable from J in the directed graph $\mathcal{R}(\mathcal{D}, a)$.

Proof. In the first place, observe that the tangent directed hypergraph

$$\mathcal{G}_L := \mathcal{G}(\Lambda \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L}, a)$$

is obtained by adding hyperarcs connecting I_l with $[n] \setminus I_l$, for $l \in L$, to the tangent directed hypergraph $\mathcal{G}(\Lambda, a)$.

Assume I_r is reachable from J in $\mathcal{R}(\mathcal{D}, a)$ for some $r \in L$. Then, I_r is also reachable from J in the hypergraph \mathcal{G}_L . Since \mathcal{G}_L contains a hyperarc connecting I_r with $[n] \setminus I_r$, we conclude that $[n]$ is reachable from J in \mathcal{G}_L . Therefore, by Proposition 18, it follows that (17) holds.

Assume now that (17) is satisfied. Then, by Proposition 18 we know that $[n]$ is reachable from J in \mathcal{G}_L . Consider a hyperpath connecting J with $[n]$ in \mathcal{G}_L . It is convenient to split the rest of the proof into two cases.

If one of the hyperarcs in the hyperpath connecting J with $[n]$ is associated with a half-space in $\{\mathcal{H}(a, I_l)\}_{l \in L}$, let $\mathcal{H}(a, I_r)$ be the half-space corresponding to the first occurrence of such a hyperarc in the hyperpath. Then, I_r is reachable from J in $\mathcal{G}(\Lambda, a)$, and consequently in $\mathcal{R}(\mathcal{D}, a)$.

If no hyperarc in the considered hyperpath is associated with a half-space in $\{\mathcal{H}(a, I_l)\}_{l \in L}$, we conclude that $[n]$ is reachable from J in $\mathcal{G}(\Lambda, a)$. Therefore, any subset of $[n]$ is reachable from J in $\mathcal{G}(\Lambda, a)$, and so any node of $\mathcal{R}(\mathcal{D}, a)$ is reachable from J in $\mathcal{R}(\mathcal{D}, a)$. This completes the proof of the proposition. \square

Definition 32. Two half-spaces $\mathcal{H}, \mathcal{H}' \in \Sigma$ are said to be *mutually redundant* with respect to $\Gamma \subset \Sigma$ if $\mathcal{H} \in \tau(\Gamma \cup \{\mathcal{H}'\})$ and $\mathcal{H}' \in \tau(\Gamma \cup \{\mathcal{H}\})$.

As an immediate corollary of Proposition 31, we obtain:

Corollary 33. *The half-spaces $\mathcal{H}(a, I)$ and $\mathcal{H}(a, J)$ are mutually redundant with respect to Λ if, and only if, I and J belong to the same strongly connected component of $\mathcal{R}(\mathcal{D}, a)$.*

The reachability relation associated with the directed graph $\mathcal{R}(\mathcal{D}, a)$ naturally induces a pre-order \preceq on the elements of E_a , i.e. $I \preceq J$ if, and only if, J is reachable from I . Considering the equivalence relation $I \sim J$ defined by $I \preceq J \preceq I$, the pre-order can be turned into a partial order (still denoted \preceq by abuse of notation) over the quotient set E_a/\sim formed by the strongly connected components of $\mathcal{R}(\mathcal{D}, a)$.

The abstract structure of half-spaces with the same apex a is thus in relation to a poset convex geometry. A *poset convex geometry* is a pair (G, σ) , where G is a ground set, and $\sigma : 2^G \rightarrow 2^G$ is the closure operator defined as

$$\sigma(X) := \{y \in G \mid y \preceq x \text{ for some } x \in X\}$$

for all $X \in 2^G$. Poset convex geometries arise from poset antimatroids [24], in the sense that the closed elements of a poset convex geometry are precisely the complements of the feasible elements of a poset antimatroid.

In our case, the poset convex geometry is associated with the partially ordered set formed by the strongly connected components of $\mathcal{R}(\mathcal{D}, a)$. We can then verify that the quotient set E_a/\sim has a unique minimal spanning set, consisting of the strongly connected components which are maximal for the order \preceq . This leads to the following characterization:

Theorem 34. *The following two properties hold:*

- (i) *The apex a is non-redundant if, and only if, the directed graph $\mathcal{R}(\mathcal{D}, a)$ is not strongly connected.*
- (ii) *When a is a non-redundant apex, L is a minimal subset of $[q]$ satisfying (13) if, and only if, $\{I_l\}_{l \in L}$ is composed of precisely one element of each maximal strongly connected component of $\mathcal{R}(\mathcal{D}, a)$.*

Proof. (i) Assume the directed graph $\mathcal{R}(\mathcal{D}, a)$ is strongly connected. Then, for each $I \in E_a$, node $[n]$ is reachable from I in $\mathcal{R}(\mathcal{D}, a)$, and consequently in the directed hypergraph $\mathcal{G}(\Lambda, a)$. By Proposition 18, we deduce that $\mathcal{H}(a, I)$ is redundant with respect to Λ for each $I \in [q]$. We conclude that Λ is an external representation of the cone \mathcal{C} , and since no half-space in Λ has apex a , $a \notin \mathcal{A}$ by Theorem 29.

Suppose now that $a \notin \mathcal{A}$. Then, $\mathcal{H}(a, I) \in \tau(\Lambda)$ for all $I \in E_a \setminus \{[n]\}$. By Proposition 18, node $[n]$ is reachable in $\mathcal{G}(\Lambda, a)$ from any node $I \in E_a \setminus \{[n]\}$, hence in $\mathcal{R}(\mathcal{D}, a)$. Since any node $I \in E_a$ is obviously reachable from node $[n]$ in $\mathcal{R}(\mathcal{D}, a)$, we conclude that the directed graph $\mathcal{R}(\mathcal{D}, a)$ is strongly connected.

(ii) First observe that by (i), the nodes of the maximal strongly connected components of $\mathcal{R}(\mathcal{D}, a)$ are all distinct from the node $[n]$.

To prove the “only if” part, let L be a minimal subset of $[q]$ satisfying (13). As a is a non-redundant apex, we know that $L \neq \emptyset$. In the first place, consider any maximal strongly connected component of $\mathcal{R}(\mathcal{D}, a)$, and let J be any node of that component. Since

$$\mathcal{H}(a, J) \in \Sigma = \tau(\Lambda \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L}),$$

then by Proposition 31, we know that for some $r \in L$, I_r is reachable from J in $\mathcal{R}(\mathcal{D}, a)$. As J belongs to a maximal strongly connected component of $\mathcal{R}(\mathcal{D}, a)$, I_r must belong to the same component. As a consequence, $\{I_l\}_{I_l \in L}$ contains at least one element of each maximal strongly connected component of $\mathcal{R}(\mathcal{D}, a)$. In the second place, by Corollary 33 the minimality of L implies $\{I_l\}_{I_l \in L}$ is composed of precisely one node of each maximal strongly connected component of the digraph $\mathcal{R}(\mathcal{D}, a)$.

Conversely, consider a (non-empty) subset $L \subset [q]$ such that $\{I_l\}_{I_l \in L}$ is composed of precisely one element of each maximal strongly connected component of $\mathcal{R}(\mathcal{D}, a)$. Then, (13) is satisfied because

$$\mathcal{H}(a, I_r) \in \tau(\Lambda \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L})$$

for all $r \in [q]$ by Proposition 31. We claim that L is a minimal subset of $[q]$ satisfying (13). Indeed, if L is a singleton, it is obviously minimal since a is a non-redundant apex. Similarly, if L has more than one element, by Proposition 31 we have

$$\mathcal{H}(a, I_r) \notin \tau(\Lambda \cup \{\mathcal{H}(a, I_l)\}_{I_l \in L \setminus \{r\}}) \quad \text{for all } r \in L.$$

This shows the “if” part of the statement. \square

In particular, the following corollary follows from Theorem 34 and Proposition 30.

Corollary 35. *Let Λ_1 and Λ_2 be two finite sets of half-spaces in Σ , whose apices are distinct from a , such that:*

$$\tau(\Lambda_1 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in [q]}) = \tau(\Lambda_2 \cup \{\mathcal{H}(a, I_l)\}_{I_l \in [q]}) = \Sigma.$$

If we define the tropical cones $\mathcal{D}_1 := \bigcap_{\mathcal{H} \in \Lambda_1} \mathcal{H}$ and $\mathcal{D}_2 := \bigcap_{\mathcal{H} \in \Lambda_2} \mathcal{H}$, then the maximal strongly connected components of the directed graphs $\mathcal{R}(\mathcal{D}_1, a)$ and $\mathcal{R}(\mathcal{D}_2, a)$ coincide.

Theorem 34 shows that, when a is a non-redundant apex, a half-space $\mathcal{H}(a, I)$ occurring in a non-redundant external representation of \mathcal{C} can be exchanged with another half-space $\mathcal{H}(a, J)$ if, and only if, I and J belong to the same (maximal) strongly connected components of $\mathcal{R}(\mathcal{D}, a)$. By Propositions 13 and 17, this implies that the equality constraint $a_I^- x = a_J^- x$ is satisfied for any $x \in \mathcal{C}$ located in a certain neighborhood of a . This can be seen as analogous to the situation of a non-fully dimensional ordinary polytope (i.e. whose affine hull is a proper subspace of \mathbb{R}^n). However, the difference here is that the exchange is due to the local shape of the polytope around a , and not to its global shape. Moreover, the vectors x satisfying

$$a_I^- x = a_J^- x \geq a_{[n] \setminus (I \cup J)}^- x$$

are included in a tropical hyperplane if, and only if, $I \cap J = \emptyset$. In contrast, in the case where I is included in J , this constraint is equivalent to the inequality $a_I^- x \geq a_{[n] \setminus I}^- x$.

We finally study the structure of maximal strongly connected components of $\mathcal{R}(\mathcal{D}, a)$. With this aim, recall that a *principal ideal* of E_a is a subset of E_a of the form $\{I \in E_a \mid I \subset J\}$, for certain $J \in E_a$, which is called the *principal element* of this ideal.

Proposition 36. *Every maximal strongly connected component C of $\mathcal{R}(\mathcal{D}, a)$ is a principal ideal of E_a . Moreover, the principal element of C is given by the biggest subset R of $[n]$ reachable from I in $\mathcal{G}(\Lambda, a)$, for any $I \in C$.*

Proof. Consider any $I \in C$, and (recalling Remark 14) let R be the biggest subset of $[n]$ reachable from I in $\mathcal{G}(\Lambda, a)$. Since $I \subset R$, we know that $R \in E_a$. We claim that for all $J \in E_a$, J belongs to C if, and only if, $J \subset R$.

If J belongs to C , then J is reachable from I in the graph $\mathcal{R}(\mathcal{D}, a)$, and consequently in the hypergraph $\mathcal{G}(\Lambda, a)$. By definition of R , it follows that $J \subset R$.

Conversely, if $J \subset R$, then J is reachable from I in $\mathcal{G}(\Lambda, a)$, thus also in $\mathcal{R}(\mathcal{D}, a)$ because we assume $J \in E_a$. As C is a maximal strongly connected component of $\mathcal{R}(\mathcal{D}, a)$ and $I \in C$, we conclude that $J \in C$. \square

As a consequence, any maximal strongly connected component C of $\mathcal{R}(\mathcal{D}, a)$ is completely determined by its principal element. Besides, when the apex a is non-redundant, the minimal elements (for inclusion) of C correspond to minimal half-spaces containing \mathcal{C} . To see this, assume that I is a minimal element of C , and let \mathcal{H} be a minimal half-space containing \mathcal{C} such that $\mathcal{C} \subset \mathcal{H} \subset \mathcal{H}(a, I)$. Let $\{I_l\}_{l \in L}$ be composed of precisely one element of each maximal strongly connected component of $\mathcal{R}(\mathcal{D}, a)$, except C . Replacing $\mathcal{H}(a, I)$ by \mathcal{H} in $\Lambda \cup \{\mathcal{H}(a, I_l)\}_{l \in L} \cup \{\mathcal{H}(a, I)\}$, we obtain another (finite) external representation of \mathcal{C} composed of half-spaces in Σ . Then, by Theorem 34, \mathcal{H} is necessarily of the form $\mathcal{H}(a, J)$, where $J \in C$. Moreover, as $\mathcal{H}(a, J) \subset \mathcal{H}(a, I)$, we have $J \subset I$, which shows that $I = J$ because I is a minimal element of C . Thus, $\mathcal{H}(a, I)$ is a minimal half-space containing \mathcal{C} .

4.4. Proof of Theorem 1 and illustrations

We have now all the ingredients to establish Theorem 1.

Given $a \in \mathcal{A}$, let $\mathcal{C}_a \subset 2^{\Sigma_a}$ be composed of the sets of half-spaces $\{\mathcal{H}(a, I) \mid I \in C\}$, where C ranges over the maximal strongly connected components of the directed graph $\mathcal{R}(\mathcal{D}, a)$, with \mathcal{D} defined as the intersection of the half-spaces with apices different from a in some external representation of \mathcal{C} . In the first place, observe that Corollary 35 shows \mathcal{C}_a is independent of the choice of the external representation of \mathcal{C} .

Now, let Γ be a non-redundant external representation of \mathcal{C} . By Theorem 29, Γ is the union, for $a \in \mathcal{A}$, of non-empty sets $\Gamma_a \subset \Gamma$ of half-spaces with apex a . Let $\Lambda := \bigcup_{b \in \mathcal{A} \setminus \{a\}} \Gamma_b$ and $\mathcal{D} := \bigcap_{\mathcal{H} \in \Lambda} \mathcal{H}$. If $\{I_l\}_{l \in [q]} \subset E_a$ is such that $\Gamma_a = \{\mathcal{H}(a, I_l)\}_{l \in [q]}$, then $L = [q]$ is a minimal subset of $[q]$ satisfying (13). Therefore, by Theorem 34, $\{I_l\}_{l \in [q]}$ is composed of precisely one element of each maximal strongly connected component of $\mathcal{R}(\mathcal{D}, a)$. Thus, according to the discussion above, Γ_a is composed of precisely one half-space of each set in the collection \mathcal{C}_a . This proves the “only if” part of the property in Theorem 1.

To prove the “if” part of the property in Theorem 1, we only need to note that, according to Corollary 33 and the discussion above, we can replace a half-space in Γ_a by any other half-space in the same set of the collection \mathcal{C}_a . By Theorem 34, we still obtain a non-redundant external representation of \mathcal{C} .

Example 37. Let us illustrate the results of this section on the 4th cyclic cone in \mathbb{P}^3 . For a given choice of four scalars $-\infty < t_1 < t_2 < t_3 < t_4$, this cone is generated by the vectors $(t_i \times 0, t_i \times 1, t_i \times 2, t_i \times 3)$ for $i = 1, \dots, 4$ (note that the product $t \times m$ corresponds to the tropical exponentiation t^m , which explains the name “cyclic” cone, see [8,3]). Here we use the scalars $t_i := i$ for $i = 1, \dots, 4$. The corresponding cone \mathcal{C} is depicted (in green) in Fig. 7. We start from a description Γ of \mathcal{C} obtained by saturating the half-spaces associated with the non-trivial extreme vectors of the polar cones of \mathcal{C} (see Section 2.3), and given by the following list of apices and sectors:

(0, 1, 2, 3), {4}	(0, 3, 7, 11), {1, 3}	(0, 3, 7, 11), {1, 4}	(0, 1, 2, 6), {3}
(0, 1, 3, 5), {1, 4}	(0, 1, 3, 6), {2, 4}	(0, 1, 3, 7), {1, 3}	(0, 1, 4, 8), {2, 4}
(0, 1, 5, 9), {2}	(0, 2, 4, 7), {1, 4}	(0, 2, 5, 8), {1, 4}	(0, 2, 5, 9), {1, 3}
(0, 3, 6, 10), {1, 4}	(0, 1, 2, 4), {1, 4}	(0, 1, 2, 4), {2, 4}	(0, 4, 8, 12), {1} (Γ)

We number the corresponding half-spaces from \mathcal{H}_1 to \mathcal{H}_{16} from left to right and top to bottom.

The set of non-redundant apices is given by the vectors which are not colored in red.² For instance, it can be verified that the apex $(0, 2, 5, 8)$ does not belong to \mathcal{A} . The tangent directed

² For interpretation of the references to color, the reader is referred to the web version of this article.




Fig. 7. Exchange of two mutually redundant half-spaces of the 4th cyclic cone in \mathbb{P}^3 (pictures have been drawn using `polymake` [16], JavaView, and JReality³). (For interpretation of the references to color, the reader is referred to the web version of this article.)

hypergraph $\mathcal{G}(\Gamma \setminus \{\mathcal{H}_{11}\}, (0, 2, 5, 8))$, generated by the set of half-spaces in Γ having an apex distinct from $(0, 2, 5, 8)$, is formed by the hyperarcs $(\{4\}, \{3\})$, $(\{2\}, \{3\})$, and $(\{1, 3\}, \{2\})$, see Fig. 8. They are respectively associated with the half-spaces \mathcal{H}_6 (or \mathcal{H}_{10}), \mathcal{H}_8 , and \mathcal{H}_{12} , which are active at $(0, 2, 5, 8)$. As a consequence, the set $\{1, 2, 3, 4\}$ is reachable from $\{1, 4\}$. It can be verified that a half-space $\mathcal{H}((0, 2, 5, 8), I)$ contains \mathcal{C} if, and only if, $I \supset \{1, 4\}$. In this case, the reachability graph $\mathcal{R}(\bigcap_{i \neq 11} \mathcal{H}_i, a)$ consists of the nodes $\{1, 4\}$, $\{1, 2, 4\}$, $\{1, 3, 4\}$, and $\{1, 2, 3, 4\}$, and is necessarily strongly connected. By the first part of Theorem 34, we conclude that the vector $(0, 2, 5, 8)$ is not a non-redundant apex.

We now illustrate with apex $(0, 3, 7, 11)$ the situation in which two half-spaces can be exchanged in a non-redundant representation. The tangent directed hypergraph $\mathcal{G}(\Gamma \setminus \{\mathcal{H}_2, \mathcal{H}_3\}, (0, 3, 7, 11))$ consists of the hyperarcs $(\{3\}, \{4\})$, $(\{2\}, \{3, 4\})$, and $(\{4\}, \{3\})$. A half-space \mathcal{H} with apex $(0, 3, 7, 11)$ contains the cone \mathcal{C} if, and only if, there exists $i \in \{2, 3, 4\}$ such that $\{1, i\} \subset \text{sect}(\mathcal{H})$. Let us denote by \mathcal{D} the cone $\bigcap_{\mathcal{H} \in \Gamma \setminus \{\mathcal{H}_2, \mathcal{H}_3\}} \mathcal{H}$ provided by the intersection of the half-spaces in Γ with apices different from $(0, 3, 7, 11)$. We deduce that the directed graph $\mathcal{R}(\mathcal{D}, (0, 3, 7, 11))$, depicted in Fig. 8, is not strongly connected. Thus, the apex $(0, 3, 7, 11)$ is non-redundant by Theorem 34. Besides, the graph $\mathcal{R}(\mathcal{D}, (0, 3, 7, 11))$ has only one maximal strongly connected component, composed of the subsets $\{1, 3\}$, $\{1, 4\}$, and $\{1, 3, 4\}$, the latter being the principal element of the component. It follows that the collection $\mathcal{C}_{(0,3,7,11)}$ of Theorem 1 is composed of a unique set of half-spaces $\{\mathcal{H}((0, 3, 7, 11), \{1, 3\}), \mathcal{H}((0, 3, 7, 11), \{1, 4\}), \mathcal{H}((0, 3, 7, 11), \{1, 3, 4\})\}$, and that only the first two half-spaces are minimal with respect to \mathcal{C} (see Fig. 7, where they are depicted in yellow and red respectively). By Theorem 34, we conclude that any non-redundant external representation obtained from Γ contains precisely one of the half-spaces $\mathcal{H}((0, 3, 7, 11), \{1, 3\})$ and $\mathcal{H}((0, 3, 7, 11), \{1, 4\})$.

³ Interactive 3D objects can be accessed at <http://www.cmap.polytechnique.fr/~allamigeon/gallery.html>.




Fig. 8. (a) The tangent directed hypergraph $\mathcal{G}(\{\mathcal{H}_i\}_{i \neq 11}, (0, 2, 5, 8))$; (b) The tangent directed hypergraph $\mathcal{G}(\{\mathcal{H}_i\}_{i \neq \{2,3\}}, (0, 3, 7, 11))$; (c) The reachability digraph $\mathcal{R}(\cap_{i \in \{2,3\}} \mathcal{H}_i, (0, 3, 7, 11))$ over the elements of $E_{(0,3,7,11)}$.

More generally, it can be verified that the non-redundant external representations which can be obtained from the representation Γ are the ones containing all the half-spaces in black, and precisely one half-space in green and one in blue.

4.5. The particular case of pure cones

We say that the real polyhedral cone \mathcal{C} is pure when it coincides with the topological closure of its interior. This terminology originates from the equivalence between this definition, and the fact that the maximal (inclusion-wise) bounded cells of the natural cell decomposition of \mathbb{P}^{n-1} induced by a generating set $\{v^r\}_{r \in [p]}$ of \mathcal{C} are all full-dimensional (the subcomplex of bounded cells is said to be pure).

Lemma 38. *If the real polyhedral cone \mathcal{C} is pure, for any $a \in \partial\mathcal{C}$ there exists at most one minimal half-space with respect to \mathcal{C} with apex a .*

Proof. In the first place, we claim that if $\mathcal{H}(a, I)$ and $\mathcal{H}(a, J)$ contain \mathcal{C} and $I \cap J \neq \emptyset$, then $\mathcal{H}(a, I \cap J)$ also contains \mathcal{C} . Otherwise, i.e. if $\mathcal{H}(a, I \cap J)$ does not contain \mathcal{C} (so I and J are not comparable), there exists $y \in \mathcal{C}$ such that

$$a_{[\overline{m}] \setminus (I \cup J)}^- y \oplus a_{J \setminus I}^- y \oplus a_{I \setminus J}^- y > a_{I \cap J}^- y.$$

Besides,

$$a_{I \setminus J}^- y \oplus a_{I \cap J}^- y \geq a_{J \setminus I}^- y \oplus a_{[\overline{m}] \setminus (I \cup J)}^- y,$$

$$a_{J \setminus I}^- y \oplus a_{I \cap J}^- y \geq a_{I \setminus J}^- y \oplus a_{[\overline{m}] \setminus (I \cup J)}^- y.$$

Thus, we have

$$a_{J \setminus I}^- y = a_{I \setminus J}^- y > a_{I \cap J}^- y.$$

Since \mathcal{C} is pure, we can find x in the interior of \mathcal{C} close enough to y so that

$$a_{J \setminus I}^- x > a_{I \cap J}^- x, \quad a_{I \setminus J}^- x > a_{I \cap J}^- x, \quad \text{and} \quad a_{J \setminus I}^- x \neq a_{I \setminus J}^- x.$$

If, for instance, $a_{J \setminus I}^- x > a_{I \setminus J}^- x$, then $a_{J \setminus I}^- x > a_I^- x$. However, this contradicts that x belongs to $\mathcal{H}(a, I)$, and so also to \mathcal{C} . Therefore, $\mathcal{H}(a, I \cap J)$ must contain \mathcal{C} , proving our claim.

According to the first part of the proof, if we suppose that $\mathcal{H}(a, I)$ and $\mathcal{H}(a, J)$ are two different minimal half-spaces with respect to \mathcal{C} , then necessarily $I \cap J = \emptyset$. Since in that case for all $x \in \mathcal{C}$ we have

$$a_I^- x \geq a_J^- x \oplus a_{[n] \setminus (I \cup J)}^- x \quad \text{and} \quad a_J^- x \geq a_I^- x \oplus a_{[n] \setminus (I \cup J)}^- x,$$

it follows that:

$$a_I^- x = a_J^- x \geq a_{[n] \setminus (I \cup J)}^- x, \quad \text{for all } x \in \mathcal{C}.$$

As $I \cap J = \emptyset$, this implies \mathcal{C} is contained in the tropical hyperplane whose apex is a , which contradicts the fact that \mathcal{C} is pure (hence, it has a non-empty interior). \square

Proposition 39. *If the real polyhedral cone \mathcal{C} is pure, every non-redundant external representation of \mathcal{C} consists of precisely one half-space $\mathcal{H}(a, I)$ for each non-redundant apex a . In particular, there exists a unique non-redundant external representation composed of minimal half-spaces.*

Proof. Assume that $\mathcal{R}(\mathcal{D}, a)$ contains two different maximal strongly connected components C and C' , and let I and J be minimal elements of C and C' . Then, as explained above, $\mathcal{H}(a, I)$ and $\mathcal{H}(a, J)$ are different minimal half-spaces containing \mathcal{C} , which is impossible by Lemma 38. Thus, the graph $\mathcal{R}(\mathcal{D}, a)$ has only one maximal strongly connected component. Besides, using the same argument, we can also conclude that this component has only one minimal element. This implies there is a unique non-redundant external representation of \mathcal{C} composed of minimal half-spaces. \square

5. Apices of non-redundant half-spaces and cell decomposition

In this section we show that the non-redundant apices associated with a real polyhedral cone come from a small set of candidates. In the sequel, \mathcal{C} denotes a real polyhedral cone generated by the set of vectors $\{v^1, \dots, v^p\} \subset \mathbb{P}^{n-1}$.

As already observed in Section 2.3, the real polyhedral cone \mathcal{C} has an external representation composed of minimal half-spaces with respect to \mathcal{C} . Besides, the apex of each minimal half-space belongs to \mathcal{C} . In consequence, the non-redundant apices \mathcal{A} associated with \mathcal{C} are apices of minimal half-spaces, and so they belong to the cells of the natural cell decomposition of \mathbb{P}^{n-1} induced by the generators $\{v^r\}_{r \in [p]}$ of \mathcal{C} characterized by the conditions of Theorem 4. We next show that the elements of \mathcal{A} satisfy a stronger property, implying they are special vertices (recall that cells composed of apices of minimal half-spaces need not be zero-dimensional).

Definition 40. Let I be a non-empty proper subset of $[n]$ and $j \in [n] \setminus I$. A vector $a \in \mathbb{P}^{n-1}$ is said to be an (I, j) -vertex of \mathcal{C} if

$$S_i(a) \cap S_j(a) \not\subset \bigcup_{k \in I \setminus \{i\}} S_k(a) \quad \text{for all } i \in I, \tag{C4}$$

and Conditions (C1) and (C2) are satisfied, where $(S_1(a), \dots, S_n(a)) = \text{type}(a)$ is the type of a relative to the generating set $\{v^r\}_{r \in [p]}$ of \mathcal{C} .

Remark 41. As in the case of Conditions (C1) and (C2), Condition (C4) above is independent of the choice of the generating set of \mathcal{C} . Indeed, assuming $\mathcal{C} \subset \mathcal{H}(a, I)$, it is equivalent to

$$\text{for each } i \in I \text{ there exists } x \in \mathcal{C} \text{ such that } a_i^- x_i = a_j^- x_j > \bigoplus_{k \in I \setminus \{i\}} a_k^- x_k. \tag{C4'}$$

This provides a geometric interpretation of Condition (C4): $\mathcal{H}(a, I)$ separates the cone \mathcal{C} from the sector $\mathcal{S}(a, j)$, and this separation is “tight”, since $\mathcal{S}(a, i) \cap \mathcal{S}(a, j)$ has a non-empty interior for all $i \in I$.

Observe that any vector a satisfying the conditions of Definition 40 is a vertex of the natural cell decomposition of \mathbb{P}^{n-1} induced by $\{v^r\}_{r \in [p]}$. Let G_S be the graph associated with the cell X_S , where $S = \text{type}(a)$ (see Section 2.1). Note that Condition (C4) above implies $S_i(a) \cap S_j(a) \neq \emptyset$ for all $i \in I$, and so node j is connected in G_S with each node of I . Moreover, by Condition (C2), each node of $[n] \setminus I$ is connected in G_S with some node of I . Therefore, the graph G_S associated with an (I, j) -vertex of \mathcal{C} is always connected. This shows that X_S is a zero-dimensional cell.

Also observe that if a is an (I, j) -vertex of \mathcal{C} , then $\mathcal{H}(a, I)$ is a minimal half-space with respect to \mathcal{C} , because Condition (C4) above is stronger than Condition (C3).

The following proposition shows that (I, j) -vertices are associated with non-trivial extreme vectors of the j th polar of \mathcal{C} .

Proposition 42. *The following three properties are equivalent:*

- (i) *the vector a is an (I, j) -vertex of \mathcal{C} ,*
- (ii) *$\mathcal{H}(a, I)$ is a minimal half-space with respect to \mathcal{C} and $a_1^- \oplus a_j^- e^j$ is a non-trivial extreme vector of the j th polar of \mathcal{C} ,*
- (iii) *$\mathcal{H}(a, I)$ can be obtained by the saturation of a half-space associated with a non-trivial extreme vector of the j th polar of \mathcal{C} .*

Proof. (i) \implies (ii): Assume the three conditions of Definition 40 are satisfied. Thus, as we observed above, $\mathcal{H}(a, I)$ is a minimal half-space with respect to \mathcal{C} . This implies $a_1^- \oplus a_j^- e^j$ belongs to the j th polar of \mathcal{C} , because the inequality $a_1^- x \geq a_j^- x_j$ holds for all $x \in \mathcal{C}$. Then, since Condition (C4) is satisfied, by Theorem 3 we conclude that $a_1^- \oplus a_j^- e^j$ is a non-trivial extreme vector of the j th polar of \mathcal{C} .

(ii) \implies (i): Conversely, if $\mathcal{H}(a, I)$ is a minimal half-space with respect to \mathcal{C} , Conditions (C1) and (C2) are satisfied. Besides, if $a_1^- \oplus a_j^- e^j$ is a non-trivial extreme vector of the j th polar of \mathcal{C} , Condition (C4'), and subsequently Condition (C4), are satisfied by Theorem 3. Thus, a is an (I, j) -vertex of \mathcal{C} .

(ii) \implies (iii): Suppose that $u := a_1^- \oplus a_j^- e^j$ is a non-trivial extreme vector of the j th polar of \mathcal{C} and $\mathcal{H}(a, I)$ is minimal with respect to \mathcal{C} . Let $\mathcal{H}(b, J)$ be the half-space obtained by the saturation of $\{x \in \mathbb{P}_{\max}^{n-1} \mid \bigoplus_{i \in [n] \setminus \{j\}} u_i x_i \geq u_j x_j\}$. From Proposition 10, it follows that $J = I$ and $u = b_1^- \oplus b_j^- e^j$, and so $a_i = b_i$ for all $i \in I$. Besides, since $\mathcal{H}(b, I)$ is minimal with respect to \mathcal{C} , we have $b \in \mathcal{C}$ by Corollary 5. Thus, $a^- b = a_1^- b = 0$, which shows that $b_h \leq a_h$ for all $h \in [n]$. The symmetric inequality can be obtained by exchanging $\mathcal{H}(a, I)$ and $\mathcal{H}(b, I)$ (since the former is also minimal), which proves $a = b$.

(iii) \implies (ii): Straightforward by Proposition 10. \square

We are now ready to prove one of the main results of this section.

Theorem 43. *For any non-redundant apex $a \in \mathcal{A}$ there exist a non-empty proper subset I of $[n]$ and $j \in [n] \setminus I$ such that a is an (I, j) -vertex of \mathcal{C} .*

Proof. Let $\Gamma \subset \Sigma$ be a finite set of minimal half-spaces with respect to \mathcal{C} such that $\mathcal{C} = \bigcap_{\mathcal{H} \in \Gamma} \mathcal{H}$. Up to extracting a subset of half-spaces, we may assume Γ is a non-redundant external representation of \mathcal{C} .

Given a non-redundant apex a , let $\mathcal{H}(a, I) \in \Gamma$ be a half-space with apex a . By assumption, $\mathcal{H}(a, I)$ is non-redundant in Γ , so there exists a vector x such that $x \notin \mathcal{H}(a, I)$ and $x \in \mathcal{H}'$ for all $\mathcal{H}' \in \Gamma \setminus \{\mathcal{H}(a, I)\}$. For $r \in [p]$, let $w^r \in \mathbb{P}^{n-1}$ be the vector defined by

$$w^r := (a_{[n] \setminus I}^- x) v^r \oplus (a_1^- v^r) x.$$

This vector w^r is located on the tropical (projective) segment joining v^r and x . Moreover, since $x \notin \mathcal{H}(a, I)$ and $v^r \in \mathcal{C} \subset \mathcal{H}(a, I)$, we have $a_1^- x < a_{[n] \setminus I}^- x$ and $a_1^- v^r \geq a_{[n] \setminus I}^- v^r$. Then,

$$a_I^- w^r = (a_{[n] \setminus I}^- x)(a_I^- v^r) \oplus (a_I^- v^r)(a_I^- x) = (a_{[n] \setminus I}^- x)(a_I^- v^r)$$

and

$$a_{[n] \setminus I}^- w^r = (a_{[n] \setminus I}^- x)(a_{[n] \setminus I}^- v^r) \oplus (a_I^- v^r)(a_{[n] \setminus I}^- x) = (a_I^- v^r)(a_{[n] \setminus I}^- x).$$

Thus, $a_I^- w^r = a_{[n] \setminus I}^- w^r$, which means that w^r lies on the boundary of $\mathcal{H}(a, I)$. It follows that w^r belongs to \mathcal{C} , because w^r also belongs to \mathcal{H}' for any $\mathcal{H}' \in \Gamma \setminus \{\mathcal{H}(a, I)\}$ (as a tropical linear combination of vectors of \mathcal{H}').

We now claim that $\arg \max(a_I^- w^r) = \arg \max(a_I^- v^r)$ and $\arg \max(a_{[n] \setminus I}^- w^r) \supset \arg \max(a_{[n] \setminus I}^- x)$ for all $r \in [p]$.

Indeed, observe that for any $i \in I$,

$$(a_I^- v^r)(a_i^- x_i) < (a_I^- v^r)(a_{[n] \setminus I}^- x) = a_I^- w^r.$$

Then, as $a_i^- w_i^r = (a_{[n] \setminus I}^- x)(a_i^- v_i^r) \oplus (a_I^- v^r)(a_i^- x_i)$, we have $i \in \arg \max(a_I^- w^r)$ if, and only if, $a_i^- v_i^r = a_I^- v^r$. This proves that $\arg \max(a_I^- w^r)$ coincides with $\arg \max(a_I^- v^r)$.

Similarly, let $j \in \arg \max(a_{[n] \setminus I}^- x)$. Since $a_I^- v^r \geq a_{[n] \setminus I}^- v^r \geq a_j^- v_j^r$ and $a_{[n] \setminus I}^- x = a_j^- x_j$, it follows that:

$$a_I^- w_j^r = (a_{[n] \setminus I}^- x)(a_j^- v_j^r) \oplus (a_I^- v^r)(a_j^- x_j) = (a_I^- v^r)(a_j^- x_j) = (a_I^- v^r)(a_{[n] \setminus I}^- x).$$

Thus, $a_j^- w_j^r$ is equal to $a_{[n] \setminus I}^- w^r$, which shows that $j \in \arg \max(a_{[n] \setminus I}^- w^r)$. This completes the proof of the claim.

Now, consider any $j \in \arg \max(a_{[n] \setminus I}^- x)$. Then, j also belongs to $\arg \max(a_{[n] \setminus I}^- w^r)$ for all $r \in [p]$. Besides, since the half-space $\mathcal{H}(a, I)$ is minimal, Condition (C3) is satisfied, so for each $i \in I$ there exists $r_i \in [p]$ such that $r_i \in S_i(a)$ and $r_i \notin \bigcup_{k \in I \setminus \{i\}} S_k(a)$. Equivalently, $\arg \max(a_I^- w^{r_i}) = \arg \max(a_I^- v^{r_i})$ is reduced to the singleton $\{i\}$. As $a_I^- w^{r_i} = a_{[n] \setminus I}^- w^{r_i}$, we conclude that w^{r_i} satisfies

$$a_i^- w_i^{r_i} = a_j^- w_j^{r_i} > \bigoplus_{k \in I \setminus \{i\}} a_k^- w_k^{r_i}.$$

This shows that Condition (C4') above is satisfied. Moreover, as $\mathcal{H}(a, I)$ is a minimal half-space with respect to \mathcal{C} , we know that Conditions (C1) and (C2) are also satisfied. In consequence, the non-redundant apex a is an (I, j) -vertex of \mathcal{C} . \square

We are now going to study a sufficient condition for an (I, j) -vertex a to be a non-redundant apex. We first show that this condition implies node j does not belong to the head of the hyperarcs associated with half-spaces different from $\mathcal{H}(a, I)$.

Lemma 44. *Let a be an (I, j) -vertex of \mathcal{C} satisfying*

$$S_i(a) \cap S_j(a) \not\subset \bigcup_{k \in [n] \setminus \{j, i\}} S_k(a) \text{ for all } i \in I. \tag{C5}$$

If b is a (K, l) -vertex of \mathcal{C} such that the half-spaces $\mathcal{H}(a, I)$ and $\mathcal{H}(b, K)$ are different, and $\mathcal{H}(b, K)$ is active at a , then $j \notin \arg \max(b_{[n] \setminus K}^- a)$.

Proof. By contradiction, assume that $j \in \arg \max(b_{[n] \setminus K}^- a)$. Then, we have $b_j^- a_j = b_K^- a$ because $\mathcal{H}(b, K)$ is active at a .

In the first place, assume $I \not\subset \arg \max(b_K^- a)$. Given $i \in I \setminus \arg \max(b_K^- a)$, since a satisfies Condition (C5), we know that there exists $r \in [p]$ such that $r \in S_i(a) \cap S_j(a)$, and $r \notin S_k(a)$ for all $k \notin \{i, j\}$. Equivalently, $a_i^- v_i^r = a_j^- v_j^r > a_k^- v_k^r$ for all $k \notin \{i, j\}$. Consider η such that

$$\bigoplus_{k \notin \{i, j\}} a_k^- v_k^r < \eta^- < a_i^- v_i^r = a_j^- v_j^r.$$

Then, the vector $x := a \oplus \eta v^r$ satisfies $x_i = \eta v_i^r$, $x_j = \eta v_j^r$, and $x_k = a_k$ for all $k \notin \{i, j\}$. In particular, we have $x_k = a_k$ for all $k \in \arg \max(b_K^- a)$ because $i \notin \arg \max(b_K^- a)$. Besides, choosing η such that η^- is close enough to $a_i^- v_i^r$, we can also suppose that $\arg \max(b_K^- x) \subset \arg \max(b_K^- a)$. Then, we have

$$b_j^- x_j = b_j^- \eta v_j^r > b_j^- a_j = b_K^- a = \bigoplus_{k \in \arg \max(b_K^- a)} b_k^- x_k \geq \bigoplus_{k \in \arg \max(b_K^- x)} b_k^- x_k = b_K^- x.$$

This shows that x does not belong to the half-space $\mathcal{H}(b, K)$. This is a contradiction because $x \in \mathcal{C}$ (as a tropical linear combination of two elements of \mathcal{C}) and $\mathcal{C} \subset \mathcal{H}(b, K)$ (by Condition (C1) applied to $\mathcal{H}(b, K)$).

Now assume $I \subset \arg \max(b_K^- a)$. Then, since $b_j^- a_j = b_K^- a$, we have $b_i^- a_i = b_j^- a_j$ for all $i \in I$. It is convenient to split the rest of the proof into two cases:

$I \subsetneq K$: Let $k \in K \setminus I$. Since b is a (K, l) -vertex of \mathcal{C} , by Condition (C4) there exists $r \in [p]$ such that $b_k^- v_k^r = b_l^- v_l^r > \bigoplus_{h \in K \setminus \{k\}} b_h^- v_h^r$. Then, we have

$$a_j^- v^r = a_j^- b_j (b_l^- v^r) \leq a_j^- b_j (b_{K \setminus \{k\}}^- v^r) < a_j^- b_j b_k^- v_k^r \leq a_k^- b_k b_k^- v_k^r = a_k^- v_k^r, \tag{18}$$

where the last inequality follows from $j \in \arg \max(b^- a)$, because it implies $b_j^- a_j \geq b_k^- a_k$. Since $k \in [n] \setminus I$, we conclude from (18) that $v^r \notin \mathcal{H}(a, I)$, which is a contradiction.

$I = K$: We know that $\mathcal{H}(a, I)$ and $\mathcal{H}(b, K) = \mathcal{H}(b, I)$ are both minimal half-spaces with respect to \mathcal{C} . Then, since $a \in \mathcal{C} \subset \mathcal{H}(b, I)$ by Corollary 5, it follows that for all $h \in [n] \setminus I$,

$$b_j^- a_j = b_l^- a \geq b_{[n] \setminus I}^- a \geq b_h^- a_h.$$

Symmetrically, it can be proved that $a_j^- b_j \geq a_h^- b_h$ for all $h \in [n] \setminus I$, using the fact that $a_j^- b_j = a_l^- b$ and $b \in \mathcal{H}(a, I)$. We conclude that a and b are identical (as elements of the projective space), which contradicts that $\mathcal{H}(a, I)$ and $\mathcal{H}(b, K)$ are different half-spaces. \square

Theorem 45. *If a is an (I, j) -vertex of \mathcal{C} satisfying Condition (C5), then a is a non-redundant apex of \mathcal{C} .*

Proof. It suffices to consider the external representation Γ of \mathcal{C} composed of the half-spaces provided by Proposition 7, when applied to the non-trivial extreme vectors of the j th polar of \mathcal{C} , for all $j \in [n]$.

By Proposition 42, a half-space $\mathcal{H}(b, K)$ belongs to Γ if, and only if, its apex b is a (K, l) -vertex of \mathcal{C} for some $l \in [n] \setminus K$. Then, since a is assumed to be an (I, j) -vertex of \mathcal{C} satisfying Condition (C5), in particular we have $\mathcal{H}(a, I) \in \Gamma$. Moreover, Lemma 44 ensures that $j \notin H$ for any hyperarc (T, H) in the tangent directed hypergraph $\mathcal{G}(\Gamma \setminus \{\mathcal{H}(a, I)\}, a)$, because such hyperarc is associated with a half-space $\mathcal{H}(b, K)$ such that b is a (K, l) -vertex of \mathcal{C} for some $l \in [n] \setminus K$. As a consequence, the set $[n]$ cannot be reachable from I in $\mathcal{G}(\Gamma \setminus \{\mathcal{H}(a, I)\}, a)$. From Proposition 18, we conclude that $\mathcal{H}(a, I)$ is not redundant in Γ , and then $a \in \mathcal{A}$ by Theorem 29. \square

Remark 46. When $\mathcal{C} \subset \mathbb{P}^2$, Theorems 43 and 45 allow us to establish that the non-redundant apices of \mathcal{C} are precisely the vectors $a \in \mathbb{P}^2$ such that a is an (I, j) -vertex of \mathcal{C} for some non-empty proper subset I of $[3]$ and $j \notin I$.

To see this, in the first place assume $I = \{i_1, i_2\}$, with $i_1 \neq i_2$. Then, Condition (C5) amounts to $S_{i_1}(a) \cap S_j(a) \not\subset S_{i_2}(a)$ and $S_{i_2}(a) \cap S_j(a) \not\subset S_{i_1}(a)$, which is equivalent to Condition (C4). Thus, Theorem 45 ensures that a is a non-redundant apex.

Assume now I consists of only one element i , and let $k \neq j$ be the second element of $[n] \setminus I$. Let Γ be a non-redundant external representation of \mathcal{C} , and assume it does not contain any half-space with apex a . Since the half-space $\mathcal{H}(a, \{i\})$ is redundant with respect to Γ , the tangent directed hypergraph $\mathcal{G}(\Gamma, a)$ must necessarily contain a hyperarc from $\{i\}$ to $\{j\}$ or to $\{k\}$ (but not to $\{j, k\}$ because no half-space in Γ has apex a). Suppose, for instance, that $\{i\}$ is connected with $\{j\}$ by a hyperarc associated with a half-space in Γ , and let b be its apex. Thus $b_i^- a_i = b_j^- a_j > b_k^- a_k$. We get a contradiction, since $a_i^- b_i < a_j^- b_j \oplus a_k^- b_k$ while $b \in \mathcal{C} \subset \mathcal{H}(a, \{i\})$.

We now exhibit a class of real polyhedral cones for which the non-redundant apices are precisely the vertices satisfying Definition 40.

Definition 47. The real polyhedral cone \mathcal{C} is said to have *generic extremities* if for each of its generators v^r there exists a non-trivial (i.e. of positive radius) Hilbert ball containing v^r and included in \mathcal{C} .

Remark 48. Definition 47 does not depend on the choice of the generating set $\{v^r\}_{r \in [p]}$. Indeed, \mathcal{C} has generic extremities if, and only if, for each $x \in \mathcal{C}$ there exists a non-trivial Hilbert ball \mathcal{B} such that $x \in \mathcal{B} \subset \mathcal{C}$.

To see this, let $x = \bigoplus_{r \in [p]} \lambda_r v^r$ be an arbitrary element of \mathcal{C} . For each $r \in [p]$, suppose that the Hilbert ball with center c^r and radius $\epsilon > 0$ is contained in \mathcal{C} and contains v^r . Without loss of generality, we can assume $\min_{i \in [n]} (v_i^r - c_i^r) = 0$ for all $r \in [p]$. Define $c := \bigoplus_{r \in [p]} \lambda_r c^r$, and let \mathcal{B} be the Hilbert ball with center c and radius ϵ .

First, let us show that $x \in \mathcal{B}$. For each $i \in [n]$, there exists $r \in [p]$ such that $x_i = \lambda_r v_i^r$. Since $c_i \geq \lambda_r c_i^r$, we have $x_i - c_i \leq v_i^r - c_i^r \leq \epsilon$. Similarly, let $s \in [p]$ such that $c_i = \lambda_s c_i^s$. Then, $x_i - c_i \geq v_i^s - c_i^s \geq 0$ by assumption. It follows that $d_H(x, c) \leq \epsilon$.

It remains to prove that $\mathcal{B} \subset \mathcal{C}$. Consider any $y \in \mathcal{B}$ and, without loss of generality, assume that $\min_{i \in [n]} (y_i - c_i) = 0$. For each $i \in [n]$, define $\mu_i := (\epsilon c_i^{r_i})^- y_i$, where $r_i \in [p]$ is such that $c_i = \bigoplus_{r \in [p]} \lambda_r c_i^r = \lambda_{r_i} c_i^{r_i}$. Observe that $\epsilon^- \lambda_{r_i} \leq \mu_i \leq \lambda_{r_i}$. We claim that

$$y = \bigoplus_{j \in [n]} \mu_j (c^{r_j} \oplus \epsilon c_j^{r_j} e^j).$$

Given $i \in [n]$, the equality $y_i = \mu_i \epsilon c_i^{r_i} = \mu_i (c^{r_i} \oplus \epsilon c_i^{r_i} e^i)_i$ holds by definition of μ_i . Besides, for $j \neq i$, we have

$$\mu_j (c^{r_j} \oplus \epsilon c_j^{r_j} e^j)_i = \mu_j c_i^{r_j} \leq \lambda_{r_j} c_i^{r_j} \leq \lambda_{r_i} c_i^{r_i} \leq \epsilon \mu_i c_i^{r_i}.$$

This shows that y_i is the maximum of the $\mu_j (c^{r_j} \oplus \epsilon c_j^{r_j} e^j)_i$ for $j \in [n]$, proving the claim. Finally, since each $c^{r_i} \oplus \epsilon c_i^{r_i} e^i$ belongs to the Hilbert ball with center c^{r_i} and radius ϵ , which is contained in \mathcal{C} , we conclude that $y \in \mathcal{C}$.

Note that, from Remark 48, we conclude that any real polyhedral cone which has generic extremities is pure.

The term *generic extremities* originates from the fact the aforementioned property holds if, and only if, each extreme vector of \mathcal{C} belongs to a non-trivial Hilbert ball contained in \mathcal{C} . This enforces that around each of its extreme vectors, the cone has the shape of a Hilbert ball, ensuring a certain “genericity”.

Remark 49. It can be shown that the cone \mathcal{C} has generic extremities as soon as the following two conditions hold:

- (i) \mathcal{C} is pure;
- (ii) the 2×2 -minors

$$v_i^r v_j^s \oplus v_j^r v_i^s = \max\{v_i^r + v_j^s, v_j^r + v_i^s\} \quad (i, j \in [n], r, s \in [p], i \neq j, r \neq s)$$

are non-singular, i.e. the maximum in the right-hand side is reached exactly once.

In particular, the latter condition is satisfied when the vectors v^1, \dots, v^p are in general position in the sense of [27].

We say that a cone \mathcal{D} approximates the cone \mathcal{C} with precision $\epsilon > 0$ if the Hausdorff distance between \mathcal{C} and \mathcal{D} (derived from the metric d_H) is bounded by ϵ . Observe that the real polyhedral




Fig. 9. Three Hilbert balls (left) and a tropical cone with generic extremities (right). (For interpretation of the references to color, the reader is referred to the web version of this article.)

cone \mathcal{C} can be approximated with an arbitrary precision by another one having generic extremities: given $\epsilon > 0$, it suffices to define the tropical cone \mathcal{C}_ϵ as the one generated by the Hilbert balls \mathcal{B}^r with center v^r and radius ϵ , for $r \in [p]$, i.e. the set of tropical linear combinations of the form:

$$\lambda_1 x^1 \oplus \dots \oplus \lambda_p x^p, \quad \text{where } \lambda_r \in \mathbb{R}_{\max} \text{ and } x^r \in \mathcal{B}^r \text{ for } r \in [p].$$

Since any Hilbert ball with center c and radius ϵ is polyhedral (its extreme vectors are the vectors $c \oplus \epsilon_i e^i$ for $i \in [n]$), \mathcal{C}_ϵ is a real polyhedral cone. Moreover, it can be shown that \mathcal{C}_ϵ approximates \mathcal{C} with precision ϵ . Also note that other deformations are possible, for instance choosing balls with different radii for each generator, or approximating each generator by a generic polytrope⁴ containing it.

Example 50. Three Hilbert balls of radius $\frac{1}{2}$ centered at the generators $v^1 = (0, 1, 3)$, $v^2 = (0, 4, 1)$ and $v^3 = (0, 9, 4)$ of the tropical cone of Fig. 1 are depicted on the left-hand side of Fig. 9. Due to the shape of these non-trivial Hilbert balls of \mathbb{P}^2 , it is geometrically clear that the tropical cone of Fig. 1 does not have generic extremities.

A tropical cone with generic extremities is shown on the right-hand side of Fig. 9. This cone is generated by the Hilbert balls on the left. Observe that letting the radii of these Hilbert balls tend to zero, the tropical cone of Fig. 1 can be approximated as much as we want.

The following theorem shows that if \mathcal{C} has generic extremities, then the vertices of the cell decomposition introduced in Definition 40 are precisely the non-redundant apices. Besides, it proves that they provide the unique non-redundant external representation composed of minimal half-spaces.

Theorem 51. *If the real polyhedral cone \mathcal{C} has generic extremities, the non-redundant apices of \mathcal{C} are precisely the vectors a for which there exist a non-empty proper subset I of $[n]$ and $j \in [n] \setminus I$ such that a is an (I, j) -vertex of \mathcal{C} .*

Moreover, each such set I is uniquely determined, and the collection of the half-spaces $\mathcal{H}(a, I)$ is the unique non-redundant external representation of \mathcal{C} composed of minimal half-spaces.

Proof. In the first place, we prove that any (I, j) -vertex of \mathcal{C} satisfies Condition (C5).

By the contrary, assume a is an (I, j) -vertex of \mathcal{C} for which Condition (C5) does not hold. Then, for some $i \in I$, given any extreme vector v of \mathcal{C} satisfying $a_i^- v_i = a_j^- v_j = a^- v = a_I^- v$, there exists $k \in [n] \setminus \{i, j\}$ such that $a_j^- v_j = a_k^- v_k = a_I^- v$. Since Condition (C4) holds, v can be chosen so that $a_i^- v_i = a_j^- v_j > a_h^- v_h$ for all $h \in I \setminus \{i\}$, and so $k \in [n] \setminus I$.

⁴ A polytrope is a tropical cone which is also convex in the classical sense, see [22] for further details. We say that it is generic when its extreme vectors are in general position.

Let \mathcal{B} be a Hilbert ball with center c and radius $\epsilon > 0$ such that $v \in \mathcal{B} \subset \mathcal{C}$. Since v is extreme in \mathcal{C} , it is also extreme in \mathcal{B} . Recalling that the extreme vectors of \mathcal{B} are the vectors $c \oplus \epsilon c_i e^i$ for $i \in [n]$, it follows that there exists $l \in [n]$ such that $v_l = \epsilon c_l$, and $v_h = c_h$ for $h \neq l$. Then, for any $0 < \eta < \epsilon$, the vector x defined by $x_l = \eta^{-1} v_l$, and $x_h = v_h$ for $h \neq l$, is in the interior of \mathcal{B} .

Now, suppose that $l \neq j$. Since x is in the interior of \mathcal{B} , there exists $\eta' > 0$ such that the vector y defined by $y_j = \eta' x_j$, and $y_h = x_h$ for $h \neq j$, belongs to \mathcal{B} . Then, we have

$$a_j^- y_j > a_j^- x_j = a_j^- v_j = a_l^- v \geq a_l^- x = a_l^- y.$$

However, this is impossible, because $y \in \mathcal{B} \subset \mathcal{C} \subset \mathcal{H}(a, I)$. Thus, l must be equal to j . The same reasoning holds with k instead of j , and leads to $l = k$. Since j and k are distinct, we obtain a contradiction. Therefore, every (I, j) -vertex of \mathcal{C} must satisfy Condition (C5). As a consequence, we conclude from Theorems 43 and 45 that the non-redundant apices are precisely those vertices a for which there exist a non-empty proper subset I of $[n]$ and $j \in [n] \setminus I$ such that a is an (I, j) -vertex of \mathcal{C} .

Finally, assume a is both an (I, j) -vertex and an (I', j') -vertex of \mathcal{C} . Then, since $\mathcal{H}(a, I)$ and $\mathcal{H}(a, I')$ are minimal half-spaces with respect to \mathcal{C} and \mathcal{C} is pure, by Lemma 38 we necessarily have $I = I'$. This proves that I is indeed uniquely determined. Moreover, by Proposition 39, there is a unique non-redundant external representation composed of minimal half-spaces. According to Lemma 38 and Proposition 42, $\mathcal{H}(a, I)$ is the only half-space with apex a appearing in such representation. \square

Example 52. The non-redundant apices of the cone of Fig. 9 (right) are the vectors $(0, \frac{1}{2}, 3)$, $(0, \frac{7}{2}, \frac{5}{2})$, $(0, 6, \frac{1}{2})$, $(0, \frac{19}{2}, \frac{9}{2})$, and $(0, 8, \frac{7}{2})$, which are respectively $(\{2\}, \cdot)$ -, $(\{2, 3\}, 1)$ -, $(\{3\}, \cdot)$ -, $(\{1\}, \cdot)$ -, and $(\{1, 2\}, 3)$ -vertices (the notation (I, \cdot) stands for any couple (I, j) with $j \notin I$). They are depicted in orange together with the corresponding minimal half-spaces, while the extreme vectors w^1, \dots, w^5 of the cone are represented in blue.

Remark 53. Theorem 51 cannot be generalized to the case of pure cones. As an example, consider the perturbation of the 4th cyclic cone in \mathbb{P}^3 generated by the following vectors:

$$\begin{matrix} (0, 1, 2, 3) & \left(0, 1, \frac{5}{2}, \frac{7}{2}\right) & \left(0, \frac{3}{2}, \frac{5}{2}, \frac{7}{2}\right) & \left(0, \frac{3}{2}, 4, 6\right) & (0, 2, 4, 6) \\ \left(0, \frac{5}{2}, 6, 9\right) & (0, 3, 6, 9) & \left(0, \frac{7}{2}, \frac{15}{2}, 12\right) & \left(0, \frac{7}{2}, 8, 12\right) & (0, 4, 8, 12) \end{matrix}$$

This cone can be verified to be pure (see Fig. 10), for instance by testing that the subcomplex of bounded cells of the natural cell decomposition induced by its generators is pure and full-dimensional. For the sake of completeness, we provide a `polymake`⁵ script allowing to check this property:

```
application "tropical";
$gen = new Matrix<Rational>([[0,1,2,3], [0,1,5/2,7/2], [0,3/2,5/2,7/2], [0,3/2,4,6],
                             [0,2,4,6], [0,5/2,6,9], [0,3,6,9], [0,7/2,15/2,12],
                             [0,7/2,8,12], [0,4,8,12]]);
$sp = new TropicalPolytope<Rational>(POINTS=>-$gen);
$strunc_vertices = $sp->PSEUDOVERTICES->minor(All, range(1, $sp->AMBIENT_DIM));
$n_vertices = scalar(@$strunc_vertices);
$all_ones = new Vector<Rational>([ (1)x$n_vertices ]);
$vertices = ($all_ones|$strunc_vertices);
$max_cells = $sp->ENVELOPE->BOUNDED_COMPLEX->MAXIMAL_POLYTOPES;
$cell_complex = new fan::PolyhedralComplex(VERTICES=>$vertices, MAXIMAL_CELLS=>$max_cells);
if ($cell_complex->FULL_DIM && $cell_complex->PURE) {
    print "The cone is pure."
} else {
    print "The cone is not pure."
};
```

⁵ Version 2.12 or later.




Fig. 10. Perturbation of the 4th cyclic cone of \mathbb{P}^3 into a pure cone.

The external representation obtained by saturating the half-spaces associated with the non-trivial extreme vectors of the polar cones contains a half-space with apex $a = (0, 1, \frac{5}{2}, \frac{9}{2})$. The type of the latter vector is $S(a) = (\{1, 2\}, \{1, 2, 3\}, \{2, 4, 5\}, \{4, 5, \dots, 10\})$, so that $S_2(a) \cap S_3(a) \not\subset S_4(a)$ and $S_4(a) \cap S_3(a) \not\subset S_2(a)$. The vector a is consequently a $(\{2, 4\}, 3)$ -vertex. However, by Theorem 29, a is not a non-redundant apex since the following list of half-spaces provides a non-redundant external representation of the cone:

- | | | |
|-------------------------------|-----------------------------|--------------------------------|
| $(0, 1, 2, 3), \{4\}$ | $(0, 1, 2, 13/2), \{3\}$ | $(0, 1, 7/2, 13/2), \{2, 4\}$ |
| $(0, 1, 9/2, 17/2), \{2, 4\}$ | $(0, 1, 11/2, 19/2), \{2\}$ | $(0, 3/2, 5/2, 9/2), \{1, 4\}$ |
| $(0, 3/2, 7/2, 8), \{1, 3\}$ | $(0, 2, 4, 7), \{1, 4\}$ | $(0, 2, 5, 19/2), \{1, 3\}$ |
| $(0, 3, 6, 10), \{1, 4\}$ | $(0, 3, 7, 23/2), \{1, 3\}$ | $(0, 4, 8, 12), \{1\}$ |

Acknowledgments

Both authors are very grateful to Stéphane Gaubert for many helpful discussions on tropical convexity, and for his suggestion to study the special case of tropical cones with generic extremities. The authors also thank the two anonymous reviewers for their comments and suggestions which helped to improve the presentation of the results.

References

[1] X. Allamigeon, S. Gaubert, É. Goubault, The tropical double description method, in: J.-Y. Marion, Th. Schwentick (Eds.), Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science (STACS 2010), in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 5, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2010, pp. 47–58.

[2] X. Allamigeon, S. Gaubert, É. Goubault, Computing the vertices of tropical polyhedra using directed hypergraphs, Discrete Comput. Geom. (2012), <http://dx.doi.org/10.1007/s00454-012-9469-6>, in press, arXiv:0904.3436.

- [3] X. Allamigeon, S. Gaubert, R.D. Katz, The number of extreme points of tropical polyhedra, *J. Combin. Theory Ser. A* 118 (1) (2011) 162–189, arXiv:0906.3492.
- [4] X. Allamigeon, S. Gaubert, R.D. Katz, Tropical polar cones, hypergraph transversals, and mean payoff games, *Linear Algebra Appl.* 435 (7) (2011) 1549–1574, arXiv:1004.2778.
- [5] X. Allamigeon, TPLib: Tropical polyhedra library in OCaml, available at <https://gforge.inria.fr/projects/tplib>, 2009.
- [6] W. Bricc, C. Horvath, \mathbb{B} -convexity, *Optimization* 53 (2004) 103–127.
- [7] W. Bricc, C. Horvath, A. Rubinov, Separation in \mathbb{B} -convexity, *Pac. J. Optim.* 1 (2005) 13–30.
- [8] F. Block, J. Yu, Tropical convexity via cellular resolutions, *J. Algebraic Combin.* 24 (1) (2006) 103–114, arXiv:math.MG/0503279.
- [9] G. Cohen, D. Dubois, J.P. Quadrat, M. Viot, A linear system theoretic view of discrete event processes and its use for performance evaluation in manufacturing, *IEEE Trans. Automat. Control* 30 (1985) 210–220.
- [10] G. Cohen, S. Gaubert, J.P. Quadrat, Max-plus algebra and system theory: where we are and where to go now, *Annual Rev. Control* 23 (1999) 207–219.
- [11] G. Cohen, S. Gaubert, J.P. Quadrat, Hahn–Banach separation theorem for max-plus semimodules, in: J.L. Menaldi, E. Rofman, A. Sulem (Eds.), *Optimal Control and Partial Differential Equations*, IOS Press, 2001, pp. 325–334.
- [12] G. Cohen, S. Gaubert, J.P. Quadrat, Duality and separation theorems in idempotent semimodules, *Linear Algebra Appl.* 379 (2004) 395–422, arXiv:math.FA/0212294.
- [13] G. Cohen, S. Gaubert, J.P. Quadrat, I. Singer, Max-plus convex sets and functions, in: G.L. Litvinov, V.P. Maslov (Eds.), *Idempotent Mathematics and Mathematical Physics*, in: *Contemp. Math.*, vol. 377, American Mathematical Society, 2005, pp. 105–129, also ESI Preprint 1341, arXiv:math.FA/0308166.
- [14] M. Develin, B. Sturmfels, Tropical convexity, *Doc. Math.* 9 (2004) 1–27 (electronic), arXiv:math.MG/0308254.
- [15] M. Develin, J. Yu, Tropical polytopes and cellular resolutions, *Experiment. Math.* 16 (3) (2007) 277–292, arXiv:math.CO/0605494.
- [16] E. Gawrilow, M. Joswig, `polymake`: a framework for analyzing convex polytopes, in: G. Kalai, G.M. Ziegler (Eds.), *Polytopes – Combinatorics and Computation*, Birkhäuser, 2000, pp. 43–74.
- [17] S. Gaubert, R.D. Katz, Max-plus convex geometry, in: R.A. Schmidt (Ed.), *Proceedings of the 9th International Conference on Relational Methods in Computer Science and 4th International Workshop on Applications of Kleene Algebra (RelMiCS/AKA 2006)*, in: *Lecture Notes in Comput. Sci.*, vol. 4136, Springer, 2006, pp. 192–206.
- [18] S. Gaubert, R.D. Katz, The Minkowski theorem for max-plus convex sets, *Linear Algebra Appl.* 421 (2–3) (2007) 356–369, arXiv:math.GM/0605078.
- [19] S. Gaubert, R.D. Katz, The tropical analogue of polar cones, *Linear Algebra Appl.* 431 (5–7) (2009) 608–625, arXiv:0805.3688.
- [20] S. Gaubert, R.D. Katz, Minimal half-spaces and external representation of tropical polyhedra, *J. Algebraic Combin.* 33 (3) (2011) 325–348, arXiv:0908.1586.
- [21] G. Gallo, G. Longo, S. Pallottino, S. Nguyen, Directed hypergraphs and applications, *Discrete Appl. Math.* 42 (2–3) (1993) 177–201.
- [22] M. Joswig, K. Kulas, Tropical and ordinary convexity combined, *Adv. Geom.* 10 (2010) 333–352, arXiv:0801.4835.
- [23] M. Joswig, Tropical halfspaces, in: *Combinatorial and Computational Geometry*, in: *Math. Sci. Res. Inst. Publ.*, vol. 52, Cambridge Univ. Press, Cambridge, 2005, pp. 409–431, arXiv:math.CO/0312068.
- [24] B. Korte, L. Lovász, R. Schrader, *Greedoids*, Springer, 1991.
- [25] G.L. Litvinov, V.P. Maslov, G.B. Shpiz, Idempotent functional analysis: an algebraic approach, *Math. Notes* 69 (5) (2001) 696–729, arXiv:math.FA/0009128.
- [26] V. Nitica, I. Singer, Max-plus convex sets and max-plus semispaces. I, *Optimization* 56 (1–2) (2007) 171–205.
- [27] J. Richter-Gebert, B. Sturmfels, T. Theobald, First steps in tropical geometry, in: *Idempotent Mathematics and Mathematical Physics*, in: *Contemp. Math.*, vol. 377, American Mathematical Society, Providence, RI, 2005, pp. 289–317, arXiv:math.AG/0306366.
- [28] I. Singer, *Abstract Convex Analysis*, Wiley, 1997.
- [29] S.N. Samborskii, G.B. Shpiz, Convex sets in the semimodule of bounded functions, in: *Idempotent Analysis*, American Mathematical Society, Providence, RI, 1992, pp. 135–137.
- [30] K. Zimmermann, A general separation theorem in extremal algebras, *Ekonom.-Mat. Obz.* 13 (2) (1977) 179–201.