On a Cone Covering Problem

Khaled Elbassioni

Max Planck Institut für Informatik
Saarbrücken, D-66123 Germany

Hans Raj Tiwary[1]

Technische Universität
MA 6-2, Institut für Mathematik
10623 Berlin, Germany

Abstract

Given a set of polyhedral cones $\mathcal{C}_1, \cdots, \mathcal{C}_k \subset \mathbb{R}^d$, and a convex set D, does the union of these cones cover the set D? In this paper we consider the computational complexity of this problem for various cases such as whether the cones are defined by extreme rays or facets, and whether D is the entire $\mathbb{R}^d$ or a given linear subspace $\mathbb{R}^t$. As a consequence, we show that it is coNP-complete to decide if the union of a given set of convex polytopes is convex, thus answering a question of Bemporad, Fukuda and Torrisi.

1. Introduction

Let $S \subseteq \mathbb{R}^d$ be a finite set of points in $\mathbb{R}^d$. The *conic hull* of S, denoted by cone(S), is the set of all non-negative linear combinations of points in S, i.e., $\text{cone}(S) = \{\sum_{p \in S} \mu_p p : \ \mu_p \geq 0 \text{ for all } p \in S\}$. It is well-known that any polyhedral cone cone(S) can be written equivalently as the intersection of finitely many half-spaces, i.e., $\text{cone}(S) = \{x \in \mathbb{R}^d : \ Ax \leq \mathbf{0}\}$, where $A \in \mathbb{R}^{m \times d}$. The two representations are called the $\mathcal{V}$- and the $\mathcal{H}$-representations, respectively.

In this note we are interested in the complexity of covering problems of the following form:

ConeCover($\mathcal{C}, D$): Given a collection of cones $\mathcal{C} = \{C_1, \ldots, C_k\}$, and a convex set D, does $\bigcup_{i=1}^k C_i \not\supseteq D$?

Email addresses: elbassio@mpi-sb.mpg.de (Khaled Elbassioni), tiwary@math.tu-berlin.de (Hans Raj Tiwary)

[1] Corresponding author

A polytope P is the convex hull of a finite set S of points in $\mathbb{R}^d$, and it can also be written in one of two equivalent forms: $P = \mathrm{conv}(S) = \{\sum_{p\in S} \mu_p p : \ \mu_p \geq 0 \text{ for all } p \in S, \ \sum_{p\in S} \mu_p = 1\}$ ($\mathcal{V}$-representation), or $P = \{x \in \mathbb{R}^d | Ax \leq \mathbf{1}\}$, where $\mathbf{1}$ is the vector in which each component is 1 ($\mathcal{H}$-representation)[2]. A polyhedron Q is the Minkowski sum of a polytope P and a cone C: $Q = P+C \stackrel{\text{def}}{=} \{x+y | x \in P, \text{ and } y \in C\}$. Similarly, one can also consider the problem PolytopeCover($\mathcal{P}, D$): Given a collection of polytopes $\mathcal{P} = P_1, \ldots, P_k$, and a convex polytope D, does $\bigcup_{i=1}^{k} P_i \not\supseteq D$?

Our motivation for studying the above covering problems comes from two other related problems on polytopes. The first is the well-known Vertex Enumeration problem of finding the vertices of a polytope given its facet defining inequalities, to be described in more details in the next section. The second problem is to check whether the union of a given set of polytopes is convex. Bemporad, Fukuda and Torrisi [3] gave polynomial-time algorithms for checking if the union of $k = 2$ polyhedra is convex, and if so finding this union, no matter whether they are given in $\mathcal{V}$ or $\mathcal{H}$ representations. They also gave necessary and sufficient conditions for the union of a finite number of convex polytopes in $\mathbb{R}^d$ to be convex, and asked whether these conditions can be used to design a polynomial time algorithm for checking if the union is convex. Bárány and Fukuda gave slightly stronger conditions in [2]. It will follow from our results that, if both d and k are part of the input, then these conditions cannot be checked in polynomial time unless P=NP.

Unless otherwise specified, all the cones considered throughout the paper will be assumed to be pointed, i.e., contain no lines, or equivalently, have a well defined apex, namely the origin. As we shall see, the complexity of the above problem depends on how the cones are represented, and whether they are disjoint or not. We consider 3 different factors, namely:

(f1) whether the cones in $\mathcal{C}$ are given in $\mathcal{V}$- or $\mathcal{H}$-representations, or both representations ($\mathcal{VH}$);

(f2) what the set D is: we consider $D = \mathbb{R}^d$ and $D = \mathbb{R}^t$ for some arbitrary $t \leq d$;

(f3) whether the cones in $\mathcal{C}$ are

- (f3)-(I): pairwise disjoint in the interior and intersect only at faces;
- (f3)-(II): pairwise disjoint in the interior , but can intersect anywhere on the boundaries; and
- (f3)-(III): not necessarily pairwise disjoint.

We denote by ConeCover[$F1, F2, F3$] the different variants of the problem, where $F1 \in \{\mathcal{V}, \mathcal{H}, \mathcal{VH}\}$, $F2 \in \{\mathbb{R}^t, \mathbb{R}^d\}$ and $F3 \in \{I, II, III\}$ describes cases (f3)-(I), (f3)-(II), and (f3)-(III).

[2] possibly after moving first the polytope so that its relative interior contains the origin

	$\mathbb{R}^d$			$\mathbb{R}^t$		
	I	II	III	I	II	III
$\mathcal{V}$	VE-hard (Thm. 1)	VE-hard (Cor. 1)	NPC (Cor. 4)	NPC (Thm. 2)	NPC (Cor. 2)	NPC (Cor. 2)
$\mathcal{H}$	P (Thm. 3)	?	NPC (Cor. 4)	P (Cor. 3)	?	NPC (Cor. 4)
$\mathcal{VH}$	P (Cor. 3)	?	NPC (Thm. 4)	P (Cor. 3)	?	NPC (Cor. 4)

Table 1: Complexity of Cone Covering problem for various input representations.

Following are the main results of this paper:

- We establish the complexity of various variants of the ConeCover problem (summarized in Table 1).
- PolytopeCover is NP-complete for $\mathcal{V}$- or $\mathcal{H}$-polytopes.
- Checking if the union of a set of polytopes is convex or not is NP-complete for $\mathcal{V}$- or $\mathcal{H}$-polytopes.

Some of the results in this paper relate the complexity of some variants of the cone cover problem to another problem, like Vertex Enumeration or Hypergraph Transversal, whose complexity status is itself not very clear. For such a comparison we use a notion similar to the notion of NP-hardness. For a problem Φ, the class Φ-hard denotes all problems Φ' such that there is a polynomial time Turing reduction of Φ to Φ'. Essentially, saying a certain problem is Φ-hard means that if this problem has a polynomial time algorithm then Φ has a polynomial time algorithm as well.

2. Results

Converting the $\mathcal{H}$-representation of a polytope to its $\mathcal{V}$-representation and vice versa, is a well studied problem. Despite years of research, it is neither known if an output-sensitive algorithm exists for this problem, nor is it known to be NP-hard. The following decision version of this problem is known to be equivalent to the enumeration problem [1].

VertEnum(P, V): Given an $\mathcal{H}$-polytope $P \subseteq \mathbb{R}^d$ and a subset of its vertices $V \subseteq \mathcal{V}(P)$, check whether $P = \text{conv}(V)$.

Let P be the polytope defined as $\{x | Ax \leq \mathbf{1}\}$, where $A \in \mathbb{R}^{m \times d}$. Every full-dimensional rational polytope can be brought into this form by moving the origin to its interior and scaling the normals of the facet-defining hyperplanes appropriately. For any vertex v of P, consider the cone of all vectors c such that v is the solution of the following linear program: $\max c^T x$ s.t. $Ax \leq \mathbf{1}$. For every vertex v of P, this cone is uniquely defined. We call this cone the *maximizer cone* of v. Such a maximizer cone can be defined for every proper face of a polytope. The union of all such cones is also known as the *normal fan* of a polytope [8]. It is easy to see that if A' is the maximal subset of rows of A such that $A'v = \mathbf{1}$, then the maximizer cone of v is the conic hull of the rows of A' treated as vectors in $\mathbb{R}^d$.

Theorem 1. *Problem* ConeCover$[\mathcal{V}, \mathbb{R}^d, I]$ *is* VertEnum*-hard.*

Proof. For an $\mathcal{H}$-polytope P and a subset of its vertices V, the $\mathcal{V}$-representation of the maximizer cone for each vertex in V can be computed easily from the facets of P. Clearly, the union of these cones covers $\mathbb{R}^d$ if and only if $P = conv(V)$. To see this, note that if $P \neq conv(V)$ then P has a vertex v not in V and any vector in the interior of the maximizer cone of v does not lie in any of the cones corresponding to the given vertices. □

Corollary 1. *Problem* ConeCover$[\mathcal{V}, \mathbb{R}^d, II]$ *is* VertEnum*-hard.*

Theorem 2. *Problem* ConeCover$[\mathcal{V}, \mathbb{R}^t, I]$ *is NP-complete.*

Proof. ConeCover$[\mathcal{V}, \mathbb{R}^t, \mathrm{I}]$ is clearly in NP. Now, given an $\mathcal{H}$-polytope $P \subset \mathbb{R}^d$, a linear subspace $\mathbb{R}^t$ and a $\mathcal{V}$-polytope $Q \subset \mathbb{R}^t$, it is NP-complete to decide whether Q is the projection of P onto the given subspace [7]. We give a polynomial reduction from this problem to ConeCover$[\mathcal{V}, \mathbb{R}^t, \mathrm{I}]$.

Every vertex v of Q is an image of some (possibly more than one) vertices of P. If this is not the case then Q clearly cannot be the projection of P. Since the vertices of Q are known this condition can be checked in polynomial time. To see why this is true, consider a vertex v of Q and consider any direction α in the affine hull of Q such that $\alpha^T x$ is maximized at v over all points in Q. If we use the same vector α as objective function over the points in P then the maximum is achieved at the face containing all vertices whose image under projection is v. Therefore from now on consider only the case in which the vertices of Q are a subset of the projected vertices of P.

Now, for any vertex v of Q pick any vertex v' of P whose projection is v. We associate the maximizer cone $\mathcal{C}(v')$ of v' with v and refer to it as $\mathcal{C}(v)$. Since $\mathcal{C}(v)$ for every vertex v of Q is actually the maximizer cone of some vertex of P, the $\mathcal{V}$-representation of $\mathcal{C}(v)$ can be easily computed from the matrix A of the normals of facet defining hyperplanes of P.

It is not difficult to see that if Q is not the projection of P onto the given subspace $\mathbb{R}^t$, then one can find a direction c parallel to the given subspace such that a vertex that maximizes $c^T x$ in P is such that its projection is a vertex of the projection of P but not of Q. Hence, the union of cones $\mathcal{C}(v)$ for each vertex v of Q covers $\mathbb{R}^t$ if and only if Q is the projection of P. Also, all these cones intersect each other only at some proper face. □

Corollary 2. *Problems* ConeCover$[\mathcal{V}, \mathbb{R}^t, II]$, *and* ConeCover$[\mathcal{V}, \mathbb{R}^t, III]$ *are NP-complete.*

For a given set of $\mathcal{H}$-cones, if the union does not cover $\mathbb{R}^d$ then there is a facet with normal $a \in \mathbb{R}^d$, of at least one of these cones and a point p in the interior of this facet such that $p + \epsilon a$ lies outside every cone, for some $\epsilon > 0$. Let us call this facet a *witness facet*, and p a *witness point* of the fact that $\mathbb{R}^d$ is not covered.

Theorem 3. $\textsc{ConeCover}[\mathcal{H}, \mathbb{R}^d, I]$ *can be solved in polynomial time.*

Proof. If the cones are allowed to intersect only at common faces, then every point in the interior of a witness facet is a witness point. Thus, one can determine in polynomial time whether the union of the given cones cover $\mathbb{R}^d$ or not as follows.

For every facet f, with normal a, of each cone C pick a point $p \in \text{relint}(f)$. For every other cone $C' \neq C$ in $\mathcal{C}$ compute the smallest $\epsilon \geq 0$ such that $p + \epsilon a$ lies in C'. This can be done via linear programming. Note that some of the linear programs might be infeasible if $p + \epsilon a$ never enters some cone for any positive ϵ, and we ignore these cones. If all the feasible linear programs output a strictly positive value of ϵ then we know that $p + \epsilon a$ does not lie in any of the cones for some value of ϵ (in particular, for a value of ϵ smaller than the smallest minimum). In this case we declare that the cones do not cover $\mathbb{R}^d$. If the minimum ϵ for some cone is zero for each f then we conclude that none of the facets is a witness facet and thus $\mathbb{R}^d$ is covered. □

Corollary 3. *The problems* $\textsc{ConeCover}[\mathcal{VH}, \mathbb{R}^d, I]$, $\textsc{ConeCover}[\mathcal{H}, \mathbb{R}^t, I]$, *and* $\textsc{ConeCover}[\mathcal{VH}, \mathbb{R}^t, I]$ *can all be solved in polynomial time.*

Proof. The polynomiality of $\textsc{ConeCover}[\mathcal{VH}, \mathbb{R}^d, I]$ is obvious from Theorem 3. Since for an $\mathcal{H}$-cone one can find the facets of its intersection with any flat in polynomial time, it follows from Theorem 3 that $\textsc{ConeCover}[\mathcal{H}, \mathbb{R}^t, I]$ is polynomially solvable too. Again, it is an obvious consequence of this that $\textsc{ConeCover}[\mathcal{VH}, \mathbb{R}^t, I]$ can be solved in polynomial time as well. □

Fact 1. *For any* $t \in \mathbb{N}$, *we can write* $\mathbb{R}^t = \bigcup_{i=1}^{t+1} R_i$, *where* $R_1, \ldots, R_{t+1}$ *are pointed cones, pairwise-disjoint in the interior, whose* $\mathcal{H}$- *and* $\mathcal{V}$-*representations can be found in polynomial time.*

Proof. Take a simplex containing the origin in its interior. There are $t+1$ facets and the conic hull of the vertices each facet defines a cone. These $t+1$ cones are pairwise-disjoint in the interior, cover $\mathbb{R}^t$, and the $\mathcal{H}$- and $\mathcal{V}$-representations of these cones can be computed from the $\mathcal{H}$- and the $\mathcal{V}$-representation of the simplex. □

Let $C_1 = \{x \in \mathbb{R}^m \mid A_1 x \leq \mathbf{0}\} = \text{cone}(S_1)$ and $C_2 = \{x \in \mathbb{R}^n \mid A_2 x \leq \mathbf{0}\} = \text{cone}(S_2)$, where $A_1 \in \mathbb{R}^{l \times m}, A_2 \in \mathbb{R}^{r \times n}$ and $S_1 \subseteq \mathbb{R}^m, S_2 \subseteq \mathbb{R}^n$, be two polyhedral cones. The *direct-sum* of C_1 and C_2, is defined as:

$$\begin{aligned} C_1 \oplus C_2 &= \{(x, y) \in \mathbb{R}^m \times \mathbb{R}^n |\ A_1 x \leq \mathbf{0},\ A_2 y \leq \mathbf{0}\} \\ &= \text{cone}\left(\left\{\begin{pmatrix} v \\ \mathbf{0} \end{pmatrix} : v \in S_1\right\} \cup \left\{\begin{pmatrix} \mathbf{0} \\ v \end{pmatrix} : v \in S_2\right\}\right). \end{aligned}$$

Theorem 4. *Problem* ConeCover$[\mathcal{VH}, \mathbb{R}^d, III]$ *is NP-complete.*

Proof. Clearly the problem is in NP since a direction exists outside the union of the given cones if they do not cover $\mathbb{R}^d$. We can easily check if such a given direction indeed lies outside each of the cones since the facets of each cone are known. For proving its NP-hardness, we use a reduction from the following problem:

Sat$(V, \mathcal{F}, \mathcal{G})$: Given a finite set V and two hypergraphs $\mathcal{F}, \mathcal{G} \subseteq 2^V$, is there a set $X \subseteq V$ such that:

$$X \not\supseteq F \text{ for all } F \in \mathcal{F} \text{ and } X \not\subseteq G \text{ for all } G \in \mathcal{G}? \tag{1}$$

When $\mathcal{F} = \mathcal{G}$, this problem is called the *saturation problem* in [4], where it is proved to be NP-complete. Given $\mathcal{F}, \mathcal{G} \subseteq 2^V$, we construct two families of cones $\mathcal{C}_\mathcal{F}$ and $\mathcal{C}_\mathcal{G}$ in $\mathbb{R}^V$, such that there is a point $x \in \mathbb{R}^V \setminus \bigcup_{C \in \mathcal{C}_\mathcal{F} \cup \mathcal{C}_\mathcal{G}} C$ if and only if the pair $(\mathcal{F}, \mathcal{G})$ is not saturated (i.e. there is a set $X \subseteq V$ satisfying (1)).

For $X \subseteq V$, denote respectively by $\mathbb{R}^X_\geq$ and $\mathbb{R}^X_\leq$ the cones $\text{cone}\{\mathbf{e}_i : i \in X\} = \{x \in \mathbb{R}^X : x \geq \mathbf{0}\}$ and $\text{cone}\{-\mathbf{e}_i : i \in X\} = \{x \in \mathbb{R}^X : x \leq \mathbf{0}\}$, where $\mathbf{e}_i$ denotes the standard ith unit vector. Let $\overline{X} = V \setminus X$, and $R_i(X)$, for $i \in [|X|+1]$ the partition of $\mathbb{R}^X$ given by Fact 1. For each $F \in \mathcal{F}$, we define $|V| - |F| + 1$ cones $C^i_F = \mathbb{R}^F_\geq \oplus R_i(\overline{F})$, for $i \in [|\overline{F}|+1]$, and for each $G \in \mathcal{G}$, we define $|G|+1$ cones $C^i_G = \mathbb{R}^{\overline{G}}_\leq \oplus R_i(G)$, for $i \in [|G|+1]$. Finally, we let $\mathcal{C}_\mathcal{F} = \{C^i_F : F \in \mathcal{F},\ i \in [|\overline{F}|+1]\}$, $\mathcal{C}_\mathcal{G} = \{C^i_G : G \in \mathcal{G},\ i \in [|G|+1]\}$, and $\mathcal{C} = \mathcal{C}_\mathcal{F} \cup \mathcal{C}_\mathcal{G}$. Then it is not difficult to see that all the cones in $\mathcal{C}$ are pointed.

Suppose that $X \subseteq V$ satisfies (1). Define $x \in \mathbb{R}^V$ by

$$x_i = \begin{cases} 1, & \text{if } i \in X, \\ -1, & \text{if } i \in V \setminus X. \end{cases}$$

Then $x \notin \bigcup_{C \in \mathcal{C}} C$. Indeed, if $x \in C^i_F$, for some $F \in \mathcal{F}$ and $i \in [|\overline{F}|+1]$, then $x_j \geq 0$ and hence $x_j = 1$, for all $j \in F$, implying that $X \supseteq F$. Similarly, if $x \in C^i_G$, for some $G \in \mathcal{G}$ and $i \in [|G|+1]$, then $x_j \leq 0$ and hence $x_j = -1$, for all $j \in \overline{G}$, implying that $X \subseteq G$.

Conversely, suppose that $x \in \mathbb{R}^V \setminus \bigcup_{C \in \mathcal{C}} C$. Let $X = \{i \in V : x_i \geq 0\}$. Then we claim that X satisfies (1). Indeed, if $X \supseteq F$ for some $F \in \mathcal{F}$, then $x_i \geq 0$ for all $i \in F$, and hence there exists an $i \in [|\overline{F}|+1]$ such that $x \in C^i_F$ (since the cones $R_1(\overline{F}), \ldots, R_{|\overline{F}|+1}(\overline{F})$ cover $\mathbb{R}^{\overline{F}}$). Similarly, if $X \subseteq G$ for some $G \in \mathcal{G}$, then $x_i < 0$ for all $i \in \overline{G}$, and hence there exists an $i \in [|G|+1]$ such that $x \in C^i_G$. In both cases we get a contradiction. $\square$

Corollary 4. *The problems* ConeCover$[\mathcal{V}, \mathbb{R}^d, III]$, ConeCover$[\mathcal{H}, \mathbb{R}^d, III]$, ConeCover$[\mathcal{VH}, \mathbb{R}^t, III]$ *and* ConeCover$[\mathcal{H}, \mathbb{R}^t, III]$ *are all NP-complete.*

Proof. NP-completeness of $\textsc{ConeCover}[\mathcal{V}, \mathbb{R}^d, \text{III}]$ and $\textsc{ConeCover}[\mathcal{H}, \mathbb{R}^d, \text{III}]$ follows from Theorem 4. NP-completeness of $\textsc{ConeCover}[\mathcal{H}, \mathbb{R}^t, \text{III}]$ is an immediate consequence of the NP-hardness of $\textsc{ConeCover}[\mathcal{H}, \mathbb{R}^d, \text{III}]$ by setting $t = d$. $\square$

An interesting special case of problem $\textsc{Sat}$ is when the hypergraphs $\mathcal{F}$ and $\mathcal{G}$ are *transversal* to each other:

$$F \not\subseteq G \text{ for all } F \in \mathcal{F} \text{ and } G \in \mathcal{G}. \tag{2}$$

In this case, the problem is known as the *hypergraph transversal problem*, denoted $\textsc{HyperTrans}$. Even though the complexity of this problem is still open, it is unlikely to be NP-hard since there exist algorithms [5] that solve the problem in quasi-polynomial time $|V|m^{o(\log m)}$, where $m = |\mathcal{F}| + |\mathcal{G}|$. Improving this to a polynomial bound is a standing open problem. We observe from our reduction in Theorem 4 that $\textsc{ConeCover}$ includes $\textsc{HyperTrans}$ as a special case.

Corollary 5. *Consider a family of cones $\mathcal{C}$ that can be partitioned into two families $\mathcal{C}_1$ and $\mathcal{C}_2$ such that*

$$\text{int}(C_1) \cap \text{int}(C_2) = \emptyset, \quad \textit{for all } C_1 \in \mathcal{C}_1 \textit{ and } C_2 \in \mathcal{C}_2. \tag{3}$$

Then $\textsc{ConeCover}(\mathcal{C}, \mathbb{R}^d)$ *is* $\textsc{HyperTrans}$*-hard.*

Proof. We note in the construction used for the proof of Theorem 4 that if the hypergraphs $\mathcal{F}$ and $\mathcal{G}$ satisfy (2), then the families of cones $\mathcal{C}_{\mathcal{F}}$ and $\mathcal{C}_{\mathcal{G}}$ satisfy (3). Indeed, if $x \in C_F^i \cap C_G^j$, for some $F \in \mathcal{F}$, $i \in [|\overline{F}| + 1]$, $G \in \mathcal{G}$, and $j \in [|G| + 1]$, then $x_k \geq 0$ for all $k \in F$ and $x_k \leq 0$ for all $k \in \overline{G}$. Thus for any $k \in F \setminus G$ (which must exist by (2)), we have $x_k = 0$, implying that x is not an interior point in either C_F^i or C_G^j. $\square$

Freund and Orlin [6] proved that, for an $\mathcal{H}$-polytope P and a $\mathcal{V}$-polytope Q, checking if $Q \supseteq P$ is NP-hard. For all other representations of P and Q, checking $P \subseteq Q$ can be done by solving a linear program. Here we can show that the union version of this problem is hard, no matter how the polytopes are represented.

Corollary 6. *Given a set of $\mathcal{H}$-polytopes $\mathcal{P} = \{P_1, \ldots, P_k\}$ and an $\mathcal{H}$-polytope Q, problem* $\textsc{PolytopeCover}(\mathcal{P}, Q)$ *is NP-hard.*

Proof. We give a reduction from problem $\textsc{ConeCover}[\mathcal{H}, \mathbb{R}^d, \text{III}]$ which is NP-hard by Corollary 4. Let S_d be a "shifted" simplex in $\mathbb{R}^d$ such that $\mathbf{0} \in \text{int}(S_d)$. Given cones $C_1, \ldots, C_k$, we define polytopes $P_1, \ldots P_k$, by $P_i = C_i \cap S_d$. Given the $\mathcal{H}$-representations of C_i, we can compute the $\mathcal{H}$-representations of P_i in polynomial time using linear programming (LP) for removing possible redundancies.

Now one can easily see that $\bigcup_{i=1}^k C_i = \mathbb{R}^d$ if and only if $\bigcup_{i=1}^k P_i = S_d$. $\square$

Corollary 7. *Given a set of $\mathcal{V}$-polytopes $\mathcal{P} = \{P_1, \ldots, P_k\}$ and a $\mathcal{V}$-polytope Q, problem* PolytopeCover$(\mathcal{P}, Q)$ *is NP-hard.*

Proof. We give a reduction from problem ConeCover$[\mathcal{V}, \mathbb{R}^d, \mathrm{III}]$ which is NP-hard by Corollary 4. Recall that in the proof of Theorem 4, for each hyperedge F we construct a set of pointed cones $C_F^i = \mathbb{R}_{\geq}^F \oplus R_i(\overline{F})$, for $i \in [|\overline{F}| + 1]$. Instead of constructing multiple cones for each hyperedge let us just consider one cone $C_F = \mathbb{R}_{\geq}^F \oplus \mathbb{R}^{|\overline{F}|}$ per hyperedge. Similarly for the cones corresponding to the hypergraph $\mathcal{G}$. It is clear that $C_F = \bigcup_{i=1}^{|\overline{F}|+1} C_F^i$. Note that each such cone is not pointed but instead has a pointed part $\mathbb{R}_{\geq}^F$ corresponding to the vertices in the hyperedge F and the affine space $\mathbb{R}^{|\overline{F}|}$ corresponding to the vertices not in F. Also, $\mathbb{R}_{\geq}^F$ is one orthant in $\mathbb{R}^{|F|}$.

For such cones checking whether the union covers $\mathbb{R}^d$ or not is NP-hard as well (see proof of Theorem 4). Now consider the d-dimensional cross-polytope $\beta_d = \mathrm{conv}(\pm e_1, \cdots, \pm e_d)$, where e_i is the i-th unit vector. Also, let $C_1, \ldots, C_k$ be the cones constructed above. The cross-polytope β_d contains the origin in its interior, and the vertices of $P_i = \beta_d \cap C_i$ for each cone constructed above can be easily computed. (Note that intersecting each of the cones with a simplex, as in the proof of previous theorem, does not work since it is not clear whether the intersection of C_i and a simplex does not have very large number of vertices, let alone computing them.)

It is easy to see that $\bigcup_{i=1}^k C_i = \mathbb{R}^d$ if and only if $\bigcup_{i=1}^k P_i = \beta_d$, thus completing the proof of the theorem. □

Theorem 5. *Given a set of rational convex polytopes $P_1, \ldots, P_k \subseteq \mathbb{R}^d$, it is coNP-complete to check whether their union is convex, for both $\mathcal{H}$ and $\mathcal{V}$-representations of the input polytopes.*

Proof. First we show that the problem is in coNP. Let $Q = \bigcup_{i=1}^k P_i$. If this union is not convex then there are two points $x, y \in Q$, such that the line segment $[x, y] \stackrel{\text{def}}{=} \{\lambda x + (1 - \lambda) y | \ \lambda \in [0, 1]\} \not\subseteq Q$. Given such a certificate line segment it is easy to verify that $[x, y] \not\subseteq Q$ by iterating the algorithm for two polytopes in [3]:

1. Let P be the polytope P_i such that $x \in P_i$;
2. Find the (last) point $z \in P$ on the ray $\{x + \lambda(y - x) | \ \lambda \geq 0\}$ such that λ is maximized;
3. If $x = y$ then output "Yes" and halt;
4. If there is another polytope P_j such that $z \in P_j$, then set $P \leftarrow P_j$, $x \leftarrow z$, and go to step 2 else output "No" and halt.

The reader can verify that all the above steps can be implemented in polynomial time no matter how the polytopes are represented.

To establish NP-hardness, consider the $\mathcal{H}$-representation first. Let $\mathcal{P} = \{P_1, \ldots, P_k\}$ and S_d be the polytope used in the construction in Corollary 6. We now reduce the problem PolytopeCover$(\mathcal{P}, S_d)$ to checking if the union of a given set of polytopes is convex. Using an algorithm for the latter problem, we can check if $P = \bigcup_{i=1}^{k} P_i$ is convex. If the answer is "No", we conclude that $P \neq S_d$. Otherwise, since $P \subseteq S_d$, either $P = S_d$, or there is hyperplane separating a vertex of S_d from P. The latter condition can be checked in polynomial time by solving k linear programs for each vertex.

For the $\mathcal{V}$-representation the same argument as above works if we use β_d instead of S_d. □

3. Conclusion and Outlook

In this paper we studied the complexity of some polyhedral covering problems. Since a polytope (polyhedral cone resp.) can be represented both by its vertices (extreme rays resp.) or by its facet-defining hyperplanes there are many variants of these problems based on the input representation. We settle the complexity of most of the variants, but the status of a few variants remain unknown. In particular, when the input cones do not intersect in the interior but neither are the intersections necessarily common faces (condition (f3)-(II)), then it is not clear whether the problem is NP-hard like case (f3)-(III) or can be solved in polynomial time.

Also, ConeCover$[\mathcal{V}, \mathbb{R}^d, \mathrm{I}]$ and ConeCover$[\mathcal{V}, \mathbb{R}^d, \mathrm{II}]$ are shown to be at least as hard as the Vertex Enumeration problem. The complexity of the latter problem itself is not known, and thus a polynomial algorithm for these variants of the ConeCover problem would be a very interesting achievement. It is also very well possible that these two variants are NP-hard irrespective of the complexity of Vertex Enumeration, but at the moment we cannot prove an independent hardness result for either of these two variants.

We also showed that checking whether the union of a set of polytopes is convex or not is NP-hard when the polytopes are given either by facets of vertices. When both the representations are known then it is not clear if the problem remains NP-hard. The corresponding version of the problem for cones remains NP-hard even in this case but since for our reductions we need to intersect the cones with a simplex or a cross-polytope depending on the representation, we cannot extend the NP-hardness to the polytope version. We believe that choosing a simplex or a cross-polytope is just a technicality that can be removed from the proofs.

References

[1] D. Avis, D. Bremner, and R. Seidel. How good are convex hull algorithms? *Comput. Geom.*, 7:265–301, 1997.

[2] I. Bárány and K. Fukuda. A case when the union of polytopes is convex. *Linear Algebra and its Applications*, 397(6):381–388, 2005.

[3] A. Bemporad, K. Fukuda, and F. D. Torrisi. Convexity recognition of the union of polyhedra. *Comput. Geom.*, 18(3):141–154, 2001.

[4] T. Eiter and G. Gottlob. Identifying the minimal transversals of a hypergraph and related problems. *SIAM J. Comput.*, 24(6):1278–1304, 1995.

[5] M. L. Fredman and L. Khachiyan. On the complexity of dualization of monotone disjunctive normal forms. 21:618–628, 1996.

[6] R. M. Freund and J. B. Orlin. On the complexity of four polyhedral set containment problems. *Mathematical Programming*, 33(2):139–145, 1985.

[7] H. R. Tiwary. On computing the shadows and slices of polytopes. *CoRR*, abs/0804.4150, 2008.

[8] G. M. Ziegler. *Lectures on Polytopes*. Graduate Texts in Mathematics, No. 152. Springer-Verlag, Berlin, 1995.