Les Cahiers du GERAD ISSN: 0711–2440

Tilted inequalities and facets of the set covering polytope: a theoretical analysis

F. Lamothe, C. Contardo, M. Gruson

G–2023–10

April 2023

Citation suggérée : F. Lamothe, C. Contardo, M. Gruson (Avril 2023). Tilted inequalities and facets of the set covering polytope: a theoretical analysis, Rapport technique, Les Cahiers du GERAD G–2023–10, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (https://www.gerad.ca/fr/papers/G-2023-10) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

Suggested citation: F. Lamothe, C. Contardo, M. Gruson (April 2023). Tilted inequalities and facets of the set covering polytope: a theoretical analysis, Technical report, Les Cahiers du GERAD G–2023–10, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (https://www.gerad.ca/en/papers/G-2023-10) to update your reference data, if it has been published in a scientific journal.

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2023
– Bibliothèque et Archives Canada, 2023

Legal deposit – Bibliothèque et Archives nationales du Québec, 2023
– Library and Archives Canada, 2023

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Tilted inequalities and facets of the set covering polytope: a theoretical analysis

François Lamothe [a, c, d]

Claudio Contardo [b, c]

Matthieu Gruson [a, d]

[a] *Department of Analytics, Operations and Information Technologies, Université du Québec à Montréal, Montréal (Qc), Canada, H2X 3X2*

[b] *Department of Mechanical, Industrial and Aerospace Engineering, Concordia University, Montréal (Qc), Canada, H3G 1M8*

[c] *GERAD, Montréal (Qc), Canada, H3T 2A7*

[d] *CIRRELT, Montréal (Qc), Canada, H3T 1J4*

lamothe.francois@uqam.ca
claudio.contardo@concordia.ca
gruson.matthieu@uqam.ca

April 2023
Les Cahiers du GERAD
G–2023–10

Abstract : Given a ground-set of elements and a family of subsets, the set covering problem consists in choosing a minimum number of elements such that each subset contains at least one of the chosen elements. This research focuses on the set covering polytope, which is the convex hull of integer solutions to the set covering problem. We investigate the connection between the study of the facets of the set covering polytope and tilting theory. This theory studies how inequalities can be rotated around their contact points with a polyhedron in order to obtain inequalities inducing higher dimensional faces. To study this connection, we introduce the concept of *tilting vectors* which characterize the degrees of freedom of rotation of an inequality. These vectors characterize facet-defining inequalities and can be used to tilt inequalities with a similar procedure to the one used for arbitrary polyhedra. Additionally, we demonstrate that the computational effort needed to tilt an inequality can be reduced when the inequality has many null coefficients. Finally, we use the tilting vectors to extend several necessary and/or sufficient conditions for facets of the set covering polytope presented by several previous works of the literature.

Keywords : Set covering, facets, tilting, valid inequalities

Acknowledgements: This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) under Discovery Grants 2020–06311 and 2021–03327. This support is gratefully acknowledged.

1 Introduction

Given a ground-set $\mathcal{E}$ of elements and a family $\mathcal{S}$ of subsets of $\mathcal{E}$, the set covering problem consists in choosing a minimum number of elements of $\mathcal{E}$ such that each subset in $\mathcal{S}$ contains at least one of the chosen elements. A feasible set of chosen elements is referred to as a cover.

To model this problem, one can use a binary matrix B whose columns are indexed by the elements $e \in \mathcal{E}$ and rows by the subsets $s \in \mathcal{S}$. The coefficient B_{se} indicates whether the subset s contains the element e. If x_e takes the value 1 when element e is chosen, then the set covering problem can be modeled using the following mixed integer linear program:

$$\min_{x} \quad \mathbb{1}x \tag{1a}$$

$$\text{subject to} \quad Bx \geqslant \mathbb{1} \tag{1b}$$

$$x \in \{0,1\}^{|\mathcal{E}|}. \tag{1c}$$

The set covering problem arises in many applications in a wide variety of fields such as logistics management (Mihelic and Robic, 2004; Cacchiani et al., 2014), crew scheduling (Caprara et al., 1999), manufacturing (Stanfel, 1989), data extraction and manipulation (Day, 1965), and medicine (Reggia et al., 1983). Thus, understanding the structural properties of this problem and their impact on solution methods has very broad implications. Mathematically, we study the properties of the associated set covering polytope $\mathcal{Q}(B)$ which is the convex hull of the set $P = \{x \in \{0,1\}^{|\mathcal{E}|} : Bx \geqslant \mathbb{1}\}$. A comprehensive polyhedral description of $\mathcal{Q}(B)$ via valid inequalities would allow solving the set covering problem using classical linear programming tools. However, even a partial description of the polytope would enable the design of fast and scalable enumeration-based methods.

Numerous families of valid inequalities have already been proposed in the literature (see the thesis of Borndörfer (1998) for an overview). A central question regarding these valid inequalities is whether they induce facets of the polytope $\mathcal{Q}(B)$, i.e., whether they are necessary to describe the polytope. There have been studies on proposing valid inequalities along with the conditions under which they induce facets of the set covering polytope (Cornuéjols and Sassano, 1989; Nobili and Sassano, 1989; Balas and Ng, 1989).

Other polyhedral studies have approached the question by describing necessary and/or sufficient conditions for known valid inequalities to induce facets of $\mathcal{Q}(B)$. In this line of work, it is worth mentioning that Balas and Ng (1989) characterized all the facets having coefficients and right-hand sides in $\{0,1,2\}$. Similarly, Sánchez-García et al. (1998) and Saxena (2004) gave necessary and sufficient conditions for the inequalities with coefficients and right-hand sides in $\{0,1,2,3\}$ to be facets. Another class of inequalities for which the characterization of facets has been studied is the rank inequalities, i.e., inequalities with binary coefficients but with arbitrary right-hand sides. Cornuéjols and Sassano (1989) proposed a necessary condition for these inequalities to be facets, while Sassano (1989) proposed a sufficient condition.

Instead of directly proposing facets of $\mathcal{Q}(B)$, a complementary approach to finding facets of $\mathcal{Q}(B)$ is to use dominance arguments to derive stronger inequalities from inequalities that are not facet defining. The approach referred to as *lifting* answers this question by posing the problem of finding the best set of coefficients for an inequality as a series of mixed-integer optimization problems (Padberg, 1973; Wolsey, 1976). More recently, Chvátal et al. (2013) introduced a more general approach called *tilting*. The authors show that a non-facet defining valid inequality can be rotated around its contact points with the polyhedron and remain valid. They also describe a procedure to perform the largest such rotation and show that the resulting inequality induces a face of higher dimension. Thus, by applying this procedure at most as many times as the dimension of the ambient space, one ends up with a facet-defining inequality. The number of linearly independent axes on which a rotation can be performed corresponds to the degrees of freedom associated with the possible rotations.

In this work, we aim to deepen the understanding of the set covering polytope $\mathcal{Q}(B)$ for inequalities with arbitrary coefficients. To that end, we introduce a new mathematical object for the study of the facets of the set covering polytope: the *tilting vectors*. The tilting vectors characterize the degrees of freedom on the possible rotations of a valid inequality. A facet-defining inequality has zero degrees of freedom. A valid inequality with $f > 0$ degrees of freedom is associated with a vector space of tilting vectors of dimension f. These tilting vectors can be used to generate tighter valid inequalities, inducing faces of higher dimensions. The main contributions of this paper are:

- We study the properties of tilting vectors and their relationship with facet-defining inequalities for the set covering polytope. In particular, we study how tilting vectors characterize facets.
- Given a non-facet defining valid inequality $\pi x \geqslant \pi_0$ and a tilting vector, we show how to obtain an inequality $\mu x \geqslant \mu_0$ inducing a face of higher dimension. Moreover, our procedure allows us to preserve the following separation property: if in addition to (π, π_0) we are also given a point x^* such that $\pi x^* < \pi_0$, then our procedure constructs (μ, μ_0) such that $\mu x^* < \mu_0$. We establish a relationship between the proposed procedure and that of Chvátal et al. (2013).
- We highlight that each null coefficient in a valid inequality may be associated to a specific tilting vector and that all the remaining tilting vectors have their support included in the support of the inequality. The implications are twofold. First, one can reduce the study of facets from general inequalities to inequalities without null coefficients. Second, the number of computations required to tilt a sparse inequality is significantly reduced as compared to the dense case.
- We use the proposed techniques to extend several results from the set covering literature. In particular, we extend to inequalities with arbitrary coefficients the necessary and sufficient conditions that Cornuéjols and Sassano (1989) and Sassano (1989) initially stated for rank inequalities. We also provide an alternative proof for the characterization introduced by Balas and Ng (1989) for the facet-defining inequalities with coefficients and right-hand side in $\{0, 1, 2\}$. Finally, we extend the results of Balas and Ng (1989) by characterizing the tilting vectors for these inequalities.

The remainder of this article is organized as follows. We begin by introducing notations and basic results in Section 2. In Section 3, we present the main results of the tilting theory introduced by Chvátal et al. (2013). Section 4 is dedicated to the introduction of the tilting vectors, their properties, and their specific use in the context of set covering. We then focus in Section 5 on how to exploit the null coefficients of set covering inequalities. In Section 6, we extend the necessary and sufficient conditions of Cornuéjols and Sassano (1989) and Sassano (1989) for rank inequalities to general inequalities. This last section also contains an extension of the work of Balas and Ng (1989) on inequalities with coefficients and right-hand sides in $\{0, 1, 2\}$. Finally, Section 7 concludes this article.

2 Notations and preliminaries

Following the notations in Sassano (1989); Cornuéjols and Sassano (1989), a set covering instance is associated with a bipartite graph $G = (\mathcal{E}, \mathcal{S}, A)$. The arc set A is composed of the pairs $(e, s) \in \mathcal{E} \times \mathcal{S}$ such that $e \in s$. In this case, a cover is a subset of nodes of $\mathcal{E}$ such that all nodes in $\mathcal{S}$ is connected to at least one node in the cover. An illustrative example can be found in Figure 1.

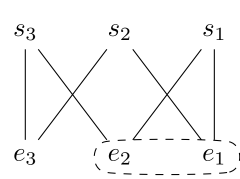

Figure 1: The bipartite graph associated to a set covering instance where $\mathcal{E} = \{1, 2, 3\}$ and $\mathcal{S} = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$. An example of cover (circled) is $\{1, 2\}$

For a given subset $E \subseteq \mathcal{E}$, $N(E) = \{s \in \mathcal{S} | \exists e \in E, (e, s) \in A\}$ denotes the neighbors of the elements in E. In this context, a cover can equivalently be defined as a set of nodes in $\mathcal{E}$ whose set of neighbors is $\mathcal{S}$. We will denote $\mathcal{Q}(G)$ the polytope of a set covering instance associated with a bipartite graph G.

Given a polyhedron $\mathcal{Q}$ and a valid inequality $\alpha x \geqslant k$ for this polyhedron, the face associated to this inequality is $F = \{x \in \mathcal{Q} | \alpha x = k\}$. A facet of a polyhedron is a face whose dimension is the dimension of the polyhedron minus one. To prevent our explanations from becoming too heavy and uneasy to read, we will not always differentiate the inequalities from the face they induce on the polyhedron $\mathcal{Q}(G)$. For the same reason, we will not always differentiate covers from their incidence vectors. This allows us, for example, to abuse the vocabulary and say that some covers are affinely independent or that an inequality is a facet.

For any vector $\alpha \in \mathbb{R}^{|\mathcal{E}|}$ and any sub-family $S \subset \mathcal{S}$, we will denote by $\gamma(\alpha, S)$ the minimum value of αx over the binary vectors $x \in \{0,1\}^{|\mathcal{E}|}$ representing a set of elements covering all the subsets in S. For simplicity, we will use $\gamma(\alpha)$ in place of $\gamma(\alpha, \mathcal{S})$. In this work, when considering an inequality $\alpha x \geqslant k$, we assume that its right-hand side is minimum, *i.e.*, $k = \gamma(\alpha) = \min_{x \in Q(G)} \alpha x$. To make this explicit, from now on we will denote $\alpha x \geqslant \gamma(\alpha)$ everywhere. Moreover, for a vector α and an integer p, we denote $E_p(\alpha) = \{e \in \mathcal{E} | \alpha_e = p\}$ the set of indices for which α has coefficient p.

A central object to the analysis of the strength of an inequality $\alpha x \geqslant \gamma(\alpha)$ is the set of covers that satisfy this inequality to equality. This set is denoted $C^{=}(\alpha)$ and contains most of the relevant information about the inequality. In particular, by definition, the inequality is a facet if and only if there are $|\mathcal{E}|$ affinely independent covers in $C^{=}(\alpha)$.

We now state a proposition taken from Nobili and Sassano (1989) that highlights some basic properties of the set covering polytope.

Proposition 1 (Nobili and Sassano (1989))**.**

1. *$\mathcal{Q}(G)$ is empty if and only if at least one subset in $\mathcal{S}$ is empty;*
2. *$\mathcal{Q}(G)$ is full-dimensional if and only if no subset in $\mathcal{S}$ is a singleton.*

If $\mathcal{Q}(G)$ is full-dimensional then, for each $e \in \mathcal{E}$:

1. *the inequality $x_e \geqslant 0$ defines a (trivial) facet of $\mathcal{Q}(G)$ unless for some subset $s \in \mathcal{S}$, $s \backslash \{e\}$ is a singleton;*
2. *the inequality $x_e \leqslant 1$ defines a (trivial) facet of $\mathcal{Q}(G)$;*
3. *every non-trivial facet of $\mathcal{Q}(G)$ is defined by an inequality of the form $\alpha x \geqslant \gamma(\alpha)$ where $\alpha \geqslant 0$ and $\gamma(\alpha) > 0$;*

From now on, we will consider that the set covering polytope $\mathcal{Q}(G)$ is full-dimensional. This happens if and only if no subset in $\mathcal{S}$ is a singleton and one can reduce any set covering instance to this case by setting to one the variable corresponding to elements present in singleton. Moreover, we will assume when considering an inequality $\alpha x \geqslant \gamma(\alpha)$ that $\alpha \geqslant 0, \gamma(\alpha) > 0$.

3 Tilting for general polyhedra

In this section, we present the tilting procedure introduced by Chvátal et al. (2013). Let $\mathcal{Q} \subset \mathbb{R}^n$ be a polyhedron and let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality separating a point from $\mathcal{Q}$ and inducing a face F. The main idea of tilting is to rotate the inequality around its contact points $\mathcal{Q}$ in order to create a separating inequality inducing a face of dimension higher than the one of F. This process can be repeated until a facet is obtained. Let us now formally introduce the concept of rotation of an inequality around a set of points.

Definition 1. Let $\pi x \geqslant \pi_0$ be an inequality satisfied at equality by a set of points X, i.e., $X \subseteq \{x \in \mathbb{R}^n : \pi x = \pi_0\}$. A rotation of $\pi x \geqslant \pi_0$ around X is another inequality $\mu x \geqslant \mu_0$ such that $X \subseteq \{x \in \mathbb{R}^n : \mu x = \mu_0\}$.

Now let $\mathcal{Q}^{\perp}$ be the vector space defining the *implicit equations* of $\mathcal{Q}$, *i.e.*, $\mathcal{Q}^{\perp}$ is the set of the vectors q for which qx is constant over $\mathcal{Q}$. Also assume that we know a set $\{x_1, ..., x_I\}$ of affinely independent vectors of the face F and a set $\{q_1, ..., q_K\}$ of linearly independent vectors of $\mathcal{Q}^{\perp}$. They can both be empty initially and their size is denoted by I and K, respectively. By definition, a facet of $\mathcal{Q}$ is a face of $\mathcal{Q}$ that contains $dim(\mathcal{Q})$ affinely independent points. Since $dim(\mathcal{Q}) + dim(\mathcal{Q}^{\perp}) = n$, if we can tilt the inequality $\alpha x \geqslant \gamma(\alpha)$ while extending the sets $\{x_1, ..., x_I\}$ and $\{q_1, ..., q_K\}$ until their sizes I and K sum to n, the resulting tilted inequality will be a facet of $\mathcal{Q}$.

Tilting the inequality $\alpha x \geqslant \gamma(\alpha)$ corresponds to rotating it around its contact points with $\mathcal{Q}$. Instead of considering all the possible degrees of freedom of rotation at once we will consider only one degree of freedom at a time and characterize it with a tilting direction. Formally, a tilting direction is a tuple $(\beta, \mu) \in \mathbb{R}^n \times \mathbb{R}$ such that

$$\beta x_i = \mu \quad \forall i \in \{1, ..., I\}. \tag{2a}$$

One can construct rotations of the inequality $\alpha x \geqslant \gamma(\alpha)$ by taking any linear combination of the inequalities $\alpha x \geqslant \gamma(\alpha)$ and $\beta x \geqslant \mu$. In particular, for a scalar λ, let us denote $I(\lambda, \alpha, \beta, \mu)$ the following rotation:

$$(\lambda\alpha + (1-\lambda)\beta)\, x \geqslant \lambda\gamma(\alpha) + (1-\lambda)\mu \tag{3}$$

When no confusion is induced, we will simply use $I(\lambda)$ instead of $I(\lambda, \alpha, \beta, \mu)$. Using the above concepts and a procedure illustrated in Algorithm 1, one can tilt an inequality until it becomes a facet. In the remainder of this section, we detail each step of the procedure.

Algorithm 1 Sketch of the tilting procedure

```
Step 0: Check whether αx ⩾ γ(α) is an implicit equation of Q;
while I + K < n do                                  ▷ I = |{x_1, ..., x_I}| and K = |{q_1, ..., q_K}
    Step 1: Obtain a tilting direction (β, μ);
    Step 2: Compute λ* the smallest scalar such that I(λ*) is valid for Q.
               At the same time a point x* ∈ Q affinely independent from
               {x_1, ..., x_I} and satisfying I(λ*) to equality will be computed;
    if I(λ*) is an implicit equation of Q then
        Step 3a: add the coefficient vector I(λ*) to {q_1, ..., q_K};
    else                                                        ▷ I(λ*) is a proper face of Q
        Step 3b: add x* to {x_1, ..., x_I} and replace the inequality
                    to be tilted by I(λ*)
```

Step 0. Check whether $\boldsymbol{\alpha x \geqslant \gamma(\alpha)}$ is an implicit equation of $\mathcal{Q}$: this can be done by maximizing αx over $\mathcal{Q}$. If the optimal value is $\gamma(\alpha)$, then $\alpha x \geqslant \gamma(\alpha)$ is an implicit equation. Otherwise, as a by-product, the maximization yields a point $\bar{x}$ of $\mathcal{Q}$ satisfying $\alpha\bar{x} > \gamma(\alpha)$. This makes $\bar{x}$ affinely independent from the face F.

Step 1. Obtaining a tilting direction: to that end, let us consider a non-null solution (β, μ) of the following system:

$$\beta q_k = 0 \quad \forall k \in \{1, ..., K\} \tag{4a}$$

$$\beta x_i = \mu \quad \forall i \in \{1, ..., I\} \tag{4b}$$

$$\beta \bar{x} = \mu. \tag{4c}$$

Note that if the system above admits the null vector as a unique solution, then the inequality $\alpha x \geqslant \gamma(\alpha)$ is already a facet (Chvátal et al., 2013).

Step 2. Tilt the inequality: Upon finding a tilting direction (β, μ), we find the smallest scalar λ^* such that the inequality $I(\lambda^*)$ remains a valid inequality for $\mathcal{Q}$. The resulting inequality is called the tilted inequality. Note that λ^* must be greater than 0 because for all $\lambda < 0$, $I(\lambda)$ is not valid for $\bar{x}$ and thus not valid for $\mathcal{Q}$. An illustration of an original inequality, a tilting direction (β, μ) and a tilted inequality is provided in Figure 2. This tilting step revolve around the following central result of the tilting routine.

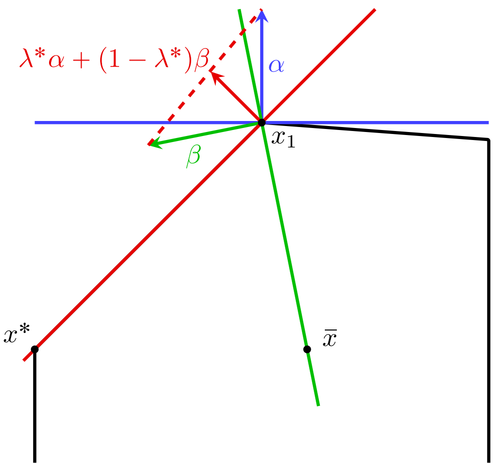

Figure 2: A valid inequality (blue) for a polyhedron with normal vector α, a tilting direction β (green) and the corresponding tilted inequality (red)

Theorem 1. *The tilted inequality $I(\lambda^*)$ is satisfied to equality by a point x^* of $\mathcal{Q}$ affinely independent from $\{x_1, ..., x_I\}$.*

Proof. Let us first exclude the case $\lambda^* = 0$ since in this case we can take $x^* = \bar{x}$.

Let $(\lambda_j)_{j\in\mathbb{N}}$ be a sequence converging towards λ^* with $\lambda^* > \lambda_j > 0$. By definition of λ^*, for every term in this sequence, the inequality $I(\lambda_j)$ is not valid for $\mathcal{Q}$. Thus, let us associate each λ_j with a point y_j of $\mathcal{Q}$ not satisfying the inequality $I(\lambda_j)$. Let us note that since the inequality $I(\lambda_j)$ is not valid for $\mathcal{Q}$, either a vertex of $\mathcal{Q}$ does not satisfy it, in which case, we set y_j equal to this vertex; or an extreme ray of $\mathcal{Q}$ exists such that $(\lambda_j\alpha + (1-\lambda_j)\beta)r < 0$, in which case, we set y_j to be $x_I + r$ (which does not satisfy $I(\lambda_j)$ since x_I satisfies it to equality).

Now, since there is only a finite number of vertices and extreme rays of $\mathcal{Q}$, the sequence $(x_j)_{j\in\mathbb{N}}$ is a converging sequence over a finite set. Thus, it contains a constant sub-sequence. Let us define the candidate point x^* as the unique point in this constant sub-sequence and let us denote $(\lambda'_j)_{j\in\mathbb{N}}$ the associated sub-sequence of λ_j.

If x^* could be written as an affine combination of the points in $\{x_1, ..., x_I\}$ then it would satisfy to equality all the inequalities $I(\lambda'_j)$. Since this is not true by construction of the sequence $(x_j)_{j\in\mathbb{N}}$ and thus by construction of x^*, the point x^* is affinely independent from $\{x_1, ..., x_I\}$.

Let us now finish the proof by showing that x^* is satisfying the tilted inequality $I(\lambda^*)$ to equality. To that end let us denote $\delta(\lambda) = (\lambda\alpha + (1-\lambda)\beta)x^* - \lambda\gamma(\alpha) + (1-\lambda)\mu$. Since the tilted inequality $I(\lambda^*)$ is valid for $\mathcal{Q}$ we have $\delta(\lambda^*) \geqslant 0$. Moreover, since the sequence $(\lambda'_j)_{j\in\mathbb{N}}$ converges towards λ^* and x^* does not satisfy $I(\lambda'_j)$, the sequence $(\delta(\lambda'_j))_{j\in\mathbb{N}}$ is a sequence of negative numbers converging towards $\delta(\lambda^*)$. Thus, $\delta(\lambda^*) = 0$ which means that the point x^* satisfies the tilted inequality $I(\lambda^*)$ to equality. □

The result of Theorem 1 relies on two properties of the point $\bar{x}$: $\bar{x} \in \mathcal{Q}$ and $\bar{x}$ is affinely independent from $\{x_1, ..., x_I\}$, which is implied by the fact that this point does not satisfy the inequality $\alpha x \geqslant \gamma(\alpha)$ to equality. To perform the computation of λ^* one can solve a non-linear program or make a sequence of calls to an oracle optimizing a linear function over $\mathcal{Q}$. These methods also provide the point x^* as a by-product. We refer to the works of Espinoza et al. (2010) and Chvátal et al. (2013) for a presentation on how to compute λ^* for general polyhedra. A tailored method for the set covering case is discussed in Section 4.4.

Step 3a. Case: $\boldsymbol{I(\lambda^*)}$ is an implicit equation of $\boldsymbol{\mathcal{Q}}$: this can be identified by maximizing $(\lambda^*\alpha + (1-\lambda^*)\beta)x$ over $\mathcal{Q}$. If the result of this maximization equals $\lambda^*\gamma(\alpha) + (1-\lambda^*)\mu$, then the tilted inequality is an implicit equation of $\mathcal{Q}$. Otherwise, it induces a proper face. Note that, when $I(\lambda^*)$ is an implicit equation, we have $\lambda^* = 0$ as otherwise $\bar{x}$ can be shown to satisfy $\alpha\bar{x} = \gamma(\alpha)$ (which is by hypothesis false) using the following identities:

$$\begin{aligned}\alpha\bar{x} =& \frac{\lambda^*\alpha\bar{x}}{\lambda^*}\\ =& \frac{(\lambda^*\alpha + (1-\lambda^*)\beta)\bar{x} - (1-\lambda^*)\beta\bar{x}}{\lambda^*}\\ =& \frac{\lambda^*\gamma(\alpha) + (1-\lambda^*)\mu - (1-\lambda^*)\mu}{\lambda^*}\\ =& \gamma(\alpha).\end{aligned}$$

Thus, in this case the tilted inequality $I(\lambda^*)$ is in fact $\beta x \geqslant \mu$. Equations (4a) ensure that β is linearly independent from the set $\{q_1, ..., q_K\}$. It can thus be added to it as a new independent implicit equation. The known dimensions of the space of implicit equations has been increased and one can retry tilting the inequality $\alpha x \geqslant \gamma(\alpha)$ with a new solution of system (4a)–(4c) once it has been updated with the new implicit equation. The point $\bar{x}$ can remain unchanged.

Step 3b. $\boldsymbol{I(\lambda^*)}$ is a proper face of $\boldsymbol{\mathcal{Q}}$: Since $I(\lambda^*)$ is a rotation of the original inequality, its induced face contains $\{x_1, ..., x_I\}$. Moreover, Theorem 1 shows it also contains a new independent point x^*. Thus, the point x^* can be added to $\{x_1, ..., x_I\}$, effectively increasing the known dimensions of the induced face. One can continue the tilting procedure by tilting the inequality $I(\lambda^*)$. The only missing piece to continue the procedure is a point $\bar{x}$ to define the system (4a)–(4c). If $\lambda^* \neq 0$, it can remain unchanged; otherwise $\bar{x}$ can be taken as the maximizer of $(\lambda^*\alpha + (1-\lambda^*)\beta)x$ over $\mathcal{Q}$ which has already been computed to verify if $I(\lambda^*)$ is an implicit equation or a proper face of $\mathcal{Q}$.

4 Tilting for set covering

In this section we start by comparing tilting for set covering to tilting in the general case. Then, we introduce the main concept of this article, the tilting vectors. After presenting how tilting vectors characterize the dimension of the face induced by an inequality, we discuss how tilting vectors can be used to tilt inequalities in the context of set covering. Finally, we discuss some computational aspects of the tilting procedure and show that, in the case of set covering polytopes, our procedure is particularly efficient to tilt inequalities $\alpha x \geqslant \gamma(\alpha)$ with α sparse.

4.1 Comparison to the general case and definitions

Compared to tilting for general polyhedra, several simplifications occur in the context of set covering. First, the set covering polytope is compact and always assumed to be full-dimensional. Thus, it has neither implicit equations, nor extreme rays. Moreover, we will change the assumptions made on the point $\bar{x}$ used in Equation (4c) to define the tilting direction (β, μ) in the general case. We will use $\bar{x} = 0$. Amongst the properties required to prove Theorem 1, the origin still satisfies the very important condition $\alpha\bar{x} \neq \gamma(\alpha)$ (since $\gamma(\alpha) > 0$; see Proposition 1). However, it does not belong to the set covering polytope. This prevents the use of Theorem 1 in tilting for set covering and we will rely on similar results presented later in this Section. Using the origin, one can simplify Equation (4c) to $\mu = 0$ and the whole system (4a)–(4c) reduces to:

$$\beta x_i = 0 \quad \forall i \in \{1, ..., I\}. \tag{6}$$

With μ being null, the tilted inequality will have the same right-hand side as the original one. This makes sense for set covering since all non-trivial valid inequalities have the form $\alpha x \geqslant \gamma(\alpha)$ with $\alpha \geqslant 0$

and $\gamma(\alpha) > 0$ (see Proposition 1). Thus, one can switch between any two inequalities by changing only the variable coefficients since, by re-scaling the inequality, one can ensure that the right-hand side remains constant.

Because we are not only interested in tilting *per se* but also in its relationship with the characterization of facets, we will assume that we have complete information about the face induced by the inequality $\alpha x \geqslant \gamma(\alpha)$. More precisely, the complete set $C^{=}(\alpha)$ is assumed to be known instead of a subset of affinely independent vectors $\{x_1, ..., x_I\}$. Thus, the system (6) will contain an equation for each point in $C^{=}(\alpha)$. This basically ensures that we have a complete description of the face induced by the inequality that is being tilted. We make this assumption to simplify the presentation but this hypothesis can be relaxed for the tilting results.

Let us now introduce the main concept of this article, the tilting vectors.

Definition 2. Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$ and let M_α be the matrix whose rows are the covers in $C^{=}(\alpha)$. A tilting vector β is a solution of the system $M_\alpha \beta = 0$.

In other words, a vector β is a tilting vector of the inequality $\alpha x \geqslant \gamma(\alpha)$ if it is orthogonal to all the covers in $C^{=}(\alpha)$. Note that, as the solution of a linear system, the set of tilting vectors of an inequality is a vector space. Moreover, although we should write "the tilting vectors of inequality $\alpha x \geqslant \gamma(\alpha)$", for the sake of brevity, we will omit the reference to the studied inequality when this creates no confusion. A more graphical way to define tilting vectors, illustrated in Figure 3, is the following. Let $\mathcal{H}(\alpha) = (\mathcal{E}, C^{=}(\alpha))$ be the hypergraph with node set $\mathcal{E}$ and with each hyper-edge representing a cover of $C^{=}(\alpha)$. A tilting vector is an assignment of weights β_e to the nodes of $\mathcal{H}(\alpha)$ such that the sum of the weights on the nodes of each hyper-edge is zero. An illustration of this graphical definition is given in Figure 3.

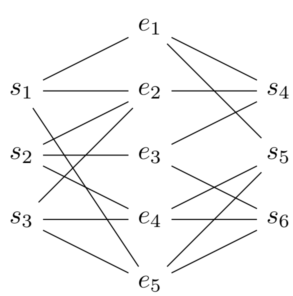

(a) An instance of set covering

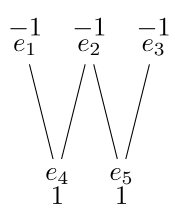

(b) The hypergraph $\mathcal{H}(\alpha)$ for the inequality $x_1 + x_2 + x_3 + x_4 + x_5 \geqslant 2$ and the coefficients of a tilting vector

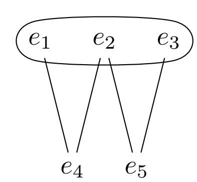

(c) The hypergraph $\mathcal{H}(\alpha)$ for the inequality $x_1 + x_2 + x_3 + 2x_4 + 2x_5 \geqslant 3$ for which the only tilting vector is null

Figure 3: A set covering instance: hypergraphs and tilting vectors corresponding to two inequalities

What makes the strength of the tilting vectors is that they can be used as tilting directions to obtain inequalities inducing higher dimensional faces and they are a useful tool to study and characterize the facets of the set covering polytope. In Section 4.2, we introduce some of their basic relations to facets and their role in tilting.

4.2 Characterizing facets with tilting vectors

In this section, we highlight the relationship between the space of tilting vectors of an inequality and the dimension of the face it induces. We start with a proposition showing that the existence of tilting vectors characterize whether an inequality is a facet or not.

Proposition 2. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$. It is a facet of $\mathcal{Q}(G)$ if and only if its only tilting vector is the null vector.*

Proposition 2 is a direct corollary of the following theorem linking the dimension of the space of tilting vectors and dimension of the face induced by an inequality.

Theorem 2. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$, let D_α be the dimension of its induced face and let D_β be the dimension of the space of its tilting vectors. Then, $D_\alpha + D_\beta = |\mathcal{E}| - 1$.*

Proof. By definition, the dimension of the induced face is $N - 1$ where N is the number of affinely independent points of $\mathcal{Q}(G)$ that satisfy $\alpha x = \gamma(\alpha)$. Let M_α be the matrix whose rows are the covers in $C^=(\alpha)$, with $Im(M_\alpha)$ being its image, *i.e.*, $\{x | \exists y, x = M_\alpha y\}$, and $Ker(M_\alpha)$ being its kernel, *i.e.*, $\{x | M_\alpha x = 0\}$. We will show that $N = dim(Im(M_\alpha))$ and, since by definition the space of tilting vectors is $Ker(M_\alpha)$, we obtain from the rank theorem of matrices that $dim(Im(M_\alpha)) + D_\beta = |\mathcal{E}|$, which implies $D_\alpha + D_\beta = |\mathcal{E}| - 1$.

Thus, let us prove that $N = dim(Im(M_\alpha))$. First, the number of affinely independent points of $\mathcal{Q}(G)$ that satisfy $\alpha x = \gamma(\alpha)$ is also the number of affinely independent covers in $C^=(\alpha)$. Second, since the null vector does not satisfy $\alpha 0 = \gamma(\alpha)$, any affinely independent points of the hyperplane $\alpha x = \gamma(\alpha)$ are equivalently linearly independent. Thus, N is the number of linearly independent covers of $C^=(\alpha)$. This is exactly the dimension of $Im(M_\alpha^T)$ which is equal to $dim(Im(M_\alpha))$. □

Proposition 2 is similar to Lemma 1 from Chvátal et al. (2013) for general tilting. The generalization from Proposition 2 to Theorem 2 can also be made for arbitrary polyhedra.

4.3 Tilting inequalities

Proposition 2 and Theorem 2 show that the existence of non-null tilting vectors certifies that the corresponding inequality is not a facet of $\mathcal{Q}(G)$. In such case, it would be helpful to derive a facet or at least a face of higher dimension. This can be achieved by tilting the inequality as presented in Algorithm 1, but we present here an adaptation for the set covering case. Let us start by defining the tilting of an inequality as follows.

Definition 3. A tilting of an inequality $\alpha x \geqslant \gamma(\alpha)$ with a tilting vector β is any valid inequality of the form $(\alpha + \epsilon\beta)x \geqslant \gamma(\alpha)$ where ϵ is a scalar.

As in the general case, the tilted inequality is a rotation of the original inequality around the set of contact points $C^=(\alpha)$ with the polytope $\mathcal{Q}(G)$. This is induced by $\beta x_c = 0$ for all $c \in C^=(\alpha)$ in the definition of the tilting vector β.

The next proposition shows that facets can be obtained in one tilting operation.

Proposition 3. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$. For each non-trivial facet containing $C^=(\alpha)$, there exists a tilting vector that can be used to tilt the valid inequality into the facet.*

Proof. Since we are considering non-trivial facets of the set covering polytope, the facet can be written $\alpha' x \geqslant \gamma(\alpha)$ through the right scaling. Take $\epsilon = 1$ and $\beta = \alpha' - \alpha$ as tilting vector. It is indeed a tilting vector since for each cover c in $C^=(\alpha)$ we have $\beta x_C = (\alpha' - \alpha)x^c = \gamma(\alpha) - \gamma(\alpha) = 0$. □

Although one can obtain any facet containing $C^=(\alpha)$ in one tilting operation, this requires the knowledge of the corresponding tilting vector which can be as difficult to obtain as the facet. However, any tilting vector can be used to obtain a face of higher dimension than the one of the original inequality. Compared to the general case, tilting using tilting vectors does not change the right-hand side of the inequality. This prevents the tilting operation from reaching the trivial facets (which do not have a positive right-hand side). Thus, two special cases are required for the two types of trivial facets ($x \geqslant 0$ and $x \leqslant 1$). We first present a theorem showing that tilting with tilting vectors does lead to higher dimensional faces outside of the two special cases. Those cases are dealt with in the subsequent two theorems. The fact that the tilted inequalities induce higher dimensional faces is to be related to with the existence the point x^* affinely independent from $\{x_1, ..., x_I\}$ in Theorem 1.

Theorem 3. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$ such that, for each element $e \in \mathcal{E}$, there is at least one cover in $C^=(\alpha)$ that contains e, and one that does not. Let β be a tilting vector for that inequality. Then, there exist positive numbers ϵ^+, ϵ^- such that the tilted inequalities*

$$(\alpha + \epsilon^+\beta)x \geqslant \gamma(\alpha) \tag{7}$$

$$(\alpha - \epsilon^-\beta)x \geqslant \gamma(\alpha) \tag{8}$$

are both valid inequalities inducing faces of strictly higher dimension.

Proof. *Sketch of the proof: we will show that, for $\epsilon \in [\epsilon_{min}, \epsilon_{max}]$, the inequality $(\alpha + \epsilon\beta)x \geqslant \gamma(\alpha)$ is valid for $\mathcal{Q}(G)$. Then we will show that with $\epsilon^- = -\epsilon_{min}$ and $\epsilon^+ = \epsilon_{max}$, we obtain faces of higher dimension.*

We will make the proof only for ϵ^+ as the ϵ^- case is similar.

The inequality $(\alpha + \epsilon^+\beta)x \geqslant \gamma(\alpha)$ is invalid for a cover c if and only if $\beta x^c < 0$ and $\epsilon^+ > \frac{\gamma(\alpha) - \alpha x^c}{\beta x^c}$. Thus, let us define ϵ_{max} as follows:

$$\epsilon_{max} = \min_{x^c \in \mathcal{Q}(G), \beta x^c < 0} \frac{\gamma(\alpha) - \alpha x^c}{\beta x^c}$$

$\boldsymbol{\epsilon_{max}}$ **positive:** there are finitely many covers and for the covers for which $\gamma(\alpha) - \alpha x^c = 0$ we also have $\beta x^c = 0$. Thus, ϵ_{max} is the minimum over finitely many positive numbers which makes it positive.

$\boldsymbol{\epsilon_{max}}$ **finite:** we now show that under the assumption that for each element $e \in \mathcal{E}$ there is a cover in $C^=(\alpha)$ that contains e and one that does not, ϵ_{max} is finite. Since $\beta \neq 0$, it has at least one non-zero coefficient. The corresponding element is contained in a cover of $C^=(\alpha)$ for which $\beta x^c = 0$. Thus, β has at least one other coefficient of the opposite sign and thus has at least one negative coefficient. Thus, for high enough values of ϵ^+, the vector $\alpha + \epsilon^+\beta$ has a negative coefficient for an element e'. There is at least one cover c in $C^=(\alpha)$ not containing e', thus, $c \cap \{e'\}$ is a cover satisfying:

$$\begin{aligned}(\alpha + \epsilon^+\beta)x_{c \cap \{e'\}} =& (\alpha + \epsilon^+\beta)x^c + (\alpha + \epsilon^+\beta)_{e'} \\ =& \gamma(\alpha) + \epsilon^+ 0 + (\alpha + \epsilon^+\beta)_{e'} \\ <& \gamma(\alpha).\end{aligned}$$

Thus, for values of ϵ^+ where $\alpha + \epsilon^+\beta$ has a negative coefficient, the tilted inequality (7) is not valid. Thus, ϵ_{max} is finite as it is upper bounded by the largest value of ϵ such that $\alpha + \epsilon\beta \geqslant 0$.

Higher dimension face: let us take $\epsilon^+ = \epsilon_{max}$ and let c^* be a cover minimizing $\frac{\gamma(\alpha) - \alpha x^c}{\beta x^c}$ when $\beta x^c < 0$. We will show that the face of the tilted contains $C^=(\alpha) \cup \{c^*\}$. First, ϵ_{max} is defined so that every minimizer such as c^* will satisfy the tilted inequality to equality. On the other hand, the covers in $C^=(\alpha)$ also satisfy the tilted inequality (7) to equality as they are orthogonal to β (remember that the goal of tilting is to rotate the original inequality around its contact points with $\mathcal{Q}(G)$ in order to keep them in the face). Moreover, $\beta x_{c*} \neq 0$ thus the cover c^* is affinely independent of the covers in $C^=(\alpha)$. Thus, the tilted inequality induces a face of strictly higher dimension. □

Theorem 3 shows that under an unique condition, any non-facet defining inequality with positive coefficients can be decomposed as the sum of two non-trivial inequalities that induce faces of higher dimension. When this condition is not satisfied, it can be decomposed as the sum of one of the trivial facets ($x_e \geqslant 0$ or $x_e \leqslant 1$) and one other non-trivial inequality. These particular cases yield similar results and are treated in the next two theorems.

Theorem 4. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality and e an element of $\mathcal{E}$. Let η be such that $\eta + \gamma(\alpha)$ is the minimum of αx^c over the set of covers c not containing e. The following propositions are equivalent:*

(i) *e is contained in all covers in $C^=(\alpha)$;*

(ii) *$\alpha - \gamma(\alpha)e_e$ is a tilting vector;*

(iii) *η is positive and $\alpha x \geqslant \gamma(\alpha)$ is the sum of $-\eta x_e \geqslant -\eta$ and $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$ with both inequalities being valid and the last one inducing a face of higher dimension.*

Proof. $(i) \Rightarrow (ii)$ If e is contained in all the covers in $C^=(\alpha)$ then for each cover in $C^=(\alpha)$, $(\alpha - \gamma(\alpha)e_e)x = 0$. Thus, the vector $\alpha - \gamma(\alpha)e_e$ is a tilting vector.

$(i) \Rightarrow (iii)$ Suppose that e is contained in all the covers in $C^=(\alpha)$. Then, all the covers c not containing e must satisfy $\alpha x_c > \gamma(\alpha)$. Thus, η is positive. Let us now show that the tilted inequality $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$ is valid for $\mathcal{Q}(G)$. For all $x \in \mathcal{Q}(G)$ we have $\alpha x \geqslant \gamma(\alpha)$. Thus, the incidence vector of each cover containing e satisfies $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$. Moreover, by definition of η, the incidence vector of each cover not containing e satisfies $\alpha x \geqslant \gamma(\alpha) + \eta$ and therefore also $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$. Thus, $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$ is valid for $\mathcal{Q}(G)$. Finally, let c^* be a cover that minimizes αx^c over the set of covers c not containing e. It must satisfy $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$ to equality. Note that all the covers c in $C^=(\alpha)$ also satisfy $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$ to equality as they contain e. Thus, the face associated with $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$ contains $C^=(\alpha) \cup \{c^*\}$. Since c^* contains e while the covers in $C^=(\alpha)$ do not, the cover c^* is affinely independent of the covers in $C^=(\alpha)$. Thus, the tilted inequality induces a face of strictly higher dimension.

$(iii) \Rightarrow (i)$ Suppose that $\alpha x \geqslant \gamma(\alpha)$ is the sum of $-\eta x_e \geqslant -\eta$ and another valid inequality with $\eta > 0$. Since both inequalities of the decomposition are valid for $\mathcal{Q}(G)$, then the incidence vector of of $\mathcal{Q}(G)$ satisfying $\alpha x = \gamma(\alpha)$ (*i.e.*, corresponding to covers of $C^=(\alpha)$) must satisfy both inequalities to equality. Thus, we have in particular $-\eta x_e = -\eta$ which implies $x_e = 1$ since $\eta > 0$. Thus, e is contained in all covers in $C^=(\alpha)$.

$(ii) \Rightarrow (i)$ Suppose that $\alpha - \gamma(\alpha)e_e$ is a tilting vector. For each cover c in $C^=(\alpha)$, we have $\alpha x^c = \gamma(\alpha)$ and $(\alpha - \gamma(\alpha)e_e)x^c = 0$. Thus, $\gamma(\alpha)e_e x^c = \gamma(\alpha)$ which means $x_e^c = 1$. □

Theorem 5. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality and e an element of $\mathcal{E}$. Let η be such that $\eta + \gamma(\alpha)$ is the minimum of αx^c over the set of covers c containing e. The following propositions are equivalent:*

(i) *e is not contained in any cover of $C^=(\alpha)$;*

(ii) *e_e is a tilting vector;*

(iii) *η is positive and $\alpha x \geqslant \gamma(\alpha)$ is the sum of $\eta x_e \geqslant 0$ and $\alpha x - \eta x_e \geqslant \gamma(\alpha)$, with both inequalities being valid and the last one inducing a face of higher dimension.*

Proof. $(i) \Leftrightarrow (ii)$ e is not contained in any cover of $C^=(\alpha)$ if and only if for each cover in $C^=(\alpha)$, $xe_e = x_e = 0$ which is the definition of e_e being a tilting vector.

$(i) \Rightarrow (iii)$ Suppose that e is not contained in any cover of $C^=(\alpha)$. Then, all the covers c containing e must satisfy $\alpha x^c > \gamma(\alpha)$, thus η is positive. Let us now show that the tilted inequality $\alpha x - \eta x_e \geqslant \gamma(\alpha)$ is valid for $\mathcal{Q}(G)$. For all $x \in \mathcal{Q}(G)$ we have $\alpha x \geqslant \gamma(\alpha)$. Thus, the incidence vector of each cover not containing e satisfies $\alpha x - \eta x_e \geqslant \gamma(\alpha)$. Moreover, by definition of η, the incidence vector of each cover containing e satisfies $\alpha x \geqslant \gamma(\alpha) + \eta$ and thus also $\alpha x - \eta x_e \geqslant \gamma(\alpha)$. Thus, $\alpha x - \eta x_e \geqslant \gamma(\alpha)$ is valid for

$\mathcal{Q}(G)$. Finally, let c^* be a cover that minimizes αx^c over the set of covers c containing e. It must satisfy $\alpha x - \eta x_e \geqslant \gamma(\alpha)$ to equality. Note that all the covers c in $C^=(\alpha)$ also satisfy $\alpha x - \eta x_e \geqslant \gamma(\alpha)$ to equality as they do not contain e. Thus, the face associated with $\alpha x - \eta x_e \geqslant \gamma(\alpha)$ contains $C^=(\alpha) \cup \{c^*\}$. Since c^* does not contain e while the covers in $C^=(\alpha)$ do, the cover c^* is affinely independent of the covers in $C^=(\alpha)$. Thus, the tilted inequality induces a face of strictly higher dimension.

(iii) $\Rightarrow$ *(i)* Suppose that $\alpha x \geqslant \gamma(\alpha)$ is the sum of $\eta x_e \geqslant 0$ and another valid inequality with $\eta > 0$. Since both inequalities of the decomposition are valid for $\mathcal{Q}(G)$ then the incidence vector of of $\mathcal{Q}(G)$ satisfying $\alpha x = \gamma(\alpha)$ (*i.e.* corresponding to covers of $C^=(\alpha)$) must satisfy both inequalities to equality. Thus, we have in particular $\eta x_e = 0$ which implies $x_e = 0$ since $\eta > 0$. Thus, e is not contained in any cover of $C^=(\alpha)$. □

In the context of the separation of a point from $Q(G)$, it is convenient to be able to decompose a separating inequality as the sum of two valid inequalities inducing faces of higher dimension. Indeed, since the inequalities of the decomposition sum to the original inequality, then at least one of them must also separate the point $\hat{x}$ from $Q(G)$. Thus, given a separating inequality, tilting guarantees the generation of another separating inequality of higher dimension. In the case of Theorem 5, only the non-trivial inequality is guaranteed to be stronger than the original one (in Theorem 4, the trivial inequality is always a facet). This is, however, not an issue in practice. Indeed, in most cases, when a point is separated from $\mathcal{Q}(G)$, it is the solution of a (linear) relaxation of the set covering model. In this case, the separated point verifies the trivial inequalities. Thus, the non-trivial inequality of the decomposition, which is guaranteed to be stronger, must be separating the point from $\mathcal{Q}(G)$.

4.4 Computation of the tilted inequalities

Let us now discuss how to compute the optimal tilted inequalities. We start by discussing the two special cases treated in Theorems 4 and 5, where one of the tilted inequalities is one of the trivial facets ($x_e \geqslant 0$ or $x_e \leqslant 1$). We then focus on the general case given by Theorem 3.

Special cases of Theorems 4 and 5: in order to be able to apply Theorem 3 and rule out the cases treated in Theorems 4 and 5, one needs to tilt the original inequality $\alpha x \geqslant \gamma(\alpha)$ so that each element $e \in \mathcal{E}$ is contained in a cover of $C^=(\alpha)$ and not contained in another cover of $C^=(\alpha)$. Assuming that at least one cover c in $C^=(\alpha)$ is known (it can be obtained by minimizing αx over $\mathcal{Q}(G)$), then for each element $e \in \mathcal{E}$, we know that either $e \in c \in C^=(\alpha)$ or $e \notin c \in C^=(\alpha)$. Thus, one needs to check only one of the two conditions of each element e. This can be done for each element by minimizing αx over $\mathcal{Q}(G)$ with the additional constraint $x_e = 1 - x_e^c$. Let us call η the minimal value and x^* the minimizer. If $x_e^c = 1$, then Theorem 4 tells us to replace the original inequality $\alpha x \geqslant \gamma(\alpha)$ by $\alpha x + \eta x_e \geqslant \gamma(\alpha) + \eta$; otherwise Theorems 4 tells us to replace the original inequality $\alpha x \geqslant \gamma(\alpha)$ by $\alpha x - \eta x_e \geqslant \gamma(\alpha)$. In both cases, the face induced by the new inequality contains x^c and the minimizer x^* which correspond to one cover containing e and one cover not containing e. Overall, the special cases can be treated in one initial call to the set covering oracle plus one call for each element.

General case of Theorem 3: the value of ϵ^+, and similarly ϵ^-, can be determined with the following linear program:

$$\max \ \epsilon \tag{9}$$

$$\text{subject to } (\alpha + \epsilon\beta)x_i \geqslant \gamma(\alpha) \quad \forall x_i \in \mathcal{Q}(G) \tag{10}$$

$$\epsilon \in \mathbb{R}. \tag{11}$$

This linear program can be solved with row generation and deciding whether there exists a row cutting of a value ϵ^* amounts to minimize $(\alpha + \epsilon^*\beta)x$ over $\mathcal{Q}(G)$. This is simply a call to a set covering oracle. At the end of the row generation, the only active row will correspond to a point that satisfies the tilted inequality to equality and is not orthogonal to β (otherwise ϵ would disappear from the constraint). This second fact implies that the point is affinely independent from the points in $C^=(\alpha)$. Note that since the above linear program has only one variable, no real linear programming machinery is required to solve it. Thus, the row generation algorithm can be replaced by a simple algorithm iteratively calling a set covering oracle. This algorithm generates a sequence of candidates $\epsilon_1^*, \epsilon_2^*, \ldots$ together with a sequence $x_1^*, x_2^*, \ldots$ where x_i^* minimizes $(\alpha + \epsilon^* \beta)x_i^*$ over $\mathcal{Q}(G)$. This iterative algorithm and sequences are illustrated in Figure 4.

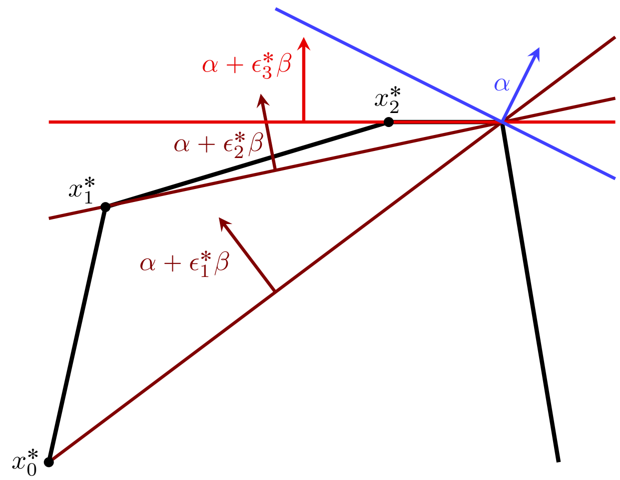

Figure 4: Illustration of the iterative algorithm solving the problem (9)–(11).

Another approach that finds ϵ^+ in only one call to a Branch & Bound algorithm is the following. For each point $y \in \mathcal{Q}(G)$ with $\beta y < 0$, one can define $\epsilon_y = \frac{\gamma(\alpha) - \alpha y}{\beta y}$ and y then satisfies the inequality $(\alpha + \epsilon_y \beta)x \geqslant \gamma(\alpha)$ to equality. This inequality is either invalid if there is a point of $\mathcal{Q}(G)$ violating it, or it is the optimal tilted inequality we are looking for. Thus, the point of $\mathcal{Q}(G)$ maximizing the smallest violation of this set of inequalities is on the optimal tilted inequality. Therefore, the following mixed integer linear program computes the optimal tilted inequality:

$$\max z \tag{12}$$

subject to

$$z \leqslant \gamma(\alpha) - (\alpha - \epsilon_y \beta)x \quad \forall y \in \mathcal{Q}(G) | \beta y < 0 \tag{13}$$

$$x \in \mathcal{Q}(G) \tag{14}$$

$$z \in \mathbb{R}. \tag{15}$$

This program can be solved with a Branch & Cut algorithm. One can start with only one of the constraints (13) and every time an integer solution y^* with positive objective value is found, one can add to the problem the constraint associated to y^*.

5 Taking advantage of the null coefficients in set covering inequalities

In this section, we show how to take advantage of the null coefficients in an inequality. In particular, we highlight that each null coefficient in a valid inequality may be associated to a specific tilting vector and that all the remaining tilting vectors have their support included in the support of the inequality. The implications are twofold. First, we will show that one can reduce the study of facets from general inequalities to inequalities without null coefficients. Second, we show that the number of computations required to tilt a sparse inequality is significantly reduced as compared to the dense case.

5.1 Tilting vectors associated with null coefficients of α

The next theorem highlights the special place that occupy the null coefficients of an inequality. In short, for an element $e_0 \in E_0(\alpha)$, *i.e.*, such that $\alpha_{e_0} = 0$, an important information is to know whether e_0 is contained in all the covers of $C^=(\alpha)$. If this is not the case, all tilting vectors have a null coefficient for e_0. Otherwise, $(\alpha - \gamma(\alpha)e_{e_0})$ is a tilting vector.

Theorem 6. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$. Let E_0^* be the set of elements of $E_0(\alpha)$ contained in all the covers of $C^=(\alpha)$. A basis of the space of tilting vectors is $(\alpha - \gamma(\alpha)e_{e_0})_{e_0 \in E_0^*} \cup (\beta^1, ..., \beta^p)$, where $(\beta^1, ..., \beta^p)$ is a basis of the tilting vectors satisfying $\beta_{e_0} = 0$ for each $e_0 \in E_0(\alpha)$.*

Proof. The vectors $(\alpha - \gamma(\alpha)e_{e_0})_{e_0 \in E_0^*}$ are tilting vectors since the elements of E_0^* are contained in all the covers of $C^=(\alpha)$.

A vector $\alpha - \gamma(\alpha)e_{e_0}$ is linearly independent from the other vector of this form and from $(\beta^1, ..., \beta^p)$ since it is the only vector in the basis that has a non-zero coefficient for e_0. Moreover, it is assumed that the vectors $(\beta^1, ..., \beta^p)$ are independent from each other so that they can be part of a basis. Thus, we only need to show that all tilting vectors can be generated using the aforementioned vectors.

Let us now consider a tilting vector β. Note that for all e_0 in $E_0(\alpha) \subset E_0^*$, there is a cover c in $C^=(\alpha)$ not containing e_0. Also note that $c \cup \{e_0\}$ belongs to $C^=(\alpha)$. Thus, $\beta_{e_0} = \beta(x^{c \cup \{e_0\}} - x^c) = 0 - 0 = 0$. Let us denote δ^{e_0} the vector $\frac{\beta_{e_0}}{\gamma(\alpha)}(\alpha - \gamma(\alpha)e_{e_0})$. It must be a tilting vector since $\alpha - \gamma(\alpha)e_{e_0}$ is one. Note that δ^{e_0} has only null coefficients for the elements in $E_0(\alpha)$ expect for e_0 for which the coefficient is β_{e_0}. Thus, the vector $\beta - \sum_{e_0 \in E_0^*} \delta^{e_0}$ is a tilting vector with null coefficients for all elements in $E_0(\alpha)$. Thus, it can be decomposed as a linear combination of the $(\beta^1, ..., \beta^p)$. Therefore, β can be decomposed as a linear combination of the $(\beta^1, ..., \beta^p)$ and the $(\alpha - \gamma(\alpha)e_{e_0})_{e_0 \in E_0^*}$. □

The above theorem shows that apart from the vectors $(\alpha - \gamma(\alpha)e_{e_0})_{e_0 \in E_0^*}$, the other tilting vectors have null coefficients for elements for which α does. As it turns out, these remaining tilting vectors correspond to the tilting vectors of $\alpha' x \geqslant \gamma(\alpha)$ for the polytope $\mathcal{Q}(G \backslash E_0(\alpha))$, where α' is the vector α without its zero coefficients. Recall that replacing $\mathcal{Q}(G)$ by $\mathcal{Q}(G \backslash \{e_0\})$ is the same as enforcing $x_{e_0} = 1$, *i.e.*, enforcing e_0 to be in all covers.

Proposition 4. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$. The tilting vectors satisfying $\beta_{e_0} = 0$ for each $e_0 \in E_0(\alpha)$ are exactly the tilting vectors of the inequality $\alpha' x \geqslant \gamma(\alpha)$ for the polytope $\mathcal{Q}(G \backslash E_0(\alpha))$ with additional zero coefficients for $E_0(\alpha)$ where α' is the vector α without its zero coefficients.*

Proof. The tilting vectors for $\alpha x \geqslant \gamma(\alpha)$ are the vectors that satisfy $\beta x^c = 0$ for each cover c of $C^=(\alpha)$. When additionally, the condition $\beta_{e_0} = 0$ for each $e_0 \in E_0(\alpha)$ is imposed, this is the same as having $\beta x^{c \subset E_0(\alpha)} = 0$ for each cover c of $C^=(\alpha)$. The sets $(c \subset E_0(\alpha))_{c \in C^=(\alpha)}$ are exactly the covers in $C^=(\alpha')$ for the polytope $\mathcal{Q}(G \backslash E_0(\alpha))$. Thus, the two sets of tilting vector mentioned in the theorem are the solution set of the same system of equations. They are thus equal (up to the additional null coefficients). □

By combining the above two results, we can show that the study of set covering facets can be reduced from arbitrary inequalities to inequalities without null coefficients. This result is crystallized in Proposition 5 which contains a necessary and sufficient condition for arbitrary inequalities to be facets. Although presented in a different manner, this result can be found in (Laurent, 1989; Cornuéjols and Sassano, 1989) for the special case of rank inequalities. It is also present in all facet characterisations for special cases we are aware of, such as those for facets with coefficients $0, 1$, and 2 by Balas and Ng (1989) and $0, 1, 2$, and 3 by Saxena (2004) and Sánchez-García et al. (1998).

Proposition 5. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$. It is a facet if and only if :*

(i) *$\alpha' x \geqslant \gamma(\alpha')$ is a facet of the polytope $\mathcal{Q}(G \backslash E_0(\alpha))$, where α' is the vector α without its zero coefficients;*

(ii) *for all $e \in E_0(\alpha)$, there is a cover of $C^{=}(\alpha)$ not containing e.*

With the above proposition, one can see that when searching for characterizations of facets, the study can be restricted to inequalities without null coefficients. Indeed, the conditions required in the general case will be the ones from the restricted case together with condition ((ii)) of Proposition 5.

5.2 Faster computation of tilting vectors for sparse inequalities

Most of the algorithms presented in Section 4.4 rely on calls to a set covering oracle for $\mathcal{Q}(G)$ to compute the tilted inequality (the arguments of this section also work for the algorithm of the last paragraph of Section 4.4 relying on Branch & Bound). When the oracle is given an objective function with many null coefficients, one can expect the oracle to run much faster in practice. Indeed, setting the variables corresponding to these null coefficients to one does not increase the cost of the solution but covers several of the subsets to be covered. As a consequence, a smaller set covering problem can be solved without compromising the optimality of the solution. In particular, oracle calls for sparse objective functions may be much faster than finding a minimum cardinality cover. In the case of tilting, the objective function given to the oracle is always $\alpha + \epsilon\beta$ for $\epsilon \in [0, 1]$. Thus, if the tilting vector β has its support (set of non-zero coefficients) included in the one of α, the objective vector is as sparse as α. Note that in this case the tilted inequality, if non-trivial, will have a support included in the support of α, allowing one to re-apply tilting to the new equality with the same sparsity properties.

One can accelerate the computation of the tilted inequality when the tilting vector has its support included in the support of α. However, some tilting vectors have a larger support than α. This is the case of the tilting vector $\alpha - \gamma(\alpha)e_e$ appearing in Theorem 4 when e is outside of the support of α. However, one can greatly increase the dimension of original inequality by using only tilting vectors as sparse as α. In fact, once all sparse tilting vectors have been exploited, there only remains the aforementioned tilting vectors of Theorem 4. This fact is apparent in the following corollary of Theorem 6.

Corollary 1. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$. Let E_0^* be the set of elements of $E_0(\alpha) = \{e \in \mathcal{E} | \alpha_e = 0\}$ contained in all the covers of $C^{=}(\alpha)$. Assume that the only tilting vector satisfying $\beta_{e_0} = 0$ for all $e_0 \in E_0(\alpha)$ is $\beta = 0$. Then, a basis of the vector space of tilting vectors is $(\alpha - \gamma(\alpha)e_{e_0})_{e_0 \in E_0^*}$.*

Thus, given a sparse initial inequality, one can tilt it into a stronger inequality by using only tilting vectors whose support is included in the initial inequality. This will require only calls to a set covering oracle with sparse objective functions which can be much quicker. Afterward, Corollary 1 shows that it only remains to check that the conditions in Theorem 4 are verified for the elements of $E_0(\alpha)$ which can be done with $|E_0(\alpha)|$ calls to the set covering oracle. Unfortunately, at each of these steps, the support of the inequality may increase its size which may slow down the last oracle calls.

6 Extending necessary and sufficient conditions for facets

In this section, we revisit several conditions from the literature that characterize facets of the set covering polytope. In particular, we extend the conditions of Cornuéjols and Sassano (1989) and Sassano (1989) from rank inequalities to arbitrary inequalities. We also give an alternative proof for the characterization of Balas and Ng (1989) on facets having coefficients and right hand side in $\{0, 1, 2\}$. We finally complement their result by characterizing the tilting vectors for these inequalities.

6.1 Necessary conditions of Cornuéjols and Sassano (1989)

Let us first restate a theorem from Cornuéjols and Sassano (1989) (Proposition 1) on rank inequalities (*i.e.*, inequalities having only binary coefficients), giving a condition under which an inequality is not a facet. We derive a new proof showing that under this condition one can pinpoint a tilting vector. The proof is not specific to rank inequalities which enables us to generalize their result to arbitrary inequalities. In this section, we consider inequalities without null coefficients but null coefficients can be taken into account by using Proposition 5.

To state the theorem, its proof, and the ensuing discussion, let us introduce a few notations and concepts. For any vector α and set of elements E, we will denote α_E the vector whose coefficient is α_e if e belongs to E and zero otherwise. Let us also introduce the concept of cutset as it is done in Cornuéjols and Sassano (1989). Let E be a set of elements of $\mathcal{E}$ and let $\bar{E} = \mathcal{E}\backslash E$. For each $E \subset \mathcal{E}$, the cutset S_E is the set of nodes adjacent to at least one node in E and one node in $\bar{E}$, *i.e.* $S_E = N(E) \cap N(\bar{E})$. Let us also recall that $\gamma(\alpha, S)$ denotes the minimum value of αx over the binary vectors representing a cover of the subsets in S.

Proposition 6 (Cornuéjols and Sassano (1989)). *If there is a non-critical cutset, i.e., for some* $E \subset \mathcal{E}$, $\gamma(\mathbb{1}) = \gamma(\mathbb{1}, \mathcal{S}\backslash S_E)$ *then the inequality* $\mathbb{1}x \geqslant \gamma(\mathbb{1})$ *is not a facet.*

Proof. We will prove that if a cutset is not critical, one can derive a tilting vector.

Let us assume that for some $E \subset \mathcal{E}$, the cutset S_E is not critical, *i.e.*, $\gamma(\mathbb{1}) = \gamma(\mathbb{1}, \mathcal{S}\backslash S_E)$. Note that the size of a cover c is also $\mathbb{1}x_c$. We show that in this case, $c \mapsto \mathbb{1}_E x_c$ is constant over all covers c of size $\gamma(\mathbb{1})$. To that end, let us consider two covers c and c' of $\mathcal{S}$ of size $\gamma(\mathbb{1})$. Note that both $c \cap E$ and $c' \cap E$ cover $N(E)\backslash N(\bar{E})$ and also that both $c' \cap \bar{E}$ and $c \cap \bar{E}$ cover $N(\bar{E})\backslash N(E)$. Thus, both $(c \cap E) \cup (c' \cap \bar{E})$ and $(c' \cap E) \cup (c \cap \bar{E})$ cover of $\mathcal{S}\backslash S_E$. If one of these two covers of $\mathcal{S}\backslash S_E$ had a size strictly higher than $\gamma(\mathbb{1})$ then the other one would have a size strictly lower than $\gamma(\mathbb{1})$ which would contradict the criticality of S_E. Thus, they both have size $\gamma(\mathbb{1})$. Thus, we have:

$$\begin{aligned}\mathbb{1}_E x_{c'} + \mathbb{1}_{\bar{E}} x_c =& \gamma(\mathbb{1}) \\ =& \mathbb{1}x_c \\ =& \mathbb{1}_E x_c + \mathbb{1}_{\bar{E}} x_c,\end{aligned}$$

which implies $\mathbb{1}_E x_{c'} = \mathbb{1}_E x_c$. Thus, $c \mapsto \mathbb{1}_E x_c$ is constant over all covers c of $\mathcal{S}$ of size $\gamma(\mathbb{1})$ which means that $c \mapsto \mathbb{1}_{\bar{E}} x_c = \gamma(\mathbb{1}) - \mathbb{1}_E x_c$ is also constant. Since the covers of size $\gamma(\mathbb{1})$ are exactly the covers in $C^{=}(\mathbb{1})$, one can see that the vector $(\mathbb{1}_{\bar{E}} x_c)\mathbb{1}_E - (\mathbb{1}_E x_c)\mathbb{1}_{\bar{E}}$ is a tilting vector for $\mathbb{1}x \geqslant \gamma(\mathbb{1})$. □

In the proof above, we have shown that when a cutset S_E is not critical, then all minimal covers have a constant size intersection with E and this enables us to highlight a tilting vector. Moreover, one can see that the proof remains valid if the vector $\mathbb{1}$ is replaced by α (except for the size of a cover c that goes from $\mathbb{1}x_c$ to αx_c). This observation leads to the following generalization of Proposition 6.

Proposition 7. *Let* $\alpha x \geqslant \gamma(\alpha)$ *be a valid inequality for* $\mathcal{Q}(G)$ *with* $\alpha > 0$. *If there is a non-critical cutset (i.e., for some* $E \subset \mathcal{E}$, $\gamma(\alpha) = \gamma(\alpha, \mathcal{S}\backslash S_E)$*), then the inequality* $\alpha x \geqslant \gamma(\alpha)$ *is not a facet.*

At this point, it is important to note that there are set covering instances where every minimal cover has a constant size intersection with a set E but where for every $E' \subset \mathcal{E}$ the cutset $S_{E'}$ is critical.

Example *Let us consider a set covering instance with* $\mathcal{E} = \{1, 2, 3, 4, 5, 6\}$ *and a set* $E = \{1, 2, 3\}$. *The family of subsets* $\mathcal{S}$ *of the set covering instance is* $\mathcal{S} = \{s \subset \mathcal{E} | |s| = 3 \text{ and } |s \cap E| \neq 1\}$. *There are two types of inclusion-wise minimal covers: the subset of* $\mathcal{E}$ *of size 3 whose intersection with* E *has size 1; and the three covers formed of* $\{1, 2, 3, e\}$ *for* $e \in \{4, 5, 6\}$. *Among these inclusion-wise minimal covers, only the first type also has minimum size. Thus, all the minimal sized covers have an intersection with* E *of size 1. For this set covering instance, every cutset is critical. For instance,* $\mathcal{S}\backslash S_E = \{\{1, 2, 3\}, \{4, 5, 6\}\}$, *which can be covered with* $\{1, 4\}$.

Thus, Proposition 7 can be extended as follows.

Proposition 8. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$ with $\alpha > 0$. If there exists a set $E \subset \mathcal{E}$ such that $\alpha_E x_c = \alpha_E x_{c'}$ for all $c, c' \in C^=(\alpha)$ then the inequality $\alpha x \geqslant \gamma(\alpha)$ is not a facet.*

6.2 Sufficient conditions of Sassano (1989)

We now consider a lemma from Sassano (1989) (Lemma 3.1) which gives a sufficient condition for the rank inequality $\mathbb{1}x \geqslant \gamma(\mathbb{1})$ to be a facet. We will discuss underlying ideas of this lemma and generalize it to arbitrary inequalities. To that end, let us denote $N(e, e') = N(\{e\}) \cap N(\{e'\})$ the common neighbors to elements e and e' and let us introduce a critical graph G^*. This critical graph has $\mathcal{E}$ as its set of nodes and contains an edge (e, e') if and only if the set of common neighbors to e and e' is critical: $\gamma(\mathbb{1}) > \gamma(\mathbb{1}, \mathcal{S}\backslash N(e, e'))$. The lemma is then the following:

Lemma 1 (Sassano (1989))**.** *If the critical graph G^* is connected, then the inequality $\mathbb{1}x \geqslant \gamma(\mathbb{1})$ is a facet of $\mathcal{Q}(G)$.*

Proof. Let us consider an edge (e, e') of the critical graph and let us consider a minimal cover c of $\mathcal{S}\backslash N(e, e')$ which has a size of $\gamma(\mathbb{1}, \mathcal{S}\backslash N(e, e'))$. Then both $c \cup \{e\}$ and $c \cup \{e'\}$ cover $\mathcal{S}$ and have a size equal to $\gamma(\mathbb{1}, \mathcal{S}\backslash N(e, e')) + 1$. Thus, we have that $\gamma(\mathbb{1}) \leqslant \gamma(\mathbb{1}, \mathcal{S}\backslash N(e, e')) + 1$. Since, by definition of the critical graph, we know that $\gamma(\mathbb{1}) > \gamma(\mathbb{1}, \mathcal{S}\backslash N(e, e'))$, then we have $\gamma(\mathbb{1}) = \gamma(\mathbb{1}, \mathcal{S}\backslash N(e, e')) + 1$. This also means that $c \cup \{e\}$ and $c \cup \{e'\}$ belong to $C^=(\mathbb{1})$ which means that for any tilting vector β, we have:

$$\begin{aligned} \beta_e - \beta_{e'} &= \beta x_{c \cup \{e\}} - \beta x_{c \cup \{e'\}} \\ &= 0 - 0 \\ &= 0. \end{aligned}$$

Hence, the presence of an edge in the critical graph implies that the two corresponding coefficients of any tilting vector are equal. From this argument, we can deduce that, if the critical graph is connected, then all the coefficients in tilting vectors are equal. This also implies that all the coefficients are zero since they must sum to zero for the covers in $C^=(\mathbb{1})$. In this case, we know by Proposition 2 that the inequality $\mathbb{1}x \geqslant \gamma(\mathbb{1})$ is a facet. □

In the proof of Theorem 1, we showed that the presence of an edge in the critical graph induces the existence of two covers $c \cup \{e\}$ and $c \cup \{e'\}$ in $C^=(\mathbb{1})$ differing only by one element. The main argument of the proof is that the existence of these two covers implies that the coefficients β_e and $\beta_{e'}$ must be equal in any tilting vector β. Unlike the previous lemma, this argument is not restricted to the rank constraint $\mathbb{1}x \geqslant \gamma(\mathbb{1})$ and is the basis of the generalization to arbitrary inequalities. Let us introduce an adequate notion of criticality for arbitrary inequalities and show that it is equivalent to the presence of two covers in $C^=(\alpha)$ differing by one element.

For an arbitrary inequality $\alpha x \geqslant \gamma(\alpha)$, the existence of two covers $c \cup \{e\}$ and $c \cup \{e'\}$ in $C^=(\alpha)$ implies that

$$\begin{aligned} \alpha_e - \alpha_{e'} &= \alpha x_{c \cup \{e\}} - \alpha x_{c \cup \{e'\}} \\ &= \gamma(\alpha) - \gamma(\alpha) \\ &= 0. \end{aligned}$$

Thus, the notion of criticality is defined only for pairs of elements that share the same coefficients in the inequality.

Definition 4. Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality and e, e' be such that $\alpha_e = \alpha_{e'}$. The common neighbors to e and e' are defined as critical when $\gamma(\alpha) = \gamma(\alpha, \mathcal{S}\backslash N(e, e')) + \alpha_e$.

In other words, the common neighbors to e and e' are critical when their removal induces a decrease of $\gamma(\alpha)$ equal to α_e. This reduction could not be greater. Indeed, one can create a cover of $\mathcal{S}$ by adding e or e' to a cover of $\mathcal{S}\backslash N(e,e')$ and thus, we always have $\gamma(\alpha) \geqslant \gamma(\alpha, \mathcal{S}\backslash N(e,e')) \geqslant \gamma(\alpha) - \min(\alpha_e, \alpha_{e'})$. As before, there is an edge (e,e') in the critical graph G^* when the common neighbors to e and e' are critical. Let us now show the equivalence between an edge in G^* and the presence of two covers in $C^=(\alpha)$ differing by one element.

Proposition 9. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality and let two elements e, e' be such that $\alpha_e = \alpha_{e'}$. The following conditions are equivalent :*

(i) *there exists two covers $c \cup \{e\}$ and $c \cup \{e'\}$ in $C^=(\alpha)$;*

(ii) *the critical graph G^* contains an edge between e and e', i.e. $\gamma(\alpha) = \gamma(\alpha, \mathcal{S}\backslash N(e,e')) + \alpha_e$.*

Proof. (i)⇒(ii) Let $c \cup \{e\}$ and $c \cup \{e'\}$ be two covers of $C^=(\alpha)$. We know that c covers $\mathcal{S}\backslash N(\{e\})$ since $(c) \cup \{e\}$ is a cover. Similarly, it covers $\mathcal{S}\backslash N(\{e'\})$. Thus, c covers $\mathcal{S}\backslash N(e,e')$ and its incidence vector satisfies $\alpha x_c = \alpha x_{c\cup\{e\}} - \alpha_e = \gamma(\alpha) - \alpha_e$. Thus, we must have $\gamma(\alpha, \mathcal{S}\backslash N(e,e')) \leqslant \gamma(\alpha) - \alpha_e$. Since we also know that we always have $\gamma(\alpha, \mathcal{S}\backslash N(e,e')) \geqslant \gamma(\alpha) - \alpha_e$, we have proved that condition (ii) is verified.

(ii)⇒(i) Let us call c a cover of $\mathcal{S}\backslash N(e,e')$ that achieves the value $\gamma(\alpha, \mathcal{S}\backslash N(e,e'))$. Then, $c \cup \{e\}$ and $c \cup \{e'\}$ both cover $\mathcal{S}$ and satisfy $\alpha x_{c\cup\{e\}} = \alpha x_{c\cup\{e'\}} = \gamma(\alpha)$, which shows that condition (i) is satisfied. □

Now, the arguments of the discussion above imply that, for an edge (e,e') in the critical graph G^*, there exists two covers $c \cup \{e\}$ and $c \cup \{e'\}$ in $C^=(\alpha)$. This, in turn, means that the coefficients corresponding to e and e' in any tilting vector β are equal since $\beta_e - \beta_{e'} = \beta x_{c\cup\{e\}} - \beta x_{c\cup\{e'\}} = 0 - 0 = 0$. Thus, the nodes in a connected component of the critical graph share the same coefficients in the tilting vector. If a connected component contains a cover c then the coefficients of the elements in c sum to zero which means all the coefficients of component are null. Thus, if each connected component contains a cover then $\alpha x \geqslant \gamma(\alpha)$ induces a facet of $\mathcal{Q}(G)$. Thus, we can generalize the lemma from Sassano (1989) as follows.

Lemma 2. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$ with $\alpha > 0$. If each connected component of its critical graph contains a cover of $C^=(\alpha)$ then it is a facet.*

Let us recall that edge (e,e') in the critical graph can exist only when $\alpha_e = \alpha_{e'}$. Therefore, the connected components of the critical graph subdivide the sets of elements with equal coefficients. In some cases, one may not be able to prove that an inequality is a facet because not all connected components of the critical graph may contain a cover of $C^=(\alpha)$. However, even if only one component does, we can show that the coefficients of its elements in all tilting vector are null. As shown in the following theorem, this has consequences on the facets obtainable from the original inequality $\alpha x \geqslant \gamma(\alpha)$. More precisely, the coefficients of the facet will be equal to the ones of α for these elements.

Theorem 7. *Let $\alpha x \geqslant \gamma(\alpha)$ be a valid inequality for $\mathcal{Q}(G)$ with $\alpha > 0$. If, for all tilting vectors β we have $\beta_e = 0$ for some element $e \in \mathcal{E}$, then all the non-trivial facets $\alpha' x \geqslant \gamma(\alpha)$ containing the face associated with $\alpha x \geqslant \gamma(\alpha)$ satisfy $\alpha'_e = \alpha_e$.*

Proof. Since we are considering non-trivial facets of the set covering polytope, the facet can be written $\alpha' x \geqslant \gamma(\alpha)$ through the right scaling. The vector $\alpha' - \alpha$ is a tilting vector since for each cover c in $C^=(\alpha)$ (which is included in $C^=(\alpha')$ by hypothesis) we have $(\alpha' - \alpha)x^c = \gamma(\alpha) - \gamma(\alpha) = 0$. Thus, we have $\alpha'_e - \alpha_e = (\alpha' - \alpha)_e = 0$ which means $\alpha'_e = \alpha_e$. □

6.3 Inequalities with coefficients and right hand side in $\{0,1,2\}$

In their seminal paper, Balas and Ng (1989) studied all the inequalities for the set covering polytope with coefficients and right hand side in $\{0, 1, 2\}$. In particular, they characterized the ones that induce facets of the set covering polytope. To understand their characterization, recall that $E_i(\alpha) = \{e \in \mathcal{E} | \alpha_e = i\}$ and let us introduce the 2-cover graph corresponding to an inequality $\alpha x \geqslant 2$. Its node set is $E_1(\alpha)$ and it contains an edge between e and e' when $E_0(\alpha) \cup \{e, e'\}$ is a cover (and thus belongs to $C^{=}(\alpha)$). Theorem 8 below states the characterization of Balas and Ng (1989) in a slightly differently manner in order to fit the notations of this article. The three conditions are indexed with 0, 1 and 2 because they can be associated to the sets of all the coefficients 0, 1 and 2, respectively.

Theorem 8 (Balas and Ng (1989)). *Let $\alpha x \geqslant 2$ be a valid inequality for $\mathcal{Q}(G)$ with $\alpha \in \{0, 1, 2\}^{\{|\mathcal{E}|\}}$. It is a facet if and only if the following three conditions hold:*

0. *for each $e \in E_0(\alpha)$, there is a cover of $C^{=}(\alpha)$ not containing e;*
1. *each connected component of the 2-cover graph contains an odd cycle (i.e. is not bipartite);*
2. *for each $e \in E_2(\alpha)$, $E_0(\alpha) \cup \{e\}$ is a cover.*

The above theorem characterizes facets corresponding to $\alpha x \geqslant 2$ with three conditions. We complement this result by characterizing the tilting vectors of these inequalities. In particular, each of the three conditions corresponds to a family of tilting vector. In order to understand the tilting vector characterization, note that a classical result from graph theory is that a graph is bipartite if and only if it does not contain an odd cycle (see Theorem 1.2 in Bondy and Murty (1976)). Moreover, if a bipartite graph is connected, then it has a unique bipartition.

Theorem 9. *Let $\alpha x \geqslant 2$ be a valid inequality for $\mathcal{Q}(G)$ with $\alpha \in \{0, 1, 2\}^{\{|\mathcal{E}|\}}$. A basis of the space of its tilting vectors is given by the juxtaposition of the following three families:*

0. *$\alpha - 2e_e$ for each $e \in E_0(\alpha)$ such that all covers of $C^{=}(\alpha)$ contain e;*
1. *$\mathbb{1}_U - \mathbb{1}_V$ for each bipartite connected component of the 2-cover graph where U, V is the bipartition of the connected component;*
2. *e_e for each $e \in E_2(\alpha)$ such that $E_0(\alpha) \cup \{e\}$ is a not cover.*

We will now prove the characterization of tilting vectors, i.e., Theorem 9 which implies the theorem of Balas and Ng (1989) through Proposition 2. For this proof, we will use a lemma from graph theory. In this lemma, we will not always differentiate an edge (u, v) from its incidence vector. This incidence vector has a null coefficient for each node of the graph except for u and v for which it has coefficient one.

Lemma 3. *Let $G = (U, E)$ be a graph. A basis of the vector space induced by the edges of G is given by the edges of the following subgraph. This subgraph contains a spanning tree of each connected component of the main graph and for each non-bipartite connected component, it also contains one additional edge so that the tree and this edge forms a non-bipartite graph (equivalently this edge adds an odd cycle to the tree).*

Note that the additional edge always exists as otherwise the connected component would be bipartite. We call the subgraph of Lemma 3 a basis subgraph.

Proof of Lemma 3. We will separate the proof into two parts. First, we will show linear independence of the edges of the basis subgraph. Second, we will show that every edge of the main graph can be written as a linear combination of the edges of the basis subgraph.

First, we have linear independence of the edges of the different connected components as the edges in these groups have disjoint supports. We will thus analyze them separately. Let us now consider a bipartite connected component of the main graph. The corresponding component in the subgraph is a spanning tree. Let us denote $e_1, \ldots, e_n$ the edges of this tree, and $x_1, \ldots, x_n$ the corresponding incidence vectors. Showing linear independence of the edges/vectors is equivalent to showing that the only assignment of coefficients $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ for which $\sum_i \lambda_i x_i = 0$ is $\lambda_1 = \ldots = \lambda_n = 0$. An equivalent formulation of having a linear combination that sums to the null vector is to assign coefficients $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ to the edges of the tree and asking, for each node, for the sum of the coefficients of the outgoing edges to be null. The edges are linearly independent if and only if the only such affectation is $\lambda_1 = \ldots = \lambda_n = 0$. Clearly, the edges connected to the leaves of the tree (nodes connected to only one edge) must have coefficient zero. Thus, these edges are linearly independent from the other ones and they can be removed from the tree (or equivalently their incidence vector can be removed from the linear combination). This removal creates new leaves in the tree which enables the process to be repeated. Since the component is initially a tree, this process can be repeated until all edges have been removed thus showing that the edges of the connected component were all independent from each other. In a similar fashion, for non-bipartite connected components, the edges of the corresponding component in the basis subgraph can be removed until the component is reduced to an odd cycle. The edges of an odd cycle are linearly independent.

Second, let us show now that every edge of the 2-cover graph can be written as a linear combination of the edges of the basis subgraph. First, the edges of the bipartite components are either part of the corresponding tree in the subgraph or they create a cycle when added to the tree (since the tree is a spanning tree). The created cycle must have an even number of edges since the considered component of the main graph is bipartite. Thus, by assigning alternating $+1$ and -1 coefficients along the edges of the cycle, one can show that the edge is linearly dependent of the edges of the tree. Second, for the non-bipartite connected components, the corresponding component in the basis subgraph has the same number of nodes and edges. Since the edges connecting the only nodes of this connected component have non zero coefficients only for these node, the vector space of these edges must have its dimension equal to the number of nodes in the connected component. Thus the corresponding edges of the basis subgraph must be a basis of this vector as we showed above they are linearly independent and they number is the number of node. Thus, they must generate all the edges connecting the only nodes of this connected component. □

We are now in a position to prove Theorem 9.

Proof of Theorem 9. Let us denote N_β the number of tilting vectors mentioned in the theorem. We will prove that 1) these vectors are indeed tilting vector of the inequality, 2) they are linearly independent and 3) they can be used to generate all tilting vectors.

1) If, for some $e \in E_2(\alpha)$, $E_0(\alpha) \cup \{e\}$ is a not cover, then, no cover of $C^=(\alpha)$ contains e. Thus, according to Theorem 5, e_e is a tilting vector. Second, the vectors $\alpha - 2e_e$ are given by the Theorem 4. Third, note that the covers in $C^=(\alpha)$ take only two forms, $E_0 \cup \{e_2\}$ and $E_0 \cup \{e_1, e_1'\}$, where E_0 is a subset of $E_0(\alpha)$, e_2 belongs to $E_2(\alpha)$, and e_1, e_1' belong to $E_1(\alpha)$. We want to show that for each of these covers c we have $(\mathbb{1}_U - \mathbb{1}_V)x_c = 0$. It is clear for the first type of covers since the support of $\mathbb{1}_U - \mathbb{1}_V$ and $x_{E_0 \cup \{e_2\}}$ are disjoint. For the second type, by definition of the 2-cover graph, (e_1, e_1') is an edge of the 2-cover graph. If this edge is not contained in the connected component associated with $\mathbb{1}_U - \mathbb{1}_V$ then again the support of $\mathbb{1}_U - \mathbb{1}_V$ and $x_{E_0 \cup \{e_1, e_1'\}}$ are disjoint. Finally, if the edge (e_1, e_1') is contained in the connected component associated with $\mathbb{1}_U - \mathbb{1}_V$ —since this connected component is bipartite— we either have $(e_1, e_1') \in U \times V$ or $(e_1, e_1') \in V \times U$. In both cases, we have $(\mathbb{1}_U - \mathbb{1}_V)x_c = -1 + 1 = 0$. Thus, $\mathbb{1}_U - \mathbb{1}_V$ is orthogonal to all covers of $C^=(\alpha)$ which makes it a tilting vector.

2) The tilting vectors in the second and third families are linearly independent as all their supports are disjoint. As for the first family of tilting vectors, they are also linearly independent of the others as each of them correspond to an element of $E_0(\alpha)$ for which they are the only one having a non-zero coefficient.

3) Let us now derive $|\mathcal{E}| - N_\beta$ linearly independent covers of $C^=(\alpha)$. Thanks to Theorem 2, this implies that the number of independent tilting vector of $\alpha x \geqslant 2$ is less than N_β. This, in turns, means that the N_β tilting vectors of the theorem can be used to generate all the tilting vectors since they are independent. The covers will be separated into three families:

0. $c \cup (E_0(\alpha) \backslash e_0)$ for each $e \in E_0(\alpha)$ for which a cover c not containing e exists;
1. $E_0(\alpha) \cup \{e_1, e_1'\}$ for each edge (e_1, e_1') of the basis subgraph of the 2-cover graph;
2. $E_0(\alpha) \cup \{e_2\}$ for each $e_2 \in E_2(\alpha)$ for which $E_0(\alpha) \cup \{e_2\}$ is a cover.

Note that a basis subgraph has a number of edges equal to its number of nodes minus its number of bipartite connected components. Thus, one can see by pairing the above families of covers and the families of tilting vectors from the Theorem, that we have characterized $|E_0(\alpha)|+|E_1(\alpha)|+|E_2(\alpha)| = |\mathcal{E}|$ objects, covers and tilting vectors, in this proof. Hence, there is indeed a total of $|\mathcal{E}| - N_\beta$ covers of $C^=(\alpha)$ in the three previous families.

Let us now discuss the linear independence of the covers in the above three families. To that end, we will consider their incidence vectors as the columns of a matrix and show that this matrix is full column rank. If the first columns are the vectors of the family 0, then those of family 1 and then family 2, the matrix can be written as follows:

$$\begin{pmatrix} 1 - J_0 & 1 & 1 \\ Y_1 & X & 0 \\ Y_2 & 0 & J_2 \end{pmatrix}$$

In the previous matrix, J_0 is composed of an identity matrix of size η_0 — the size of the family 0 — on top of $|E_0(\alpha)| - \eta_0$ rows of zeros. J_2 is composed of an identity matrix of size η_2 — the size of the family 2 — on top of $|E_2(\alpha)| - \eta_2$ rows of zeros. Finally, X is the edge-node incidence matrix of the basis subgraph of the 2-cover graph. For each cover c associated to $e_0 \in E_0(\alpha)$ in family 0, the cover $c \cup \{e_0\}$ is either a cover of family 2 or it is equal to $E_0(\alpha) \cup \{e_1, e_1'\}$ for some edge (e_1, e_1') of the 2-cover graph. In this second case, Lemma 3 tells us that the edge (e_1, e_1') can be written as a linear combination of the edges in the basis subgraph. Note that the sum of the coefficients in the combination is 1. Thus, in both cases, the cover $c \cup \{e_0\}$ can be written as an affine combination of the covers in families 2 and 1. Therefore, by column manipulations, we can replace the previous matrix by the following matrix that has the same rank:

$$\begin{pmatrix} -J_0 & 1 & 1 \\ 0 & X & 0 \\ 0 & 0 & J_2 \end{pmatrix}$$

Since the matrices J_0, J_2 and X are full column rank, the complete matrix must also be. □

7 Concluding remarks

In this work, we introduced a new mathematical object, the tilting vectors, which are derived from a variation of the tilting concepts introduced by Chvátal et al. (2013). These vectors can be used to tilt set covering inequalities and provide tools to derive properties and proofs for facets of the set covering polytope. In particular, thanks to the tilting vectors, we were able to generalize some facet characterizations from rank inequalities to arbitrary inequalities. We also showed that the null coefficients in a set covering inequality can be treated separately. Indeed, one can study/tilt a set covering inequality by first ignoring its null coefficients. Special properties or tilting procedures can then be used to take them into account.

Although the study of the structure of the set covering problem has not received much recent attention, we believe that it remains an important topic of research. The set covering problem can be used to model any problem whose set of solutions X is monotonic ($x \in X$ and $x \leqslant y$ implies $y \in X$). This includes a wide variety of problems, such as covering problems (such as vertex covering, or feedback sets, for instance), packing problems (such as set packing, node packing, or independent sets), knapsack problems (single knapsack, multiple knapsack), and others. Thus, advances in understanding the structure of the set covering problem can be directly applied to multiple other optimization problems.

References

Egon Balas and Shu Ming Ng. On the set covering polytope: I. all the facets with coefficients in {0, 1, 2}. Mathematical Programming, 43(1):57–69, 1989.

John Adrian Bondy and Uppaluri Siva Ramachandra Murty. Graph theory with applications. Macmillan, London, 1976.

Ralf Borndörfer. Aspects of set packing, partitioning, and covering. PhD thesis, 1998.

Valentina Cacchiani, Vera C Hemmelmayr, and Fabien Tricoire. A set-covering based heuristic algorithm for the periodic vehicle routing problem. Discrete Applied Mathematics, 163:53–64, 2014.

Alberto Caprara, Matteo Fischetti, and Paolo Toth. A heuristic method for the set covering problem. Operations research, 47(5):730–743, 1999.

Vašek Chvátal, William Cook, and Daniel Espinoza. Local cuts for mixed-integer programming. Mathematical Programming Computation, 5(2):171–200, 2013.

Gérard Cornuéjols and Antonio Sassano. On the 0, 1 facets of the set covering polytope. Mathematical Programming, 43(1):45–55, 1989.

Richard H Day. On optimal extracting from a multiple file data storage system: An application of integer programming. Operations Research, 13(3):482–494, 1965.

Daniel Espinoza, Ricardo Fukasawa, and Marcos Goycoolea. Lifting, tilting and fractional programming revisited. Operations research letters, 38(6):559–563, 2010.

Monique Laurent. A generalization of antiwebs to independence systems and their canonical facets. Mathematical Programming, 45(1-3):97–108, 1989.

Jurij Mihelic and Borut Robic. Facility location and covering problems. In Proc. of the 7th International Multiconference Information Society, volume 500, 2004.

Paolo Nobili and Antonio Sassano. Facets and lifting procedures for the set covering polytope. Mathematical Programming, 45:111–137, 1989.

Manfred W Padberg. On the facial structure of set packing polyhedra. Mathematical programming, 5(1): 199–215, 1973.

James A Reggia, Dana S Nau, and Pearl Y Wang. Diagnostic expert systems based on a set covering model. International journal of man-machine studies, 19(5):437–460, 1983.

Miguel Sánchez-García, María Inés Sobrón, and Begoña Vitoriano. On the set covering polytope: Facets with coefficients in {0, 1, 2, 3}. Annals of Operations Research, 81:343–356, 1998.

Antonio Sassano. On the facial structure of the set covering polytope. Mathematical Programming, 44(1): 181–202, 1989.

Anureet Saxena. On the set-covering polytope: I. all facets with coefficients in {0, 1, 2, 3}. Technical report, GSIA Working Paper# 2004-E29, 2004.

Larry E Stanfel. A successive approximations method for a cellular manufacturing problem. Annals of Operations Research, 17:13–30, 1989.

Laurence A Wolsey. Facets and strong valid inequalities for integer programs. Operations research, 24(2): 367–372, 1976.