

Circuits in Extended Formulations^{*}

Steffen Borgwardt^a, Matthias Brugger^b

^aUniversity of Colorado Denver, 1201 Larimer Street, Denver, CO, 80204, USA

^bTechnical University of Munich, Arcisstr. 21, 80333 München, Germany

Abstract

Circuits and extended formulations are classical concepts in linear programming theory. The circuits of a polyhedron are the elementary difference vectors between feasible points and include all edge directions. We study the connection between the circuits of a polyhedron P and those of an extended formulation of P , i.e., a description of a polyhedron Q that linearly projects onto P . It is well known that the edge directions of P are images of edge directions of Q . We show that this ‘inheritance’ under taking projections does not extend to the set of circuits, and that this non-inheritance is quite generic behavior. We provide counterexamples with a provably minimal number of facets, vertices, and extreme rays, including relevant polytopes from clustering, and show that the difference in the number of circuits that are inherited and those that are not can be exponentially large in the dimension. We further prove that counterexamples exist for any fixed linear projection map, unless the map is injective. Finally, we characterize those polyhedra P whose circuits are inherited from all polyhedra Q that linearly project onto P . Conversely, we prove that every polyhedron Q satisfying mild assumptions can be projected in such a way that the image polyhedron P has a circuit with no preimage among the circuits of Q . Our proofs build on standard constructions such as homogenization and disjunctive programming.

Keywords: circuits, extended formulations, polyhedral theory, linear programming

MSC: 52B05, 52B12, 90C05

1. Introduction

Extended formulations have been widely studied in polyhedral combinatorics and the theory of linear programming, with many recent advances (see, e.g., [2, 21, 27, 28, 29, 31, 37, 38, 43] as well as [22, Chapter 4] and references therein). An *extended formulation* of a polyhedron P is a linear system $Ay = b, By \leq d$ in variables y defining a polyhedron Q that can be affinely projected onto P . The polyhedron Q , along with the affine projection that maps Q to P , is called an *extension* of P . When the projection is clear from the context, for a simple wording, we may refer to Q itself as an extension or an *extension polyhedron*. The relevance of extended formulations for optimization comes from the fact that one may optimize any linear functional over P by solving a linear program (LP) with feasible region Q instead, which may yield a more compact formulation with fewer constraints. Especially for many problems in combinatorial

^{*}The first author was supported by Air Force Office of Scientific Research grant FA9550-21-1-0233 and NSF grant 2006183, Algorithmic Foundations, Division of Computing and Communication Foundations. The second author was partially supported by the Alexander von Humboldt Foundation with funds from the German Federal Ministry of Education and Research (BMBF).

Email addresses: `steffen.borgwardt@ucdenver.edu` (Steffen Borgwardt), `matthias.brugger@tum.de` (Matthias Brugger)

optimization, and LP-based approaches in the form of relaxations, extended formulations have become a powerful tool since the associated polyhedra typically have an exponential number of facets. In some of these cases, significantly smaller extensions with only a polynomial number of facets have been shown to exist [2, 31, 40, 47, 48]. In other situations, such as for the fixed-shape partition polytopes that we discuss below, there does not even exist a known linear formulation [5, 13, 20, 34].

In this paper, we study the connection between the set of circuits of a polyhedron and the set of circuits of an extension. The *circuits* or elementary vectors [42] of a polyhedron are the minimal linear dependence relations in its constraint matrix and include all edge directions (we provide a formal definition in Section 1.2). Circuit walks, circuit augmentation schemes [25], and circuit diameters [11] generalize the classical concepts of edge walks, the Simplex method, and combinatorial diameters by following steps along the more general set of circuits, so in particular through the interior of a polyhedron. Circuits and circuit walks are of particular interest for polyhedra associated with combinatorial optimization problems, so precisely in the setting where extended formulations are especially useful. This is because the circuits of polyhedra in combinatorial optimization can often be readily interpreted in terms of the underlying application [36]. For example, the circuits for network flow problems are cycles and well-known results like flow decomposition [1, 10] are immediate consequences of such interpretations.

We are interested in whether the set of circuits of a polyhedron, which may be difficult to describe directly, can be ‘accessed’ through the set of circuits of an extension. We ask the following fundamental question:

When are all circuits of a polyhedron P projections of circuits of a given extension Q ?

As we will see, the connection between the sets of circuits is far weaker than it is for the edge directions. We will be able to quantify this weaker behavior in several ways. Before we explain our main contributions in Section 1.3, we describe our original motivation for this work in Section 1.1, and recall some necessary concepts and introduce convenient notation in Section 1.2.

1.1. Motivation

Fixed-shape clustering is the task of partitioning a data set X of n items into k clusters C_1, \dots, C_k such that the number of items in each C_i equals a fixed number $\kappa_i \in \mathbb{N}$, where $\sum_{i=1}^k \kappa_i = n$. For a data set $X = \{x^{(1)}, \dots, x^{(n)}\} \subseteq \mathbb{R}^d$, popular clustering objectives like least-squares assignments can be found through linear optimization over the so-called *fixed-shape partition polytopes* [5, 8, 15, 20, 30, 33, 34]. These are formed as the convex hull of all feasible *clustering vectors* $(c^{(1)}, \dots, c^{(k)}) \in (\mathbb{R}^d)^k$, where each cluster C_i is represented through a vector $c^{(i)} = \sum_{x \in C_i} x$ of items assigned to it. For given X , k and a vector of cluster sizes $\kappa := (\kappa_1, \dots, \kappa_k)$, we denote these polytopes by $P(X, k, \kappa)$.

Our original motivation for the work in this paper was based on our interest in the circuits of the polytopes $P(X, k, \kappa)$. For example, a characterization of the set of circuits would lead to improved methods for gradual transitions between separable clusterings [14] and could provide insight into possible efficient approximations of the hard-to-compute normal cones of its vertices, whose volume can be used to measure robustness or quality of the associated clusterings [13].

An explicit inequality description of the fixed-shape partition polytope $P(X, k, \kappa)$ is not known. Instead, computations are performed over a certain transportation polytope, which is

given by the following system in variables $y \in \mathbb{R}^{k \times n}$ and which we denote by $T(n, k, \kappa)$:

$$\begin{aligned} \sum_{j=1}^n y_{ij} &= \kappa_i \quad \forall i \in [k] \\ \sum_{i=1}^k y_{ij} &= 1 \quad \forall j \in [n] \\ y &\geq \mathbf{0} \end{aligned}$$

Via the linear map $\pi_X: y \mapsto (c^{(1)}, \dots, c^{(k)})$ with $c^{(i)} = \sum_{j=1}^n y_{ij} \cdot x^{(j)}$ for all $i \in [k]$, the polytope $T(n, k, \kappa)$ projects to $P(X, k, \kappa)$.

We became interested in whether the set of circuits of $P(X, k, \kappa)$ could be characterized via projecting from $T(n, k, \kappa)$. There are a couple of favorable properties of the two polyhedra that made such an approach especially promising [17, 18]; for example, the edges of both polyhedra have near-identical characterizations and all circuit walks in $T(X, k, \kappa)$ are, in fact, edge walks. Nonetheless, somewhat surprisingly, in Section 3.4 we exhibit that new circuits may appear in the projection onto $P(X, k, \kappa)$.

Our interest in the behavior of circuits and circuit walks under taking projections is further motivated by a well-known fact about *edge* walks: for every edge walk in the original polyhedron, there is an edge walk in the extension that projects onto it. In particular, all edge directions of the original polyhedron are images of edge directions of the extension polyhedron (see Section 2 for a proof). In the context of linear programming, there is a pivot rule for the Simplex method that relies precisely on that relationship: the shadow vertex pivot rule [7]. This pivot rule constructs a Simplex path by following an edge walk in a two-dimensional projection (shadow) of the feasible region of the LP. The shadow vertex pivot rule and modifications thereof play an important role in the probabilistic analysis of the Simplex method [23, 45, 46] and in recent work on strong bounds for the performance of the Simplex method on 0/1 polytopes [6]. We believe that a deeper understanding of the behavior of circuits in the context of projections and lifting could also provide new tools for the analysis of circuit augmentation schemes and corresponding pivot rules.

1.2. Notation and Definitions

We begin with a formal definition of the set of circuits of a polyhedron, following [11, 12, 16, 25, 42], and then introduce some new terminology for our purposes.

Circuits. Let $P = \{x \in \mathbb{R}^n: Ax = b, Bx \leq d\}$ be a polyhedron in \mathbb{R}^n where $A \in \mathbb{R}^{p \times n}$, $B \in \mathbb{R}^{q \times n}$ and $b \in \mathbb{R}^p$, $d \in \mathbb{R}^q$. The *circuits* of P with respect to its linear description are the vectors $g \in \ker(A) \setminus \{\mathbf{0}\}$ such that Bg is support-minimal (w.r.t. inclusion) in $\{By: y \in \ker(A) \setminus \{\mathbf{0}\}\}$. Here, $\ker(A)$ denotes the kernel of A , $\mathbf{0}$ denotes the all-zero vector in appropriate dimension, and the *support* of a vector $z \in \mathbb{R}^p$ is the set $\{i \in [p]: z_i \neq 0\}$. Circuits correspond to directions in the underlying space \mathbb{R}^n , and any form of normalization leads to a finite set of unique representatives for the directions. A standard normalization scheme is to assume co-prime integer components (which assumes rational data). For our purposes, it will often be convenient to view the set of circuits as a finite union of one-dimensional linear subspaces. Each of these subspaces is generated by a pair of circuits $\pm g$, and we call these pairs of circuits the *unique circuit directions* of P .

The set of circuits depends on the description of P . We denote it as $\mathcal{C}(A, B)$ when we refer to a specific description; when a system defining P is clear from the context, we use $\mathcal{C}(P)$ in place of $\mathcal{C}(A, B)$. For any description, $\mathcal{C}(P)$ always contains all edge directions of P , where an

edge direction is either the direction of an extreme ray of P or a nonzero multiple of $u - v$ for some pair of adjacent vertices $u, v \in \mathcal{V}(P)$. Here, $\mathcal{V}(P)$ denotes the set of vertices of P .

Unless stated otherwise, we assume a minimal (or irredundant) description for P when it is not provided explicitly; this implies that each inequality constraint defines a facet of P . This is not a restriction for A , as $\ker(A)$ is not affected by the existence of redundant equalities. For B , it is standard to assume irredundancy for a different reason: the addition of redundant inequalities to B may lead to a larger set of circuits that are not tied to the geometry of the underlying polyhedron. The benefit of an irredundant description can most easily be seen through an equivalent ‘geometric’ definition of the set of circuits: $\mathcal{C}(A, B)$ consists of all potential edge directions as the right-hand side vectors b and d vary [32]. With an irredundant description of P , all facets for any choice of b and d correspond to an original facet of P . In this case, the entire set $\mathcal{C}(P)$ can be constructed through forming all possible one-dimensional intersections of (translated) facets of P [18]; for a redundant description of P , the non-facet-defining constraints also have to be considered. Finally, note that any two irredundant descriptions of P yield the same set of circuits; see, e.g., [35, Lemma 3].

Inheritance under affine projections. Let $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ be polyhedra such that $P = \pi(Q)$ for some affine map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$. We say that a circuit $g \in \mathcal{C}(P)$ is *inherited from Q under π* if $g \in \pi(\mathcal{C}(Q)) - \pi(\mathbf{0})$, i.e., if g is the image of a circuit of Q under the ‘linearized’ map $x \mapsto \pi(x) - \pi(\mathbf{0})$. If π is clear from the context, we use the simpler wording $g \in \mathcal{C}(P)$ is inherited from Q . It is easy to see that $\mathcal{C}(P)$ is unaffected by translations of the polyhedron P (cf. Lemma 2 in Section 2). In our discussion of extensions of P , we may therefore restrict ourselves to polyhedra that *linearly* project to P (unless stated otherwise).

Recall that P and Q are either given through an explicit description or are assumed to have a *minimal* description. In the latter case, the assumption of minimality of the implicit linear descriptions is crucial: if we add suitable redundant inequalities to the description of P , we may always generate a circuit that is not inherited from Q . Likewise, by introducing redundancy to the description of Q , one could blow up the set of possible circuits in Q such that all circuits of P would ultimately be inherited from this expanded extended formulation. Thus, our setting is well-defined to only obtain the strongest statements about the inheritance of circuits.

Most of our results will assume *pointed* polyhedra, i.e., polyhedra given by $Ax = b, Bx \leq d$ whose lineality space $\ker(A) \cap \ker(B)$ is trivial, i.e., the matrix $\binom{A}{B}$ has full column rank. This is appropriate for our purposes since $\mathcal{C}(A, B) = \ker(A) \cap \ker(B) \setminus \{\mathbf{0}\}$ if $\ker(A) \cap \ker(B) \neq \{\mathbf{0}\}$ and, hence, $\mathcal{C}(\pi(Q)) = \pi(\mathcal{C}(Q))$ is trivially satisfied for any polyhedron Q with a nontrivial lineality space, for all linear maps π .

1.3. Contributions and Outline

We show that polyhedra do not necessarily inherit their circuits from extended formulations. This illustrates a fundamental difference between edge directions and circuits. We demonstrate this in two different ways. After collecting some tools in Section 2 that simplify our discussion, we first prove in Section 3 that polytopes with sufficiently many facets in general position have extensions from which not all circuits are inherited. For a polyhedron P in \mathbb{R}^n given by an irredundant description $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$, we say that its facets are *in general position* if, for every row submatrix B' of B with $\dim(P)$ rows, the matrix $\binom{A}{B'}$ has full rank n . If P is full-dimensional, this reduces to the usual definition of general position for the facet normals of P (the rows of B), which requires all subsets of n facet normals to be linearly independent.

Theorem 1. *Let $P \subseteq \mathbb{R}^n$ be a polytope whose facets are in general position, and suppose that P has more facets than vertices. Then there exist a simple polytope $Q \subseteq \mathbb{R}^m$ for some $m \in \mathbb{N}$*

and a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$ such that $\mathcal{C}(P) \not\subseteq \pi(\mathcal{C}(Q))$. Moreover, if P has at least twice as many facets as vertices, then the number of unique circuit directions in $\mathcal{C}(P)$ exceeds the number of those in $\mathcal{C}(Q)$ by a factor of at least $2^{\Omega(\dim(P))}$.

The proof of Theorem 1 is based on a simple (non-constructive) counting argument: if P has more unique circuit directions than its extension Q , not all circuits can be inherited from Q .

Second, we provide an alternative, constructive approach and show how to construct counterexamples with a minimal number of facets, vertices, and extreme rays. Our construction yields both bounded and unbounded polyhedra in every dimension greater than 2, with corresponding extensions just one dimension higher.

Theorem 2. *For all $m, n \in \mathbb{N}$ with $m > n \geq 3$, there exist full-dimensional pointed polyhedra $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ and a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$ such that $\mathcal{C}(P) \not\subseteq \pi(\mathcal{C}(Q))$ and $\mathcal{C}(P) \cap \pi(\mathcal{C}(Q))$ consists precisely of the edge directions of P . Moreover, Q can be chosen to be simple, and P can be chosen to be either a polytope with $n+2$ facets and $n+2$ vertices, or a pointed polyhedral cone with $n+1$ facets and $n+1$ extreme rays.*

Next, we show that an exponential gap in the number of unique directions between the subset of circuits that are inherited and the entire set of circuits, as given by the second part of Theorem 1, can also be established in a constructive manner.

Theorem 3. *For all $n \geq 3$, one can construct full-dimensional pointed polyhedra $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ and a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$ such that $\mathcal{C}(P)$ contains $2^{\Omega(n)}$ unique circuit directions while the number of unique circuit directions in $\mathcal{C}(Q)$ is $O(n^2)$.*

Theorems 1 to 3 suggest that polyhedra and projection maps need to satisfy special properties in order for circuits to be inherited. We conclude Section 3 with a transfer to the fixed-shape partition polytopes from Section 1.1, where an inheritance fails despite these polyhedra exhibiting a number of favorable properties.

We are interested in understanding which properties of the three ‘ingredients’ – the original polyhedron P , the extension polyhedron Q , and the projection map π from Q to P – can guarantee the inheritance of circuits by themselves. This leads us to the following three questions, which we study in Section 4:

- (Q1) *Which linear maps π have the property that, for every polyhedron Q , all circuits of $\pi(Q)$ are inherited from Q ?*
- (Q2) *Which polyhedra P inherit their circuits from every extension?*
- (Q3) *For which polyhedra Q does every polyhedron P that is a linear projection of Q inherit its circuits from Q ?*

It is not hard to exhibit two sufficient properties for membership in the classes stated in questions (Q1) and (Q2), respectively. We show that injective maps belong to the class of maps for question (Q1), and we show that polyhedra in which all circuits are edge directions belong to the class of polyhedra for (Q2). As these observations are useful multiple times for our discussion, the brief arguments can already be found in the preliminaries in Section 2.

What is more interesting is that these properties are, in fact, both sufficient and *necessary*, thus completely resolving questions (Q1) and (Q2). More precisely, in Section 4, we obtain the following two results. First, we strengthen Theorem 2 and show that, for any linear projection, one can construct a (bounded) counterexample in which no circuits other than the edge directions are inherited, unless the projection map is injective.

Theorem 4. *Let $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear map such that $\dim(\pi(\mathbb{R}^m)) \geq 3$. Then $\mathcal{C}(\pi(Q)) \subseteq \pi(\mathcal{C}(Q))$ for all polyhedra $Q \subseteq \mathbb{R}^m$ if and only if π is injective. In particular, if π is not injective, there exists a full-dimensional simple polytope $Q \subseteq \mathbb{R}^m$ such that $\mathcal{C}(\pi(Q)) \not\subseteq \pi(\mathcal{C}(Q))$ and $\mathcal{C}(\pi(Q)) \cap \pi(\mathcal{C}(Q))$ consists precisely of the edge directions of $\pi(Q)$.*

Second, we provide a formal answer to question (Q2).

Theorem 5. *Let $P \subseteq \mathbb{R}^n$ be a pointed polyhedron. All circuits in $\mathcal{C}(P)$ are edge directions of P if and only if $\mathcal{C}(P) \subseteq \pi(\mathcal{C}(Q))$ for all polyhedra $Q \subseteq \mathbb{R}^m$ and all linear maps $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$.*

Our proofs of Theorems 4 and 5 in Section 4 are constructive, and build upon observations and tools collected in Sections 2 and 3. Finally, we give a partial answer to question (Q3), showing that no polyhedron with a non-degenerate vertex has the property that (Q3) asks for.

Theorem 6. *Let $Q \subseteq \mathbb{R}^m$ be a polyhedron with $\dim(Q) \geq 4$. If Q has a non-degenerate vertex, then there exists a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^{\dim(Q)-1}$ such that $\pi(Q)$ is full-dimensional and $\mathcal{C}(\pi(Q)) \not\subseteq \pi(\mathcal{C}(Q))$.*

In summary, Theorems 4 to 6 show that, whenever a polyhedron P inherits all of its circuits from another polyhedron Q (with a non-degenerate vertex) under some affine projection π , this is not a property of any single one of the three ‘ingredients’ P , Q , and π – unless inheritance is immediate because π defines an affine isomorphism between P and Q or because P has no circuits that are not edge directions. This means that the inheritance of circuits, beyond these simple cases, can only be a property of specific *combinations* of the three ingredients. We conclude Section 4 with a family of examples of such nontrivial combinations that guarantee inheritance, and provide some final remarks in Section 5.

2. Preliminaries

We first note that every pointed polyhedron P has some circuits that are naturally inherited from any extension: the edge directions of P . This is a well-known fact about projections of polyhedra. As a service to the reader, we include a brief proof.

Lemma 1. *Let $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ be pointed polyhedra and let $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear map with $\pi(Q) = P$. For every edge direction g of P , there exists an edge direction f of Q such that $\pi(f) = g$.*

Proof. First, recall that any face F of P is the image of a face of Q under π , namely the face $\pi^{-1}(F) := \{y \in Q: \pi(y) \in F\}$ (see, e.g., [27, Proposition 2.1] or [49, Lemma 7.10] for the polytopal case).

Now let e be an edge of P . Since P is pointed by hypothesis, e has a vertex v . Then both $\pi^{-1}(e)$ and $\pi^{-1}(v)$ are faces of Q with $\pi^{-1}(v) \subsetneq \pi^{-1}(e)$. Since Q is pointed, $\pi^{-1}(v)$ is pointed, too, and therefore has a vertex v' . Consider the edges of $\pi^{-1}(e)$ that are incident with v' . Not all of them can be contained in $\pi^{-1}(v)$, for otherwise $\pi^{-1}(v)$ contains the inner cone of $\pi^{-1}(e)$ at v' , and therefore $\pi^{-1}(v) \supseteq \pi^{-1}(e)$, a contradiction. Thus, $\pi^{-1}(e)$ has an edge that contains both v' and some point w' with $\pi(w') \neq v$. For this point, $v' - w'$ is an edge direction of Q such that $\pi(v' - w') = v - \pi(w') \neq \mathbf{0}$ and $\pi(w') \in e$. \square

Lemma 1 shows that, if all circuits of P are edge directions, then $\mathcal{C}(P) \subseteq \pi(\mathcal{C}(Q))$ for any extension of P specified by Q and π . This special case includes hypercubes and simplices,

and more complicated polytopes such as Birkhoff polytopes and fractional matching polytopes [26, 44].

In general, the set of circuits of a polyhedron may of course be much larger than the set of its edge directions; c.f., Theorem 1. However, we can make a simple a priori observation: two polyhedra that are affinely isomorphic trivially are extensions of one another. Recall that two polyhedra $P \subseteq \mathbb{R}^p$ and $Q \subseteq \mathbb{R}^q$ are *affinely (linearly) isomorphic* if there exists an affine (linear) map $\pi: \mathbb{R}^q \rightarrow \mathbb{R}^p$ such that $\pi(Q) = P$ and, for all $x \in P$, there exists a unique $y \in Q$ with $\pi(y) = x$. The sets of circuits of affinely isomorphic polyhedra are isomorphic, too, as the next lemma states. Since translations are special affine isomorphisms, this justifies our assumption made in Section 1.2 that all projection maps are *linear* maps.

Lemma 2. *Let $Q \subseteq \mathbb{R}^m$ be a polyhedron and let $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be an affine map. If Q and $\pi(Q)$ are affinely isomorphic, then $\mathcal{C}(\pi(Q)) = \pi(\mathcal{C}(Q)) - \pi(\mathbf{0})$.*

Proof. Let π be an affine isomorphism between $P := \pi(Q)$ and Q . It is easy to see that the affine hulls $\text{aff}(P)$ and $\text{aff}(Q)$ are affinely isomorphic as well, and that there is a one-to-one correspondence between the facets of P and the facets of Q . Using the geometric interpretation of the set of circuits (see Section 1.2), we obtain the statement. \square

Another setting that is easy to resolve is when the polyhedra are low-dimensional. Consider a polyhedron P and an extension Q of P with $\dim(Q) \leq 3$. Then either $\dim(P) = 3$ and thus $\dim(Q) = 3$, in which case P and Q must be affinely isomorphic, or $\dim(P) \leq 2$. In the latter case, every facet of P (if one exists) is an edge of P . Hence, every circuit of P trivially is an edge direction. In summary, we obtain the following corollary to Lemmas 1 and 2.

Corollary 1. *Let $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ be pointed polyhedra and let $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear map such that $\pi(Q) = P$. If $\dim(P) \leq 2$ or $\dim(Q) \leq 3$, then $\mathcal{C}(P) \subseteq \pi(\mathcal{C}(Q))$.*

Lemma 2 will be one of the key ingredients for proving Theorem 4 (see Section 4.1). Further, the lemma has several interesting implications for certain types of extended formulations and projections: one of the simplifying assumptions commonly made in the study of extended formulations is that the projection is an orthogonal projection onto a subspace of the variables. In our context, such a projection is just as general as any other linear projection.

Corollary 2. *Let $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ be polyhedra and let $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear map such that $\pi(Q) = P$. Further let $Q' := \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^m : x = \pi(y), y \in Q\}$ and $\pi': \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ defined by $(x, y) \mapsto x$. Then $\pi'(Q') = P$ and $\pi(\mathcal{C}(Q)) = \pi'(\mathcal{C}(Q'))$.*

Proof. The claim follows immediately from Lemma 2, since the map $\tau: y \mapsto (\pi(y), y)$ defines a linear isomorphism between Q and $Q' = \tau(Q)$. \square

Another consequence of Lemma 2 that was already observed in [24, Lemma 7.1] is that every pointed polyhedron is affinely isomorphic to a polyhedron in standard form whose set of circuits is isomorphic to the set of circuits of the original polyhedron. We include a brief proof of this fact. Recall that a polyhedron P is in *standard form* if it is given by a linear system of the form $Ax = b, x \geq \mathbf{0}$.

Corollary 3 ([24]). *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a pointed polyhedron where $B \in \mathbb{R}^{m \times n}$. Define the affine map $\sigma: \mathbb{R}^n \rightarrow \mathbb{R}^m, x \mapsto d - Bx$. Then $\sigma(P)$ is a polyhedron with a standard form description such that $\mathcal{C}(\sigma(P)) = B \cdot \mathcal{C}(P) = \{Bg : g \in \mathcal{C}(P)\}$. In other words, $\mathcal{C}(\sigma(P))$ is the set of support-minimal vectors in $B \cdot \ker(A)$.*

Proof. Let $x, y \in P$ such that $\sigma(x) = \sigma(y)$. Then $A(x - y) = \mathbf{0}$ and $B(x - y) = \mathbf{0}$. Since P is pointed, it follows that $x = y$. Hence, σ is an isomorphism between P and $\sigma(P)$. Note that $\text{aff}(\sigma(P)) = \sigma(\text{aff}(P))$. We further claim that $\sigma(P) = \text{aff}(\sigma(P)) \cap \mathbb{R}_{\geq 0}^m$ (see also [27, Section 2.2]). Clearly, $\sigma(P) \subseteq \text{aff}(\sigma(P)) \cap \mathbb{R}_{\geq 0}^m$. To see that the converse inclusion also holds, let $s \in \text{aff}(\sigma(P)) \cap \mathbb{R}_{\geq 0}^m$, i.e., $s = \sigma(z) \geq \mathbf{0}$ for some $z \in \text{aff}(P)$. In particular, we have that $Az = b$ and $Bz \leq d$, which implies that $z \in P$ as claimed. Thus, the description of $\sigma(P)$ as $\text{aff}(\sigma(P)) \cap \mathbb{R}_{\geq 0}^m$ is in standard form. By applying Lemma 2, we obtain $\mathcal{C}(\sigma(P)) = \sigma(\mathcal{C}(P)) - d = B \cdot \mathcal{C}(P)$. \square

We point out that Corollary 3 contrasts with the behavior of circuits under the standard conversion of a polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ to standard form: in addition to introducing slack variables $s \geq \mathbf{0}$ to obtain $Bx + s = d$, one splits each variable x into a positive and a negative part $x = x^+ - x^-$, both of which are constrained to be nonnegative. It is shown in [19] that this conversion may introduce exponentially many new circuits. Corollary 3 suggests that this behavior is a consequence of splitting the variables and not of introducing slack variables, which is what applying the slack map σ defined in Corollary 3 implicitly does as well. More precisely, $\sigma(P)$ is the projection of $P' := \{(x, s) \in \mathbb{R}^n \times \mathbb{R}^m : Ax = b, Bx + s = d, s \geq \mathbf{0}\}$ onto the slack variables s . By Corollary 2, P' and P (and, hence, P' and $\sigma(P)$) are affinely isomorphic. Characterizing $\sigma(P)$ via P' adds the benefit that one can derive an explicit standard form representation of $\sigma(P)$ from the description of P' , using a projection technique found, e.g., in [22, Theorem 3.46]: for a basis $\{(u^{(1)}, v^{(1)}), \dots, (u^{(l)}, v^{(l)})\}$ of $\ker((B^\top \ A^\top))$, we have that $\sigma(P) = \{s \in \mathbb{R}^m : s \geq \mathbf{0}, (u^{(i)})^\top s = (u^{(i)})^\top d \ \forall i \in [l]\}$.

We conclude these preliminaries with a final simple tool that will be useful in the next sections. Every circuit of the Cartesian product $P_1 \times P_2$ of polyhedra P_1 and P_2 is a circuit of one of the product terms, suitably padded with zeros. This corresponds to a standard fact for the direct sum of two matroids (see [41, p. 131]) and transfers immediately to circuits of polyhedra in canonical form, as shown in [16, Lemma 3.9]. We provide the result in all generality here.

Proposition 1 ([16, 41]). *Let $P_1 \subseteq \mathbb{R}^{n_1}, P_2 \subseteq \mathbb{R}^{n_2}$ be pointed polyhedra. Then $\mathcal{C}(P_1 \times P_2) = (\mathcal{C}(P_1) \times \{\mathbf{0}\}) \cup (\{\mathbf{0}\} \times \mathcal{C}(P_2))$.*

Proof. Let $P_i = \{x \in \mathbb{R}^{n_i} : A^{(i)}x = b^{(i)}, B^{(i)}x \leq d^{(i)}\}$ for $i \in \{1, 2\}$. Then $\mathcal{C}(P_1 \times P_2)$ consists precisely of those nonzero vectors $(g^{(1)}, g^{(2)}) \in \ker(A^{(1)}) \times \ker(A^{(2)})$ for which the support of $(B^{(1)}g^{(1)}, B^{(2)}g^{(2)})$ is inclusion-minimal. Let $(g^{(1)}, g^{(2)}) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ be a nonzero vector in $\ker(A^{(1)}) \times \ker(A^{(2)})$. W.l.o.g., we may assume that $g^{(1)} \neq \mathbf{0}$. Then $(g^{(1)}, \mathbf{0}) \in \ker(A^{(1)}) \times \ker(A^{(2)})$ and the support of $(B^{(1)}g^{(1)}, \mathbf{0})$ is contained in the support of $(B^{(1)}g^{(1)}, B^{(2)}g^{(2)})$. Since P_2 is pointed, it follows that $(g^{(1)}, g^{(2)}) \in \mathcal{C}(P_1 \times P_2)$ if and only if $g^{(2)} = \mathbf{0}$ and $g^{(1)} \in \mathcal{C}(P_1)$. \square

3. Counterexamples for the Inheritance of Circuits

In this section, we prove that, in general, circuits of polyhedra are not inherited from extended formulations. This contrasts with the behavior for edge directions stated in Lemma 1. We begin by showing that the weaker behavior of the set of circuits is rather generic.

3.1. A Generic Counting Argument

Given two pointed polyhedra P and Q such that $P = \pi(Q)$ for some linear map π , a necessary condition for the circuits of P to be inherited from Q under π is that Q has at least as many unique circuit directions as P . Under a general position assumption, this condition is not satisfied for polytopes P with many facets and suitable extensions Q .

Theorem 1. *Let $P \subseteq \mathbb{R}^n$ be a polytope whose facets are in general position, and suppose that P has more facets than vertices. Then there exist a simple polytope $Q \subseteq \mathbb{R}^m$ for some $m \in \mathbb{N}$ and a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$ such that $\mathcal{C}(P) \not\subseteq \pi(\mathcal{C}(Q))$. Moreover, if P has at least twice as many facets as vertices, then the number of unique circuit directions in $\mathcal{C}(P)$ exceeds the number of those in $\mathcal{C}(Q)$ by a factor of at least $2^{\Omega(\dim(P))}$.*

To prove Theorem 1, we will use the following well-known characterization of the set of circuits [19, 32, 36]: for a pointed polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$, the set $\mathcal{C}(A, B)$ consists precisely of the nonzero solutions of systems of the form $Ax = \mathbf{0}, B'x = \mathbf{0}$ where B' is a row submatrix of B such that the rank of $\binom{A}{B'}$ is $n - 1$.

Proof of Theorem 1. Suppose that P is given by an irredundant description with equality and inequality constraint matrices A and B , respectively. Further let $d := \dim(P)$, and let f_0 and f_{d-1} denote the number of vertices and facets of P , respectively. In other words, f_{d-1} is the number of rows of B and $f_0 = |\mathcal{V}(P)|$, where $\mathcal{V}(P)$ denotes the set of vertices of P .

We first claim that each subset of $d - 1$ rows of B induces a unique circuit direction of P , and that no two distinct subsets induce the same direction. To see this, take an arbitrary row submatrix B' of B consisting of $d - 1$ rows. By the hypothesis that the facets of P are in general position, the matrix $\binom{A}{B'}$ must have rank $n - 1$ (recall that A has $n - d$ rows in an irredundant description). Its kernel is therefore generated by a circuit $g \in \mathcal{C}(P)$. Further, B' is the maximal row submatrix of B such that $B'g = \mathbf{0}$. Otherwise, we could add a row to B' to obtain a matrix B'' with d rows such that the rank of $\binom{A}{B''}$ is still $n - 1$, contradicting the general position assumption.

It follows from the claim that P has $\binom{f_{d-1}}{d-1}$ unique circuit directions (pairs of circuits $\pm g$), each induced by a subset of $d - 1$ rows of B . We now show that there is an extension of P with fewer unique circuit directions, which then implies the first part of the statement. Indeed, since P is a polytope, it is the image of the simplex

$$Q := \left\{ \lambda \in \mathbb{R}^{\mathcal{V}(P)} : \lambda \geq \mathbf{0}, \sum_{v \in \mathcal{V}(P)} \lambda_v = 1 \right\}$$

under the linear map $\lambda \mapsto \sum_{v \in \mathcal{V}(P)} \lambda_v v$. All circuits of Q are edge directions (see Section 2). Thus, the number of unique circuit directions in $\mathcal{C}(Q)$ is equal to $\binom{f_0}{2}$. Note that $f_0 \geq d + 1$ because P is a polytope. As $f_{d-1} = f_0$ for all polytopes with $d \leq 2$, the hypothesis that $f_{d-1} > f_0$ implies that $d \geq 3$. It follows that

$$\binom{f_{d-1}}{d-1} > \binom{f_0}{d-1} \geq \binom{f_0}{f_0-2} = \binom{f_0}{2}$$

where the first inequality is due to $f_{d-1} > f_0$, and the second one follows from $2 \leq d - 1 \leq f_0 - 2$ since $d \geq 3$ and $f_0 \geq d + 1$.

To prove the second part of the statement, now suppose that $f_{d-1} \geq 2f_0$. Then we have that

$$\binom{f_{d-1}}{d-1} \geq \binom{2f_0}{d-1} > 2^{d-1} \binom{f_0}{d-1} \geq 2^{d-1} \binom{f_0}{2}$$

where the last inequality follows as before. \square

We note that the facets of simple or simplicial polytopes can always be brought into general position by slightly tilting the facets or perturbing the vertices, respectively. In general, such

perturbations must preserve the combinatorial structure (in particular, retain the number of facets and vertices) for Theorem 1 to apply. In the next two sections, we show that counterexamples to the inheritance of circuits (including ones with an exponential blow-up in the number of unique circuit directions) can alternatively be obtained in a constructive way, as opposed to the non-constructive counting argument of Theorem 1.

3.2. A Family of Minimal Counterexamples

We construct a family of provably minimal counterexamples. These will also serve as the basis of our proofs in Sections 3.4 and 4. The essential building block for our constructions is a carefully chosen family of linear projections $\pi_{n,m}: \mathbb{R}^m \rightarrow \mathbb{R}^n$ for all $m, n \in \mathbb{N}$ with $m > n \geq 3$. We define a matrix

$$\Pi_{n,m} = \left(\begin{array}{ccc|c|c} 2 & 1 & 0 & 0 & \mathbf{0} \\ 0 & 0 & 2 & 1 & \mathbf{0} \\ 0 & 1 & 0 & 1 & \mathbf{0} \\ \hline \mathbf{0} & & & 2\mathbf{I}_{n-3} & \mathbf{0} \end{array} \right) \in \mathbb{R}^{n \times m},$$

where \mathbf{I}_d denotes the $d \times d$ identity matrix whose rows are the standard unit vectors \mathbf{e}_i . For reasons that will become clear later, the nonzero columns of $\Pi_{n,m}$ are scaled to sum up to precisely 2. Let $\pi_{n,m}$ be the map defined by $x \mapsto \Pi_{n,m}x$. This map allows us to state our first family of (unbounded) counterexamples, which all are projections of the nonnegative orthant, a simplicial cone:

Lemma 3. *Let $m > n \geq 3$ and $\pi := \pi_{n,m}$. Then $\pi(\mathbb{R}_{\geq 0}^m)$ is a full-dimensional pointed polyhedral cone with $n+1$ facets and $n+1$ extreme rays. Further, $\mathcal{C}(\pi(\mathbb{R}_{\geq 0}^m)) \not\subseteq \pi(\mathcal{C}(\mathbb{R}_{\geq 0}^m))$ where $\mathcal{C}(\pi(\mathbb{R}_{\geq 0}^m)) \cap \pi(\mathcal{C}(\mathbb{R}_{\geq 0}^m))$ is equal to the set of edge directions of $\pi(\mathbb{R}_{\geq 0}^m)$.*

Proof. Let $R_n := \pi(\mathbb{R}_{\geq 0}^m)$. As a projection of a cone, R_n is a cone again, spanned by the first $n+1$ column vectors of the matrix $\Pi_{n,m}$. Further, $R_n \subseteq \mathbb{R}^n$, so it is a pointed cone with vertex $\mathbf{0}$. Since $\Pi_{n,m}$ has full row rank, we have that $\dim(R_n) = n$. We claim that each of these vectors generates an extreme ray of R_n , and that R_n is defined by the following $n+1$ inequalities, all of which are facet-defining:

$$\begin{aligned} x &\geq \mathbf{0} \\ x_1 + x_2 - x_3 &\geq 0 \end{aligned}$$

In order to prove the claim, we proceed by induction on n . The case $n = 3$ is easily verified (see Figure 1). Now let $n \geq 4$. Observe that $\{x \in R_n : x_n = 0\}$ is a face of R_n which is isomorphic to R_{n-1} and, thus, is a facet. The unique column of $\Pi_{n,m}$ not contained in this facet is the vector $2\mathbf{e}_n$, which must therefore generate an extreme ray of R_n . All other inequalities except $x_n \geq 0$ define facets of $\{x \in R_n : x_n = 0\}$ by the induction hypothesis.

Since n of the $n+1$ facets of R_n are defined by nonnegativity constraints, no circuit of R_n can be supported in more than two components. This implies that the vectors in $\mathcal{C}(R_n)$ are multiples of $\mathbf{e}_1 - \mathbf{e}_2$, \mathbf{e}_3 , or of one of the nonzero column vectors of $\Pi_{n,m}$ (which capture all edge directions of R_n). It is easy to see that the set $\pi(\mathcal{C}(\mathbb{R}_{\geq 0}^m))$, in turn, consists of all multiples of column vectors of $\Pi_{n,m}$. The set of edge directions of R_n is therefore given by $\pi(\mathcal{C}(\mathbb{R}_{\geq 0}^m)) \setminus \{\mathbf{0}\} = \mathcal{C}(R_n) \cap \pi(\mathcal{C}(\mathbb{R}_{\geq 0}^m))$. \square

The construction in Lemma 3 readily generalizes to the bounded case if we replace the nonnegative orthant $\mathbb{R}_{\geq 0}^m$ with the standard hypercube in \mathbb{R}^m , which we denote by $C_m = [0, 1]^m$.

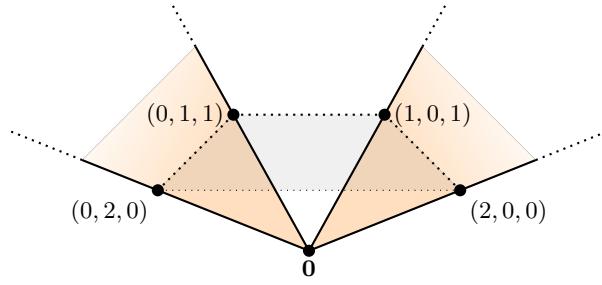


Figure 1: The cone $R_3 = \pi_{3,4}(\mathbb{R}_{\geq 0}^4)$ from Lemma 3, shown here intersected with the hyperplane $x_1 + x_2 + x_3 = 2$. The two highlighted facets are defined by $x_1 \geq 0$ and $x_2 \geq 0$, respectively. Their intersection yields the circuit $e_3 \in \mathcal{C}(R_3)$.

Lemma 4. *Let $m > n \geq 3$ and $\pi := \pi_{n,m}$. Then $\pi(C_m)$ is a full-dimensional polytope and $\mathcal{C}(\pi(C_m)) \not\subseteq \pi(\mathcal{C}(C_m))$. Moreover, $\mathcal{C}(\pi(C_m)) \cap \pi(\mathcal{C}(C_m))$ consists precisely of the edge directions of $\pi(C_m)$.*

Proof. First, note that each facet-defining hyperplane for C_m is a (shifted) coordinate hyperplane, which implies that $\mathcal{C}(C_m) = \mathcal{C}(\mathbb{R}_{\geq 0}^m)$. So to prove the first part of the statement, it suffices to show that $\mathcal{C}(\pi(C_m)) \supseteq \mathcal{C}(\pi(\mathbb{R}_{\geq 0}^m))$ since $\mathcal{C}(\pi(\mathbb{R}_{\geq 0}^m)) \not\subseteq \pi(\mathcal{C}(\mathbb{R}_{\geq 0}^m)) = \pi(\mathcal{C}(C_m))$ by Lemma 3. Indeed, $\mathbf{0} \in \pi(C_m) \subseteq \mathbb{R}_{\geq 0}^n$. Thus, $\mathbf{0}$ is a vertex of $\pi(C_m)$. Observe that $\pi(\mathbb{R}_{\geq 0}^m)$ is the inner cone of $\pi(C_m)$ at $\mathbf{0}$. Hence, all $n+1$ facet-defining inequalities of $\pi(\mathbb{R}_{\geq 0}^m)$ also define facets of $\pi(C_m)$. Since both polyhedra are full-dimensional, it follows from the geometric interpretation of the set of circuits that $\mathcal{C}(\pi(C_m)) \supseteq \mathcal{C}(\pi(\mathbb{R}_{\geq 0}^m))$ as desired.

To prove the second part of the statement, note that $\pi(C_m)$ is the linear projection of a hypercube and, as such, a zonotope (see Figure 2). Equivalently, $\pi(C_m)$ can be written as the Minkowski sum of $n+1$ line segments $[\mathbf{0}, \pi(\mathbf{e}_i)]$ for every $i \in [n+1]$. Since every edge of a zonotope is parallel to one of the line segments from which it is generated, it follows that each nonzero vector in $\mathcal{C}(\pi(C_m))$ is an edge direction of $\pi(C_m)$. \square

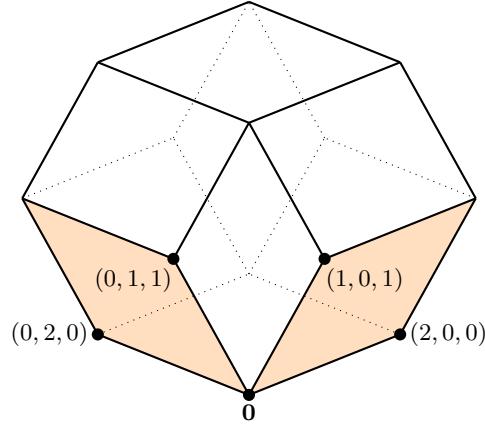


Figure 2: The zonotope $\pi_{3,4}(C_4)$ from Lemma 4. The two facets that yield the circuit direction e_3 are highlighted.

We note that the proof of the second statement of Lemma 4 readily extends to any other linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ by replacing the line segments $[\mathbf{0}, \pi_{n,m}(\mathbf{e}_i)]$ for $i \in [n+1]$ with $[\mathbf{0}, \pi(\mathbf{e}_i)]$

for all $i \in [m]$ with $\pi(\mathbf{e}_i) \neq \mathbf{0}$. We thus obtain a more general statement about the inheritance of circuits in zonotopes: the only circuits that a zonotope inherits from its hypercube extension are the edge directions.

Corollary 4. *Let $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear map. Then $\mathcal{C}(\pi(C_m)) \cap \pi(\mathcal{C}(C_m))$ consists precisely of the edge directions of the zonotope $\pi(C_m)$.*

Let us point out another implication of the previous results. Let $P_1, P_2 \subseteq \mathbb{R}^n$ be polyhedra. Then the Minkowski sum $P_1 + P_2$ is the image of $P_1 \times P_2$ under the map $\sigma: (x, y) \mapsto x + y$. The proof of Lemma 4 implies that, in general, $\mathcal{C}(P_1 + P_2) \not\subseteq \mathcal{C}(P_1) \cup \mathcal{C}(P_2) = \sigma(\mathcal{C}(P_1 \times P_2))$, where the last identity follows from Proposition 1. This means that taking the Minkowski sum of polyhedra may create a circuit that is not a circuit of any of the summands.

Recall that, by Corollary 1, 3 is the minimal dimension of any polyhedron which does not inherit its circuits from an extension. Indeed, the lowest-dimensional counterexamples given in Lemmas 3 and 4 are 3-dimensional. The family of cones given in Lemma 3 is minimal in yet another sense: any n -dimensional *unbounded* pointed polyhedron with n facets (and, thus, n extreme rays) is a simplicial cone. Therefore, all circuit directions in such a polyhedron are edge directions, which are naturally inherited from any extension (see Lemma 1). In the case of *bounded* polyhedra, any counterexample needs one more facet (and one more vertex), i.e., at least $n + 2$ facets and vertices each – otherwise, it is a simplex and has no circuits that are not also edge directions. Even though the zonotopes from Lemma 4 do not satisfy this additional minimality requirement, we can obtain a family of polytopes that are minimal in this sense with a little extra work, using the same projections $\pi_{n,m}$. To this end, let S_m denote the simplex $S_m = \{x \in \mathbb{R}^m: x \geq \mathbf{0}, \sum_{i=1}^m x_i \leq 1\}$ for $m \in \mathbb{N}$.

Lemma 5. *Let $m > n \geq 3$ and $\pi := \pi_{n,m}$. Then $\pi(S_m)$ is a full-dimensional polytope with $n + 2$ facets and $n + 2$ vertices, and $\mathcal{C}(\pi(S_m)) \not\subseteq \pi(\mathcal{C}(S_m))$. Moreover, $\mathcal{C}(\pi(S_m)) \cap \pi(\mathcal{C}(S_m))$ consists precisely of the edge directions of $\pi(S_m)$.*

Proof. Let $R_n := \pi(\mathbb{R}_{\geq 0}^m)$ and $P_n := \pi(S_m)$. First observe that $P_n = \{x \in R_n: \sum_{i=1}^n x_i \leq 2\}$ since the entries of each nonzero column of $\Pi_{n,m}$ sum to 2 (see Figure 1). In particular, this implies that the inequality $\sum_{i=1}^n x_i \leq 2$ is facet-defining for P_n , and that all $n + 1$ nonzero column vectors of $\Pi_{n,m}$, along with the origin $\mathbf{0}$, are vertices of P_n . Further note that $\dim(P_n) = \dim(R_n) = n$ and therefore $\mathcal{C}(P_n) \supseteq \mathcal{C}(R_n)$.

Up to rescaling, $\pi(\mathcal{C}(S_m))$ is the set of all difference vectors between pairs of vertices of P_n . We claim that every such difference vector that belongs to $\mathcal{C}(P_n)$ is the difference of two adjacent vertices and therefore an edge direction of P_n . Indeed, for any $i \geq 4$, the facet $\{x \in P_n: x_i = 0\}$ contains all vertices of P_n but $2\mathbf{e}_i$. Hence, $2\mathbf{e}_i$ must be adjacent to all other vertices. This implies that the only candidate pairs of non-adjacent vertices are contained in the face of P_n defined by $x_i = 0$ for all $i \geq 4$. Since this face is isomorphic to P_3 , there are exactly two pairs of non-adjacent vertices, as can be seen in Figure 1. The corresponding difference vectors are $2\mathbf{e}_1 - \mathbf{e}_2 - \mathbf{e}_3$ and $2\mathbf{e}_2 - \mathbf{e}_1 - \mathbf{e}_3$, respectively. Neither of them is a circuit in $\mathcal{C}(P_n)$ because $\mathbf{e}_1 - \mathbf{e}_2$ has strictly smaller support. This implies that every vector in $\mathcal{C}(P_n) \cap \pi(\mathcal{C}(S_m))$ is an edge direction of P_n .

It remains to show that P_n has a circuit which is not an edge direction. Since $\mathcal{C}(P_n) \supseteq \mathcal{C}(R_n)$, the proof of Lemma 3 implies that $\mathbf{e}_3 \in \mathcal{C}(P_n)$. Observe that P_n has no pair of vertices that only differ in the third coordinate and, hence, $\mathbf{e}_3 \notin \pi(\mathcal{C}(S_m))$. \square

Combining Lemmas 3 and 5, we obtain the following theorem.

Theorem 2. *For all $m, n \in \mathbb{N}$ with $m > n \geq 3$, there exist full-dimensional pointed polyhedra $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ and a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$ such that $\mathcal{C}(P) \not\subseteq \pi(\mathcal{C}(Q))$ and*

$\mathcal{C}(P) \cap \pi(\mathcal{C}(Q))$ consists precisely of the edge directions of P . Moreover, Q can be chosen to be simple, and P can be chosen to be either a polytope with $n+2$ facets and $n+2$ vertices, or a pointed polyhedral cone with $n+1$ facets and $n+1$ extreme rays.

Even though the counterexamples discussed above may seem pathological, they have an interesting property: for $m = n+1$, the simplicial extensions in Lemmas 3 and 5 have the same number of vertices and extreme rays as their projections. Indeed, such a canonical ‘simplex extension’ exists for every pointed polyhedron [27] and will be the starting point for proving Theorem 5 in Section 4.

Next, we will see that one can not only construct counterexamples that fail the inheritance of circuits, but there is a construction principle that guarantees that the difference in the number of circuits that are inherited and those that are not is exponentially large in the dimension. This will yield a constructive analogue to the second part of Theorem 1.

3.3. Constructing Counterexamples with Many Non-Inherited Circuits

Our goal is to prove Theorem 3, which we restate for convenience.

Theorem 3. *For all $n \geq 3$, one can construct full-dimensional pointed polyhedra $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ and a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$ such that $\mathcal{C}(P)$ contains $2^{\Omega(n)}$ unique circuit directions while the number of unique circuit directions in $\mathcal{C}(Q)$ is $O(n^2)$.*

Our proof of Theorem 3 relies on the interplay between the basic solutions of a polytope and the circuits in its homogenization. We first recall the relevant concepts.

A *basic solution* of a linear system $Ax = b, Bx \leq d$ in variables $x \in \mathbb{R}^n$ is a vector $\bar{x} \in \mathbb{R}^n$ with $A\bar{x} = b$ such that the row submatrix of $\begin{pmatrix} A \\ B \end{pmatrix}$ obtained by taking all rows for which \bar{x} satisfies the corresponding constraint at equality has full column rank. Note that \bar{x} need not be feasible. If P is the polyhedron in \mathbb{R}^n defined by the above system, we will also call \bar{x} a basic solution of P . We denote the set of all such basic solutions of P by $\mathcal{B}(P)$. Note that $\mathcal{B}(P)$ contains the set of vertices $\mathcal{V}(P)$ of P . For our purposes, the following equivalent characterization of $\mathcal{B}(P)$ will be convenient.

Lemma 6. *Let $P = \{x \in \mathbb{R}^n: Ax = b, Bx \leq d\}$ be a pointed polyhedron. Then $\mathcal{B}(P)$ is equal to the set of all vectors $g \in \mathbb{R}^n$ such that $Ag = b$ and $Bg - d$ is support-minimal in $\{By - d: y \in \mathbb{R}^n, Ay = b\}$.*

Proof. Let $g \in \mathbb{R}^n$ such that $Ag = b$. Note that the matrix $\begin{pmatrix} A \\ B \end{pmatrix}$ has rank n since P is pointed. Therefore, $Bg - d$ is support-minimal in $\{By - d: y \in \mathbb{R}^n, Ay = b\}$ if and only if $\begin{pmatrix} A \\ B' \end{pmatrix}$ has rank n , where B' is obtained from B by deleting all rows in the support of $Bg - d$. \square

Note the similarity between the above characterization of $\mathcal{B}(P)$ and the definition of the set of circuits $\mathcal{C}(P)$ (cf. Section 1.2). In this sense, one can think of the basic solutions (which include all vertices) as the zero-dimensional analogue of the circuits (which include all edge directions). In fact, we can state this connection more precisely as follows.

Following [49], we define the *homogenization* of a polyhedron $P = \{x \in \mathbb{R}^n: Ax = b, Bx \leq d\}$ as

$$\text{hom}(P) := \{(t, x) \in \mathbb{R} \times \mathbb{R}^n: t \geq 0, Ax - bt = \mathbf{0}, Bx - dt \leq \mathbf{0}\}. \quad (\text{hom})$$

Observe that $P = \{x \in \mathbb{R}^n: (1, x) \in \text{hom}(P)\}$. If P is pointed, then $\text{hom}(P)$ is a pointed polyhedral cone whose extreme rays are generated by all vectors $(0, g)$, where g is the direction of an extreme ray of P , and $(1, v)$ for all vertices $v \in \mathcal{V}(P)$. The circuits of $\text{hom}(P)$ are in correspondence with the basic solutions and circuits of P , as shown next.

Lemma 7. Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a pointed polyhedron. Up to rescaling, the circuits of $\text{hom}(P)$ w.r.t. the system (hom) are the nonzero vectors $(\gamma, g) \in \mathbb{R} \times \mathbb{R}^n$ for which one of the following holds:

- (i) $\gamma = 0$ and $g \in \mathcal{C}(P)$,
- (ii) $\gamma = 1$ and $g \in \mathcal{B}(P)$.

Proof. Let $(\gamma, g) \in \mathbb{R} \times \mathbb{R}^n$ be a nonzero vector with $Ag - b\gamma = \mathbf{0}$. If $\gamma = 0$, then $(\gamma, g) \in \mathcal{C}(\text{hom}(P))$ if and only if Bg is support-minimal in $\{By : y \neq \mathbf{0}, Ay = \mathbf{0}\}$, i.e., if and only if $g \in \mathcal{C}(P)$. If $\gamma \neq 0$, we may assume after rescaling that $\gamma = 1$. Suppose that $(1, g) \in \mathcal{C}(\text{hom}(P))$. Then we must have in particular that $Bg - d$ is support-minimal in $\{By - d : y \in \mathbb{R}^n, Ay = b\}$. Hence, $g \in \mathcal{B}(P)$ by Lemma 6. Conversely, $(1, g) \in \mathcal{C}(\text{hom}(P))$ if $g \in \mathcal{B}(P)$. To see this, suppose for the sake of contradiction that there exists some vector $y \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ such that $Ay = \mathbf{0}$ and the support of By is contained in the support of $Bg - d$. Let B' be the matrix obtained from B by deleting all rows in the support of $Bg - d$. Then $B'y = \mathbf{0}$. Since $\begin{pmatrix} A \\ B' \end{pmatrix}$ has rank n , we must have that $y = \mathbf{0}$, a contradiction. \square

The crucial observation for proving Theorem 3 now is the following: if P is a polytope with vertex set $\mathcal{V}(P)$, then $\text{hom}(P)$ is the image of the nonnegative orthant $\mathbb{R}_{\geq 0}^{\mathcal{V}(P)}$ under the projection $x \mapsto \sum_{v \in \mathcal{V}(P)} x_v(1, v)$. This projection maps the circuits of the nonnegative orthant to the edge directions of $\text{hom}(P)$. In particular, the number of unique circuit directions of $\mathbb{R}_{\geq 0}^{\mathcal{V}(P)}$ equals $|\mathcal{V}(P)|$ while $\mathcal{C}(\text{hom}(P)) \supseteq \{1\} \times \mathcal{B}(P)$ by Lemma 7. Here it is important to stress that both $\text{hom}(P)$ and $\mathcal{B}(P)$ depend on the particular inequality description of P . If we assume a minimal description, then every inequality in (hom) (possibly except $t \geq 0$) defines a facet of $\text{hom}(P)$ and all circuits in Lemma 7(ii) are indeed circuits. To prove Theorem 3, it therefore suffices to exhibit a family of polytopes with polynomially many (in the dimension) vertices but exponentially many basic solutions (w.r.t. a minimal description). The corresponding homogenizations will then have an exponential number of unique circuit directions of which only a polynomial number are inherited from the associated nonnegative orthant extension. We show that for all $n \geq 2$, the standard cross-polytope $Q_n := \{x \in \mathbb{R}^n : x^\top y \leq 1 \ \forall y \in \{-1, 1\}^n\}$, suitably cropped by intersecting it with a hypercube, satisfies all the desired properties. This will complete the proof of Theorem 3.

Lemma 8. Let $n \geq 2, \delta \in (\frac{1}{2}, 1)$ and $Q'_n := Q_n \cap [-\delta, \delta]^n$. Then $|\mathcal{V}(Q'_n)| = 4n(n-1)$ and $\mathcal{B}(Q'_n) \supseteq \{-\delta, \delta\}^n$.

Proof. We first argue that no face of $[-\delta, \delta]^n$ of dimension $n-2$ or less intersects Q_n . Indeed, let F be a face of $[-\delta, \delta]^n$ with $\dim(F) \leq n-2$. By symmetry, we may assume that $F \subseteq \{x \in [-\delta, \delta]^n : x_1 = x_2 = \delta\}$. Then the inequality $x_1 + x_2 \leq 1$, which is valid for Q_n , separates F from Q_n since $\delta > \frac{1}{2}$.

This means that each vertex of Q'_n is contained in at most one facet of $[-\delta, \delta]^n$. In fact, it must be contained in exactly one: none of the vertices of Q_n (which are the positive and negative unit vectors) is contained in $[-\delta, \delta]^n$ since $\delta < 1$. Hence, each vertex of Q'_n is the intersection of exactly one facet of $[-\delta, \delta]^n$ with an edge of Q_n . Observe that every edge of Q_n intersects exactly two distinct facets of $[-\delta, \delta]^n$. Again, this is due to the choice of $\delta < 1$. Hence, the number of vertices of Q'_n equals twice the number of edges of Q_n , i.e., $|\mathcal{V}(Q'_n)| = 4n(n-1)$. Moreover, none of the vertices of $[-\delta, \delta]^n$ is a vertex of Q'_n and for all $i \in [n]$, the inequalities $-\delta \leq x_i \leq \delta$ are facet-defining for Q'_n . Therefore, $\mathcal{B}(Q'_n) \supseteq \mathcal{V}([- \delta, \delta]^n) = \{-\delta, \delta\}^n$. \square

Although the above construction is phrased in terms of unbounded polyhedra, one may readily obtain the same asymptotic gap for the bounded case, as in Theorem 1. One only needs

to replace the nonnegative orthant $\mathbb{R}_{\geq 0}^{\mathcal{V}(P)}$ above with the simplex $S_{|\mathcal{V}(P)|}$. Its projection then is the polytope $\{(t, x) \in \text{hom}(P) : t \leq 1\}$, in which all circuits in $\mathcal{C}(\text{hom}(P))$ still are circuits. Further, the number of unique circuit directions of $S_{|\mathcal{V}(P)|}$ is quadratic in the number of those of $\mathbb{R}_{\geq 0}^{\mathcal{V}(P)}$.

A careful analysis of the counterexamples in Section 3.2 shows that the constructions are, in fact, homogenizations, too: consider again the linear map $\pi_{n,m} : \mathbb{R}^m \rightarrow \mathbb{R}^n$ from Section 3.2. Define a linear transformation of \mathbb{R}^n which maps $x \in \mathbb{R}^n$ to the vector $x' \in \mathbb{R}^n$ defined by $x'_3 = \frac{1}{2} \sum_{i=1}^n x_i$ and $x'_i = x_i$ for all $i \neq 3$. Under this transformation, the cone $R_n = \pi_{n,m}(\mathbb{R}_{\geq 0}^m)$ from Lemma 3 can be viewed as the homogenization of some polytope $P \subseteq \mathbb{R}^{n-1}$ whose vertices are the nonzero column vectors of the matrix $\Pi_{n,m}$ after projecting out the third coordinate. This coordinate takes over the role of the homogeneous coordinate t (recall that all nonzero column vectors of $\Pi_{n,m}$ satisfy $\sum_{i=1}^n x_i = 2$). Then $\mathbf{0}$ is a basic solution of P which is not a vertex. In the homogenization, $\mathbf{0}$ yields the circuit \mathbf{e}_3 by Lemma 7.

3.4. Fixed-Shape Partition Polytopes

In all constructions seen so far, the extensions may seem specifically designed so as to fail the inheritance of circuits. We conclude this section with an example from combinatorial optimization that exhibits this undesirable behavior despite a number of favorable properties.

Recall the clustering application and the definitions of the associated fixed-shape partition polytopes $P(X, k, \kappa)$ and their extensions $T(n, k, \kappa)$ from Section 1.1. There are several reasons why one might hope that the set of circuits of $P(X, k, \kappa)$ is inherited from $T(n, k, \kappa)$. Most importantly, for any choice of n, k, κ and X , edges of both $P(X, k, \kappa)$ and $T(n, k, \kappa)$ have a near-identical characterization in terms of the underlying application in the form of ‘cyclic exchanges’ of items between clusters [8, 9, 17, 30, 34, 39]: a subset of clusters are ordered in a cycle, and one item from each cluster is transferred to the next along the cycle. Further, $T(n, k, \kappa)$ exhibits a number of interesting properties: it is a 0/1 polytope whose constraint matrix is totally unimodular and in which all circuits appear as edge directions; in fact, any circuit walk in $T(n, k, \kappa)$ is an edge walk [18]. The special case $k = n$ and $\kappa_i = 1$ for all $i \in [n]$ yields the n th Birkhoff polytope. Finally, the projection π_X is highly symmetric, using the same $x^{(j)}$ in combination with y_{ij} for all $i \in [k]$.

Despite the combination of these many favorable properties, somewhat surprisingly, new circuits may appear in the projection onto $P(X, k, \kappa)$, even for small d, n , and k .

Lemma 9. *For all $n \geq 5$, there exist $k \in \mathbb{N}$, $\kappa = (\kappa_1, \dots, \kappa_k) \in \mathbb{N}^k$, and $X \subseteq \mathbb{R}^{n-2}$ with $|X| = \sum_{i=1}^k \kappa_i = n$ such that $\mathcal{C}(P(X, k, \kappa)) \not\subseteq \pi_X(\mathcal{C}(T(n, k, \kappa)))$, where $\pi_X : \mathbb{R}^{k \times n} \rightarrow (\mathbb{R}^d)^k$ is defined as above.*

Proof. For $n \geq 5$, let $X := \mathcal{V}(P_{n-2}) \subseteq \mathbb{R}^{n-2}$ be the set of vertices of the polytope $P_{n-2} = \pi_{n-2, n-1}(S_{n-1})$ from Lemma 5, where we may assume without loss of generality that $x^{(n)} = \mathbf{0}$ and the remaining $n-1$ vertices are labelled in arbitrary order. (For given X , the set of all possible clustering vectors is invariant under reordering the data points $x^{(j)} \in X$.)

Consider the fixed-shape partition polytope $P(X) := P(X, k, \kappa)$ for $k = 2$ and cluster sizes $\kappa_1 = 1$ and $\kappa_2 = n-1$. $P(X)$ is the convex hull of all vectors $(c^{(1)}, c^{(2)}) \in \mathbb{R}^{n-2} \times \mathbb{R}^{n-2}$ such that $c^{(1)} = \sum_{j=1}^{n-1} y_{1j} \cdot x^{(j)}$ and $c^{(2)} = \sum_{j=1}^{n-1} x^{(j)} - c^{(1)}$ for some vector $y \in \mathbb{R}^{2 \times n}$ in the

corresponding transportation polytope $T := T(n, 2, (1, n-1))$, which is described by

$$\begin{aligned} \sum_{j=1}^n y_{1j} &= 1 \\ y_{2j} &= 1 - y_{1j} \quad \forall j \in [n] \\ y &\geq \mathbf{0} \end{aligned}$$

Eliminating the variables y_{2j} , it is easy to see that T and the simplex S_{n-1} are affinely isomorphic. Moreover, $P(X)$ is affinely isomorphic to its projection onto the first half $c^{(1)}$ of the clustering vector, which equals $\pi_{n-2, n-1}(S_{n-1}) = P_{n-2}$. The statement then follows immediately from Lemmas 2 and 5. \square

While our proof of Lemma 9 uses the explicit constructions of Section 3.2, it is not difficult to see that one may replace the particular data set X with the set of vertices of any sufficiently generic polytope that satisfies the conditions of Theorem 1. Indeed, as noted in the proof of Lemma 9, for the particular choice of parameters k and κ above, the transportation polytope degenerates to a simplex with $|X|$ vertices. Via the projection used in the proof of Theorem 1, it follows that fixed-shape partition polytopes fail the inheritance of circuits for sufficiently generic data sets X .

Despite these negative results, we stress that there do exist classes of fixed-shape partition polytopes in which all circuits are inherited from the transportation-style extensions, even though this happens for one of the two trivial reasons stated in Lemmas 1 and 2. For instance, $T(n, k, \kappa)$ is a fixed-shape partition polytope itself for any n and k , using the standard unit vectors in \mathbb{R}^n as item locations, as already observed in [18]. Similarly, suppose that we augment a given data set $X \subseteq \mathbb{R}^d$ of size $|X| = n$ with the unit vectors in \mathbb{R}^n , i.e., we replace each item location $x^{(i)}$ with $(x^{(i)}, \mathbf{e}_i) \in \mathbb{R}^d \times \mathbb{R}^n$ for all $i \in [n]$. Then the fixed-shape partition polytope resulting from this augmented embedding can equivalently be derived using the construction in Corollary 2. In particular, the resulting polytope is affinely isomorphic to $T(n, k, \kappa)$.

4. The Role of Projection Maps and Polyhedra for the Inheritance of Circuits

In the previous section, we saw that there exist polyhedra that do not inherit all their circuit directions from an extension. In this section, we explore the role that the individual ‘ingredients’ of those counterexamples – the original polyhedron, the extension polyhedron, and the projection map between them – play for the inheritance of circuits. Our discussion is driven by the following three questions, first stated in Section 1.3:

- (Q1) *Which linear maps π have the property that, for every polyhedron Q , all circuits of $\pi(Q)$ are inherited from Q ?*
- (Q2) *Which polyhedra P inherit their circuits from every extension?*
- (Q3) *For which polyhedra Q does every polyhedron P that is a linear projection of Q inherit its circuits from Q ?*

We provide a complete characterization of the maps for question (Q1) in Section 4.1 and of the polyhedra for question (Q2) in Section 4.2. As we will see, they correspond to restrictive properties that make the inheritance of circuits trivial. We further provide a partial answer to question (Q3) in Section 4.3. In Section 4.4, we explain why our characterizations are best possible. We do so through the discussion of some combinations of polyhedra and maps that lead to an inheritance of circuits, but where neither of them exhibits the aforementioned properties.

4.1. Inheritance Based on the Projection Map

In Lemma 2, we saw that linear isomorphisms essentially preserve the set of circuits. We first show that no other type of linear map guarantees inheritance for *all* polyhedra, thus resolving (Q1).

Theorem 4. *Let $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be a linear map such that $\dim(\pi(\mathbb{R}^m)) \geq 3$. Then $\mathcal{C}(\pi(Q)) \subseteq \pi(\mathcal{C}(Q))$ for all polyhedra $Q \subseteq \mathbb{R}^m$ if and only if π is injective. In particular, if π is not injective, there exists a full-dimensional simple polytope $Q \subseteq \mathbb{R}^m$ such that $\mathcal{C}(\pi(Q)) \not\subseteq \pi(\mathcal{C}(Q))$ and $\mathcal{C}(\pi(Q) \cap \pi(\mathcal{C}(Q)))$ consists precisely of the edge directions of $\pi(Q)$.*

Recall from Theorem 2 that in every dimension greater than 2, there are polyhedra that do not inherit their circuits from all extensions. The key observation for proving Theorem 4 will be that, in any fixed dimension, the particular projection used to obtain Theorem 2 can be exchanged for any other one after a suitable linear transformation of the domain space.

Proof of Theorem 4. By Lemma 2, it suffices to show the ‘if’ part of the statement and we may assume that $\pi(\mathbb{R}^m) = \mathbb{R}^n$. If π is not injective, then $m > n$. By Theorem 2, there exists a full-dimensional simple polytope $Q \subseteq \mathbb{R}^m$ and a linear map $\sigma: \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that the polytope $P := \sigma(Q) \subseteq \mathbb{R}^n$ is full-dimensional, $\mathcal{C}(P) \not\subseteq \sigma(\mathcal{C}(Q))$, and the set of edge directions of P is precisely the set $\mathcal{C}(P) \cap \sigma(\mathcal{C}(Q))$. Since $\dim(P) = n$, we have that $\sigma(\mathbb{R}^m) = \mathbb{R}^n = \pi(\mathbb{R}^m)$. Hence, there exists a linear transformation $\tau: \mathbb{R}^m \rightarrow \mathbb{R}^m$ such that $\pi = \sigma \circ \tau$. Now consider the polytope $\tilde{Q} := \tau^{-1}(Q)$. Clearly, \tilde{Q} is simple again with $\dim(\tilde{Q}) = m$ and

$$\pi(\tilde{Q}) = (\sigma \circ \tau \circ \tau^{-1})(Q) = \sigma(Q) = P.$$

Using Lemma 2, we conclude that

$$\pi(\mathcal{C}(\tilde{Q})) = (\sigma \circ \tau)(\mathcal{C}(\tilde{Q})) = \sigma(\mathcal{C}(Q)) \not\supseteq \mathcal{C}(P)$$

and, thus, $\mathcal{C}(P) \cap \pi(\mathcal{C}(\tilde{Q})) = \mathcal{C}(P) \cap \sigma(\mathcal{C}(Q))$. □

4.2. Inheritance for all Extensions

Next, we resolve (Q2) by showing that any polyhedron which inherits its circuits from *every* extension cannot have a circuit that is not an edge direction already.

Theorem 5. *Let $P \subseteq \mathbb{R}^n$ be a pointed polyhedron. All circuits in $\mathcal{C}(P)$ are edge directions of P if and only if $\mathcal{C}(P) \subseteq \pi(\mathcal{C}(Q))$ for all polyhedra $Q \subseteq \mathbb{R}^m$ and all linear maps $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$.*

In fact, we will prove a slightly stronger statement that clearly implies Theorem 5:

Theorem 7. *Let $P \subseteq \mathbb{R}^n$ be a pointed polyhedron and $g \in \mathbb{R}^n \setminus \{\mathbf{0}\}$. g is an edge direction of P if and only if $g \in \pi(\mathcal{C}(Q))$ for all polyhedra $Q \subseteq \mathbb{R}^m$ and all linear maps $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$.*

For any pointed polyhedron P and a nonzero vector g which is not among the edge directions of P , we construct an extension of P none of whose circuits projects to g . The construction is based on a classical result of Balas [3, 4] on the union of polyhedra, which we recall next. We denote the union over a family \mathcal{P} of sets as $\bigcup \mathcal{P}$.

Proposition 2 ([3, 4]). *Let $P \subseteq \mathbb{R}^n$ be a polyhedron, and let $\mathcal{P} = \{P_1, \dots, P_p\}$ be a family of nonempty polyhedra $P_i = \{x \in \mathbb{R}^n : A^{(i)}x = b^{(i)}, B^{(i)}x \leq d^{(i)}\}$, $i \in [p]$, such that $P = \text{conv}(\bigcup \mathcal{P})$. Consider the polyhedron $Q_{\mathcal{P}} \subseteq \mathbb{R}^p \times (\mathbb{R}^n)^p$ defined by the following linear system in variables $\lambda \in \mathbb{R}^p$ and $x^{(i)} \in \mathbb{R}^n$ for all $i \in [p]$:*

$$\begin{aligned} \lambda &\geq \mathbf{0} \\ \sum_{i=1}^p \lambda_i &= 1 && \text{(disj)} \\ A^{(i)}x^{(i)} &= b^{(i)}\lambda_i && \forall i \in [p] \\ B^{(i)}x^{(i)} &\leq d^{(i)}\lambda_i && \forall i \in [p] \end{aligned}$$

Then $P = \{\sum_{i=1}^p x^{(i)} : (\lambda, x^{(1)}, \dots, x^{(p)}) \in Q_{\mathcal{P}}\}$.

Next, we give a characterization of the circuits of the extension $Q_{\mathcal{P}}$ defined in Proposition 2.

Lemma 10. *Let $P \subseteq \mathbb{R}^n$ be a pointed polyhedron and let \mathcal{P} and $Q_{\mathcal{P}}$ be defined as in Proposition 2. Up to rescaling, the circuits of $Q_{\mathcal{P}}$ w.r.t. the system (disj) are the nonzero vectors $(f, g^{(1)}, \dots, g^{(p)}) \in \mathbb{R}^p \times (\mathbb{R}^n)^p$ for which one of the following holds:*

- (i) $f = \mathbf{0}$; $g^{(i)} \in \mathcal{C}(P_i)$ for some $i \in [p]$ and $g^{(k)} = \mathbf{0}$ for all $k \neq i$,
- (ii) $f = \mathbf{e}_i - \mathbf{e}_j$ for some $i, j \in [p], i \neq j$; $g^{(i)} \in \mathcal{B}(P_i), g^{(j)} \in \mathcal{B}(P_j)$, and $g^{(k)} = \mathbf{0}$ for all $k \neq i, j$.

Proof. Let $(f, g^{(1)}, \dots, g^{(p)}) \in \mathbb{R}^p \times (\mathbb{R}^n)^p$ be a circuit of $Q_{\mathcal{P}}$. If $f = \mathbf{0}$, then $(g^{(1)}, \dots, g^{(p)}) \in \mathcal{C}(P_1 \times \dots \times P_p)$. Statement (i) immediately follows from an inductive application of Proposition 1.

Now suppose that $f \neq \mathbf{0}$. Since $\sum_{i=1}^p f_i = 0$, f must be supported in at least two components, say, $f_1 \neq 0$ and $f_2 \neq 0$. We claim that these are the only nonzero components of f , and that $g^{(3)} = \dots = g^{(p)} = \mathbf{0}$. Then, after rescaling, we may assume that $f_1 = -f_2 = 1$, and statement (ii) follows from Lemma 6. In order to prove the claim, observe that the vector $(\mathbf{e}_1 - \mathbf{e}_2, \frac{1}{f_1}g^{(1)}, -\frac{1}{f_2}g^{(2)}, \mathbf{0}, \dots, \mathbf{0}) \in \mathbb{R}^p \times (\mathbb{R}^n)^p$ belongs to $\mathcal{C}(Q_{\mathcal{P}})$, too. Then f and $\mathbf{e}_1 - \mathbf{e}_2$ must have the same support. Further, by support-minimality of $B^{(k)}g^{(k)} - d^{(k)}f_k$, we cannot have that $g^{(k)} \neq \mathbf{0}$ for some $k \geq 3$ as all P_k are pointed. \square

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

Proof of Theorem 7. The ‘only if’ part immediately follows from Lemma 1. For the converse implication, suppose that g is not an edge direction of P . We first show the statement for the case that P is a polytope.

Let $\mathcal{U} := \{\{u, v\} : u, v \in \mathcal{V}(P), g \in \mathbb{R}(u - v)\}$ be the set of unordered pairs $\{u, v\}$ of vertices of P whose difference is in direction of g (possibly $\mathcal{U} = \emptyset$). Observe that the pairs in \mathcal{U} are pairwise disjoint: if $\{u, v\}, \{v, w\} \in \mathcal{U}$ then u, v, w are collinear. Since all three of them are vertices, it follows that $u = w$. For every pair $\{u, v\} \in \mathcal{U}$, let $F_{\{u, v\}}$ be the minimal face of P containing both u and v . Since $u - v$ is not an edge direction of P , we have $\dim(F_{\{u, v\}}) \geq 2$. Hence, there exists a vector $z \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ orthogonal to $u - v$ such that $\frac{u+v}{2} \pm \varepsilon z \in F_{\{u, v\}}$ for some small $\varepsilon > 0$. This means that the parallelogram

$$P_{\{u, v\}} := \text{conv} \left\{ u, v, \frac{u+v}{2} + \varepsilon z, \frac{u+v}{2} - \varepsilon z \right\}$$

is contained in $F_{\{u, v\}}$.

Moreover, for all $u, v \in \mathcal{V}(P)$ with $u \neq v$ and $\{u, v\} \notin \mathcal{U}$, there exist $a \in \mathbb{R}^n \setminus \{\mathbf{0}\}$, $\beta \in \mathbb{R}$ such that $a^\top g = 0$ and $a^\top u < \beta < a^\top v$, i.e., the hyperplane $\{x \in \mathbb{R}^n : a^\top x = \beta\}$ strictly separates u and v . If ε is sufficiently small, then also $P_{\{u,v\}}$ and $P_{\{u',v'\}}$, and $P_{\{u,v\}}$ and w can be strictly separated by a hyperplane whose normal vector is orthogonal to g , for any distinct $\{u, v\}, \{u', v'\} \in \mathcal{U}$ and $w \in \mathcal{V}(P) \setminus \bigcup \mathcal{U}$.

Now define

$$\mathcal{P} := \left\{ P_{\{u,v\}} : \{u, v\} \in \mathcal{U} \right\} \cup \left\{ \{v\} : v \in \mathcal{V}(P) \setminus \bigcup \mathcal{U} \right\}.$$

Since \mathcal{P} is a family of polytopes (singletons and parallelograms) contained in P which covers $\mathcal{V}(P)$, we have that $P = \text{conv}(\bigcup \mathcal{P})$. For this family \mathcal{P} , we consider the extension $Q_{\mathcal{P}}$ as defined in Proposition 2. We claim that under the projection given in Proposition 2, none of the circuits of $Q_{\mathcal{P}}$ maps to a multiple of g . By Lemma 10, the circuits of $Q_{\mathcal{P}}$ either map to (i) edge directions of some member of the family \mathcal{P} or to (ii) (scaled) differences of two vertices of different members of \mathcal{P} . This is because $\dim(Q) \leq 2$ for all $Q \in \mathcal{P}$ and all basic solutions of parallelograms are vertices. By construction, none of the parallelograms $P_{\{u,v\}}$ has an edge in direction g , ruling out case (i). For case (ii), recall that any two distinct members of \mathcal{P} can be strictly separated by some hyperplane whose normal vector is orthogonal to g . We conclude that g is not inherited from $Q_{\mathcal{P}}$.

Now suppose that P is unbounded. Then $P = P' + \text{rec}(P)$ where $P' := \text{conv}(\mathcal{V}(P))$ and $\text{rec}(P)$ denotes the recession cone of P . We first define an extension for each Minkowski summand individually and then combine the two. Indeed, since the first summand P' is a polytope, the first part of the proof yields an extension $Q_{P'}$ of P' none of whose circuits is sent to g . Suppose that the second summand $\text{rec}(P)$ is generated by q extreme rays in directions $\{r^{(1)}, \dots, r^{(q)}\} \subseteq \mathbb{R}^n$. Then it is the image of the nonnegative orthant $\mathbb{R}_{\geq 0}^q$ under the map $\mathbb{R}^q \ni y \mapsto \sum_{i=1}^q y_i r^{(i)} \in \mathbb{R}^n$. By assumption, g is not a multiple of $r^{(i)}$ for any $i \in [q]$. Now consider the polyhedron $Q_{P'} \times \mathbb{R}_{\geq 0}^q$. It is an extension of P , where the corresponding projection first maps $Q_{P'} \times \mathbb{R}_{\geq 0}^q$ to $P' \times \text{rec}(P)$ and then applies the map $(x, y) \mapsto x + y$. By Proposition 1, g is not inherited from $Q_{P'} \times \mathbb{R}_{\geq 0}^q$ under this combined map. This concludes the proof. \square

We stress that the extension $Q_{\mathcal{P}}$ constructed in the above proof is not necessarily given by an irredundant system if we follow Proposition 2. However, for the purpose of proving a negative result about the *non*-inheritance of a particular direction, this is not a restriction.

Before we focus on the last of the three ingredients, the extension polyhedron Q , let us remark that the simplex extension that we saw in Lemma 5 and the proof of Theorem 1 is, in fact, a special case of the more general extension $Q_{\mathcal{P}}$ used in the proof of Theorem 5: for a polytope P and the decomposition $\mathcal{P} := \{\{v\} : v \in \mathcal{V}(P)\}$, the polyhedron $Q_{\mathcal{P}}$ is affinely isomorphic to the simplex $S_{|\mathcal{V}(P)|-1}$. Proposition 2 also generalizes another result from Section 3: let P be a polytope and consider $\mathcal{P} := \{\{\mathbf{0}\}, \{1\} \times P\}$; then $Q_{\mathcal{P}}$ as defined in Proposition 2 equals $\{(t, x) \in \text{hom}(P) : t \leq 1\}$.

4.3. (No) Inheritance Based on the Extension Polyhedron

In Section 3.2, we saw that there exist polyhedra – simplices, simplicial cones, and hypercubes in dimension 4 and greater – that can be projected in such a way that not all circuits of the image polyhedron are inherited from the original one. In this section, we prove that these polyhedra can essentially be exchanged for any other polyhedron Q (of the same dimension), provided that Q has a non-degenerate vertex.

Theorem 6. *Let $Q \subseteq \mathbb{R}^m$ be a polyhedron with $\dim(Q) \geq 4$. If Q has a non-degenerate vertex, then there exists a linear map $\pi : \mathbb{R}^m \rightarrow \mathbb{R}^{\dim(Q)-1}$ such that $\pi(Q)$ is full-dimensional and $\mathcal{C}(\pi(Q)) \not\subseteq \pi(\mathcal{C}(Q))$.*

Before we give a detailed proof of this result, let us take a closer look at the proofs of Lemmas 3 to 5 and identify a common theme: all polyhedra that we projected from in Section 3.2 have a non-degenerate vertex at the origin $\mathbf{0}$, and their inner cone at $\mathbf{0}$ equals the nonnegative orthant. So in any fixed dimension, they are all identical *locally* at $\mathbf{0}$. We then applied a carefully chosen linear projection map which preserves this local resemblance. This allowed us to always generate a particular circuit direction \mathbf{e}_3 , for which we were then able to establish non-inheritance. In this last step, however, knowledge of the set of circuits of the extension polyhedron was crucial. This will be the major technical challenge when applying the above proof strategy to an arbitrary polyhedron Q : neither do we know the other facets of Q that are not incident with $\mathbf{0}$ nor is $\mathcal{C}(Q)$ given explicitly. We address this challenge by defining an infinite family of linear projections such that every member of the family maps Q to a polyhedron with vertex $\mathbf{0}$ in which the non-inherited circuit direction \mathbf{e}_3 from the results in Section 3.2 still appears as a circuit. Moreover, the family will have the property that no nonzero vector is sent to \mathbf{e}_3 (or a multiple thereof) under more than one of the projections in the family. Since our family is *infinite* but $\mathcal{C}(Q)$ is *finitely* generated, there must be some member of the family which does not send any of the circuits of Q to \mathbf{e}_3 . This will be the map that we can apply to Q and obtain the same non-inheritance result as in Section 3.2. The remainder of this section is dedicated to the proof details.

Proof of Theorem 6. After an affine transformation, we may assume that Q is full-dimensional, $\mathbf{0}$ is a non-degenerate vertex of Q , and the inner cone of Q at $\mathbf{0}$ equals $\mathbb{R}_{\geq 0}^m$. For all $\alpha \in \mathbb{N} \setminus \{1\}$, we define the matrix

$$\Pi_\alpha := \left(\begin{array}{cccc|c} \alpha & 1 & 0 & 0 & \mathbf{0} \\ 0 & 0 & \alpha & 1 & \\ 0 & \alpha - 1 & 0 & \alpha - 1 & \\ \hline \mathbf{0} & & & & \alpha \mathbf{I}_{n-3} \end{array} \right) \in \mathbb{R}^{(m-1) \times m}$$

and a corresponding linear map $\pi_\alpha: \mathbb{R}^m \rightarrow \mathbb{R}^{m-1}, x \mapsto \Pi_\alpha x$. Note that $\pi_2 = \pi_{m-1,m}$ where $\pi_{m-1,m}$ is the projection used in Section 3.

Consider the cone $\pi_\alpha(\mathbb{R}_{\geq 0}^m) \subseteq \mathbb{R}^{m-1}$. It is defined by the m inequalities

$$\begin{aligned} x &\geq \mathbf{0} \\ (\alpha - 1)x_1 + (\alpha - 1)x_2 - x_3 &\geq 0 \end{aligned}$$

all of which are facet-defining. This can be seen using the same arguments as in the proof of Lemma 3. In fact, the cone above is obtained from $\pi_2(\mathbb{R}_{\geq 0}^m)$ by rescaling it along the third coordinate. In particular, $\mathbf{e}_3 \in \mathcal{C}(\pi_\alpha(\mathbb{R}_{\geq 0}^m))$. Now consider $\pi_\alpha^{-1}(\mathbb{R}\mathbf{e}_3) =: K_\alpha$, i.e., K_α is the set of all vectors in \mathbb{R}^m that π_α sends to a multiple of $\mathbf{e}_3 \in \mathbb{R}^{m-1}$. Since Π_α has full row rank, K_α is a two-dimensional linear subspace of \mathbb{R}^m , spanned by the vectors $\alpha\mathbf{e}_2 - \mathbf{e}_1$ and $\alpha\mathbf{e}_4 - \mathbf{e}_3$. For any $\alpha \neq \beta$, we have that $K_\alpha \cap K_\beta = \{\mathbf{0}\}$ because the four basis vectors of K_α and K_β are linearly independent. Now recall from the definition of the set of circuits that $\mathcal{C}(Q)$ consists of a *finite* number of one-dimensional linear subspaces of \mathbb{R}^m . Hence, there must exist some $\alpha \neq 1$ such that $\mathcal{C}(Q) \cap K_\alpha = \emptyset$. For this choice of α , we conclude that $\mathbf{e}_3 \notin \pi_\alpha(\mathcal{C}(Q))$ while $\mathbf{e}_3 \in \mathcal{C}(\pi_\alpha(\mathbb{R}_{\geq 0}^m)) \subseteq \mathcal{C}(\pi_\alpha(Q))$, where the last inclusion follows from the fact that $\pi_\alpha(\mathbb{R}_{\geq 0}^m)$ is the inner cone of $\pi_\alpha(Q)$ at the vertex $\mathbf{0}$. \square

4.4. Inheritance in Nontrivial Instances

Theorems 4 to 6 imply that, beyond the trivial cases that we saw in Section 2, inheritance of circuits cannot be a property of a polyhedron, of a specific extension polyhedron, or of the map between the two by itself. We conclude our discussion by showing that there do exist instances for which a combination of these three ingredients leads to the desired inheritance of circuits while each individual ingredient does not satisfy the restrictive assumptions of Theorems 4 and 5. In this sense, Theorems 4 and 5 are the best possible statements.

Lemma 11. *For all $m, n \in \mathbb{N}$ with $n \geq 3$ and $m \geq n + 3$, there exist full-dimensional polytopes $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ and a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ with $\pi(Q) = P$ such that $\mathcal{C}(P) \subseteq \pi(\mathcal{C}(Q))$, P and Q are not linearly isomorphic, not all circuits of P are edge directions, and Q has a non-degenerate vertex.*

Proof. We again modify the construction from Lemma 5. For $n \geq 3$ and $m \geq n + 3$, we define the matrix

$$\Pi'_{n,m} := \left(\begin{array}{cc|c|c} 1 & 1 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 2 \\ 0 & 1 & 1 & 0 & 1 & 1 \\ \hline 0 & & & & & \\ \end{array} \right) \in \mathbb{R}^{n \times m}$$

and let $\pi': \mathbb{R}^m \rightarrow \mathbb{R}^n$ be the linear map $x \mapsto \Pi'_{n,m}x$. Consider the polytope $P'_n := \pi'(S_m)$ (see Figure 3). We claim that P'_n is defined by the following irredundant system of inequalities:

$$\begin{aligned} x &\geq \mathbf{0} \\ x_3 &\leq 1 \\ x_1 + x_2 - x_3 &\geq 0 \\ x_1 + x_2 - x_3 &\leq 1 \end{aligned}$$

All inequalities are valid for P'_n . To see that they are facet-defining, we proceed by induction on n , similar to the proof of Lemma 5. For the case $n = 3$, we refer to Figure 3. If $n \geq 4$, then $\{x \in P'_n: x_n = 0\}$ is a face of P'_n which contains all column vectors of $\Pi'_{n,m}$ but \mathbf{e}_n . Hence, it is isomorphic to P'_{n-1} and therefore is a facet of P'_n .

In particular, since P'_n has more than $n + 1$ facets, it is not a simplex and, thus, cannot be isomorphic to S_m . Consider again the cone $R_n \subseteq \mathbb{R}^n$ defined in the proof of Lemma 3. It is easy to see that $\mathcal{C}(P'_n) = \mathcal{C}(R_n)$. Hence, after rescaling, every circuit direction of P'_n appears as one of the column vectors of $\Pi'_{n,m}$ or as the difference of two of them. This implies that $\mathcal{C}(P'_n) \subseteq \pi'(\mathcal{C}(S_m))$. Further, $\mathbf{e}_3 \in \mathcal{C}(P'_n)$ is not an edge direction of P'_n , and S_m clearly has a non-degenerate vertex. \square

5. Final Remarks

We showed that the connection between the sets of circuits of polyhedra and their extensions is much weaker than the connection between their edge directions: in general, circuits are not inherited under affine projections. Whenever this does happen for a nontrivial combination of two polyhedra and a projection map between them, it is due to the specific combination of the three ingredients and not due to any single one of them by itself. Therefore, a natural direction of future work would be to identify properties of these combinations that are sufficient for the inheritance of circuits (beyond our characterizations in Section 4).

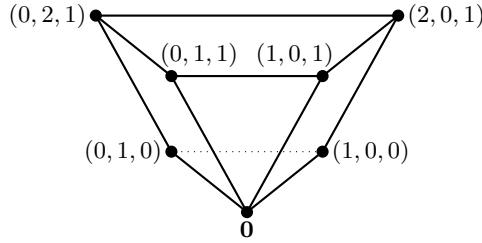


Figure 3: The polytope P'_3 from the proof of Lemma 11.

References

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows*. Prentice Hall, Englewood Cliffs, NJ, 1993.
- [2] M. Aprile and S. Fiorini. Regular matroids have polynomial extension complexity. *Mathematics of Operations Research*, 47(1):540–559, 2022.
- [3] E. Balas. Disjunctive programming and a hierarchy of relaxations for discrete optimization problems. *SIAM Journal on Algebraic Discrete Methods*, 6(3):466–486, 1985.
- [4] E. Balas. Disjunctive programming: Properties of the convex hull of feasible points. *Discrete Applied Mathematics*, 89(1):3–44, 1998.
- [5] E. R. Barnes, A. J. Hoffman, and U. G. Rothblum. Optimal partitions having disjoint convex and conic hulls. *Mathematical Programming*, 54(1):69–86, 1992.
- [6] A. Black, J. A. De Loera, S. Kafer, and L. Sanità. On the simplex method for 0/1 polytopes. *eprint arXiv:2111.14050*, 2021.
- [7] K. H. Borgwardt. *The Simplex Method: A Probabilistic Analysis*, volume 1 of *Algorithms and Combinatorics*. Springer Berlin, Heidelberg, 1987.
- [8] S. Borgwardt. *A Combinatorial Optimization Approach to Constrained Clustering*. 2010. PhD Thesis.
- [9] S. Borgwardt. On the diameter of partition polytopes and vertex-disjoint cycle cover. *Mathematical Programming, Series A*, 141(1):1–20, 2013.
- [10] S. Borgwardt, J. A. De Loera, E. Finhold, and J. Miller. The hierarchy of circuit diameters and transportation polytopes. *Discrete Applied Mathematics*, 240:8–24, 2018.
- [11] S. Borgwardt, E. Finhold, and R. Hemmecke. On the circuit diameter of dual transportation polyhedra. *SIAM Journal on Discrete Mathematics*, 29(1):113–121, 2016.
- [12] S. Borgwardt, E. Finhold, and R. Hemmecke. Quadratic diameter bounds for dual network flow polyhedra. *Mathematical Programming*, 159(1):237–251, 2016.
- [13] S. Borgwardt and F. Happach. Good Clusterings Have Large Volume. *Operations Research*, 67(1):215–231, 2019.
- [14] S. Borgwardt, F. Happach, and S. Zirkelbach. An Algorithm for the Separation-Preserving Transition of Clusterings. *INFORMS Journal on Optimization*, 5(1):1–26, 2023.

- [15] S. Borgwardt and S. Onn. Efficient solutions for weight-balanced partitioning problems. *Discrete Optimization*, 21(C):71–84, 2016.
- [16] S. Borgwardt, T. Stephen, and T. Yusun. On the circuit diameter conjecture. *Discrete and Computational Geometry*, 60(3):558–587, 2018.
- [17] S. Borgwardt and C. Viss. Constructing Clustering Transformations. *SIAM Journal on Discrete Mathematics*, 35(1):152–178, 2021.
- [18] S. Borgwardt and C. Viss. Circuit Walks in Integral Polyhedra. *Discrete Optimization*, 44(1):100566, 2022.
- [19] S. Borgwardt and C. Viss. A polyhedral model for enumeration and optimization over the set of circuits. *Discrete Applied Mathematics*, 308:68–83, 2022.
- [20] A. Brieden and P. Gritzmann. On optimal weighted balanced clusterings: Gravity bodies and power diagrams. *SIAM Journal on Discrete Mathematics*, 26:415–434, 2012.
- [21] M. Conforti, G. Cornu  jols, and G. Zambelli. Extended formulations in combinatorial optimization. *Annals of Operations Research*, 204(1):97–143, 2013.
- [22] M. Conforti, G. Cornu  jols, and G. Zambelli. *Integer Programming*, volume 271 of *Graduate Texts in Mathematics*. Springer, Cham, 2014.
- [23] D. Dadush and S. Huijberts. A friendly smoothed analysis of the simplex method. *SIAM Journal on Computing*, 49(5):STOC18–449, 2020.
- [24] D. Dadush, Z. K. Koh, B. Natura, and L. A. V  gh. On circuit diameter bounds via circuit imbalances. *eprint arXiv:2111.07913*, 2021.
- [25] J. A. De Loera, R. Hemmecke, and J. Lee. On augmentation algorithms for linear and integer-linear programming: from Edmonds-Karp to Bland and beyond. *SIAM Journal on Optimization*, 25(4):2494–2511, 2015.
- [26] J. A. De Loera, S. Kafer, and L. Sanit  . Pivot rules for circuit-augmentation algorithms in linear optimization. *SIAM Journal on Optimization*, 32(3):2156–2179, 2022.
- [27] S. Fiorini, V. Kaibel, K. Pashkovich, and D. O. Theis. Combinatorial bounds on nonnegative rank and extended formulations. *Discrete Mathematics*, 313(1):67–83, 2013.
- [28] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf. Exponential lower bounds for polytopes in combinatorial optimization. *Journal of the ACM*, 62(2):1–23, 2015.
- [29] S. Fiorini, T. Rothvo  , and H. R. Tiwary. Extended formulations for polygons. *Discrete and Computational Geometry*, 48(3):658–668, 2012.
- [30] K. Fukuda, S. Onn, and V. Rosta. An adaptive algorithm for vector partitioning. *Journal of Global Optimization*, 25:305–319, 2003.
- [31] M. X. Goemans. Smallest compact formulation for the permutohedron. *Mathematical Programming*, 153(1):5–11, 2015.
- [32] J. E. Graver. On the foundation of linear and integer programming I. *Mathematical Programming*, 9:207–226, 1975.

- [33] F. K. Hwang, S. Onn, and U. G. Rothblum. Representations and characterizations of vertices of bounded-shape partition polytopes. *Linear Algebra and its Applications*, 278(1-3):263–284, 1998.
- [34] F. K. Hwang, S. Onn, and U. G. Rothblum. A polynomial time algorithm for shaped partition problems. *SIAM Journal on Optimization*, 10(1):70–81, 1999.
- [35] S. Kafer. *Polyhedral diameters and applications to optimization*. PhD thesis, University of Waterloo, 2022.
- [36] S. Kafer, K. Pashkovich, and L. Sanit  . On the Circuit Diameter of Some Combinatorial Polytopes. *SIAM Journal on Discrete Mathematics*, 33(1):1–25, 2019.
- [37] V. Kaibel. Extended formulations in combinatorial optimization. *Optima*, 85:2–7, 2011.
- [38] V. Kaibel and S. Weltge. A short proof that the extension complexity of the correlation polytope grows exponentially. *Discrete and Computational Geometry*, 53(2):397–401, 2015.
- [39] V. Klee and C. Witzgall. Facets and vertices of transportation polyhedra. In G. B. Dantzig and A. F. Veinott, editors, *Mathematics of the decision sciences*, volume 1, pages 257–282. American Mathematical Society, Providence, RI, 1968.
- [40] R. K. Martin. Using separation algorithms to generate mixed integer model reformulations. *Operations Research Letters*, 10(3):119–128, 1991.
- [41] J. G. Oxley. *Matroid Theory*. Oxford Graduate Texts in Mathematics. Oxford University Press, 2011.
- [42] R. T. Rockafellar. The elementary vectors of a subspace of R^N . In *Combinatorial Mathematics and its Applications*, pages 104–127. University of North Carolina Press, Chapel Hill, NC, 1969.
- [43] T. Rothvo  . The matching polytope has exponential extension complexity. *Journal of the ACM*, 64(6):1–19, 2017.
- [44] L. Sanit  . The diameter of the fractional matching polytope and its hardness implications. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 910–921. IEEE, 2018.
- [45] D. A. Spielman and S.-H. Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51(3):385–463, 2004.
- [46] R. Vershynin. Beyond Hirsch conjecture: Walks on random polytopes and smoothed complexity of the simplex method. *SIAM Journal on Computing*, 39(2):646–678, 2009.
- [47] R. T. Wong. Integer programming formulations of the traveling salesman problem. In *Proceedings of the IEEE International Conference on Circuits and Computers*, pages 149–152, 1980.
- [48] M. Yannakakis. Expressing combinatorial optimization problems by linear programs. *Journal of Computer and System Sciences*, 43(3):441–466, 1991.
- [49] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer, New York, NY, 1995.