

Circuits, Cutting Planes, and Compact Formulations

Matthias Brugger

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz: Prof. Dr. Gregor Kemper

Prüfende der Dissertation:

1. Prof. Dr. Andreas S. Schulz
2. Prof. Dr. Volker Kaibel
3. Prof. Dr. Steffen Borgwardt

Die Dissertation wurde am 26.09.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 20.02.2024 angenommen.

Abstract

Linear programs (LPs) are well-solved optimization problems, both in theory and in practice. This is also exploited by fundamental techniques for solving integer programs such as the cutting plane method. In this thesis, we study three fundamental aspects of the theory of integer and linear programming.

In the first part, we study $\{0, 1/2\}$ -cuts, which are Gomory-Chvátal cuts that are derived from a linear system using multipliers 0 or $1/2$ only. They play a central role in polyhedral combinatorics, notably in Edmonds' pioneering work on the matching polytope. We prove that recognizing rational polyhedra for which all $\{0, 1/2\}$ -cuts suffice to determine the integer hull is strongly NP-hard. We further investigate structural properties of the family of $\{0, 1/2\}$ -cuts and their implications for the separation problem.

Next, we consider extended formulations of polyhedra, which allow for compact LPs for many combinatorial optimization problems. An extended formulation describes a higher-dimensional polyhedron that linearly projects onto the original one. Rothvoß (2017) proved that any extended formulation for the matching polytope needs an exponential number of constraints. We show that the main tool in his proof, the hyperplane separation bound, does not directly improve upon the best known lower bounds for sizes of extended formulations for spanning tree and completion time polytopes.

In the final two chapters, we study the geometric underpinning of circuit augmentation schemes. These are LP algorithms that generalize the Simplex method by moving between solutions along circuits of the feasible region instead of only edges. We first show that edges and circuits behave in fundamentally different ways under projections of polyhedra, as they arise in extended formulations. We then study circuit diameters of polyhedra, which lower bound combinatorial diameters. The upper bound for the combinatorial diameter predicted by the Hirsch conjecture turned out to be false, with counterexamples found by Klee and Walkup (1967) and Santos (2012). The circuit analogue of the Hirsch conjecture, the so-called circuit diameter conjecture, is open. Previously, only the unbounded Klee-Walkup polyhedron was studied in the circuit setting. We consider Santos' counterexample to the Hirsch conjecture and prove that the key combinatorial property of the polytopes underlying his construction is no longer true when using circuits.

Zusammenfassung

Lineare Optimierungsprobleme (LPs) lassen sich effizient lösen, sowohl aus theoretischer Sicht als auch in der Praxis. Das machen sich auch grundlegende Techniken zum Lösen ganzzahliger Optimierungsprobleme wie Schnittebenenverfahren zunutze. Diese Arbeit widmet sich drei theoretischen Aspekten der ganzzahligen und linearen Optimierung.

Im ersten Teil untersuchen wir $\{0, 1/2\}$ -Schnitte, die sich ausgehend von einem System linearer Ungleichungen als Gomory-Chvátal-Schnitte mit Koeffizienten 0 oder $1/2$ herleiten lassen. Diese spielen eine wichtige Rolle in der polyedrischen Kombinatorik, insbesondere in Edmonds' wegweisender Arbeit zum Matching-Polytop. Wir zeigen unter anderem, dass es NP-schwer ist, rationale Polyeder zu erkennen, deren ganzzahlige Hülle sich nur mit $\{0, 1/2\}$ -Schnitten erreichen lässt. Außerdem untersuchen wir strukturelle Eigenschaften der Familie aller $\{0, 1/2\}$ -Schnitte und was sich daraus für das Separierungsproblem folgern lässt.

Als Nächstes betrachten wir erweiterte Formulierungen von Polyedern. Diese erlauben es, viele kombinatorische Optimierungsprobleme als kompakte LPs zu formulieren, und beschreiben dabei höherdimensionale Polyeder, die sich auf das ursprüngliche Polyeder projizieren lassen. Wie von Rothvoß (2017) bewiesen, benötigt man für jede erweiterte Formulierung des Matching-Polytops exponentiell viele Ungleichungen. Wir zeigen, dass die Technik, auf der Rothvoß' Beweis fußt, für Polytope im Zusammenhang mit Spannbäumen und Scheduling-Problemen nicht unbedingt bessere untere Schranken an die Größe von erweiterten Formulierungen liefert als das, was bereits bekannt ist.

Der dritte und letzte Teil der Arbeit widmet sich den geometrischen Grundlagen einer Familie von Verfahren zum Lösen von LPs, die den Simplex-Algorithmus verallgemeinern, indem sie LP-Lösungen nicht nur entlang von Kanten, sondern entlang von sogenannten Circuits verbessern. Zunächst zeigen wir, dass sich Circuits unter Projektionen von Polyedern grundsätzlich anders verhalten als Kanten. Dann betrachten wir den Circuit-Durchmesser von Polyedern. Dieser verallgemeinert den kombinatorischen Durchmesser, für den die Hirsch-Vermutung eine obere Schranke nahelegte. Diese Vermutung stellte sich als falsch heraus, wie von Klee und Walkup (1967) sowie von Santos (2012) widerlegt. Ob die Schranke aus der Hirsch-Vermutung für den Circuit-Durchmesser von Polyedern gilt, ist offen. Bisher wurde nur das unbeschränkte Gegenbeispiel von Klee und Walkup auf seinen Circuit-Durchmesser hin untersucht. Wir betrachten die Polytope, auf denen Santos' Gegenbeispiel basiert, und zeigen, dass deren wesentliche kombinatorische Eigenschaft im Circuit-Fall nicht mehr erfüllt ist.

Contents

1	Introduction	1
1.1	Organization of This Thesis	9
1.2	Notation	11
2	When the $\{0, 1/2\}$-Closure Coincides with the Integer Hull	13
2.1	Introduction	13
2.2	Computational Complexity of Integrality Properties of the $\{0, 1/2\}$ -Closure	18
2.2.1	Proof of Theorem 2.3	18
2.2.2	Further Hardness Results	20
2.3	Binary Clutters and the $\{0, 1/2\}$ -Closure	22
2.3.1	Preliminaries	22
2.3.2	On the Hardness of Finding Most Violated $\{0, 1/2\}$ -Cuts	30
2.3.3	Set Packing, Idealness, and the MFMC Property	32
2.3.4	Implications for Fractional Matching and Stable Set Polytopes	35
2.4	Further Notes and Open Questions	49
3	Limitations of the Hyperplane Separation Bound for the Extension Complexity of Polytopes	55
3.1	Introduction	55
3.2	Preliminaries	58
3.2.1	Slack Matrices and Nonnegative Factorizations	58
3.2.2	The Hyperplane Separation Bound	59
3.3	Limitations of the Hyperplane Separation Bound	62
3.3.1	Spanning Tree Polytopes	62
3.3.2	Graphic Zonotopes	63
3.4	Diagonal Scalings and Redundancy	66
3.5	Further Notes and Open Questions	69
4	Circuits under Projections of Polyhedra	71
4.1	Introduction	71
4.2	Preliminaries	73
4.2.1	Circuits of Polyhedra	73
4.2.2	Circuits under Affine Equivalence of Polyhedra	74
4.2.3	Inheritance of Edge Directions	77
4.3	Counterexamples for the Inheritance of Circuits	79
4.3.1	A Family of Minimal Counterexamples	79
4.3.2	Counterexamples by Counting Circuits	82
4.3.3	Fixed-Shape Partition Polytopes	85

4.4	The Role of Projection Maps and Polyhedra for the Inheritance of Circuits	87
4.4.1	Inheritance Based on the Projection Map	87
4.4.2	Inheritance for all Extensions	88
4.4.3	(No) Inheritance Based on the Extension Polyhedron	90
4.4.4	Inheritance in Nontrivial Instances	92
4.5	Further Notes and Open Questions	93
5	Hirsch Counterexamples and the Circuit Diameter Conjecture	95
5.1	Introduction	95
5.2	Preliminaries	98
5.2.1	Circuits Revisited	98
5.2.2	Conformal Sums of Circuits	100
5.2.3	Short Circuit Walks on Polygons	102
5.3	Bounded Hirsch Counterexamples	104
5.3.1	The Circuit Length of the 5-Dimensional Spindles	104
5.3.2	The Circuit Length of the 20- and 23-Dimensional Hirsch Counterexamples Satisfies the Hirsch Bound	114
5.4	Further Notes on Circuit Lengths and Circuit Diameters	117
A	Polyhedral Theory	121
	References	125

Chapter 1

Introduction

From Problems to Integer Programs. We begin with an example. A *matching* in an undirected graph is a set of edges no two of which share an endpoint. A matching is *maximum* if it is of maximum cardinality. For example, the graph in Figure 1.1 has a matching of cardinality two. This is maximum because the graph has five nodes, so in every subset of three edges, two of them must intersect in a node.




Figure 1.1

Finding a maximum matching in a given undirected graph $G = (V, E)$ is a classical combinatorial optimization problem known as the *maximum matching problem*. To solve this problem, one approach is the following. The *incidence vector* $\chi(M)$ of a matching $M \subseteq E$ is a 0/1 vector with an entry for each edge $e \in E$ which equals 1 if $e \in M$ and 0 otherwise. By the definition of a matching, all incidence vectors of matchings in G satisfy the linear inequalities

$$\begin{aligned} x(\delta(v)) &\leq 1 && \text{for all } v \in V \\ x_e &\geq 0 && \text{for all } e \in E \end{aligned} \tag{1.1}$$

where $\delta(v)$ denotes the set of all edges incident with node v and $x(\delta(v)) = \sum_{e \in \delta(v)} x_e$. Conversely, every integral solution $x \in \mathbb{Z}^E$ of (1.1) can only have components in $\{0, 1\}$ and must therefore be the incidence vector of some matching in G . This means that one can solve the maximum matching problem by finding an integral solution x of the system (1.1) that maximizes the linear objective function $\mathbf{1}^\top x = \sum_{e \in E} x_e$. Optimization problems of this type are *integer programs (IPs)*. More generally, we consider IPs of the form

$$\max \left\{ c^\top x : Ax \leq b, x \in \mathbb{Z}^n \right\} \tag{1.2}$$

for a matrix $A \in \mathbb{Q}^{m \times n}$ and vectors $b \in \mathbb{Q}^m$ and $c \in \mathbb{Q}^n$.

Many combinatorial optimization problems, not only the maximum matching problem, can be modelled as IPs. Suppose that we switch the roles of nodes and edges in our example above. Instead of all subsets M of edges for which no node of G is an endpoint of two edges of M , we now consider all subsets S of nodes such that no edge of G has both of its endpoints in S . These subsets S are the *stable sets* in the graph G . As with

the maximum matching problem, the problem of finding a maximum stable set also has a succinct IP formulation: Assuming that G has no isolated nodes, the incidence vectors of stable sets in G are precisely the integral solutions of

$$\begin{aligned} x_u + x_v &\leq 1 && \text{for all } uv \in E \\ x_v &\geq 0 && \text{for all } v \in V \end{aligned} \tag{1.3}$$

where we write uv for the edge with endpoints u and v . So a maximum stable set corresponds to an integral solution $x \in \mathbb{Z}^V$ of (1.3) that maximizes $\mathbf{1}^\top x$. This means that general integer programming can be no easier than finding a maximum stable set, which is a well-known NP-hard problem [139]. So to find an algorithm with good theoretical running time guarantees for the maximum matching problem, one likely needs more than just access to a black-box IP solver.

Linear Programs. If we ignore the integrality constraints in (1.2), the problem becomes

$$\max \left\{ c^\top x : Ax \leq b, x \in \mathbb{R}^n \right\}. \tag{1.4}$$

This is an optimization problem of a different type: a *linear program (LP)*. LPs of the form (1.4) for rational A, b , and c can be solved in polynomial time, and there are algorithms that perform very well in practice. (We will discuss one later in this chapter.) The feasible region of (1.2) is a polyhedron, and, given rational data, is said to be a *rational polyhedron*. (In fact, we may even assume that all coefficients of A and b are integers after rescaling; whenever we speak of rational polyhedra in the following, we will make this assumption.)

Given that LPs are polynomially solvable, we can efficiently compute a vector $x \in \mathbb{R}^E$ that satisfies (1.1) and maximizes the linear objective function $\mathbf{1}^\top x$, for any given graph $G = (V, E)$. Does this imply that we can find maximum matchings efficiently? Unfortunately, no. For the graph in Figure 1.1, the maximum value of $\mathbf{1}^\top x$ over (1.1) is $\frac{5}{2}$, and a fractional solution x^* that attains this maximum is shown in Figure 1.2. However, the maximum possible value of $\mathbf{1}^\top x$ attained by any integral solution x of (1.1) is $2 < \frac{5}{2}$, as we saw above. In fact, this is no coincidence: For any non-bipartite graph G , such as the one in Figure 1.1, one can find an integral objective function vector $w \in \mathbb{Z}^E$ for which the maxima of $w^\top x$ over (1.1) and over its integer points do not coincide. This means that the polyhedron determined by (1.1) is not necessarily an integral polyhedron. For a given graph G , it is therefore called the *fractional matching polytope* of G . Note that it is indeed a polytope, i.e., a bounded polyhedron, since one can add to (1.1) the valid inequalities $x_e \leq 1$ for each edge $e \in E$ without changing the set of feasible solutions.

Nonetheless, the maximum matching problem may be reduced to solving an LP – though one with a different feasible region. Namely, as any graph G has only finitely many matchings, the number of integral solutions of the system (1.1) (which we know to be precisely the incidence vectors of matchings in G) is finite. Hence, if we take the convex hull of those integral solutions, we obtain a polytope again, the *integer hull* of the fractional matching polytope. Let us call this polytope the *matching polytope* of G . If we now maximize our linear objective function $\mathbf{1}^\top x$ over the matching polytope instead




Figure 1.2: A fractional solution x^* of the system (1.1) for the graph in Figure 1.1. Only nonzero entries are shown.

of the fractional matching polytope of G , then we know that an optimal vertex must be the incidence vector of a matching in G whose cardinality is the optimal objective function value. This approach naturally extends to any other linear objective function $w^\top x$ given by a vector of edge weights $w \in \mathbb{R}^E$: Finding matchings of maximum weight (solving the *maximum-weight matching problem*) is equivalent to linear optimization over the matching polytope. However, to formulate this problem as an LP of the form (1.4), we need a complete linear description of the matching polytope, i.e., a system of linear inequalities whose set of solutions is precisely the matching polytope. How do we find such a system?

Cutting Planes and Closures. We saw above that the system (1.1) is not sufficient to describe the matching polytope – not even for the small graph in Figure 1.1. In this example, however, it is easy to find an inequality that we may add to (1.1): Since no matching in the graph in Figure 1.1 can have more than two edges, the inequality $\mathbf{1}^\top x \leq 2$ is valid for its matching polytope. Moreover, it cuts off the fractional point x^* in Figure 1.2. Such an inequality is called a *cutting plane* (or simply *cut*). Of course, to derive this cut, we exploited our knowledge of what the integer points of the fractional matching polytope encode. We may, however, also do this entirely without any problem-specific knowledge as follows.

Namely, for the graph in Figure 1.1, summing over all but the nonnegativity constraints in (1.1) and dividing by two yields the inequality $\mathbf{1}^\top x \leq \frac{5}{2}$. By construction, this inequality is valid for the fractional matching polytope. As all coefficients on the left-hand side are integers, we know that for any integral vector x , the inner product $\mathbf{1}^\top x$ must be integral, too. This means that all integral points in the fractional matching polytope of the graph in Figure 1.1 – and hence all points in its matching polytope – must satisfy the stronger inequality $\mathbf{1}^\top x \leq \lfloor \frac{5}{2} \rfloor = 2$. More generally, take any nonnegative linear combination of inequalities from (1.1) such that the resulting valid inequality only has integer coefficients, possibly except for the right-hand side. Rounding down the right-hand side to the next integer will not cut off any integral points and therefore yields a valid inequality for the matching polytope. This recipe for generating cuts, which applies to any rational polyhedron, originated in the work of Gomory [118] and Chvátal [53]. The resulting inequalities are therefore called *Gomory-Chvátal cuts*.

Using these types of cuts, Gomory [118] showed how to reduce any IP to solving a finite sequence of LPs whose feasible region is gradually refined by adding cutting planes. Such *cutting plane methods* are one of the backbones of modern IP solvers. Chvátal [53] formalized Gomory’s cutting plane method and proved that it can be used to characterize the integer hull of polyhedra (see also [173]): All possible Gomory-

Chvátal cuts that can be derived from the linear description of a rational polyhedron P determine a rational polyhedron again, the (*first*) Gomory-Chvátal closure of P (called the *elementary closure* in Chvátal’s paper [53]). Taking the Gomory-Chvátal closure of this polyhedron, in turn, yields the *second* Gomory-Chvátal closure of P , and so on. Chvátal proved that for every rational polyhedron P , one eventually obtains the integer hull after a finite number of rounds of this procedure. For polytopes that, like the fractional matching polytope, are contained in the 0/1 cube (i.e., for which all variables can only take values between 0 and 1), this number is essentially quadratically bounded in the number of variables, up to a logarithmic factor [90, 168].

For the fractional matching polytope, in fact, already the first Gomory-Chvátal closure yields the integer hull: To see this, let us note that the cut $\mathbf{1}^\top x \leq 2$ in our small example belongs to a more general family of Gomory-Chvátal cuts for the fractional matching polytope. Indeed, for any graph $G = (V, E)$, take a subset of nodes $S \subseteq V$ of odd cardinality, and add the inequality $x(\delta(v)) \leq 1$ in (1.1) for all nodes $v \in S$ and the nonnegativity constraint $-x_e \leq 0$ for all edges $e \in \delta(S)$, i.e., all edges with exactly one endpoint in S . The resulting valid inequality is

$$\sum_{v \in S} x(\delta(v)) - x(\delta(S)) \leq |S|.$$

Denoting the set of edges with both endpoints in S by $E(S)$, we may rewrite this inequality as $2x(E(S)) \leq |S|$ since $\sum_{v \in S} x(\delta(v)) = 2x(E(S)) + x(\delta(S))$ for all $x \in \mathbb{R}^E$. Dividing by two and rounding down the right-hand side yields the Gomory-Chvátal cut

$$x(E(S)) \leq \frac{|S| - 1}{2} \tag{1.5}$$

where we used that $|S|$ is odd. For the graph in Figure 1.1, for instance, choosing all five nodes for S results in the known inequality $\mathbf{1}^\top x \leq 2$.

In a seminal paper from the 1960s, Edmonds [79] proved that the system (1.1) together with the inequalities (1.5) for all odd subsets $S \subseteq V$ is sufficient to describe the matching polytope of any graph $G = (V, E)$. Chvátal [53] reframed this in terms of the Gomory-Chvátal closure, concluding that the Gomory-Chvátal closure of the fractional matching polytope of any graph indeed coincides with its integer hull (i.e., the matching polytope).

Polyhedra with this property are especially interesting from an IP standpoint since linear optimization over their integer points is likely “easier” than general integer programming in the following sense. Consider the *decision version* of an IP of the general form (1.2), which is the problem of deciding whether the given rational linear system $Ax \leq b$ has an integral solution of objective function value at least some given number. This problem is known to be in $\text{NP} \cap \text{coNP}$ when given the promise that the Gomory-Chvátal closure of the polyhedron $\{x: Ax \leq b\}$ is integral (see, e.g., [41]). Belonging to the complexity class $\text{NP} \cap \text{coNP}$ is believed to be an indication against being a hard problem, since $\text{NP} \cap \text{coNP}$ contains no NP-hard problems unless $\text{NP} = \text{coNP}$. The latter hypothesis, in the words of Karp and Papadimitriou [140], “*is weaker than* $\text{P} = \text{NP}$, *but is generally considered almost as improbable*”.

Can one efficiently test whether the Gomory-Chvátal closure and the integer hull of a given rational polyhedron coincide? Unfortunately, recognizing this property is NP-hard

– even for polyhedra that, like the fractional matching polytope, are contained in the 0/1 cube [62, 64]. However, the fractional matching polytope satisfies an even stronger condition. Not only can we obtain its integer hull by adding the Gomory-Chvátal cuts (1.5), those cuts are also special Gomory-Chvátal cuts. Recall that any Gomory-Chvátal cut for a given polyhedron is derived from a nonnegative linear combination of inequalities of its linear description. To derive the cuts (1.5) from (1.1) as we did above, we only used linear combinations with coefficients 0 or $\frac{1}{2}$. This property is shared with several other relevant classes of facet-defining inequalities for polytopes in combinatorial optimization, e.g., the odd-cycle inequalities for the *stable set polytope* of a graph G [114], which is the convex hull of all integral solutions of (1.3). The odd-cycle inequalities, too, are such special Gomory-Chvátal cuts, which we call $\{0, \frac{1}{2}\}$ -cuts for short [49]. One of the results obtained in Chapter 2 is to prove that for a given rational polyhedron, deciding whether all $\{0, \frac{1}{2}\}$ -cuts (which yield the $\{0, \frac{1}{2}\}$ -closure) suffice to determine the integer hull is NP-hard. In fact, we obtain this hardness result by showing that one could solve (the decision version of) the maximum stable set problem in polynomial time, assuming that the above recognition problem is polynomially solvable.

Optimization and Separation. For the maximum-weight matching problem, the fact that its decision version is in $\text{NP} \cap \text{coNP}$ is, of course, not all that is known about its complexity: Edmonds’ famous blossom algorithm [79, 80] finds a matching of maximum weight in polynomial time. In LP terms, this means that one can efficiently optimize any linear objective function over the matching polytope of a given graph $G = (V, E)$. Interestingly, to do this, it suffices to have an efficient algorithm that, given a point x^* , decides whether x^* is in the matching polytope of G and, if not, finds a valid inequality for the matching polytope that is violated by x^* . This follows from a fundamental result of Grötschel, Lovász, and Schrijver [124], which states that for families of polyhedra such as the matching polytope, the above so-called *separation problem* and linear optimization are roughly of the same complexity (commonly termed the *equivalence of optimization and separation*).

To solve the separation problem for the matching polytope, we first check whether the given vector x^* is in the fractional matching polytope (there is only a small number of inequalities in (1.1) to check). If x^* satisfies all of (1.1) but is not in the matching polytope, x^* must violate one of the $\{0, \frac{1}{2}\}$ -cuts (1.5). It is well known how to find such a separating cut in polynomial time: Padberg and Rao [161] showed how to do this by means of solving an auxiliary combinatorial optimization problem, which calls for an odd cut of minimum weight in a suitably defined edge-weighted graph. Here, an *odd cut* is a subset of edges of the form $\delta(S)$ for some subset of nodes S of odd cardinality.

The feasible solutions of this auxiliary problem, the odd cuts (as subsets of edges), form a special set family: a *binary clutter*. Binary clutters are well-studied objects in polyhedral combinatorics, graph theory, and optimization, and they arise naturally in the study of $\{0, \frac{1}{2}\}$ -cuts, as shown by Caprara and Fischetti [49]. In fact, with every rational polyhedron, one can associate an auxiliary combinatorial optimization problem over a binary clutter. If this auxiliary problem is polynomially solvable, then, via the equivalence of optimization and separation, one can optimize over the integer hull in polynomial time, provided that the $\{0, \frac{1}{2}\}$ -closure coincides with the integer hull. As

we will prove in Chapter 2, the converse implication is false unless $P = NP$: There are rational polyhedra with an integral $\{0, \frac{1}{2}\}$ -closure over which one can optimize in polynomial time, yet solving the associated binary clutter problem is NP-hard. We will also show how the binary clutter framework for $\{0, \frac{1}{2}\}$ -cuts introduced in [49] can be exploited to make conclusions about integrality properties of the $\{0, \frac{1}{2}\}$ -closure.

Compact Formulations. Even though Padberg and Rao’s combinatorial separation algorithm for the matching polytope is simpler than Edmonds’ blossom algorithm, the equivalence of optimization and separation relies on the ellipsoid method, which is considered slow for practical purposes. Simply feeding the linear description of the matching polytope to an efficient LP solver will not yield an efficient algorithm either since there is an exponential number of inequalities (1.5) and all of them are facet-defining in general, so none of them may be dropped. This phenomenon is not uncommon for polyhedra in combinatorial optimization. Also the *spanning tree polytope* of a connected graph $G = (V, E)$, which is the convex hull of the incidence vectors of spanning trees in G , generally has a number of facets that is exponential in $|V|$. (Incidentally, its linear description is also due to Edmonds [82].) However, one can obtain a significantly more compact LP formulation (a so-called *extended formulation*) by introducing few additional variables, which reduces the number of required constraints to $O(|V||E|)$ [153, 204]. Using this extended formulation, minimum-weight spanning trees can be found efficiently by means of solving a single LP of polynomial size, confirming the well-known polynomial solvability of this problem (see [177]).

Can one do the same for matchings? Unfortunately, no. Rothvoß [167] proved that any extended formulation for the matching polytope of a complete graph needs exponentially many inequalities (in the number of nodes). Previously, similar exponential lower bounds had been established for other well-known polytopes in combinatorial optimization, including the stable set polytope [101] (see also [136]). Given that the stable set problem is NP-complete, the fact that there is no polynomial-size extended formulation may not come as a surprise, even though the bound of [101] does not condition on $P \neq NP$ or other complexity-theoretic hypotheses. The exponential lower bound for the matching polytope, however, is more surprising in this regard – after all, one can optimize over it in polynomial time.

It should be noted that there *are* ways to find maximum-weight matchings efficiently by compact linear programming: By a standard reduction, it suffices to be able to find maximum-weight *perfect* matchings in a graph twice the size of the input graph (see, e.g., [177]). Here, a matching is *perfect* if every node is contained in some edge of the matching. Now the maximum-weight perfect matching problem can, for example, be reduced to solving a polynomial number of polynomial-size LPs [15]. Another LP-based approach is via an equivalent variant of the separation problem, called *primal separation* (see [92, 178]), which, in the case of perfect matchings, amounts to solving a small number of minimum *s-t*-cut problems [92]. These have compact LP formulations. Moreover, there is a polynomial-size LP for testing optimality of a given perfect matching [196]. Also this is sufficient to optimize over the (perfect) matching polytope in polynomial time [178]. Finally, a careful implementation of the cutting plane method converges in polynomial time, as proved in [52].

While all of the above approaches rely on LPs of polynomial size, Rothvoß’ result in [167] implies that one cannot obtain a *single* compact LP for the maximum-weight (perfect) matching problem by means of an extended formulation. The key tool in his proof is the so-called *hyperplane separation bound* [99]. This lower bound, like many other bounding techniques, builds on a beautiful result of Yannakakis [205] that links the *extension complexity* of a polyhedron – which is the minimum number of inequalities of any extended formulation – to certain factorizations of an associated slack matrix. A *slack matrix* records the slack of a vertex in an inequality, for all pairs of vertices of the polyhedron and inequalities in a given linear description. The flexibility in choosing different but equivalent linear descriptions does not affect Yannakakis’ result, but it may have a substantial impact on the quality of the hyperplane separation bound, as we show in Chapter 3. In particular, when applied to the slack matrix of the spanning tree polytope based on Edmonds’ description in [82], we prove that the best lower bound one can hope to achieve via the hyperplane separation technique is no better than a rather immediate and well-known lower bound of $\Omega(|E|)$ (see, e.g., [100]). This is particularly interesting because the question whether the bound of $\Omega(|E|)$ can be improved upon for the extension complexity of spanning tree polytopes of complete graphs is a notoriously difficult problem [8, 141, 199].

Augmentation and Circuits. As mentioned above, LPs of the form (1.4) can be solved in polynomial time. Yet how does one actually solve an LP? A simple idea for solving *any* optimization problem, in fact, is the following. Find some feasible solution and check whether it is optimal. If not, find a better solution and repeat. This idea leads to an *augmentation scheme*. In fact, being able to find improving solutions efficiently usually suffices to obtain polynomial-time algorithms for combinatorial optimization problems [178, 180]. Several classical polynomial-time algorithms for well-known problems, especially for network flow problems, are augmentation schemes (also known as a *primal algorithms* in this context). Two famous examples are the Edmonds-Karp-Dinic algorithm for maximum flow problems [77, 87] and Goldberg and Tarjan’s cycle cancelling algorithm for computing minimum-cost flows [117].

The concept of augmentation is also very prominent in combinatorial algorithms for matching problems, including Edmonds’ blossom algorithm. For example, to find a maximum matching in the graph in Figure 1.1, we may start with an arbitrarily chosen edge. This clearly is a matching. Given this initial matching, we then check whether there is a path along which matching and non-matching edges alternate, and whose endpoints are not matched. See Figure 1.3 for an example. Flipping matching and non-matching edges along such a path yields a matching with one more edge. It is well known that a matching is maximum if no such augmenting path exists (see [177]).

Geometrically, augmenting a matching along an alternating path means moving along an edge of the matching polytope [54]. So we may rephrase the combinatorial augmentation scheme sketched above as follows: Start by finding some vertex of the matching polytope. Check whether moving along any of its incident edges improves the objective function value. If so, follow such an edge to an adjacent vertex and repeat. If there is no improving edge, we found an optimal vertex.

Replacing the matching polytope with the feasible region of any given LP, this is




Figure 1.3: An alternating path in the graph in Figure 1.1. The solid edge is the initial matching; the two dashed edges indicate a matching obtained after a single augmentation.

precisely the geometric idea behind arguably the most famous algorithm for solving LPs: the Simplex method, introduced by George B. Dantzig in the 1940s [70]. Of course, there are many details to be filled in – for example, how does one select an improving edge if more than one is available? This is governed by a *pivot rule*. To this day, no pivot rule is known for which the Simplex method is guaranteed to run in polynomial time. The existence of such a polynomial pivot rule is a long-standing open question in the theory of linear programming (see, e.g., [190]).

It is therefore natural to wonder whether good bounds on the worst-case running time of augmentation schemes for LPs may potentially be easier to obtain if the augmenting directions are drawn from a larger set than just the (incident) edge directions. Such a set must have the following property: For any linear objective function and any non-optimal solution, the set contains a direction along which the objective function value strictly increases. A set of vectors with this property is called a *test set* for linear programming. A particularly nicely structured test set that includes the edge directions is the set of *circuits* of a polyhedron [120]. This set can be thought of to consist of all potential edge directions that can arise by translating facets of the polyhedron [98]. Using circuits as augmenting directions, one can solve LPs with a polynomial number of augmentations [72, 73, 127, 132]. However, computing an augmenting circuit direction may be NP-hard [73] or require solving an auxiliary LP [34, 73, 132], although one of potentially simpler structure [31].

In rare cases, the circuits of a polyhedron correspond to the actual edge directions. This is true, for example, for the fractional matching polytope [73, 169]. In general, though, the set of circuits may be much larger than the set of edge directions. This is what one might expect especially for polyhedra in combinatorial optimization, which typically have an exponential number of facets. To describe their circuits, a tempting idea is to use an extended formulation, which may be significantly more compact, as we saw above. Given that an extended formulation determines a polyhedron that linearly projects onto the original one (by projecting out the extra variables), may it be the case that circuits are also images of circuits under the same projection map? While this is known to be true for the edge directions, we give a negative answer for circuits in Chapter 4. We prove that, for any polyhedron and a given circuit that is not an edge direction, one can construct an extended formulation none of whose circuits projects onto the given circuit. Even worse, linear images of polyhedra with no non-edge circuits may have exponentially many non-edge circuits.

Diameters of Polyhedra. As described above, the Simplex method traces a path from an initial vertex to an optimal one along edges of the polyhedron that is the feasible region. So the minimum length of such a path, measured by the number of its edges, is a lower bound for the number of steps of the Simplex method with any pivot rule. How short can these paths be? The maximum length of a shortest path along edges between any pair of vertices of a polyhedron is the (*combinatorial*) *diameter* of the polyhedron. In Dantzig’s book on linear programming [70] from 1963, Warren M. Hirsch conjectured that the diameter of a d -dimensional polyhedron with f facets is at most $f - d$. This became known as the *Hirsch conjecture*. It was soon disproved for unbounded polyhedra by Klee and Walkup [145] in 1967. The Hirsch conjecture for *polytopes*, however, remained open until 2012, when Santos [170] found a counterexample in dimension 43. This was later improved to lower dimensions in [155].

Of course, for a polynomial pivot rule to exist for the Simplex method, diameters of polyhedra do not necessarily have to satisfy the Hirsch bound of $f - d$. Still, rather surprisingly, it is not even known whether there is an upper bound on the diameter that is polynomial in f and d (see [170]). Given that good bounds on the number of augmenting steps seem easier to obtain when augmenting along circuits, an immediate question is whether the Hirsch bound holds in the circuit setting, i.e., for the length of paths between vertices that are constructed by moving maximally along circuits instead of edges. This question is known as the *circuit diameter conjecture* [35]. While this conjecture is open, it has been verified for special cases: For example, the Klee-Walkup polyhedron from [145], which violates the unbounded Hirsch conjecture, does satisfy the circuit diameter conjecture [185]. Is this also true for the known *bounded* Hirsch counterexamples from Santos’ work [155, 170] or does any of them give rise to a counterexample to the circuit diameter conjecture? This question is the subject of Chapter 5. We give a partial answer and show that the key combinatorial property for the polytopes in [155, 170] to be counterexamples to the Hirsch conjecture is no longer true when considering circuits.

1.1 Organization of This Thesis

We begin in Chapter 2 by introducing $\{0, \frac{1}{2}\}$ -cuts and the $\{0, \frac{1}{2}\}$ -closure of rational polyhedra more formally. In Section 2.2, we then explain our hardness result for testing whether the $\{0, \frac{1}{2}\}$ -closure and the integer hull coincide. As we will see, the proof also has a number of interesting consequences. Section 2.3 is largely self-contained and explores the binary clutter framework of [49] for the family of $\{0, \frac{1}{2}\}$ -cuts. We give a brief introduction to the relevant concepts from the theory of clutters in Section 2.3.1. We also review the precise relationship between binary clutters and $\{0, \frac{1}{2}\}$ -cuts in detail, providing characterizations of the binary clutters associated with the fractional matching polytope and its close relative, the fractional stable set polytope. In Section 2.3.2, we prove our hardness result for optimizing over binary clutters derived from polyhedra in the 0/1 cube whose $\{0, \frac{1}{2}\}$ -closure is integral. Section 2.3.3 contains both new and previously known observations on integrality properties of binary clutters for the $\{0, \frac{1}{2}\}$ -closure of set packing polytopes. Specialized to fractional matching and stable set polytopes, we obtain interesting connections to the strong Chvátal rank of their

coefficient matrices. Our results suggest many avenues for future work, which are summarized in Section 2.4.

In Chapter 3, we first explain the hyperplane separation bound, which operates on slack matrices of polytopes. In Section 3.3, we apply it to slack matrices of spanning tree polytopes and another class of polytopes that generalizes completion time polytopes from scheduling. In both cases, we will see that the hyperplane separation bound performs rather poorly, compared to the best known bounds. Section 3.4 mitigates our negative results of Section 3.3 by showing ways to improve the hyperplane separation bound by choosing different slack matrices.

Section 4.2 in Chapter 4 introduces the circuits of a polyhedron more formally and provides several equivalent characterizations from the literature. In Section 4.3, we then give a family of linear projections under which circuits are not necessarily images of circuits, when applied to carefully chosen polyhedra, including certain polytopes from clustering. We then extend these results in Section 4.4 to characterize for which linear projection maps or polyhedra circuits always are images of circuits under a projection.

In Chapter 5, we start in Section 5.1 by introducing diameters of polyhedra and the Hirsch conjecture in greater detail than we did above. In particular, we explain the key combinatorial parameter that Santos' construction of bounded counterexamples relies on. In Section 5.2, we prepare the arguments needed for our analysis of these counterexamples: We first describe an alternative perspective on the set of circuits of polyhedron. This will allow us to derive a well-known property of circuits, which motivates why the precise upper bound given by the Hirsch conjecture is of interest for bounding circuit diameters. We then explain our main technical tool in Section 5.2.3. This is used in the proofs in Section 5.3. Section 5.3.2 includes a detailed account of Santos' construction. Finally, Section 5.4 concludes the chapter by investigating possible extensions and implications of our results.

Reading this thesis requires some familiarity with the elementary concepts of linear programming, graph theory, and complexity theory – although at a very basic level. For instance, the reader should be familiar with LP duality; the complexity classes P, NP, and coNP; reductions between computational problems; and standard notions in graph theory. These concepts are covered in most standard textbooks on linear programming, integer programming, or combinatorial optimization, e.g., in [57, 174, 177] (to a much greater extent than what is needed here, of course). Since polyhedra and polyhedral techniques are prominent throughout this thesis, some knowledge of polyhedral theory is helpful. A brief summary of some basic concepts that we will need, especially in Chapters 4 and 5, can be found in Appendix A.

Bibliographic Notes. The results of Section 2.2 are joint work with Andreas S. Schulz and appear in [45]. Those of Chapter 3 appear in [44]. The results of Chapter 4 are joint work with Steffen Borgwardt and appear in [28]. The results of Chapter 5 are joint work with Alexander E. Black and Steffen Borgwardt, and appear in [21].

1.2 Notation

Throughout this thesis, we denote the i th standard basis vector by e_i . Further, I_n denotes the $n \times n$ identity matrix (where we omit the subscript if the dimension is clear from the context). We use $\mathbb{R}_{\geq 0}$ to denote the set of nonnegative reals and \mathbb{Z}_2 to denote the integers modulo 2. $\mathbf{0}$ and $\mathbf{1}$ are the vectors of all zeros and all ones, respectively, in appropriate dimension. For $n \in \mathbb{N}$, we write $[n]$ for $\{1, \dots, n\}$. Vectors with superscripts such as $a^{(i)}$ always denote the i th vector in a list of vectors, whereas subscripts as in a_i indicate the i th component of a vector a . For a vector $x \in \mathbb{R}^n$ and $S \subseteq [n]$, we write $x(S) = \sum_{i \in S} x_i$.

All graphs in this thesis are undirected and simple (i.e., without loops or parallel edges), unless stated otherwise. K_n denotes the complete graph on n nodes. We will stick to this terminology in order to avoid confusion with the *vertices* of polyhedra. Nonetheless, to denote the sets of nodes and edges of a graph G , we use the standard symbols V and E (or $V(G)$ and $E(G)$, respectively, whenever G is not clear from the context).

Chapter 2

When the $\{0, 1/2\}$ -Closure Coincides with the Integer Hull

The results of Section 2.2 of this chapter are joint work with Andreas S. Schulz [45]. The presentation of those results in Section 2.2 and the corresponding parts of Sections 2.1 and 2.4 is largely identical to our paper.

2.1 Introduction

Let $P = \{x \in \mathbb{R}^n : Ax \leq b\}$ with $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$ be a rational polyhedron. For all $u \in \mathbb{R}_{\geq 0}^m$ with $u^\top A \in \mathbb{Z}^n$, the inequality

$$u^\top Ax \leq \lfloor u^\top b \rfloor \quad (2.1)$$

is satisfied by all $x \in P \cap \mathbb{Z}^n$ and is therefore valid for the integer hull of P , denoted by $P_I = \text{conv}(P \cap \mathbb{Z}^n)$. Inequalities of the form (2.1) are called *Gomory-Chvátal cuts* for P [53, 118]. The intersection of all halfspaces corresponding to Gomory-Chvátal cuts yields the *Gomory-Chvátal closure* P' of P . Note that P' does not depend on the linear description $Ax \leq b$ of P (see [174]). Moreover, P' is a rational polyhedron again, which follows from the fact that all Gomory-Chvátal cuts with $[0, 1]$ -valued multipliers u suffice (see, e.g., [57]), i.e.,

$$P' = \left\{ x \in P : u^\top Ax \leq \lfloor u^\top b \rfloor, u \in [0, 1]^m, u^\top A \in \mathbb{Z}^n \right\}.$$

It even suffices to consider Gomory-Chvátal cuts for which $u^\top b \notin \mathbb{Z}$: If $u^\top b \in \mathbb{Z}$, the inequality (2.1) is redundant for P and is therefore called a *trivial* cut.

If we restrict to multipliers $u \in \{0, \frac{1}{2}\}^m$, we obtain a special family of Gomory-Chvátal cuts, first introduced by Caprara and Fischetti [49]. We refer to these special Gomory-Chvátal cuts as $\{0, \frac{1}{2}\}$ -cuts. The $\{0, \frac{1}{2}\}$ -closure of P is defined as

$$P_{\frac{1}{2}}(A, b) := \left\{ x \in P : u^\top Ax \leq \lfloor u^\top b \rfloor, u \in \{0, \frac{1}{2}\}^m, u^\top A \in \mathbb{Z}^n \right\}.$$

It is easily seen that $P_{\frac{1}{2}}(A, b)$ is a rational polyhedron. From the definition, it follows that $P_I \subseteq P' \subseteq P_{\frac{1}{2}}(A, b) \subseteq P$. Further note that, unlike P' , the $\{0, \frac{1}{2}\}$ -closure $P_{\frac{1}{2}}(A, b)$ depends on the system $Ax \leq b$ defining the polyhedron P . For instance, $2Ax \leq 2b$ clearly defines the same polyhedron P as $Ax \leq b$, while $P_{\frac{1}{2}}(2A, 2b) = P$. This contrasts with what is known for the Gomory-Chvátal closure, where $P' = P$ if and only if $P = P_I$ (see [174]).

$\{0, \frac{1}{2}\}$ -cuts are prominent in polyhedral combinatorics. We already saw an example in Chapter 1:

Example 2.1 (Matchings). Let $G = (V, E)$ be a graph. Recall from Chapter 1 that the *fractional matching polytope* of G is given by

$$\begin{aligned} x(\delta(v)) &\leq 1 \quad \text{for all } v \in V \\ x &\geq \mathbf{0} \end{aligned} \tag{2.2}$$

The *blossom inequalities* [53, 79]

$$x(E(S)) \leq \frac{|S| - 1}{2} \quad \text{for all } S \subseteq V \text{ with } |S| \text{ odd}, \tag{2.3}$$

are $\{0, \frac{1}{2}\}$ -cuts for the fractional matching polytope of G . As mentioned in Chapter 1, Edmonds [79] showed that the integer hull, the *matching polytope* of G , is determined by (2.2) and (2.3). In particular, the $\{0, \frac{1}{2}\}$ -closure of the fractional matching polytope is integral. \diamond

Let us look at another example.

Example 2.2 (Stable sets). Given a graph $G = (V, E)$ without isolated nodes, we call the polyhedron

$$\begin{aligned} x_u + x_v &\leq 1 \quad \text{for all } uv \in E \\ x &\geq \mathbf{0} \end{aligned} \tag{2.4}$$

the *fractional stable set polytope* of G . (It is readily checked that (2.4) indeed defines a polytope.) Its integer hull is the *stable set polytope* of G , the convex hull of all incidence vectors of stable sets in G (see Chapter 1).

For a cycle C of odd length in G , let us add the edge constraints in (2.4) along the edges of C . The resulting inequality is $2x(V(C)) \leq |V(C)|$, where $V(C)$ denotes the set of nodes of C . We divide by 2, round down the right-hand side, and obtain the family of $\{0, \frac{1}{2}\}$ -cuts

$$x(V(C)) \leq \frac{|V(C)| - 1}{2} \quad \text{for all odd cycles } C \tag{2.5}$$

These are the so-called *odd-cycle inequalities*. Gerards and Schrijver [114] proved that they suffice to describe the Gomory-Chvátal closure of the fractional stable set polytope of G (and therefore also its $\{0, \frac{1}{2}\}$ -closure).

Unlike the $\{0, \frac{1}{2}\}$ -closure of the fractional matching polytope, the $\{0, \frac{1}{2}\}$ -closure of the fractional stable set polytope, given by (2.4) and (2.5), is not integral in general. For instance, consider K_4 , the complete graph on 4 nodes. The point $x^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ satisfies all constraints in (2.4) and (2.5) but $\mathbf{1}^\top x^* > 1$. Clearly, a stable set in K_4 cannot contain more than one node, so x^* is not in the stable set polytope of K_4 . Following [54], a graph is called *t-perfect* if its stable set polytope coincides with the $\{0, \frac{1}{2}\}$ -closure of its fractional stable set polytope. With this terminology, K_4 is not *t-perfect*. \diamond

In both examples above, the family of $\{0, \frac{1}{2}\}$ -cuts can be separated in polynomial time [114, 122, 161]. So by the equivalence of optimization and separation [124], one can solve the stable set problem (which is NP-hard in general, as we saw in Chapter 1) on t -perfect graphs in polynomial time. In general, though, the following *separation problem* for the $\{0, \frac{1}{2}\}$ -closure is strongly NP-hard, as shown by Caprara and Fischetti [49]:

Given $A \in \mathbb{Z}^{m \times n}$, $b \in \mathbb{Z}^m$ and $x^* \in \mathbb{Q}^n$ such that $x^* \in P = \{x \in \mathbb{R}^n : Ax \leq b\}$, find a $\{0, \frac{1}{2}\}$ -cut for P that is violated by x^* or conclude that $x^* \in P_{\frac{1}{2}}(A, b)$.

We note that Caprara and Fischetti's original statement in [49] does not explicitly mention *strong* NP-hardness, even though this can be deduced from their proof that the following subproblem, which can be thought of as the decision version of the separation problem, is (strongly) NP-complete (see also [88, 152]):

Given $A \in \mathbb{Z}^{m \times n}$, $b \in \mathbb{Z}^m$ and $x^* \in \mathbb{Q}^n$ such that $x^* \in P = \{x \in \mathbb{R}^n : Ax \leq b\}$, decide whether $x^* \notin P_{\frac{1}{2}}(A, b)$.

For technical reasons, we call this subproblem the *membership problem* for the $\{0, \frac{1}{2}\}$ -closure, as in [88, 152], even though the more adequate term would be *non-membership*. The membership problem remains strongly NP-complete even when $Ax \leq b$ defines a polytope in the 0/1 cube, as shown by Letchford, Pokutta, and Schulz [152].

The fact that the membership problem is in NP is not difficult to see: An obvious NP certificate is a separating $\{0, \frac{1}{2}\}$ -cut, succinctly represented by its vector of multipliers. Now suppose that we are given the additional promise that $P_{\frac{1}{2}}(A, b) = P_I$ for an instance specified by A and b . In this case, there is also a short coNP certificate for membership in $P_{\frac{1}{2}}(A, b)$, since the vertices of P_I and (integral) directions of its extreme rays can be shown to have polynomially bounded encoding length (see Section 17.1 of [174] for details). This implies that the membership problem for the $\{0, \frac{1}{2}\}$ -closure is in $\text{NP} \cap \text{coNP}$ when the $\{0, \frac{1}{2}\}$ -closure coincides with the integer hull. This is a special case of the well-known fact that testing membership in the Gomory-Chvátal closure belongs to $\text{NP} \cap \text{coNP}$ when restricted to polyhedra P with $P' = P_I$ (see, e.g., [41]).

We saw in Example 2.1 that the fractional matching polytope has this property. Moreover, the linear system (2.2) and (2.3) that determines the $\{0, \frac{1}{2}\}$ -closure is even totally dual integral (TDI) [66]. This motivates the following questions that are the subject of the first part of this chapter: What is the computational complexity of recognizing rational polyhedra whose $\{0, \frac{1}{2}\}$ -closure coincides with the integer hull, and of deciding whether adding all $\{0, \frac{1}{2}\}$ -cuts produces a TDI system? We prove in Section 2.2 that both properties are strongly NP-hard to recognize. As a byproduct of our proof, we also obtain a hardness result for testing whether the Gomory-Chvátal closure and the $\{0, \frac{1}{2}\}$ -closure of a given rational polyhedron coincide. Recall that the polyhedra in both Examples 2.1 and 2.2 satisfy this property.

In the second part of this chapter (Section 2.3), we shall be exploring the structure of the family of $\{0, \frac{1}{2}\}$ -cuts and its implications for the separation problem in greater detail. Caprara and Fischetti's proof of their hardness result for separating $\{0, \frac{1}{2}\}$ -cuts in [49] relies on an interpretation of the family of $\{0, \frac{1}{2}\}$ -cuts for a given polyhedron as

a so-called *binary clutter*. Clutters are special set families; *binary* clutters are those for which the incidence vectors of their sets (the *members*) are solutions of systems of congruences modulo 2 (see Section 2.3.1 for a formal definition). Many combinatorial optimization problems whose feasible solutions are inherently parity-constrained can be formulated as linear optimization problems over suitable binary clutters. A classical example is the shortest path problem: The set of all s - t -paths in a graph with two distinguished nodes s and t is a binary clutter (see [60]). It is well known that for nonnegative edge weights, shortest s - t -paths can be computed in polynomial time (see, e.g., [177]). Yet also NP-hard combinatorial optimization problems such as the max-cut problem can be expressed in terms of a binary clutter (see, e.g., [111, 122, 126]; see also Section 2.4). Caprara and Fischetti [49] observed that the decision versions of all combinatorial optimization problems of this type, provided a nonnegative linear objective function, can be reduced to carefully chosen instances of the membership problem for the $\{0, \frac{1}{2}\}$ -closure, which led them to conclude NP-hardness of the membership problem.

On the other hand, if the binary clutter associated with the $\{0, \frac{1}{2}\}$ -cuts for a given polyhedron has the property that minimum-weight members can be found in polynomial time, for all nonnegative weights, then one immediately obtains an efficient separation routine for the $\{0, \frac{1}{2}\}$ -closure [49]. As we will see in Section 2.3.1, this is true, e.g., for the binary clutters associated with the fractional matching and stable set polytopes [114, 122, 161], as observed in [49]. Letchford [150] later generalized these polynomially solvable special cases of the separation problem even further, using results of [123, 193] from the theory of binary matroids, which are very closely related to binary clutters (see Section 2.3.1).

At this point, we stress that all of the above special cases of the separation problem for $\{0, \frac{1}{2}\}$ -cuts are polynomially solvable because minimum-weight members of the associated binary clutters can be found efficiently, for *arbitrary* nonnegative weights. As we will see in Section 2.3.1, this implies that, in these cases, one can not only find *some* violated $\{0, \frac{1}{2}\}$ -cut (if one exists) in polynomial time but even a *most* violated one. Here, we say that a $\{0, \frac{1}{2}\}$ -cut for $P = \{x: Ax \leq b\}$ induced by some multiplier u^* is *most violated* by $x^* \in P$ if it maximizes the amount of violation across all $\{0, \frac{1}{2}\}$ -cuts for P , i.e., if

$$(u^*)^\top Ax^* - \lfloor (u^*)^\top b \rfloor \geq u^\top Ax^* - \lfloor u^\top b \rfloor$$

for all $\{0, \frac{1}{2}\}$ -valued multipliers u for which $u^\top A$ is integral.

Note that the maximum possible amount by which any point in P may violate a $\{0, \frac{1}{2}\}$ -cut is $\frac{1}{2}$. Such *maximally violated* $\{0, \frac{1}{2}\}$ -cuts can be separated in polynomial time [50]. The problem of finding a *most* (not necessarily *maximally*) violated $\{0, \frac{1}{2}\}$ -cut, however, is at least as hard as the separation problem for the $\{0, \frac{1}{2}\}$ -closure: If there is a separating $\{0, \frac{1}{2}\}$ -cut for a given point x^* , there is also a most violated one that separates x^* from the $\{0, \frac{1}{2}\}$ -closure. So finding a most violated $\{0, \frac{1}{2}\}$ -cut is strongly NP-hard. In Section 2.3.2, we prove that, rather surprisingly, this problem remains strongly NP-hard even when the $\{0, \frac{1}{2}\}$ -closure and the integer hull coincide, and even for instances where the separation problem for the $\{0, \frac{1}{2}\}$ -closure is easy.

The theory of binary clutters (and clutters, more generally) in polyhedral combinatorics and optimization is extremely rich. So the binary clutter perspective on the family of $\{0, \frac{1}{2}\}$ -cuts does not only prove useful in studying the complexity of the separation

problem – one can even gain structural insights about the $\{0, \frac{1}{2}\}$ -closure from it. We demonstrate this in Sections 2.3.3 and 2.3.4. For instance, binary clutters always come in pairs: Each s - t -path in a graph intersects each s - t -cut in at least one edge. The set of all minimal s - t -cuts is a binary clutter again (see [60]). So from a polyhedral viewpoint, optimizing a nonnegative linear objective function over the binary clutter of minimal s - t -cuts in a graph with edge set E can be done by linear optimization over the integral points $z \in \mathbb{Z}^E$ of the set covering polyhedron

$$\begin{aligned} z(P) &\geq 1 && \text{for all } s\text{-}t\text{-paths } P \\ z &\geq \mathbf{0} \end{aligned} \tag{2.6}$$

Such a set covering polyhedron can be associated with any (binary) clutter by replacing the s - t -paths P above with the members of the clutter. In the case of s - t -paths, (2.6) is, in fact, an integral polyhedron. Whenever this happens, the binary clutter is said to be *ideal*, and it is further said to have the *max-flow min-cut (MFMC) property* if the system (2.6) is TDI. Both the clutters of s - t -paths and s - t -cuts in a graph are not only ideal but also have the MFMC property (see [60]).

Especially the MFMC property is key to many famous min-max theorems in combinatorial optimization (see Chapter 80 of [177]). For example, Menger’s theorem [157] states that in a graph with two distinguished nodes s and t , the minimum number of edges of an s - t -cut is equal to the maximum number of edge-disjoint s - t -paths. This follows directly from the fact that the clutter of s - t -paths has the MFMC property, which implies that the dual of the LP $\min\{\mathbf{1}^\top z : z \in \mathbb{R}^E, z \text{ satisfies (2.6)}\}$ has an integral optimal solution. (The MFMC property is sometimes also referred to as the *Mengerian property*.)

For binary clutters associated with the $\{0, \frac{1}{2}\}$ -closure of set packing polyhedra, we will see in Section 2.3.3 that idealness and the MFMC property imply integrality and total dual integrality of the $\{0, \frac{1}{2}\}$ -closure, respectively. Recall that a *set packing polyhedron* is given by a linear system of the form $Ax \leq \mathbf{1}, x \geq \mathbf{0}$ for a 0/1 matrix A . The fractional matching and stable set polytopes from Examples 2.1 and 2.2 both are of set packing type, where A is the node-edge or edge-node incidence matrix of a graph. In Section 2.3.4, we characterize exactly when the associated binary clutters are ideal and provide necessary conditions for them to have the MFMC property. Interestingly, our results imply that, for these special set packing systems, one can relate both idealness and the MFMC property to the strong Chvátal rank of the coefficient matrix A . Moreover, both integrality properties of the associated binary clutters can be recognized in polynomial time. To obtain our characterizations in Section 2.3.4, we combine celebrated and deep results of Gerards and Schrijver [114], Guenin [126], and Seymour [183]. Further implications of these results are discussed in the final part of this chapter, in Section 2.4.

2.2 Computational Complexity of Integrality Properties of the $\{0, 1/2\}$ -Closure

We start by settling the computational complexity of recognizing when the $\{0, \frac{1}{2}\}$ -closure of a rational polyhedron coincides with the integer hull. The related recognition problem for the Gomory-Chvátal closure was studied by Cornuéjols and Li [62]. They proved that, given a rational polyhedron P with $P_I = \emptyset$, deciding whether $P' = \emptyset$ is weakly NP-complete. This immediately implies weak NP-hardness of testing whether $P' = P_I$. Cornuéjols, Lee, and Li [64] extended these hardness results to the case when P is contained in the 0/1 cube. Moreover, they showed that deciding whether a constant number of Gomory-Chvátal inequalities is sufficient to obtain the integer hull is weakly NP-hard, even for polytopes in the 0/1 cube.

In this section, we establish analogous hardness results for the $\{0, \frac{1}{2}\}$ -closure. Our main result is the following theorem.

Theorem 2.3. *Given $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$ with $P = \{x \in \mathbb{R}^n : Ax \leq b\} \subseteq [0, 1]^n$, deciding whether $P_{\frac{1}{2}}(A, b) = P_I$ is strongly NP-hard, even when the inequalities $-x \leq \mathbf{0}$ and $x \leq \mathbf{1}$ are part of the system $Ax \leq b$.*

We give a proof of this theorem in Section 2.2.1. Our proof implies several further hardness results, which we explain in Section 2.2.2. In particular, deciding whether adding all $\{0, \frac{1}{2}\}$ -cuts to a given linear system $Ax \leq b$ produces a TDI system, is strongly NP-hard. We also establish strong NP-hardness of the following problems: deciding whether the $\{0, \frac{1}{2}\}$ -closure coincides with the Gomory-Chvátal closure; deciding whether a constant number of $\{0, \frac{1}{2}\}$ -cuts suffices to obtain the integer hull; computing the dimension of the $\{0, \frac{1}{2}\}$ -closure. Finally, we give a hardness result for the membership problem for the $\{0, \frac{1}{2}\}$ -closure, which is slightly stronger than the one of Letchford, Pokutta, and Schulz [152].

2.2.1 Proof of Theorem 2.3

We reduce from STABLE SET:

Let $G = (V, E)$ be a graph and $k \in \mathbb{N}, k \geq 2$. Does G have a stable set of size at least k ?

It is well known that STABLE SET is strongly NP-complete [139]. Note that the problem remains strongly NP-complete if restricted to graphs with minimum degree at least 2: Given an instance of STABLE SET specified by G and k , we construct a new graph G' by adding two dummy nodes to G as well as all edges with at least one endpoint being a dummy node. Every node in G' has degree at least 2, and every stable set in G' of size $k \geq 2$ is a stable set in G of the same size.

Consider an instance of STABLE SET given by $G = (V, E)$ and $k \geq 2$. By the above observation, we may assume that every node in V has degree at least 2. Note that $|V| =: n \geq 3$ and $|E| =: m \geq 3$ in this case. Let $A := 2 \cdot \mathbf{1}\mathbf{1}^\top - M^\top$ where $M \in \{0, 1\}^{m \times n}$ denotes the edge-node incidence matrix of G . We define a polytope $P \subseteq \mathbb{R}^m$ by the

following system of inequalities:

$$\mathbf{0} \leq x \leq \mathbf{1} \quad (2.7)$$

$$Ax \leq 2 \cdot \mathbf{1} \quad (2.8)$$

$$(2k-3)\mathbf{1}^\top x \geq 2k-3 \quad (2.9)$$

Claim 1. $P_I = \{x \in P : \mathbf{1}^\top x = 1\}$.

Proof of Claim 1. If we add all inequalities in (2.8), we obtain the valid inequality $2(n-1)\mathbf{1}^\top x \leq 2n$. Every integral point x in P therefore satisfies $\mathbf{1}^\top x = 1$. Since $A \in \{1, 2\}^{n \times m}$, it is easy to check that every standard basis vector is indeed contained in P . We conclude that

$$P_I = \{x \in [0, 1]^m : \mathbf{1}^\top x = 1\} \supseteq \{x \in P : \mathbf{1}^\top x = 1\} \supseteq P_I. \quad \diamond$$

The $\{0, \frac{1}{2}\}$ -cuts that can be derived from (2.7)–(2.9) are all the inequalities of the following two types with $u \in \{0, \frac{1}{2}\}^n$ and $v \in \{0, \frac{1}{2}\}^m$:

$$\sum_{i=1}^m \left(2u^\top \mathbf{1} + \lfloor v_i - (Mu)_i \rfloor \right) x_i \leq 2u^\top \mathbf{1} + \lfloor v^\top \mathbf{1} \rfloor \quad (2.10)$$

$$\sum_{i=1}^m \left(2u^\top \mathbf{1} - (k-1) + \lfloor \frac{1}{2} + v_i - (Mu)_i \rfloor \right) x_i \leq 2u^\top \mathbf{1} - (k-1) + \lfloor \frac{1}{2} + v^\top \mathbf{1} \rfloor \quad (2.11)$$

The first type (2.10) defines all cuts that are derived only from (2.7) and (2.8), whereas the second type (2.11) also uses inequality (2.9). The vector u is the vector of multipliers for inequalities (2.8) while v collects the multipliers for the upper bounds in (2.7).

In what follows, $P_{\frac{1}{2}}$ denotes the $\{0, \frac{1}{2}\}$ -closure of P defined by (2.7)–(2.9) together with (2.10) and (2.11) for all $u \in \{0, \frac{1}{2}\}^n$ and $v \in \{0, \frac{1}{2}\}^m$.

Claim 2. $P_{\frac{1}{2}} = P_I$ if and only if there is a $\{0, \frac{1}{2}\}$ -cut equivalent to $\mathbf{1}^\top x \leq 1$.

Proof of Claim 2. If there is such a cut, then $P_{\frac{1}{2}} \subseteq \{x \in P : \mathbf{1}^\top x \leq 1\} = P_I$ by Claim 1. To see the “only if” part, consider the vector $y = (\frac{1}{m} + \varepsilon) \mathbf{1}$ for some small $\varepsilon > 0$. Clearly, $y \notin P_I$ since $\mathbf{1}^\top y > 1$. We claim that there is a choice for ε such that $y \in P$ and y satisfies all $\{0, \frac{1}{2}\}$ -cuts except those that are equivalent to $\mathbf{1}^\top x \leq 1$. First observe that every cut (of either type (2.10) or (2.11)) as well as every inequality in (2.8) and (2.9) may be written as $a^\top x \leq \alpha$ for some $a \in \mathbb{Z}^m, \alpha \in \mathbb{Z}$ where $a_i \leq \alpha$ for all $i \in [m]$ and $\alpha \leq m+n$. If $\alpha \leq 0$, we clearly have $a^\top y \leq \alpha$ since $y \geq \frac{1}{m} \mathbf{1}$. If $\alpha > 0$ and $a^\top x \leq \alpha$ is not equivalent to $\mathbf{1}^\top x \leq 1$, then $a_i < \alpha$ for at least one $i \in [m]$. It follows that $a^\top y \leq \alpha - \frac{1}{m} + \varepsilon(m\alpha - 1)$. For instance, taking $\varepsilon := \frac{1}{m^2(m+n)}$ yields $a^\top y \leq \alpha$ as desired. \diamond

In particular, the proof of Claim 2 shows that the inequality $\mathbf{1}^\top x \leq 1$ is not valid for P and, hence, $P \neq P_I$.

Claim 3. No cut of type (2.10) is equivalent to $\mathbf{1}^\top x \leq 1$.

Proof of Claim 3. Let $u \in \{0, \frac{1}{2}\}^n$ and $v \in \{0, \frac{1}{2}\}^m$. If $u = \mathbf{0}$, (2.10) reduces to a trivial inequality. If $v = \mathbf{0}$, the cut (2.10) is a trivial cut which is only derived from inequalities in the description of P with even right-hand sides. Hence, we may assume that both $u \neq \mathbf{0}$ and $v \neq \mathbf{0}$. It suffices to show that $\lfloor v_i - (Mu)_i \rfloor < \lfloor v^\top \mathbf{1} \rfloor$ for at least one $i \in [m]$. If $v^\top \mathbf{1} \geq 1$, there is nothing to show. Now let $v^\top \mathbf{1} = \frac{1}{2}$ and suppose for the sake of contradiction that $\lfloor v_i - (Mu)_i \rfloor \geq 0$ for all $i \in [m]$. It follows that $Mu \leq v$. Since every column of M has at least two nonzero entries by assumption, we obtain $u = \mathbf{0}$, a contradiction. \diamond

Claim 4. A cut of type (2.11) induced by $u \in \{0, \frac{1}{2}\}^n$ and $v \in \{0, \frac{1}{2}\}^m$ is equivalent to $\mathbf{1}^\top x \leq 1$ if and only if $v = \mathbf{0}$, $2Mu \leq \mathbf{1}$, and $2u^\top \mathbf{1} \geq k$.

Proof of Claim 4. Suppose first that $v \neq \mathbf{0}$. Then, for every $i \in [m]$, we have $\lfloor \frac{1}{2} + v_i - (Mu)_i \rfloor \leq 1 \leq \lfloor \frac{1}{2} + v^\top \mathbf{1} \rfloor$. This holds with equality for all $i \in [m]$ simultaneously only if $v_i = \frac{1}{2}$ and $v^\top \mathbf{1} \leq 1$, contradicting $m \geq 3$. Thus, no inequality of the form (2.11) with $v \neq \mathbf{0}$ has identical coefficients that coincide with the right-hand side. We may therefore assume that $v = \mathbf{0}$.

If $2u^\top \mathbf{1} \leq k - 1$, inequality (2.11) is redundant: It is the sum of the inequalities $(2u^\top \mathbf{1} - (k - 1))\mathbf{1}^\top x \leq 2u^\top \mathbf{1} - (k - 1)$ and $\lfloor \frac{1}{2} - (Mu)_i \rfloor x_i \leq 0$ for all $i \in [m]$, all of which are valid for P . Assuming that $2u^\top \mathbf{1} \geq k$, inequality (2.11) is equivalent to $\mathbf{1}^\top x \leq 1$ if and only if $(Mu)_i \leq \frac{1}{2}$ for all $i \in [m]$. \diamond

Putting together Claims 2 to 4, we conclude that $P_{\frac{1}{2}} = P_I$ if and only if there exists some $u \in \{0, \frac{1}{2}\}^n$ such that $2u$ is the incidence vector of a stable set in G of size at least k . This concludes the proof of Theorem 2.3.

2.2.2 Further Hardness Results

A careful analysis of the proof of Theorem 2.3 shows that, if the polytopes P constructed in the reduction satisfy $P_{\frac{1}{2}} = P_I$, there is a single $\{0, \frac{1}{2}\}$ -cut that certifies this (see Claim 2). This observation immediately implies the following corollary.

Corollary 2.4. Let $k \in \mathbb{N}$ be a fixed constant. Given $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$ with $P = \{x \in \mathbb{R}^n : Ax \leq b\} \subseteq [0, 1]^n$, deciding whether one can obtain P_I by adding at most k $\{0, \frac{1}{2}\}$ -cuts is strongly NP-hard, even when $k = 1$, and $-x \leq \mathbf{0}$ and $x \leq \mathbf{1}$ are part of the system $Ax \leq b$.

Moreover, let us remark that $P' = P_I$ for the polytopes P arising from the reduction. This follows from the fact that for $n \geq 3$, the inequality $\mathbf{1}^\top x \leq \lfloor 2n/2(n-1) \rfloor = 1$ is a Gomory-Chvátal cut for P , see the proof of Claim 1.

Corollary 2.5. Given $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$ with $P = \{x \in \mathbb{R}^n : Ax \leq b\} \subseteq [0, 1]^n$, deciding whether $P_{\frac{1}{2}}(A, b) = P'$ is strongly NP-hard, even when $-x \leq \mathbf{0}$ and $x \leq \mathbf{1}$ are part of the system $Ax \leq b$.

The linear systems arising from our reduction have another interesting property. The inequality description (2.7)–(2.11) of $P_{\frac{1}{2}}$ in the proof of Theorem 2.3 is a TDI system if

and only if $P_{\frac{1}{2}} = P_I$. Recall that a rational linear system $Bx \leq d$ is *totally dual integral* (TDI) if the dual LP

$$\min \left\{ d^\top y : B^\top y = c, y \geq \mathbf{0} \right\}$$

has an integral optimal solution y for every integral vector c for which the minimum is finite. Further recall the following result of Edmonds and Giles [84] which states that total dual integrality of certain linear systems implies polyhedral integrality.

Proposition 2.6 ([84]). *Let $P = \{x : Bx \leq d\}$ be a rational polyhedron where d is an integral vector. If the linear system $Bx \leq d$ is TDI, then P is integral.*

It therefore suffices to show that if $P_{\frac{1}{2}} = P_I$ in the proof of Theorem 2.3, then (2.7)–(2.11) defines a TDI system. So suppose that $P_{\frac{1}{2}} = P_I$. By the proof of Theorem 2.3, there exist vectors $u', u'' \in \{0, \frac{1}{2}\}^n$ such that $2Mu' \leq \mathbf{1}$, $2Mu'' \leq \mathbf{1}$, $2(u')^\top \mathbf{1} = k$, and $2(u'')^\top \mathbf{1} = k - 2 \geq 0$ (see Claim 4). The cuts of type (2.11) derived with u' and u'' (where we take $v = \mathbf{0}$) are the inequalities $\mathbf{1}^\top x \leq 1$ and $-\mathbf{1}^\top x \leq -1$, respectively. The system defined by these two inequalities and $x \geq \mathbf{0}$ is a subsystem of (2.7)–(2.11) that is sufficient to describe $P_{\frac{1}{2}}$ (see Claim 1) and that is TDI. To see this, let $c \in \mathbb{Z}^m$. We may assume without loss of generality that c_1 is the largest coefficient of c . It follows that $\max \{c^\top x : x \in P_{\frac{1}{2}}\} = c_1$. It suffices to show that the inequality $c^\top x \leq c_1$ is a nonnegative integer linear combination of the selected subsystem. Indeed, it is the sum of $c_1 \mathbf{1}^\top x \leq c_1$ (which is a nonnegative integer multiple of $\mathbf{1}^\top x \leq 1$ or $-\mathbf{1}^\top x \leq -1$) and $-(c_1 - c_i)x_i \leq 0$ for all $i \in [m]$. The above argument shows the following result.

Corollary 2.7. *Let $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$. Deciding whether the system given by $Ax \leq b$ and all $\{0, \frac{1}{2}\}$ -cuts derived from it is TDI, is strongly NP-hard, even when $-x \leq \mathbf{0}$ and $x \leq \mathbf{1}$ are part of the system $Ax \leq b$.*

Another consequence of the proof of Theorem 2.3 is that computing the dimension of the $\{0, \frac{1}{2}\}$ -closure of a given rational polyhedron is strongly NP-hard, since the integer hull of any polytope P constructed in our reduction is a facet of P (see Claim 1 and the proof of Claim 2):

Corollary 2.8. *Given $A \in \mathbb{Z}^{m \times n}$, $b \in \mathbb{Z}^m$, and $d \in \mathbb{Z}$, deciding whether the dimension of $P_{\frac{1}{2}}(A, b)$ is at most d , is strongly NP-complete, even when $-x \leq \mathbf{0}$ and $x \leq \mathbf{1}$ are part of the system $Ax \leq b$.*

Further note that the presence of the constraints $x \leq \mathbf{1}$ in (2.7) is not essential for our reduction in the proof of Theorem 2.3. In fact, the upper bounds are redundant: For every $i \in [m]$, consider a row of A such that the entry in column i is equal to 2. Such a row exists because $n \geq 3$. The corresponding inequality in (2.8) together with the nonnegativity constraints $-x_j \leq 0$ (possibly twice) for all $j \neq i$ yields $2x_i \leq 2$ for all $x \in P$. As the only relevant cuts among (2.10) and (2.11) are those with $v = \mathbf{0}$, we conclude that all of the above results still hold true when the upper bounds $x \leq \mathbf{1}$ are not part of the input.

The final byproduct of our proof of Theorem 2.3 is that the membership problem for the $\{0, \frac{1}{2}\}$ -closure of polytopes in the 0/1 cube is strongly NP-complete. This has already

been shown by Letchford, Pokutta, and Schulz [152], strengthening the NP-completeness result of Caprara and Fischetti [49]. However, neither of the two different reductions given in [152] constructs linear systems that include both nonnegativity constraints and upper bounds on every variable. When these constraints are required to be part of the input, membership testing remains strongly NP-complete, as the following result shows.

Corollary 2.9. *The membership problem for the $\{0, \frac{1}{2}\}$ -closure of polytopes contained in the 0/1 cube is strongly NP-complete, even when the inequalities $-x \leq \mathbf{0}$ and $x \leq \mathbf{1}$ are part of the input.*

Proof. The problem clearly belongs to NP. To show hardness, we use the same reduction from STABLE SET as in the proof of Theorem 2.3. The vector y defined in the proof of Claim 2 satisfies $y \notin P_{\frac{1}{2}}$ if and only if the instance of STABLE SET is a “yes” instance. The encoding length of y is polynomial in m and n if we choose ε as in Claim 2. \square

The hardness result for the membership problem first given by Caprara and Fischetti [49] relies on the fact that the family of $\{0, \frac{1}{2}\}$ -cuts for a rational polyhedron admits a nice combinatorial interpretation as a binary clutter. In the next section, we will explore this connection to binary clutters in greater detail.

2.3 Binary Clutters and the $\{0, 1/2\}$ -Closure

The main contributions of this section are presented in Sections 2.3.2 to 2.3.4: a hardness result for separating most violated $\{0, \frac{1}{2}\}$ -cuts that is stronger than what can be derived from [49, 152]; and structural insights into integrality properties of the $\{0, \frac{1}{2}\}$ -closure of special set packing polyhedra. As mentioned above, our results rely on the close relationship between $\{0, \frac{1}{2}\}$ -cuts and binary clutters observed by Caprara and Fischetti [49]. To understand this relationship, we first need to introduce the relevant concepts in Section 2.3.1. As the entire section is intended to be largely self-contained, we then review some of Caprara and Fischetti’s results from [49]. While our presentation in parts of Section 2.3.1 follows their paper, we strongly emphasize the clutter perspective.

2.3.1 Preliminaries

There is a vast body of literature on clutters. Excellent introductions are [60] and, with a focus on ideal clutters, [1]. We follow these two references for all basic definitions. Binary clutters, idealness, and the MFMC property, more specifically, are also covered in Chapters 78 to 80 of [177]. Many of the results on binary clutters in the literature are phrased in terms of *binary matroids*, which are closely related. One of the ways in which they are related will be briefly sketched below (and revisited later in this chapter). However, all results presented in this chapter will be stated in terms of clutters, not matroids. The interested reader is referred to [160] for the elementary concepts of matroid theory, and to [1, 60, 183] for more details on connections between binary clutters and binary matroids.

Clutter Basics

Let E be a finite set. A *clutter* over ground set E is a collection \mathcal{F} of subsets of E with the property that no set in \mathcal{F} is contained in another one. In particular, the (inclusion-)minimal sets among any collection of subsets of E form a clutter. For a given clutter \mathcal{F} , the sets in \mathcal{F} are called the *members* of the clutter. \mathcal{F} is *binary* if, for any three members $S_1, S_2, S_3 \in \mathcal{F}$, their symmetric difference $S_1 \Delta S_2 \Delta S_3$ contains a member of \mathcal{F} again [148].

Given a clutter \mathcal{F} over ground set E , a *cover* (or *transversal*) of \mathcal{F} is a subset of E that intersects each member of \mathcal{F} . The set of all minimal covers of \mathcal{F} is a clutter over ground set E again: the *blocker* of \mathcal{F} , denoted by $b(\mathcal{F})$. Note that $b(b(\mathcal{F})) = \mathcal{F}$ [83]. Using the notion of blockers, binary clutters may also be characterized as follows. A clutter \mathcal{F} is binary if and only if members of \mathcal{F} and $b(\mathcal{F})$ intersect in an odd number of elements, i.e., $|S \cap C|$ is odd for all $S \in \mathcal{F}$ and all $C \in b(\mathcal{F})$ [148] (see also [182]). In particular, this implies that blockers of binary clutters are binary again.

Another well-known characterization of binary clutters is the following (see, e.g., [1]). As usual, we define the *support* of a vector $z \in \mathbb{R}^n$ as $\text{supp}(z) := \{i \in [n] : z_i \neq 0\}$.

Proposition 2.10 (see [1]). *A clutter \mathcal{F} over ground set $[q]$ is binary if and only if \mathcal{F} consists of the minimal sets in*

$$\{\text{supp}(y) : y \in \{0, 1\}^q, Qy \equiv d \pmod{2}\} \quad (2.12)$$

for a matrix $Q \in \{0, 1\}^{p \times q}$ and a vector $d \in \{0, 1\}^p$.

Proof. Let \mathcal{F} be the collection of all minimal sets in (2.12) for some $Q \in \{0, 1\}^{p \times q}$ and $d \in \{0, 1\}^p$. First note that \mathcal{F} is indeed a clutter as we only take minimal sets. For any three solutions $y_1, y_2, y_3 \in \{0, 1\}^q$ of the system of linear congruences $Qy \equiv d \pmod{2}$, we have that $Q(y_1 + y_2 + y_3) \equiv 3d \equiv d \pmod{2}$. Moreover, the support of $y_1 + y_2 + y_3$ over \mathbb{Z}_2 equals $\text{supp}(y_1 + y_2 + y_3) = \text{supp}(y_1) \Delta \text{supp}(y_2) \Delta \text{supp}(y_3)$. So for any three sets in \mathcal{F} , the incidence vector of their symmetric difference is a solution of $Qy \equiv d$ again, which means that \mathcal{F} must be binary.

Now suppose that \mathcal{F} is binary. Let Q be the matrix whose rows are the incidence vectors of the members of $b(\mathcal{F})$. Since \mathcal{F} is binary by hypothesis, it follows that $Q\chi(S) \equiv \mathbf{1} \pmod{2}$ for all $S \in \mathcal{F}$. Conversely, let y be a 0/1 vector with $Qy \equiv \mathbf{1} \pmod{2}$. Then $\text{supp}(y)$ is a cover of $b(\mathcal{F})$ and must therefore contain a member of $b(b(\mathcal{F})) = \mathcal{F}$. \square

Proposition 2.10 implies that binary clutters can be thought of as affine vector spaces over \mathbb{Z}_2 . With this characterization, one of the relationships with binary matroids becomes apparent. For a binary clutter \mathcal{F} given as in Proposition 2.10, let us “homogenize” the congruence system in (2.12) by introducing an additional variable $\eta \in \{0, 1\}$: Then a 0/1 vector y is a solution of the original system (2.12) if and only if $\begin{pmatrix} y \\ \eta \end{pmatrix}$ is a solution of $(Q \mid d) \begin{pmatrix} y \\ \eta \end{pmatrix} \equiv \mathbf{0} \pmod{2}$. Each solution of the homogenized system corresponds to a subset of columns of the 0/1 matrix $(Q \mid d)$ that are linearly dependent over \mathbb{Z}_2 . These are precisely the *cycles* of the binary matroid represented by $(Q \mid d)$. Minimal cycles are called *circuits*. So for the binary clutter \mathcal{F} given as $Qy \equiv d \pmod{2}$, there is a binary matroid \mathcal{M} over a ground set with one more element l (corresponding to the

rightmost column of $(Q \mid d)$ such that $\mathcal{F} = \{C \setminus \{l\} : C \text{ is a circuit of } \mathcal{M}, l \in C\}$. For a given element l , the clutter \mathcal{F} is called the l -port of \mathcal{M} .

This shows that binary clutters may be characterized in yet another way: as ports of binary matroids. We will briefly revisit this perspective later in this chapter, in Section 2.3.4. But first, let us introduce the binary clutters that are relevant for separating $\{0, \frac{1}{2}\}$ -cuts.

Separation of $\{0, 1/2\}$ -Cuts is a Binary Clutter Problem

Following [49], we associate a binary clutter with a given matrix $A \in \mathbb{Z}^{m \times n}$ and a vector $b \in \mathbb{Z}^m$ as follows. Let $\mathcal{F}(A, b)$ be the collection of all minimal sets in

$$\left\{ \text{supp}(y) : y \in \{0, 1\}^m, \begin{pmatrix} \bar{A}^\top \\ \bar{b}^\top \end{pmatrix} y \equiv \begin{pmatrix} \mathbf{0} \\ 1 \end{pmatrix} \pmod{2} \right\} \quad (2.13)$$

where $\bar{A} := A \pmod{2}$ and $\bar{b} := b \pmod{2}$. If we consider *all* sets in the collection (2.13) (not just the *minimal* ones), then these are precisely the supports of all multiplier vectors of nontrivial $\{0, \frac{1}{2}\}$ -cuts for the polyhedron $P = \{x \in \mathbb{R}^n : Ax \leq b\}$, as observed in [49]. Indeed, for all $u \in \{0, \frac{1}{2}\}^m$, we have that $u^\top A \in \mathbb{Z}^n$ and $u^\top b \notin \mathbb{Z}$ if and only if $y = 2u$ satisfies the congruences in (2.13), since $y^\top A \equiv y^\top \bar{A}$ and $y^\top b \equiv y^\top \bar{b} \pmod{2}$. Restricting to minimal sets in the definition of $\mathcal{F}(A, b)$ above therefore yields a subfamily of all nontrivial $\{0, \frac{1}{2}\}$ -cuts for P . We refer to them as *support-minimal* cuts, where the *support* of a given $\{0, \frac{1}{2}\}$ -cut induced by u is defined as $\text{supp}(u)$.

Here, it is important to note that Caprara and Fischetti [49] work with the entire collection (2.13) and do not distinguish between $\{0, \frac{1}{2}\}$ -cuts and support-minimal $\{0, \frac{1}{2}\}$ -cuts at all. We only make this distinction in order to be consistent with the definition of clutters in the literature. For all our purposes, ignoring non-minimal supports is no restriction: The next result, which is implicit in [49] (see also [150]), shows that the support-minimal nontrivial $\{0, \frac{1}{2}\}$ -cuts for $P = \{x : Ax \leq b\}$ are sufficient to describe $P_{\frac{1}{2}}(A, b)$.

Lemma 2.11 (see [49, 150]). *For a polyhedron $P = \{x \in \mathbb{R}^n : Ax \leq b\}$ with $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$, the $\{0, \frac{1}{2}\}$ -closure $P_{\frac{1}{2}}(A, b)$ is the projection of the following polyhedron onto the space of the x variables:*

$$\begin{aligned} Ax + s &= b \\ s(S) &\geq 1 \quad \text{for all } S \in \mathcal{F}(A, b) \\ s &\geq \mathbf{0} \end{aligned} \quad (2.14)$$

Proof. First, note that P is obtained from the polyhedron $\{(s, x) : Ax + s = b, s \geq \mathbf{0}\}$ by projecting out the slack variables s . We now argue that each nontrivial $\{0, \frac{1}{2}\}$ -cut for P (trivial cuts are redundant) may be expressed in the slack variables s only. To this end, let $u \in \{0, \frac{1}{2}\}^m$ such that $u^\top A \in \mathbb{Z}^n$ and $u^\top b \notin \mathbb{Z}$. Since $u^\top b$ is half-integral (i.e., $2u^\top b \in \mathbb{Z}$), we have that $\lfloor u^\top b \rfloor = u^\top b - \frac{1}{2}$. So multiplying the cut $u^\top A \leq \lfloor u^\top b \rfloor$ by 2, rearranging and substituting $s = b - Ax$ yields the equivalent inequality

$$2u^\top s \geq 1.$$

If $\text{supp}(u) \notin \mathcal{F}(A, b)$, then there exists some $v \in \{0, 1\}^m$ with $\text{supp}(v) \subsetneq \text{supp}(u)$ such that v also induces a nontrivial $\{0, \frac{1}{2}\}$ -cut for P . Since $v \leq u$ and $s \geq \mathbf{0}$, the inequality $2u^\top s \geq 1$ is dominated by $2v^\top s \geq 1$. We may therefore restrict to support-minimal multipliers u , i.e., those for which $2u = \chi(S)$ for some $S \in \mathcal{F}(A, b)$. \square

Lemma 2.11 allows for a reformulation of the membership and separation problems for the $\{0, \frac{1}{2}\}$ -closure, as observed by Caprara and Fischetti [49]. Given a rational point x^* in a polyhedron $P = \{x: Ax \leq b\}$, Lemma 2.11 implies that $x^* \in P_{\frac{1}{2}}(A, b)$ if and only if (s^*, x^*) is in the polyhedron defined by (2.14) where $s^* := b - Ax^* \geq \mathbf{0}$. So the membership problem amounts to deciding whether there is a member $S \in \mathcal{F}(A, b)$ of weight $s^*(S) < 1$, where the (nonnegative) weights are given by s^* . Finding such a member immediately gives a separating $\{0, \frac{1}{2}\}$ -cut. Similarly, a most violated $\{0, \frac{1}{2}\}$ -cut is induced by a member $S \in \mathcal{F}(A, b)$ that minimizes $s^*(S)$. To see this, note that for all $S, S^* \in \mathcal{F}(A, b)$, we have that $s^*(S^*) \leq s^*(S)$ if and only if

$$\frac{1}{2}\chi(S^*)^\top Ax^* - \left\lfloor \frac{1}{2}\chi(S^*)^\top b \right\rfloor \geq \frac{1}{2}\chi(S)^\top Ax^* - \left\lfloor \frac{1}{2}\chi(S)^\top b \right\rfloor.$$

Strictly speaking, this only means that $\{0, \frac{1}{2}\}$ -cuts induced by members of $\mathcal{F}(A, b)$ of minimum s^* -weight are most violated *support-minimal* cuts for P . However, support-minimality is no restriction here as $s^* \geq \mathbf{0}$.

By the definition of the clutters $\mathcal{F}(A, b)$, the membership problem for the $\{0, \frac{1}{2}\}$ -closure is therefore a special case of the following problem:

Given $Q \in \{0, 1\}^{p \times q}$, $d \in \{0, 1\}^p$ and a nonnegative weight vector $w \in \mathbb{Q}_{\geq 0}^q$,
decide whether the binary clutter \mathcal{F} over ground set $[q]$ given in the form
 $Qy \equiv d \pmod{2}$ (as in Proposition 2.10) has a member S of weight $w(S) < 1$.

Following [150], we call this problem the *binary clutter problem*. In fact, the binary clutter problem and the membership problem for the $\{0, \frac{1}{2}\}$ -closure are polynomially equivalent: Caprara and Fischetti [49] proved that every binary clutter is essentially of the form $\mathcal{F}(A, b)$ for integral A and b such that any given nonnegative weight vector can be obtained as the slack vector of a point in the polyhedron $\{x: Ax \leq b\}$. Moreover, A, b , and this point can be computed efficiently from a given instance of the binary clutter problem [49]. The stronger result of Letchford, Pokutta, and Schulz [152] mentioned in Section 2.1 follows from a more careful choice of A and b so that $Ax \leq b$ defines a polytope in the 0/1 cube.

It is well known that the binary clutter problem is strongly NP-complete; various hardness proofs have been given time and again, e.g., in [49, 152, 193] (see also [111, 122, 126]). Therefore, via the reduction of [49], also the membership problem for the $\{0, \frac{1}{2}\}$ -closure is strongly NP-complete. We remark that this reduction also implies that the separation problem for the Gomory-Chvátal closure, more generally, is strongly NP-hard, as observed by Eisenbrand [88].

On the positive side, suppose that for a rational polyhedron $P = \{x: Ax \leq b\}$, the binary clutter problem for $\mathcal{F}(A, b)$ and any given nonnegative weight vector w can be solved in polynomial time. Then, using standard reductions from combinatorial optimization (see, e.g., [125, 178]), one can construct an efficient separation routine

for the $\{0, \frac{1}{2}\}$ -closure $P_{\frac{1}{2}}(A, b)$ as follows. First, note that by scaling the weights, an algorithm for the binary clutter problem over $\mathcal{F}(A, b)$ can also decide whether $\min\{w(S) : S \in \mathcal{F}(A, b)\} < k$ for any given rational k . This already suffices to compute the minimum weight of any member of $\mathcal{F}(A, b)$ in polynomial time (by binary search) and to efficiently find a member $S \in \mathcal{F}(A, b)$ that attains the minimum (see, e.g., [178] for details). As explained above, this solves the separation problem for the $\{0, \frac{1}{2}\}$ -closure of P , for any given vector $x^* \in P$. Even stronger, one always obtains a most violated $\{0, \frac{1}{2}\}$ -cut for P in this way.

As mentioned in Section 2.1, the binary clutters associated with the fractional matching and stable set polytopes in Examples 2.1 and 2.2 belong to such a class of clutters for which the binary clutter problem is known to be solvable in polynomial time, for all nonnegative weights (see [49]). In the remainder of this section, we will look at each of the two examples in turn. As the associated binary clutters will play a central role in Section 2.3.4, we provide a combinatorial characterization of their members. The characterizations that we obtain are not new; they already appear in [49], although less explicitly and in a more general form. They also follow readily from the matroid perspective on binary clutters and using matroid duality (see, e.g., [60, 150, 160]). However, for the sake of clarity, we state and prove them explicitly in the language of binary clutters, using elementary combinatorial arguments.

Binary Clutters for Fractional Matching and Stable Set Polytopes

Let $G = (V, E)$ be a graph. The *edge-node incidence matrix* of G is the matrix $M_G \in \{0, 1\}^{E \times V}$ whose rows are the incidence vectors of the edges of G . In other words, the entry of M_G in row $e \in E$ and column $v \in V$ is 1 if $v \in e$ and 0 otherwise. The transpose of M_G is the *node-edge incidence matrix* of G .

With this definition, the fractional matching polytope of G (see Example 2.1) can equivalently be written as

$$\left\{ x \in \mathbb{R}^E : M_G^\top x \leq \mathbf{1}, x \geq \mathbf{0} \right\}. \quad (2.15)$$

Similarly, the fractional stable set polytope of G (see Example 2.2) is given by

$$\left\{ x \in \mathbb{R}^V : M_G x \leq \mathbf{1}, x \geq \mathbf{0} \right\}. \quad (2.16)$$

Both linear systems in (2.15) and (2.16) include nonnegativity constraints. Let us first derive a useful equivalent formulation of the binary clutters associated with linear systems of this form. More generally, we consider polyhedra of the form $\{x : Ax \leq b, x \geq \mathbf{0}\}$ where $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$. For the sake of brevity, we define $\mathcal{F}_+(A, b) := \mathcal{F}\left(\begin{pmatrix} A \\ -\mathbf{I} \end{pmatrix}, \begin{pmatrix} b \\ \mathbf{0} \end{pmatrix}\right)$, and label the ground set of $\mathcal{F}_+(A, b)$ by $[m] \dot{\cup} [n]$, the sets of rows and columns of A . Given a subset of rows $S \subseteq [m]$, we define $\text{odd}_A(S)$ as the support of $\sum_{i \in S} a^{(i)} \pmod{2}$, where $a^{(i)}$ denotes the i th row of A . Put differently, the set $\text{odd}_A(S) \subseteq [n]$ consists of all column indices $j \in [n]$ for which the number of rows in S that has an odd entry in column j is odd. With this notation, the members of $\mathcal{F}_+(A, b)$ can be easily described as follows.

Lemma 2.12. *For $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$, the members of $\mathcal{F}_+(A, b)$ are the minimal sets in $\{S \dot{\cup} \text{odd}_A(S) : S \subseteq [m], b(S) \text{ odd}\}$.*

Proof. By definition of the clutter $\mathcal{F}_+(A, b)$, its members correspond to the support-minimal solutions $(y, t) \in \{0, 1\}^m \times \{0, 1\}^n$ of the system of congruences

$$\begin{aligned} \bar{A}^\top y + t &\equiv \mathbf{0} \pmod{2} \\ \bar{b}^\top y &\equiv 1 \pmod{2} \end{aligned}$$

Since $\text{supp}(t) = \text{odd}_A(\text{supp}(y))$ for any such solution (y, t) , the statement follows. \square

Let us now apply Lemma 2.12 to the binary clutters associated with the fractional matching polytope of a graph G , in which case the matrix A is the node-edge incidence matrix M_G^\top . We will need the following graph terminology. A *cut* in a given graph $G = (V, E)$ is a subset of edges of the form $\delta(U)$ for some $U \subseteq V$. (We will also write $\delta_G(U)$ for $\delta(U)$ if the graph G is not clear from the context.) The sets U and $V \setminus U$ are the *shores* of the cut. For a given subset of nodes $T \subseteq V$ (the *terminals*), the cut $\delta(U)$ is called a *T-cut* if $|U \cap T|$ is odd. The *join* of two graphs G and H , denoted by $G + H$, is the graph on nodes $V(G) \dot{\cup} V(H)$ whose edge set consists of $E(G) \dot{\cup} E(H)$ and all edges between nodes of G and nodes of H (see [23]).

We are now ready to describe the members of $\mathcal{F}_+(M_G^\top, b)$, for any given integral vector b . As mentioned above, this description is a special case of a result of [49], and is implicitly used in [161]. Stated in matroid terms, it can also be found in [1, Section 9.2].

Proposition 2.13 (see [1, 49, 161]). *Let M_G be the edge-node incidence matrix of a graph $G = (V, E)$ and let $b \in \mathbb{Z}^V$. Then $\mathcal{F}_+(M_G^\top, b)$ is the clutter of minimal T -cuts in $G + K_1$ where*

$$T = \begin{cases} \{v \in V : b_v \text{ odd}\} \cup V(K_1) & \text{if } b(V) \text{ odd} \\ \{v \in V : b_v \text{ odd}\} & \text{otherwise} \end{cases}$$

Proof. Let us denote the unique node of K_1 by v^+ . For each node v of G , we identify the edge between v and v^+ with v . Thus, the edge set of $G + K_1$ is identified with $V \dot{\cup} E$, the ground set of the clutter $\mathcal{F}_+(M_G^\top, b)$. We claim that the T -cuts in $G + K_1$ are precisely the subsets of $V \dot{\cup} E$ of the form $S \dot{\cup} \text{odd}_{M_G^\top}(S)$ for some $S \subseteq V$ such that $b(S)$ is odd. The statement then follows immediately from Lemma 2.12.

To prove the claim, let $S \subseteq V$. Since the columns of M_G^\top are the incidence vectors of edges, $\text{odd}_{M_G^\top}(S)$ is the set of all edges of G with an odd number of endpoints in S . This implies that $\text{odd}_{M_G^\top}(S) = \delta_G(S)$. By our labelling of the edges of $G + K_1$, it follows that $S \dot{\cup} \text{odd}_{M_G^\top}(S) = S \dot{\cup} \delta_G(S) = \delta_{G+K_1}(S)$. This cut is a T -cut if and only if $b(S)$ is odd, as $|S \cap T| \equiv b(S) \pmod{2}$ by the definition of T .

Further note that for any T -cut in $G + K_1$, *both* its shores contain an odd number of nodes from T . To see this, recall from the definition of T that $v^+ \in T$ if and only if $b(V)$ is odd. Hence, for all $S \subseteq V$, we have that

$$|(S \cup \{v^+\}) \cap T| + |(V \setminus S) \cap T| \equiv b(V) + \underbrace{|S \cap T| + |(V \setminus S) \cap T|}_{=|V \cap T|} \equiv 0 \pmod{2}$$

Since $\delta_{G+K_1}(S \cup \{v^+\}) = \delta_{G+K_1}(V \setminus S)$, we may therefore assume that each T -cut in $G + K_1$ is induced by a subset of nodes not containing v^+ . This concludes the proof of the claim. \square

Padberg and Rao [161] gave a polynomial-time algorithm for finding a minimum-weight T -cut in a graph with terminals T and nonnegative edge weights (see also [151, 164]), and showed how to use this for separating the family of $\{0, \frac{1}{2}\}$ -cuts for the fractional matching polytope (or, more generally, the fractional b -matching polytope [79]). We note that their reduction of the separation problem to a minimum-weight T -cut problem relies on exactly the same graph construction as Proposition 2.13.

To characterize the clutters that arise for edge-node incidence matrices as in Example 2.2, we need one other notion. A *signed graph* is a tuple (G, Σ) where $G = (V, E)$ is an undirected graph, possibly with loops and parallel edges, and $\Sigma \subseteq E$ (the *signature*). An edge of G is *odd* if it is in Σ , and *even* otherwise. A subset of edges $F \subseteq E$ is called *odd* if F contains an odd number of odd edges, and *even* otherwise.

Proposition 2.14 (see [49]). *Let M_G be the edge-node incidence matrix of a graph $G = (V, E)$ and let $b \in \mathbb{Z}^E$. Then $\mathcal{F}_+(M_G, b)$ is the clutter of odd cycles in the signed graph $(G + K_1, \Sigma)$ where $\Sigma = \{e \in E : b_e \text{ odd}\}$.*

In fact, Proposition 2.14 is a special case of the following well-known observation, which may be proved using folklore arguments. For the sake of completeness, we provide a proof below.

Proposition 2.15 (see [49]). *Let M_G be the edge-node incidence matrix of a graph $G = (V, E)$ and let $b \in \mathbb{Z}^E$. Then $\mathcal{F}(M_G, b)$ is the clutter of odd cycles in the signed graph (G, Σ) where $\Sigma = \{e \in E : b_e \text{ odd}\}$.*

Using Proposition 2.15, we may prove Proposition 2.14 as follows.

Proof of Proposition 2.14. First, recall from its definition that, for any given integral A and b , the clutter $\mathcal{F}(A, b)$ does not depend on the actual entries of A and b but only on their parity (see [49]). This means that the clutter $\mathcal{F}_+(M_G, b)$, which is $\mathcal{F}(\begin{pmatrix} M_G \\ -\mathbf{I} \end{pmatrix}, \begin{pmatrix} b \\ \mathbf{0} \end{pmatrix})$ by definition, is the same as $\mathcal{F}(\begin{pmatrix} M_G \\ \mathbf{I} \end{pmatrix}, \begin{pmatrix} b \\ \mathbf{0} \end{pmatrix})$. Since the rows of M_G are the incidence vectors of the edges of G , the sum of all columns of the matrix $\begin{pmatrix} M_G \\ \mathbf{I} \end{pmatrix}$ is the vector $\begin{pmatrix} \mathbf{0} \\ \mathbf{1} \end{pmatrix}$. Next, observe that we may add this vector as a new column to $\begin{pmatrix} M_G \\ \mathbf{I} \end{pmatrix}$ without changing the set of solutions of the congruence system defining $\mathcal{F}(\begin{pmatrix} M_G \\ \mathbf{I} \end{pmatrix}, \begin{pmatrix} b \\ \mathbf{0} \end{pmatrix})$. Let us denote the matrix obtained in this way by M_G^+ . This matrix is, in fact, the edge-node incidence matrix of $G + K_1$, where the additional column corresponds to the unique node of K_1 . Since $\mathcal{F}_+(M_G, b) = \mathcal{F}(M_G^+, \begin{pmatrix} b \\ \mathbf{0} \end{pmatrix})$, Proposition 2.15 applied to $G + K_1$ and $\begin{pmatrix} b \\ \mathbf{0} \end{pmatrix}$ yields the statement. \square

It remains to prove Proposition 2.15.

Proof of Proposition 2.15. By the definition of the clutter $\mathcal{F}(M_G, b)$, the incidence vectors of its members are the minimal 0/1 solutions of the system of congruences

$$M_G^\top y \equiv \mathbf{0}, \quad b^\top y \equiv 1 \pmod{2}. \quad (2.17)$$

Let $y \in \{0, 1\}^E$ be a solution of (2.17). Since the rows of M_G^\top are the incidence vectors of all cuts of the form $\delta(v)$ for $v \in V$, the subgraph $(V, \text{supp}(y))$ is Eulerian, i.e., all nodes have even degree. Moreover, $b^\top y \equiv |\Sigma \cap \text{supp}(y)| \pmod{2}$. It follows that the

members of $\mathcal{F}(M_G, b)$ are the edge sets of minimal Eulerian subgraphs of G that contain an odd number of edges from Σ . Let us call such a subgraph *odd*. Using standard arguments (that are also used, e.g., in [114]), it can be shown that the minimal odd Eulerian subgraphs of (G, Σ) are precisely the odd cycles. For the sake of completeness, we include a brief proof of this fact.

It suffices to show that every odd Eulerian subgraph of a signed graph contains an odd cycle. Indeed, let H be an odd Eulerian subgraph. Since H is Eulerian, it contains a cycle C . If C is odd, we are done. Otherwise, we delete all edges of C from H . The resulting subgraph H' of H is Eulerian again, and since C is even by hypothesis, H' is odd (and thus nonempty). By induction on the number of edges, we know that H' contains an odd cycle, and therefore H contains an odd cycle. \square

Given a signed graph and nonnegative edge weights, one can find a shortest odd cycle in polynomial time using the algorithm described by Grötschel and Pulleyblank [122] or Gerards and Schrijver [114]. So the separation problem for the $\{0, \frac{1}{2}\}$ -closure can be solved efficiently also in this case.

Interestingly, to find a separating $\{0, \frac{1}{2}\}$ -cut for the fractional stable set polytope of a graph $G = (V, E)$, it suffices to consider all odd cycles in $G + K_1$ that do not contain the unique node of K_1 : By virtue of Lemma 2.11, it is easily verified that those cycles (as members of $\mathcal{F}_+(M_G, \mathbf{1})$) induce precisely the odd-cycle inequalities (2.5), which suffice to describe the $\{0, \frac{1}{2}\}$ -closure (see Example 2.2). This is true, more generally, for arbitrary integral right-hand sides $b \in \mathbb{Z}^E$ [114]. However, we know from Proposition 2.14 that the set of *all* support-minimal nontrivial $\{0, \frac{1}{2}\}$ -cuts is a superset of the odd-cycle inequalities (or their counterparts for the *fractional b -stable set polytope* for arbitrary integral b). Indeed, if we label the edges of $G + K_1$ in the same way as in the proof of Proposition 2.13, then a cycle in $G + K_1$ whose node set includes $V(K_1)$ corresponds to the edge set of a path in G plus its two endpoints. For a given vector $b \in \mathbb{Z}^E$ and Σ defined as in Proposition 2.14, we therefore obtain a $\{0, \frac{1}{2}\}$ -cut of the form

$$x(V_{\text{int}}(P)) \leq \frac{b(P) - 1}{2} \quad (2.18)$$

for each odd path $P \subseteq E$, where $V_{\text{int}}(P)$ denotes the set of internal nodes of P (i.e., excluding the two endpoints). The above discussion shows that these *odd-path inequalities* must be redundant for the $\{0, \frac{1}{2}\}$ -closure.

In fact, they are already redundant for the fractional b -stable set polytope. To see this, let P be a path in G for which $b(P)$ is odd. Now decompose P into two matchings P^+ and P^- . Since $b(P)$ is odd and $b(P) = b(P^+) + b(P^-)$, we must have that $b(P^+) \neq b(P^-)$; say, $b(P^+) > b(P^-)$. Note that each internal node of P is contained in exactly one edge from each matching P^\pm , as edges of P^\pm alternate along P . So if we take the sum of all edge inequalities $x_i + x_j \leq b_{ij}$ for edges $ij \in P^-$ and the nonnegativity constraint for each endpoint of P incident with an edge of P^- , we obtain the inequality

$$x(V_{\text{int}}(P)) \leq b(P^-).$$

Since $2b(P^-) < b(P)$ and b is integral, it follows that $b(P^-) \leq \frac{b(P)-1}{2}$. We have therefore derived an inequality that dominates the odd-path inequality (2.18).

Put differently, the fact that the odd-path inequalities are redundant for the fractional stable set polytope of a given graph G shows that the binary clutter problem over $\mathcal{F}_+(M_G, \mathbf{1})$, in fact, reduces to the binary clutter problem over $\mathcal{F}(M_G, \mathbf{1})$ – provided that the weights on the edges of $G + K_1$ are given by the slack of a point in the fractional stable set polytope.

Recall that in both special cases above, one can find *minimum-weight* members of the associated binary clutters in polynomial time. As explained in Section 2.1, this implies that most violated $\{0, \frac{1}{2}\}$ -cuts can be found efficiently – in particular in those cases in which the $\{0, \frac{1}{2}\}$ -closure is integral. In general, though, this is too strong a property to ask for (unless $P = NP$), as we will see in the next section.

2.3.2 On the Hardness of Finding Most Violated $\{0, 1/2\}$ -Cuts

Our goal in this section is to prove the following theorem.

Theorem 2.16. *Given $A \in \mathbb{Z}^{m \times n}$, $b \in \mathbb{Z}^m$, and $x^* \in \mathbb{Q}^n$ such that $x^* \in P = \{x \in \mathbb{R}^n : Ax \leq b\} \subseteq [0, 1]^n$, finding a $\{0, \frac{1}{2}\}$ -cut for P that is most violated by x^* is strongly NP-hard, even when $P_{\frac{1}{2}}(A, b) = P_I$ and $x^* \notin P_I$.*

Proof. We reduce from EXACT 3-COVER, which is known to be strongly NP-complete [107]:

Let $n \in \mathbb{N}$ and let \mathcal{S} be a collection of subsets of $[3n]$ with $|S| = 3$ for all $S \in \mathcal{S}$. Is there a subcollection $\mathcal{I} \subseteq \mathcal{S}$ such that $|\mathcal{I}| = n$ and $\bigcup_{S \in \mathcal{I}} S = [3n]$?

We may restrict to instances for which $\bigcup_{S \in \mathcal{S}} S = [3n]$, since otherwise the instance trivially is a “no” instance. Given such an instance of EXACT 3-COVER, let $P \subseteq \mathbb{R}^{3n}$ be the polyhedron defined by the following $3|\mathcal{S}| + 3n + 1$ inequalities:

$$\left. \begin{array}{l} x_i - x_j - x_k \leq 0 \\ -x_i + x_j - x_k \leq 0 \\ -x_i - x_j + x_k \leq 0 \end{array} \right\} \text{ for all } \{i, j, k\} \in \mathcal{S} \quad (2.19)$$

$$\mathbf{1}^\top x \leq 1 \quad (2.20)$$

$$x \geq \mathbf{0} \quad (2.21)$$

Constraints (2.20) and (2.21) imply that $P \subseteq [0, 1]^{3n}$. Further note that $\mathbf{0} \in P$ and therefore also $\mathbf{0} \in P_I$. Now consider an arbitrary element $l \in [3n]$. Since $\bigcup_{S \in \mathcal{S}} S = [3n]$ by hypothesis, l is contained in some set $S \in \mathcal{S}$. Among the three inequalities (2.19) for S , there is a unique one for which the coefficient of x_l is $+1$. The cut derived from this inequality, together with (2.20) and nonnegativity constraints $-x_i \leq 0$ (2.21) for all $i \notin S$, is the inequality $x_l \leq 0$. Since the choice of l was arbitrary, it follows that $P_{\frac{1}{2}} \subseteq P \cap \{x \in \mathbb{R}^{3n} : x \leq \mathbf{0}\} = \{\mathbf{0}\} \subseteq P_I$ and therefore $P_{\frac{1}{2}} = P_I = \{\mathbf{0}\}$. Here, we use $P_{\frac{1}{2}}$ to denote the $\{0, \frac{1}{2}\}$ -closure of P with respect to (2.19)–(2.21).

Let $x^* = \frac{1}{3n} \mathbf{1} \in \mathbb{R}^{3n}$. It is easily verified that $x^* \in P \setminus P_I$. Since (2.20) is the only inequality with an odd right-hand side, each nontrivial $\{0, \frac{1}{2}\}$ -cut for P must be supported in (2.20). Note that this inequality is tight for x^* , and the slack of x^* in all

other inequalities (2.19) and (2.21) is the same (namely, $\frac{1}{3n}$ each). Hence, the $\{0, \frac{1}{2}\}$ -cuts that are most violated by x^* are induced by the minimum-cardinality members of the binary clutter \mathcal{F} associated with $P_{\frac{1}{2}}$.

We will now argue that the minimum cardinality of a member of \mathcal{F} is $n + 1$ if and only if the given instance of EXACT 3-COVER is a “yes” instance. First note that for all $S \in \mathcal{S}$, any two of the three inequalities in (2.19) add up to an inequality with all-even coefficients. This means that no support-minimal $\{0, \frac{1}{2}\}$ -cut for P can be derived using more than one inequality from (2.19) for any fixed $S \in \mathcal{S}$. Hence, we may identify each member of \mathcal{F} with a tuple (\mathcal{I}, J) where $\mathcal{I} \subseteq \mathcal{S}$, $J \subseteq [3n]$ and, by Lemma 2.12, J consists precisely of those elements of $[3n]$ that are covered an even number of times by the sets in \mathcal{I} . Note here that we require an even parity because (2.20) is always included in any cut derivation. In particular, the cardinality of a member of \mathcal{F} specified by (\mathcal{I}, J) is $|\mathcal{I}| + |J| + 1$. We claim that this quantity is at least $n + 1$ with equality if and only if \mathcal{I} is an exact cover of $[3n]$. Indeed, for any member of \mathcal{F} given by (\mathcal{I}, J) , the sets in \mathcal{I} and the singletons $\{i\}$ for all $i \in J$ define a cover of $[3n]$. Since any cover of $[3n]$ by sets of cardinality at most 3 must have cardinality at least n , it follows that $|\mathcal{I}| + |J| \geq n$ with equality if and only if $J = \emptyset$ and \mathcal{I} is an exact cover of $[3n]$, as desired. \square

We would like to point out that a very similar reduction from EXACT 3-COVER was given in [193] to prove that finding minimum-weight members of binary clutters is NP-hard. (To be precise, the result of [193] concerns the matroid analogue: finding a minimum-weight circuit of a binary matroid that contains a fixed element of the ground set; see Section 2.3.1.) While we arrived at the statement and proof above independently, we note that our hardness result is stronger than that of [193]. Indeed, Theorem 2.16 implies that finding minimum-weight members is NP-hard even for instances of the binary clutter problem in which the weights are given by the slack vector of a point inside a polyhedron whose $\{0, \frac{1}{2}\}$ -closure is integral.

Here, we must stress that Theorem 2.16 does *not* imply that the membership problem for the $\{0, \frac{1}{2}\}$ -closure is NP-hard even when the $\{0, \frac{1}{2}\}$ -closure is integral, via scaling the weights and binary search as discussed in Section 2.3.1. Scaling the weights cannot be done here without changing the point x^* that is provided as part of the input. It can even be shown that for polytopes, no two distinct slack vectors are scalings of one another. For the particular polytopes and weights used in the proof of Theorem 2.16, this is easy to see: The only point which satisfies (2.20) at equality and whose slack is uniformly nonzero elsewhere is the point x^* from the reduction.

In fact, solving the binary clutter problem for the instances in the proof of Theorem 2.16, i.e., finding *any* member of the associated clutter of weight less than 1, is trivial. Namely, we saw above that for any $l \in [3n]$, the inequality $x_l \leq 0$ is a $\{0, \frac{1}{2}\}$ -cut that separates x^* from $P_{\frac{1}{2}}$. The weight of the corresponding member of the associated binary clutter is equal to $\frac{3n-2}{3n} < 1$.

Further note that in our reduction in the proof of Theorem 2.16, there is a nontrivial $\{0, \frac{1}{2}\}$ -cut with support of cardinality $n + 1$ (which is minimum) if and only if there is a solution to EXACT 3-COVER. This implies that finding a nontrivial $\{0, \frac{1}{2}\}$ -cut of minimum support is strongly NP-hard even when the $\{0, \frac{1}{2}\}$ -closure coincides with the integer hull, strengthening a result of Eisenbrand [89, Proposition 5.8].

Corollary 2.17. *Given $A \in \mathbb{Z}^{m \times n}$, $b \in \mathbb{Z}^m$, and $k \in \mathbb{N}$, deciding whether there is a nontrivial $\{0, \frac{1}{2}\}$ -cut for $P = \{x \in \mathbb{R}^n : Ax \leq b\}$ with support of cardinality at most k is strongly NP-complete, even when $P_{\frac{1}{2}}(A, b) = P_I$ and $P \subseteq [0, 1]^n$.*

Given the hardness result of Theorem 2.16, a natural question is to ask whether there is a class of binary clutters for which the binary clutter problem can be solved in polynomial time (for all nonnegative weights), and which contains clutters associated with a relevant family of polyhedra whose $\{0, \frac{1}{2}\}$ -closure is integral. In fact, we already saw such a class of clutters in Section 2.3.1, namely those associated with T -cuts and odd cycles in (signed) graphs. Can such binary clutters be described structurally? In the next two sections, we will be investigating two structural properties of clutters that can be regarded as “integrality” properties. We will show that they are closely related to integrality and total dual integrality of the $\{0, \frac{1}{2}\}$ -closure of set packing polyhedra, which generalize fractional matching and stable set polytopes.

2.3.3 Set Packing, Idealness, and the MFMC Property

Consider a polyhedron of the form $\{x : Ax \leq \mathbf{1}, x \geq \mathbf{0}\}$ where A is a 0/1 matrix. We call such a polyhedron a *set packing polyhedron*. (In fact, if A has no column of all zeros, then we can even speak of a set packing *polytope*.) Similarly, any polyhedron of the form $\{x : Ax \geq \mathbf{1}, x \geq \mathbf{0}\}$ with a 0/1 matrix A is a *set covering polyhedron*.

Let \mathcal{F} be a clutter over ground set E . We say that \mathcal{F} is *ideal* if the following set covering polyhedron $Q(\mathcal{F}) \subseteq \mathbb{R}^E$ associated with \mathcal{F} is integral [63]:

$$\begin{aligned} z(S) &\geq 1 \quad \text{for all } S \in \mathcal{F} \\ z &\geq \mathbf{0} \end{aligned} \tag{2.22}$$

If the linear system (2.22) is TDI, the clutter \mathcal{F} is said to have the *max-flow min-cut property* (or *MFMC property* for short) [183]. By Proposition 2.6, the MFMC property implies idealness. The two properties are also known under different names: Ideal clutters are sometimes said to have the *weak MFMC property* [183], and the MFMC property is also referred to as the *Mengerian property*.

In this section, we will see that if the clutter $\mathcal{F}_+(A, \mathbf{1})$ associated with the $\{0, \frac{1}{2}\}$ -closure of a set packing polyhedron $\{x : Ax \leq \mathbf{1}, x \geq \mathbf{0}\}$ is ideal, then the $\{0, \frac{1}{2}\}$ -closure coincides with the integer hull. Even stronger, we will prove that the linear description of the $\{0, \frac{1}{2}\}$ -closure is TDI if $\mathcal{F}_+(A, \mathbf{1})$ has the MFMC property.

The observation that idealness implies integrality of the $\{0, \frac{1}{2}\}$ -closure is implicit in [150]; we give a direct proof here.

Proposition 2.18 (see [150]). *Let $P = \{x : Ax \leq \mathbf{1}, x \geq \mathbf{0}\}$ for a 0/1 matrix A . If the clutter $\mathcal{F}_+(A, \mathbf{1})$ is ideal, then the $\{0, \frac{1}{2}\}$ -closure of P coincides with P_I .*

Proof. Let $P_{\frac{1}{2}}$ denote the $\{0, \frac{1}{2}\}$ -closure of P with respect to the system $Ax \leq \mathbf{1}, x \geq \mathbf{0}$. By Lemma 2.11, $P_{\frac{1}{2}}$ is obtained from the polyhedron

$$\begin{aligned} Ax + s &= \mathbf{1} \\ (s, x) &\in Q(\mathcal{F}_+(A, \mathbf{1})) \end{aligned} \tag{2.23}$$

by projecting out the slack variables s . (Note here that we did not introduce slack variables for the nonnegativity constraints $x \geq \mathbf{0}$ since they can be identified with x .) By Lemma 2.12, the support of each row of the matrix $(\mathbf{I} | A)$ is a member of $\mathcal{F}_+(A, \mathbf{1})$. In particular, this implies that the inequalities $Ax + s \geq \mathbf{1}$ are valid for $Q(\mathcal{F}_+(A, \mathbf{1}))$. Hence, the polyhedron given by (2.23) is a face of $Q(\mathcal{F}_+(A, \mathbf{1}))$. Now suppose that the clutter $\mathcal{F}_+(A, \mathbf{1})$ is ideal. Then the polyhedron $Q(\mathcal{F}_+(A, \mathbf{1}))$ is integral, and so is every face of $Q(\mathcal{F}_+(A, \mathbf{1}))$. As orthogonal projections onto coordinate subspaces preserve integrality, the statement follows. \square

As mentioned above, Proposition 2.18 can also be derived from a result of Letchford [150] as follows. The starting point of [150] is the following observation. For $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$, every integral point in the polyhedron

$$P = \{(s, x) \in \mathbb{R}^m \times \mathbb{R}^n : s \geq \mathbf{0}, x \geq \mathbf{0}, Ax + s = b\} \quad (2.24)$$

satisfies the system of congruences $Ax + s \equiv b \pmod{2}$. Hence, the polyhedron

$$\text{conv}\{(s, x) \in \mathbb{Z}^m \times \mathbb{Z}^n : s \geq \mathbf{0}, x \geq \mathbf{0}, Ax + s \equiv b \pmod{2}\} \quad (2.25)$$

is a relaxation of the integer hull P_I , called the *binary group relaxation* in [150]. By Proposition 2.10, the minimal 0/1 solutions of $\bar{A}x + s \equiv \bar{b} \pmod{2}$ induce a binary clutter. In fact, this can be shown to be the blocker of $\mathcal{F}_+(A, b)$, as mentioned in [150]. To see this, note that for all integral vectors $(y, t), (s, x) \in \mathbb{Z}^m \times \mathbb{Z}^n$ with

$$\begin{aligned} A^\top y + t &\equiv \mathbf{0}, \quad b^\top y \equiv 1 \pmod{2} \\ \text{and } Ax + s &\equiv b \pmod{2}, \end{aligned}$$

we have that

$$y^\top s + x^\top t \equiv y^\top (b + Ax) + x^\top t = \underbrace{y^\top b}_{\equiv 1} + x^\top \underbrace{(A^\top y + t)}_{\equiv \mathbf{0}} \equiv 1 \pmod{2}.$$

Since this holds true in particular for all 0/1 vectors (y, t) and (s, x) , we thus obtain

Proposition 2.19 (see [150]). *Given $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$, the blocker $b(\mathcal{F}_+(A, b))$ is the collection of all minimal sets in*

$$\{\text{supp}(s) \dot{\cup} \text{supp}(x) : s \in \{0, 1\}^m, x \in \{0, 1\}^n, \bar{A}x + s \equiv \bar{b} \pmod{2}\}.$$

Letchford went on by showing that, first, if $b(\mathcal{F}_+(A, b))$ is ideal, then the binary group relaxation (2.25) coincides with $Q(\mathcal{F}_+(A, b))$. Second, he observed that for set packing systems with a 0/1 matrix A and $b = \mathbf{1}$, the integer hull of P in (2.24) is the face of (2.25) defined by $Ax + s = \mathbf{1}$. Combining these two results for the set packing case, it follows that the polyhedron (2.23) is integral if $b(\mathcal{F}_+(A, \mathbf{1}))$ is ideal. Since a clutter is ideal if and only if its blocker is [149], the statement of Proposition 2.18 follows after projecting out the slack variables s .

Let us remark here that in our terms, the binary group relaxation (2.25) introduced in [150] is, in fact, the integer hull of $Q(\mathcal{F}_+(A, b))$. (This directly implies Letchford's first observation above.) Indeed, if an integral vector (s, x) satisfies $Ax + s \equiv b \pmod{2}$,

so do $(s + 2e_i, x)$ and $(s, x + 2e_j)$ for every $i \in [m]$ and $j \in [n]$. This means that the vertices of (2.25) must be the minimal 0/1 solutions of $Ax + s \equiv b \pmod{2}$, which are the incidence vectors of the members of $b(\mathcal{F}_+(A, b))$ by Proposition 2.19. The next folklore result states that these are precisely the integral vertices of $Q(\mathcal{F}_+(A, b))$ (see, e.g., Proposition 1.24 in [1] or Remark 1.16 in [60]).

Proposition 2.20 (see [1, 60]). *For every clutter \mathcal{F} , the integral vertices of $Q(\mathcal{F})$ are precisely the incidence vectors of the members of $b(\mathcal{F})$.*

Since the recession cone of both the binary group relaxation (2.25) and $Q(\mathcal{F}_+(A, b))_I$ is the nonnegative orthant, this implies that (2.25) and $Q(\mathcal{F}_+(A, b))_I$ coincide.

We next consider the case that the binary clutter associated with the $\{0, \frac{1}{2}\}$ -closure of a set packing polyhedron is not only ideal but also has the MFMC property. In this case, one may strengthen Proposition 2.18 and show that adding all $\{0, \frac{1}{2}\}$ -cuts yields a TDI system.

Theorem 2.21. *Let $P = \{x: Ax \leq \mathbf{1}, x \geq \mathbf{0}\}$ for a 0/1 matrix A . If the clutter $\mathcal{F}_+(A, \mathbf{1})$ has the MFMC property, then the linear system given by $Ax \leq \mathbf{1}, x \geq \mathbf{0}$ and all support-minimal nontrivial $\{0, \frac{1}{2}\}$ -cuts for P is TDI.*

The key ingredient to prove this result is the observation made in the proof of Proposition 2.18: If we rewrite the inequalities $Ax \leq \mathbf{1}$ using slack variables, then the resulting polyhedron is a face of $Q(\mathcal{F}_+(A, \mathbf{1}))$. Recall that this was a consequence of Lemmas 2.11 and 2.12. However, total dual integrality is a property of linear systems and not of polyhedra, so we need to analyze the transformations between the given linear systems carefully. The following well-known facts will be useful.

Proposition 2.22 ([59], see also [174]). *Let A, B, B' be rational matrices; a, b, d, d' rational vectors of appropriate dimension; and $\beta \in \mathbb{Q}$.*

- (i) *If $a^\top x \leq \beta, Bx \leq d$ is a TDI system, then $a^\top x = \beta, Bx \leq d$ is TDI again.*
- (ii) *Suppose that each inequality in $B'x \leq d'$ is a nonnegative integer linear combination of inequalities from $Bx \leq d$. If $B'x \leq d'$ is TDI, so is $Bx \leq d$.*
- (iii) *If $Ax + s = b, Bx \leq d, s \geq \mathbf{0}$ is TDI, so is $Ax \leq b, Bx \leq d$.*

We are now ready to prove Theorem 2.21.

Proof of Theorem 2.21. Suppose that $\mathcal{F}_+(A, \mathbf{1})$ has the MFMC property. Then the linear system

$$\begin{aligned} Ax + s &= \mathbf{1} \\ s(S) + x(T) &\geq 1 \quad \text{for all } S \dot{\cup} T \in \mathcal{F}_+(A, \mathbf{1}) \\ s, x &\geq \mathbf{0} \end{aligned} \tag{2.26}$$

defines a face of $Q(\mathcal{F}_+(A, \mathbf{1}))$ (see the proof of Proposition 2.18) and is therefore TDI by Proposition 2.22(i). Recall from Section 2.3.1 that for each member $S \dot{\cup} T \in \mathcal{F}_+(A, \mathbf{1})$, the corresponding support-minimal nontrivial $\{0, \frac{1}{2}\}$ -cut for P is given by

$$u^\top Ax + v^\top x \leq u^\top \mathbf{1} - \frac{1}{2} \tag{2.27}$$

where $u = \frac{1}{2}\chi(S)$ and $v = \frac{1}{2}\chi(T)$. Now we multiply (2.27) by 2 and subtract the equation $2u^\top(Ax + s) = 2u^\top \mathbf{1}$. The resulting inequality is

$$-2u^\top s - 2v^\top x \leq -1$$

By definition of u and v , this is precisely the set covering inequality in (2.26) for the member $S \dot{\cup} T$. Further note that $2u$ is 0/1 vector, which means that we obtained this inequality as a nonnegative integer linear combination of the $\{0, \frac{1}{2}\}$ -cut (2.27) and the equations $Ax + s = \mathbf{1}$. Thus, it follows from Proposition 2.22(ii) that also the following system is TDI:

$$\begin{aligned} Ax + s &= \mathbf{1} \\ \frac{1}{2}\chi(S)^\top Ax + \frac{1}{2}\chi(T)^\top x &\leq \frac{1}{2}\chi(S)^\top \mathbf{1} - \frac{1}{2} \quad \text{for all } S \dot{\cup} T \in \mathcal{F}_+(A, \mathbf{1}) \\ s, x &\geq \mathbf{0} \end{aligned} \quad (2.28)$$

Finally, the linear description of the $\{0, \frac{1}{2}\}$ -closure of P is obtained from the system (2.28) by projecting out the slack variables s . This operation preserves total dual integrality by Proposition 2.22(iii). \square

2.3.4 Implications for Fractional Matching and Stable Set Polytopes

In the previous section, we discussed two integrality properties of clutters: idealness and the MFMC property. In this section, we will look at each of the two properties in turn for the particular clutters associated with the fractional matching and stable set polytopes in Examples 2.1 and 2.2. Recall that both families of polyhedra are of set packing type. As we will see, idealness and the MFMC property are in these cases closely related to two well-known properties of integral matrices, the Edmonds-Johnson property and total unimodularity.

Idealness and the Edmonds-Johnson Property

We begin with the binary clutters associated with fractional matching polytopes. Recall from Section 2.3.1 that they are of the form $\mathcal{F}_+(M_G^\top, \mathbf{1})$, where M_G^\top is the node-edge incidence matrix of a graph G . By Proposition 2.13, each such clutter is the clutter of minimal T -cuts in a certain graph with terminals T . Edmonds and Johnson [86] proved that all clutters of minimal T -cuts are ideal.

Proposition 2.23 ([86]). *The clutter of minimal T -cuts in a graph $G = (V, E)$ is ideal for every $T \subseteq V$.*

It thus follows from Proposition 2.18 that the $\{0, \frac{1}{2}\}$ -closure of the fractional matching polytope is integral, providing yet another proof of Edmonds' characterization of the matching polytope [79].

We next consider the clutters associated with fractional stable set polytopes. In this case, Proposition 2.18 states that graphs G for which $\mathcal{F}_+(M_G, \mathbf{1})$ is ideal are t -perfect. We already saw in Example 2.2 that not all graphs are t -perfect; K_4 is not. Hence, we know by Proposition 2.18 that $\mathcal{F}_+(M_{K_4}, \mathbf{1})$ cannot be ideal. This can also be seen directly as follows.

By Proposition 2.14, the clutter $\mathcal{F}_+(M_{K_4}, \mathbf{1})$ is the clutter of odd cycles in $(K_4 + K_1, E(K_4))$. Since cuts and cycles always intersect in an even number of edges, an odd cycle $C \subseteq E(G)$ in a signed graph (G, Σ) remains odd after replacing the signature Σ with $\Sigma \Delta \delta(U)$ for some $U \subseteq V$. This operation is called a *signature exchange*. Thus, signature exchanges leave the clutter of odd cycles invariant. This means that the clutters of odd cycles of $(K_4 + K_1, E(K_4))$ and of $(K_4 + K_1, E(K_4 + K_1))$ are identical, since we may obtain one signed graph from the other by performing a signature exchange given by the cut $\delta_{K_4+K_1}(V(K_4))$. Moreover, $K_4 + K_1$ and K_5 are isomorphic by definition of the graph join operator. So after relabelling the ground set, $\mathcal{F}_+(M_{K_4}, \mathbf{1})$ is precisely the clutter of odd cycles of $(K_5, E(K_5))$. For all $n \in \mathbb{N}$, let us call the signed graph $(K_n, E(K_n))$ an *odd- K_n* for short. It is a well-known fact that the clutter of odd cycles of odd- K_5 , which we denote by \mathcal{O}_5 , is non-ideal (see, e.g., [60, 126, 183]). For the sake of completeness, we include a short proof that follows [111].

Proposition 2.24 (see [60, 126, 183]). *\mathcal{O}_5 is non-ideal.*

Proof. By the definition of idealness, it suffices to show that the set covering polyhedron $Q(\mathcal{O}_5)$ as defined in (2.22) is not integral. As the length of any odd cycle in odd- K_5 is at least 3, the vector $z^* = \frac{1}{3}\mathbf{1} \in \mathbb{R}^{E(K_5)}$ is in $Q(\mathcal{O}_5)$. One may verify by a direct computation that z^* is a (fractional) vertex of $Q(\mathcal{O}_5)$. An alternative, computation-free argument is as follows.

Since $z^* \in Q(\mathcal{O}_5)$, the minimum of the linear objective function $\mathbf{1}^\top z$ over $Q(\mathcal{O}_5)$ is at most $\mathbf{1}^\top z^* = \frac{10}{3}$. By Proposition 2.20, the minimum value of $\mathbf{1}^\top z$ across all *integral* points $z \in Q(\mathcal{O}_5)$ is the minimum cardinality of a cover of \mathcal{O}_5 . Any cover of \mathcal{O}_5 must intersect each of the 10 cycles in K_5 of length 3. Since each edge of K_5 is contained in exactly three of those 10 cycles, any cover must therefore contain at least $\lceil \frac{10}{3} \rceil = 4 > \frac{10}{3}$ edges by the pigeonhole principle. Hence, $Q(\mathcal{O}_5)$ is not an integral polyhedron, meaning that \mathcal{O}_5 is non-ideal. \square




Figure 2.1

While graphs G for which $\mathcal{F}_+(M_G, \mathbf{1})$ is ideal are t -perfect by Proposition 2.18, the converse is false: The graph G in Figure 2.1 is t -perfect [114, 115] but the clutter of odd cycles of $(G + K_1, E(G))$ can be shown to be non-ideal, as we will see later. Our goal for the remainder of this section is to derive a converse to Proposition 2.18 for the special case that A is the (edge-node or node-edge) incidence matrix of a graph. What property of A is equivalent with idealness of $\mathcal{F}_+(A, \mathbf{1})$ in this case? We will see that this property is the so-called *Edmonds-Johnson property*. An integral matrix A has the

Edmonds-Johnson property (or *EJ property* for short) if, for all integral vectors l, u, b, d of appropriate dimension, the Gomory-Chvátal closure of the polyhedron

$$\begin{aligned} d &\leq Ax \leq b \\ l &\leq x \leq u \end{aligned} \tag{2.29}$$

coincides with the integer hull. Matrices with the EJ property are also said to have *strong Chvátal rank* at most 1 in the literature (see [174]). With this definition, we may state our result as follows.

Theorem 2.25. *Let A be the node-edge or edge-node incidence matrix of a graph. Then $\mathcal{F}_+(A, \mathbf{1})$ is ideal if and only if A has the EJ property.*

Note that when A is the node-edge or edge-node incidence matrix of a graph, Theorem 2.25 indeed implies Proposition 2.18. This is because node-edge incidence matrices (and their transposes) are *2-regular* (or *totally half-modular*), that is, each nonsingular square submatrix has a half-integral inverse (see, e.g., [4, 114, 128]). It follows from Theorem 20 in [4] that for polyhedra $P = \{x: Ax \leq \mathbf{1}, x \geq \mathbf{0}\}$ with an integral 2-regular matrix A , the $\{0, \frac{1}{2}\}$ -closure $P_{\frac{1}{2}}$ and the Gomory-Chvátal closure P' coincide. Notice that P can be expressed in the form (2.29) by taking $l = \mathbf{0}, u = \mathbf{1}$ and $d = \mathbf{0}, b = \mathbf{1}$ (in appropriate dimensions). Hence, if A has the EJ property, then $P_{\frac{1}{2}} = P' = P_I$.

The EJ property is named after Edmonds and Johnson, who established the property for all *node-edge* incidence matrices of graphs [85, 86]. Combining this with Propositions 2.13 and 2.23, we therefore only need to prove Theorem 2.25 for *edge-node* incidence matrices.

To this end, take a graph G and, as above, let M_G denote its edge-node incidence matrix. We already know from Proposition 2.18 that the clutter $\mathcal{F}_+(M_G, \mathbf{1})$ is non-ideal if G is not t -perfect – in which case M_G cannot have the EJ property either since the EJ property implies t -perfection. For example, as K_4 is not t -perfect, its edge-node incidence matrix

$$M_{K_4} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{pmatrix} \tag{2.30}$$

does not have the EJ property. Recall that the associated clutter $\mathcal{F}_+(M_{K_4}, \mathbf{1})$ is isomorphic to \mathcal{O}_5 , which is non-ideal by Proposition 2.24. In fact, M_{K_4} and \mathcal{O}_5 are the only minimal such counterexamples, as proved by Gerards and Schrijver [114] and Guenin [126], respectively. By “minimal”, we mean the following. For both incidence matrices and clutters of odd cycles, there are certain operations that produce smaller matrices or clutters while maintaining the EJ property or idealness, respectively. So to characterize these properties, it suffices to list all smallest incidence matrices and clutters of odd cycles without the respective property. Showing that this list consists of just one counterexample each is the achievement of [114, 126].

In order to prove Theorem 2.25, we will argue that the two sets of operations that preserve the EJ property and idealness, respectively, can be simulated by each other

in such a way that they essentially commute with constructing the clutter $\mathcal{F}_+(M_G, \mathbf{1})$ from an edge-node incidence matrix M_G . To make this more precise, we need some further terminology.

In addition to signature exchanges, which we discussed earlier in this section, there are two other operations that can be performed on a signed graph (G, Σ) to obtain another signed graph. First, we can *delete* an edge $e \in E(G)$, which means removing it from G and replacing Σ with $\Sigma \setminus \{e\}$. Secondly, we can *contract* an edge e by first performing a signature exchange if necessary so that $e \notin \Sigma$, then removing e from G and identifying its two endpoints. See Figure 2.2 for an example of the contraction operation. A signed graph obtained from (G, Σ) by performing a sequence of deletions, contractions, and signature exchanges is called a (*signed*) *minor* of (G, Σ) . The order in which the operations are applied does not matter up to signature exchanges (see [60]).




Figure 2.2: Starting from the graph in Figure 2.1 with all edges odd (a), we perform a signature exchange using the cut with shore U and obtain the signed graph in (b). Solid edges are odd, dashed edges are even. Contracting the four even edges results in K_4 with all edges odd (c).

Recall that doing a signature exchange leaves the clutter of odd cycles of a signed graph (G, Σ) invariant. The effect of the other two operations on the clutter of its odd cycles is as follows. The clutter of odd cycles of a minor of (G, Σ) obtained after deleting an edge e consists of all odd cycles in (G, Σ) that do not contain e . When contracting e , however, we may not simply delete e from each odd cycle in (G, Σ) since the resulting sets of edges may not be cycles in the graph obtained from G after contracting e . This happens, for example, when we contract a chord of a cycle, i.e., a non-loop edge that is not part of the cycle but both its endpoints are on the cycle. To fix this, we first delete e from each odd cycle in (G, Σ) and then drop all sets which are not cycles anymore.

The following simple observation is an immediate consequence and will be useful later.

Remark 2.26. *Let C be an odd cycle in a signed graph (G, Σ) and let (G', Σ') be a minor of (G, Σ) . If C is a cycle in G' , then C is odd in (G', Σ') .*

The effect of taking minors of signed graphs on the clutter of their odd cycles can be equivalently described in polyhedral terms: For the set covering polyhedron (2.22) associated with the clutter of odd cycles, contracting an edge e corresponds to taking the face defined by $z_e = 0$. Deleting e means projecting the polyhedron onto the

coordinate hyperplane defined by $z_e = 0$ (see, e.g., [1]). Both polyhedral operations preserve integrality. This means that if the clutter of odd cycles of a signed graph is ideal, then the clutter of odd cycles of any minor is ideal, too [183]. In particular, if a signed graph has an odd- K_5 minor, the clutter of its odd cycles is non-ideal. Guenin [126] proved that the converse is also true:

Proposition 2.27 ([126], see also [176]). *The clutter of odd cycles of a signed graph (G, Σ) is ideal if and only if (G, Σ) has no odd- K_5 minor.*

This will be the first ingredient of our proof of Theorem 2.25. The second ingredient is a result of Gerards and Schrijver [114] that provides a characterization of incidence matrices with the EJ property. Their result is of similar flavour as Guenin's result above. As with signed minor operations and idealness, the EJ property is preserved under several matrix operations, as observed in [114]:

- (i) permuting rows or columns,
- (ii) deleting a row or column,
- (iii) multiplying a row or column by -1 ,
- (iv) replacing the matrix $\begin{pmatrix} 1 & g^\top \\ f & D \end{pmatrix}$ with $D - fg^\top$.

Proposition 2.28 ([114]). *The edge-node incidence matrix of a graph has the EJ property if and only if it cannot be transformed to M_{K_4} by a sequence of operations (i) to (iv).*

Notice that the class of edge-node incidence matrices of graphs is not closed under the operations (i) to (iv). In fact, Gerards and Schrijver's original result in [114] is stated in terms of *bidirected graphs*. These can be thought of as simultaneous generalizations of directed, undirected, and signed graphs. Every integral matrix for which the sum of the absolute values of the entries in each row is at most 2 is the edge-node incidence matrix of a bidirected graph. Here, a row with a single nonzero entry defines a loop. Rows with two nonzeros define edges with distinct endpoints; each endpoint receives a sign from $\{+, -\}$ depending on the sign of the corresponding entry in the incidence matrix. We will not go into detail here but refer the reader to Chapter 68 of [177] or [74] for further details and more background on bidirected graphs.

For our purposes, we only need the following fact (see [177]): Each bidirected graph has an underlying signed graph whose odd edges are those edges with endpoints of the same sign. Loops corresponding to rows of the incidence matrix with a single ± 1 entry may be ignored. Then each of the operations (i) to (iv) corresponds to a minor operation on the underlying signed graph (see also [114]): Clearly, (i) has no effect on the underlying signed graph; neither has multiplying rows by -1 in (iii). Multiplying a column corresponding to node v by -1 is a signature exchange with the cut $\delta(v)$. Operation (iv) is an edge contraction, possibly with one included signature exchange if the contracted edge is odd. For example, performing (iv) on the top left entry of M_{K_4}

as given in (2.30) yields the matrix

$$\begin{pmatrix} -1 & 1 & 0 \\ -1 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}$$

This is the incidence matrix of a bidirected graph. The underlying signed graph is obtained from odd- K_4 after contracting the edge corresponding to the first row of M_{K_4} (see Figure 2.3). Finally, (ii) corresponds to an edge or node deletion. Here, deleting a node v means deleting all edges incident with v and then removing v from the graph.

Note that we did not discuss node deletions when we introduced signed graph minors earlier in this section. This is motivated by the observation that any clutter whose ground set is the set of edges of a (signed) graph is unaffected by the presence of isolated nodes, and we may assume that, possibly after a sequence of edge deletions, any deleted node is an isolated node indeed.




Figure 2.3: Illustration of the contraction operation (iv): The underlying signed graph of K_4 , viewed as a bidirected graph, is odd- K_4 (a). Nodes are labelled by which column of M_{K_4} as given in (2.30) they correspond to; solid edges are odd, dashed edges are even. After the signature exchange with $\delta(1)$, we obtain the signed graph in (b). Contracting edge 12 yields (c).

The above graph-theoretic interpretation of operations (i) to (iv) in terms of signed graph minors suggests the following equivalent version of Proposition 2.28 given in [177].

Corollary 2.29 (Corollary 68.6b in [177], see also Corollary 1 in [114]). *Let M_G be the edge-node incidence matrix of a graph G . Then M_G has the EJ property if and only if $(G, E(G))$ has no odd- K_4 minor (possibly after node deletions).*

We now make the following key observation that links Corollary 2.29 with Proposition 2.27.

Lemma 2.30. *Let G be a graph, possibly with loops and parallel edges. For all $n \in \mathbb{N}$, the signed graph $(G, E(G))$ has an odd- K_n minor (possibly after node deletions) if and only if $(G + K_1, E(G + K_1))$ has an odd- K_{n+1} minor.*

Using this key lemma, we may readily prove Theorem 2.25.

Proof of Theorem 2.25. As remarked above, we may suppose that A is the edge-node incidence matrix M_G of a graph G . By Proposition 2.14, the clutter $\mathcal{F}_+(M_G, \mathbf{1})$ is the clutter of odd cycles of $(G + K_1, E(G + K_1))$, up to the signature exchange with $\delta(V(K_1))$. We know from Proposition 2.27 that this clutter is ideal if and only if $(G + K_1, E(G + K_1))$ has no odd- K_5 minor. This, in turn, is equivalent with $(G, E(G))$ not having an odd- K_4 minor by Lemma 2.30. Finally, it follows from Corollary 2.29 that $(G, E(G))$ has no odd- K_4 minor if and only if M_G has the EJ property. \square

For example, we saw in Figure 2.2 that the signed graph $(G, E(G))$ in Figure 2.1 has an odd- K_4 minor. Thus, by Theorem 2.25, the clutter of odd cycles of the signed graph $(G + K_1, E(G + K_1))$ is non-ideal, which means that $(G + K_1, E(G + K_1))$ must have an odd- K_5 minor (by Proposition 2.27). This is easily verified: In fact, one can use the same sequence of edge deletions and contractions that turned $(G, E(G))$ into an odd- K_4 in order to produce a minor of $(G + K_1, E(G + K_1))$ that is “almost” an odd- K_5 , up to edge deletions (see Figure 2.4).




Figure 2.4: The signed graph in (a) is $(G + K_1, E(G + K_1))$ for the graph G in Figure 2.1. Solid edges are odd, dashed edges are even. After the same sequence of minor operations performed on $(G, E(G))$ in Figure 2.2, we obtain the minor in (c). Deleting all but one odd edge from v^+ to each node distinct from v^+ yields an odd- K_5 .

This motivates our proof of Lemma 2.30.

Proof of Lemma 2.30. If $n = 1$, there is nothing to prove. So suppose that $n \geq 2$. Again, let us identify the edges of $G + K_1$ that are incident with the unique node of K_1 with their other endpoint in $V(G)$. That is, $E(G + K_1) = E(G) \dot{\cup} V(G)$. With this labelling, any sequence of signature exchanges and deletions or contractions of edges in $E(G)$ can be performed on both $(G, E(G))$ and $(G + K_1, E(G + K_1))$ simultaneously. Here, doing a signature exchange on $(G, E(G))$ using the cut $\delta_G(U)$

for some $U \subseteq V(G)$ corresponds to a signature exchange on $(G + K_1, E(G + K_1))$ with $\delta_{G+K_1}(U)$. Conversely, we may assume that for any signature exchange on $(G + K_1, E(G + K_1))$, we only use cuts of the form $\delta_{G+K_1}(U)$ for some subset U not containing $V(K_1)$, since $\delta_{G+K_1}(U) = \delta_{G+K_1}(V(G + K_1) \setminus U)$.

We call two minors (H, Σ) and (H^+, Σ^+) of $(G, E(G))$ and $(G + K_1, E(G + K_1))$, respectively, *simultaneous* if they result from the same sequence of signature exchanges and deletions or contractions of edges in $E(G)$. Our first observation is that simultaneous minors are closely related in the following sense.

Claim 1. *Let (H, Σ) and (H^+, Σ^+) be simultaneous minors of $(G, E(G))$ and $(G + K_1, E(G + K_1))$, respectively. Then H^+ has a node v^+ that is adjacent to all other nodes of H^+ such that $\delta_{H^+}(v^+) = V(G)$ and (H, Σ) is obtained from (H^+, Σ^+) after deleting v^+ and all incident edges.*

Proof of Claim 1. The statements clearly holds true for the trivial pair of simultaneous minors $(G, E(G))$ and $(G + K_1, E(G + K_1))$, where we take v^+ to be the unique node of K_1 . Now suppose that a pair of simultaneous minors (H, Σ) and (H^+, Σ^+) satisfies all the claimed properties. We will show that these properties are maintained under performing any one minor operation (signature exchange, deletion or contraction of an edge in $E(G)$) on both (H, Σ) and (H^+, Σ^+) simultaneously.

Let v^+ be the node of H^+ whose deletion from (H^+, Σ^+) (together with its incident edges) yields (H, Σ) . In particular, we have that $E(H) = E(H^+) \setminus V(G)$ and $\Sigma = \Sigma^+ \setminus V(G) = \Sigma^+ \cap E(G)$. This means that, after a simultaneous signature exchange using a cut with shore $U \subseteq V(G)$, the new signatures still coincide on all edges from $E(G)$. Clearly, signature exchanges leave the graphs themselves invariant, so the claim holds true after a simultaneous signature exchange.

Now let $e \in E(H)$. Since $E(H) \subseteq E(G)$ and $\delta_{H^+}(v^+) \cap E(G) = \emptyset$ by hypothesis, v^+ is not an endpoint of e . So deleting or contracting e in (H, Σ) is the same as doing so in (H^+, Σ^+) and then removing v^+ and all incident edges. In particular, v^+ is still a node of the resulting minor of (H^+, Σ^+) whose incident edges $V(G)$ are untouched. It is easily seen that v^+ is adjacent to all other nodes of this minor, no matter whether e was deleted or contracted. See Figure 2.4 for an example. \diamond

Now suppose that we can obtain an odd- K_n from $(G, E(G))$ by performing a sequence of edge deletions and contractions, and node deletions. As argued above, we may assume that all deleted nodes are isolated by first deleting their incident edges. So $(G, E(G))$ has a minor (H, Σ) obtained after edge deletions and contractions only such that (H, Σ) is the union of an odd- K_n with a (possibly empty) set of isolated nodes. Let (H^+, Σ^+) be the corresponding minor of $(G + K_1, E(G + K_1))$ such that (H, Σ) and (H^+, Σ^+) are a simultaneous pair of minors. For such a special pair of minors, we may strengthen the statement of Claim 1 and show that the node v^+ can even be assumed to have an *odd* edge to all other nodes (up to a signature exchange):

Claim 2. *Let (H, Σ) and (H^+, Σ^+) be a pair of simultaneous minors of $(G, E(G))$ and $(G + K_1, E(G + K_1))$, respectively, and let v^+ be the node of H^+ as per Claim 1. Suppose that each connected component of H is complete and $\Sigma = E(H)$. Then (H^+, Σ^+) has an odd edge between v^+ and each node distinct from v^+ .*

Proof of Claim 2. Suppose for the sake of contradiction that H^+ has two nodes $u_0, u_1 \neq v^+$ that belong to the same connected component of H such that all edges between u_0 and v^+ are even and all edges between u_1 and v^+ are odd. By Claim 1, there is at least one edge from each of the two nodes to v^+ , so $u_0 \neq u_1$. Since the connected component of u_0 and u_1 is complete by hypothesis, the two nodes must therefore be connected by an edge $e \in E(H^+)$. It follows from Claim 1 that $e \in E(H) \subseteq E(G)$ since no edge from $V(G)$ can be incident to both u_0 and u_1 . Let $v, w \in V(G)$ be the original endpoints of e in G . Note that $v, w \in E(H^+)$ because $V(G) = \delta_{H^+}(v^+) \subseteq E(H^+)$ by Claim 1. Now observe that $\{e, v, w\}$, as a subset of edges of $G + K_1$, is a cycle in both $G + K_1$ and H^+ . Since this cycle is odd in $(G + K_1, E(G + K_1))$, it must also be odd in (H^+, Σ^+) by Remark 2.26. In particular, since one of the two edges v, w is incident with u_0 and the other one with u_1 , exactly one of v, w is an odd edge in (H^+, Σ^+) by hypothesis. So for the cycle $\{e, v, w\}$ to be odd in (H^+, Σ^+) , e must be even. However, $e \in E(H) = \Sigma$ by hypothesis and therefore also $e \in \Sigma^+$ (by Claim 1), a contradiction.

Hence, for each node of (H^+, Σ^+) distinct from v^+ , we can select an edge that connects the node to v^+ , such that all selected edges to the same connected component of H have the same parity. Since H^+ has no edge between nodes of different connected components of H , we can perform signature exchanges using $\delta_{H^+}(W)$ for the nodes W of a connected component if necessary to make all selected edges odd. \diamond

Thus, if $(G, E(G))$ has a minor (H, Σ) whose connected components all are odd complete graphs, then $(G + K_1, E(G + K_1))$ must have a minor isomorphic to $(H + K_1, E(H + K_1))$ by Claim 2. In particular, if (H, Σ) is an odd- K_n , possibly together with a number of isolated nodes, one can obtain $(K_n + K_1, E(K_n + K_1))$ from this $(H + K_1, E(H + K_1))$ minor by contracting all edges incident with nodes v that are isolated in H . Note that each such edge can be made even by a signature exchange with $\delta(v)$, which leaves the parity of all other edges invariant. As $K_n + K_1$ and K_{n+1} are isomorphic, we have thus shown that if $(G, E(G))$ has an odd- K_n minor (possibly after node deletions), then $(G + K_1, E(G + K_1))$ has an odd- K_{n+1} minor.

To prove the converse implication, it suffices to show the following statement.

Claim 3. *Let (H^+, Σ^+) be a minor of $(G + K_1, E(G + K_1))$. Then there is a node v^+ such that deleting v^+ and all incident edges from (H^+, Σ^+) yields a minor of $(G, E(G))$ obtained by edge deletions and contractions and possibly node deletions.*

Proof of Claim 3. Suppose first that (H^+, Σ^+) is obtained from $(G + K_1, E(G + K_1))$ by deleting or contracting edges from $E(G)$ only. Then we can perform the same operations on $(G, E(G))$ and obtain a simultaneous pair of minors. By Claim 1, the statement therefore holds for (H^+, Σ^+) .

Since we may perform deletions and contractions in arbitrary order (modulo signature exchanges), we may assume without loss of generality that any minor of $(G + K_1, E(G + K_1))$ is obtained by first deleting or contracting edges from $E(G)$ and only then deleting or contracting edges from $V(G)$. So to prove the claim for all minors, it suffices to show that, starting from a pair of simultaneous minors (H, Σ) and (H^+, Σ^+) , deleting or contracting edges in $V(G)$ preserves the claimed property. By Claim 1, all edges of H^+ that come from $V(G)$ have a common endpoint v^+ whose deletion results in a minor of (H, Σ) (namely, (H, Σ) itself). We will argue that this property is maintained under

deleting or contracting edges from $V(G)$. The invariant clearly holds after deleting an edge from $V(G)$.

Now suppose that we contract some edge $v \in V(G)$. The invariant guarantees that one of the endpoints of v in (H^+, Σ^+) is v^+ . Let w denote the other endpoint. (Note that possibly $w \neq v$ since v , as a node of G , may not exist any more after contracting edges.) By definition, contracting edge v means that its two endpoints v^+ and w are identified, say, with v^+ , and all edges that were incident with w now also become incident with v^+ . Deleting v^+ and all incident edges therefore results in (H, Σ) with node v and all incident edges deleted. This means that we may simulate the contraction of edge v in (H^+, Σ^+) by deleting node v in (H, Σ) . Since node deletions are permitted, we obtain a minor of (H, Σ) as desired.

Recall that all subsequent minor operations performed on (H^+, Σ^+) after this edge contraction only touch edges from $V(G)$ by hypothesis. Deleting nodes and incident edges in (H, Σ) is therefore safe to do: We will never delete any edge from $V(G)$ because $E(H) \subseteq E(G)$ by Claim 1. \diamond

Deleting an arbitrary node of an odd- K_{n+1} and all incident edges yields an odd- K_n . So if $(G + K_1, E(G + K_1))$ has an odd- K_{n+1} minor, then $(G, E(G))$ must have an odd- K_n minor by Claim 3. This concludes the proof of Lemma 2.30. \square

Lemma 2.30 will also prove extremely useful in characterizing the edge-node incidence matrices M_G for which the clutter $\mathcal{F}_+(M_G, \mathbf{1})$ has the MFMC property, as we will see next.

The MFMC Property and Total Unimodularity

As mentioned in Section 2.1, the linear description of the matching polytope (2.2)–(2.3) is TDI, as shown by Cunningham and Marsh [66]. So the statement of Theorem 2.21, when applied to the node-edge incidence matrix of a graph, is little surprising. However, one cannot directly derive the result of [66] from Theorem 2.21 since the clutter of minimal T -cuts in a graph with terminals T does not necessarily have the MFMC property:

Proposition 2.31 (see [60, 183]). *The clutter of minimal T -cuts in K_4 with $T = V(K_4)$ does not have the MFMC property.*

Proof. For $T = V(K_4)$, every T -cut in K_4 is of the form $\delta(v)$ for some $v \in V(K_4)$. All four such cuts are minimal. Now consider the following LP in variables $z \in \mathbb{R}^{E(K_4)}$:

$$\begin{aligned} \min \quad & \mathbf{1}^\top z \\ & z(\delta(v)) \geq 1 \quad \text{for all } v \in V(K_4) \\ & z \geq \mathbf{0} \end{aligned} \tag{2.31}$$

Its dual (in variables $y \in \mathbb{R}^{V(K_4)}$) is

$$\begin{aligned} \max \quad & \mathbf{1}^\top y \\ & y_u + y_v \leq 1 \quad \text{for all } uv \in E(K_4) \\ & y \geq \mathbf{0} \end{aligned} \tag{2.32}$$

The vectors $z^* = (\frac{1}{3}, \dots, \frac{1}{3})$ and $y^* = (\frac{1}{2}, \dots, \frac{1}{2})$ are primal and dual feasible, respectively, as is easily checked. Since $\mathbf{1}^\top z^* = \mathbf{1}^\top y^* = 2$, both are, in fact, optimal solutions. Note that the feasible region of the dual (2.32) is the fractional stable set polytope of K_4 (cf. Example 2.2). So the maximum possible objective function value of an integral solution of (2.32) is $1 < \mathbf{1}^\top y^*$. This implies that the system in (2.31) is not TDI. \square

For each edge e of K_4 , there is a unique edge disjoint from e . Now relabel the set of edges of K_4 by swapping the labels of each such pair of edges. Then, for every node $v \in V(K_4)$, the cut $\delta(v)$ becomes the unique cycle of K_4 that does not contain v . So the clutter of odd *cuts* in K_4 is also the clutter of odd *cycles* of odd- K_4 , up to relabelling the ground set as described above. Let us denote this clutter by Q_6 (its ground set has 6 elements, the 6 edges of K_4). Among the binary clutters without the MFMC property, Q_6 plays a special role, as shown by Seymour [183]. To formally state Seymour's result, we first need to define *minors* of clutters. When we introduced the deletion and contraction operations for signed graphs earlier in this section, we implicitly showed how deleting or contracting an edge of a signed graph is mirrored by one of two operations performed on the clutter of its odd cycles. These operations are well-defined for any clutter, more generally. Let us make this more explicit here.

For a clutter \mathcal{F} over ground set E , *deleting* an element $e \in E$ from \mathcal{F} means dropping all members from \mathcal{F} that contain e . The resulting collection $\{S \in \mathcal{F} : e \notin S\}$ is a clutter again, this time over ground set $E \setminus \{e\}$. The clutter obtained from \mathcal{F} after *contracting* e is defined as the collection of all minimal sets in $\{S \setminus \{e\} : S \in \mathcal{F}\}$. Again, its ground set is $E \setminus \{e\}$. Any clutter obtained from \mathcal{F} after a sequence of deletions and contractions is called a *minor* of \mathcal{F} . Both idealness and the MFMC property are preserved under taking minors of clutters [183]. So any clutter with a Q_6 minor cannot have the MFMC property. Seymour [183] proved that for binary clutters, the converse is also true:

Proposition 2.32 ([183]). *A binary clutter has the MFMC property if and only if it has no Q_6 minor.*

If \mathcal{F} is the clutter of odd cycles of a signed graph (G, Σ) , the (clutter) minors of \mathcal{F} are precisely the clutters of odd cycles of (signed graph) minors of (G, Σ) , as we saw earlier in this section. Thus, Seymour's theorem (Proposition 2.32), specialized to clutters of odd cycles, states that the clutter of odd cycles of a signed graph (G, Σ) has the MFMC property if and only if (G, Σ) has no odd- K_4 minor (see also [60]). Using this and the tools that we developed to characterize for which edge-node incidence matrices M_G the clutter $\mathcal{F}_+(M_G, \mathbf{1})$ is ideal, we may even characterize exactly when $\mathcal{F}_+(M_G, \mathbf{1})$ has the MFMC property.

Theorem 2.33. *Let M_G be the edge-node incidence matrix of a graph G . Then $\mathcal{F}_+(M_G, \mathbf{1})$ has the MFMC property if and only if G is bipartite.*

We will need the following simple fact.

Lemma 2.34. *A graph G is bipartite if and only if $(G, E(G))$ has no odd- K_3 minor (possibly after node deletions).*

Proof. Suppose first that G is bipartite. Then G has no odd cycle, so the clutter of odd cycles of $(G, E(G))$ is the trivial clutter \emptyset whose only minor is \emptyset . Note, however, that

the clutter of odd cycles of odd- K_3 is nontrivial since it consists of a unique member, namely $E(K_3)$. This implies that odd- K_3 cannot be a minor of $(G, E(G))$.

Now suppose that $G = (V, E)$ has an odd cycle $C \subseteq E$. Consider the signed graph obtained from $(G, E(G))$ after deleting all edges in $E \setminus C$. Since C is a cycle in this minor, it follows from Remark 2.26 that C must still be odd. So after contracting all but 3 edges of C and possibly deleting isolated nodes, we are left with an odd- K_3 . \square

Now the proof of Theorem 2.33 is immediate.

Proof of Theorem 2.33. By Propositions 2.14 and 2.32, the clutter $\mathcal{F}_+(M_G, \mathbf{1})$ has the MFMC property if and only if $(G + K_1, E(G + K_1))$ (after a signature exchange) has no odd- K_4 minor. By Lemma 2.30, this is equivalent with $(G, E(G))$ not having an odd- K_3 minor (where deleting nodes is permitted). The statement then follows from Lemma 2.34. \square

Combined with Theorem 2.21, Theorem 2.33 provides a sufficient condition for the linear system (2.4)–(2.5) that determines the $\{0, \frac{1}{2}\}$ -closure of the fractional stable set polytope of a graph G to be TDI. Graphs for which (2.4)–(2.5) is a TDI system are called *strongly t -perfect*. We remark that, strictly speaking, the statement of Theorem 2.21 only concerns the system including all support-minimal nontrivial $\{0, \frac{1}{2}\}$ -cuts, in particular including the redundant odd-path inequalities (2.18). However, a careful analysis of our argument in Section 2.3.1 shows that the odd-path inequalities are nonnegative integer linear combinations of other inequalities, so we may apply Proposition 2.22(ii) to obtain the desired statement. Thus, Theorems 2.21 and 2.33 equivalently state that bipartite graphs are strongly t -perfect. This is a trivial statement: A bipartite graph G has no odd cycles, and its edge-node incidence matrix M_G is totally unimodular (see, e.g., [177]). By a well-known fact (see [174]), this implies that the system $M_G x \leq \mathbf{1}, x \geq \mathbf{0}$ is TDI.

A stronger statement about t -perfection may be derived from our characterization of graphs G for which $\mathcal{F}_+(M_G, \mathbf{1})$ is ideal (see Theorem 2.25). It is known that a graph G is strongly t -perfect if $(G, E(G))$ has no odd- K_4 minor [112, 175]. This is the case if and only if $\mathcal{F}_+(M_G, \mathbf{1})$ is ideal, as we know from combining Theorem 2.25 and Corollary 2.29. Since the MFMC property implies idealness, t -perfection follows for a strictly larger class of graphs than bipartite graphs.

The combination of our characterization of idealness in Theorem 2.25 with Proposition 2.32 also yields another interesting byproduct. Namely, recall from Proposition 2.15 that $\mathcal{F}(M_G, \mathbf{1})$ is precisely the clutter of odd cycles of $(G, E(G))$, which has the MFMC property if and only if $(G, E(G))$ has no odd- K_4 minor by Proposition 2.32. Thus, we obtain the following corollary.

Corollary 2.35. *Let M_G be the edge-node incidence matrix of a graph G . Then $\mathcal{F}_+(M_G, \mathbf{1})$ is ideal if and only if $\mathcal{F}(M_G, \mathbf{1})$ has the MFMC property.*

We next explore in greater detail when the clutters associated with fractional matching polytopes have the MFMC property. It turns out that, also in this case, bipartite graphs and their node-edge incidence matrices M_G^\top are the only ones for which the clutter $\mathcal{F}_+(M_G^\top, \mathbf{1})$ can have the MFMC property. To prove this, we first note that – as with the clutter of odd cycles – taking minors of the clutter of minimal T -cuts in a graph

also has a graphic interpretation using edge deletions and contractions. However, the two operations behave in the opposite way here. For an edge e with endpoints u and v , contracting e in the clutter yields the clutter of minimal T -cuts in G with e deleted from G . Now suppose that we delete e from the clutter. Then the resulting minor is the clutter of minimal T' -cuts in the graph G' , where G' is obtained from G by contracting edge e into node w , and T' consists of all nodes in $T \setminus e$, plus w if exactly one of the endpoints of e was in T (see [60]).

Given a graph G with terminals T , it thus follows from Proposition 2.31 that the clutter of its minimal T -cuts does not have the MFMC property if one can transform G into K_4 with terminals $V(K_4)$ by a sequence of edge deletions and contractions as defined above. The next result shows that this is possible for the join of any non-bipartite graph with K_1 as in Proposition 2.13.

Theorem 2.36. *Let M_G be the edge-node incidence matrix of a graph G . If $\mathcal{F}_+(M_G^\top, \mathbf{1})$ has the MFMC property, then G is bipartite.*

Proof. By Proposition 2.13, the clutter $\mathcal{F}_+(M_G^\top, \mathbf{1})$ is the clutter of minimal T -cuts in $G + K_1$, where the set of terminals T is defined as in Proposition 2.13 for $b = \mathbf{1}$. As usual, we denote the unique node of K_1 by v^+ and identify its incident edges with $V(G)$. We will show that if G is non-bipartite, then there is a sequence of edge deletions and contractions as defined above that transforms $G + K_1$ with terminals T into K_4 with terminals $V(K_4)$. The statement then follows from Proposition 2.31.

So suppose that G has an odd cycle $C \subseteq E(G)$. Consider the minor obtained from the clutter of minimal T -cuts in $G + K_1$ after contracting all edges in $E(G) \setminus C$ and deleting all edges in $V(G) \setminus V(C)$. In graphic terms, this minor is the clutter of minimal T' -cuts (for some new set of terminals T') in the graph obtained from $G + K_1$ after deleting $E(G) \setminus C$ and contracting $V(G) \setminus V(C)$.

We claim that T' consists of all nodes in the resulting graph. To see this, recall that the set of terminals may only be modified by edge contractions. All edges that we contract are from $V(G) \setminus V(C)$. In particular, no such edge has an endpoint on the cycle C , which implies that $V(C) \subseteq T'$. Instead, each contracted edge is incident with v^+ . Let us denote the node into which each such edge is contracted by v^+ again. Since $T \supseteq V(G)$, contracting an edge incident with v^+ makes v^+ a terminal if it was not a terminal before the contraction and vice versa. Now recall that, by the definition of T in Proposition 2.13, we have that $v^+ \in T$ if and only if $|V(G)|$ is odd. Since $|V(C)| = |C|$ is odd, $|V(G) \setminus V(C)|$ is even if and only if $v^+ \in T$. So after contracting $|V(G) \setminus V(C)|$ of its incident edges, v^+ must be a terminal in T' .

The final step of the proof is to contract C into a cycle of length 3. If $|C| = 3$, we are done. So suppose that $|C| > 3$, and let $e_1, \dots, e_{|C|}$ denote the edges of C ordered along the cycle. Now contract $e_4, \dots, e_{|C|}$ in this order, and let C_k be the cycle obtained from C after contracting e_4, \dots, e_k , for $4 \leq k \leq |C|$. Since any two consecutive edges in the sequence $e_4, \dots, e_{|C|}$ share an endpoint, the nodes of C_k are all terminals if k is odd. In particular, $|C|$ is odd, so all 3 nodes of the final cycle $C_{|C|}$ are terminals. Also note that v^+ is still a terminal because no edge of C contains v^+ . After deleting all but one edge from each node of $C_{|C|}$ to v^+ , we thus obtain K_4 with all nodes being terminals. \square

The converse implication of Theorem 2.36, however, is false: The graph G in Figure 2.5

is bipartite but after contracting two edges of $G + K_1$ and deleting multiple copies of parallel edges, we obtain K_4 with all nodes being terminals.




Figure 2.5: (a) A bipartite graph G , (b) its join with K_1 , and (c) the graph obtained from $G + K_1$ after contracting the two edges of the subgraph shaded in (b). Assuming that all nodes of G are terminals, the same is true for $G + K_1$, following Proposition 2.13, and therefore the node in (c) into which the two edges are contracted is a terminal again. After two edge deletions, one obtains K_4 from (c).

As remarked above, the linear systems in Examples 2.1 and 2.2 describing the fractional matching and stable set polytopes of bipartite graphs are TDI since (node-edge or edge-node) incidence matrices of bipartite graphs are totally unimodular. So from a graph-theoretic standpoint, the necessary conditions of Theorems 2.33 and 2.36 for the clutters $\mathcal{F}_+(M_G, \mathbf{1})$ or $\mathcal{F}_+(M_G^\top, \mathbf{1})$ to have the MFMC property for edge-node incidence matrices M_G are not particularly interesting. However, let us rephrase those conditions in terms of total unimodularity:

Corollary 2.37. *Let A be the node-edge or edge-node incidence matrix of a graph. If $\mathcal{F}_+(A, \mathbf{1})$ has the MFMC property, then A is totally unimodular.*

In this form, the statement now provides an interesting analogue to Theorem 2.25. Namely, by a famous result of Hoffman and Kruskal [129], an integral matrix A is totally unimodular if and only if

$$\begin{aligned} d \leq Ax \leq b \\ l \leq x \leq u \end{aligned}$$

defines an integral polyhedron for all integral vectors l, u, b, d of appropriate dimension. Matrices A with this property are also said to have strong Chvátal rank 0 (see [174]). Comparing this with the definition of the Edmonds-Johnson property, both Corollary 2.37 and Theorem 2.25 therefore relate idealness and the MFMC property of $\mathcal{F}_+(A, \mathbf{1})$ to the strong Chvátal rank of A for the special case that A is the node-edge or edge-node incidence matrix of a graph.

Using the results described in this section, we may also show that idealness of the clutters associated with fractional matching and stable set polytopes can be recognized in polynomial time. In fact, it follows from a result of Truemper [193] that binary clutters (represented as congruence systems) can be tested for having a Q_6 minor in polynomial time. To be precise, the results of [193] are phrased in the language of binary matroids. Using earlier matroid decomposition results of [194], Truemper showed how to efficiently detect a certain matroid minor (the dual of the Fano matroid; see, e.g., [60]) among all minors that contain a fixed element l of the ground set. We will not go into detail here but only mention that, for any choice of l , the l -port of this matroid minor is Q_6 [183] (see Section 2.3.1). So in our terms, the result of [193] may be restated as follows (see also [114]).

Proposition 2.38 ([193]). *There is a polynomial-time algorithm that, given a matrix $Q \in \{0, 1\}^{p \times q}$ and a vector $d \in \{0, 1\}^p$, decides whether the binary clutter \mathcal{F} over ground set $[q]$ represented as $Qy \equiv d \pmod{2}$ (as in Proposition 2.10) has a Q_6 minor.*

Combining this with Proposition 2.32 and Corollary 2.35, we obtain the following corollary.

Corollary 2.39. *The 0/1 matrices A for which $\mathcal{F}_+(A, \mathbf{1})$ is ideal can be recognized in polynomial time when A is the node-edge or edge-node incidence matrix of a graph.*

2.4 Further Notes and Open Questions

We saw in Section 2.2 that, even though the membership problem for the $\{0, \frac{1}{2}\}$ -closure of rational polyhedra $P = \{x: Ax \leq b\}$ with $P_{\frac{1}{2}}(A, b) = P_I$ is likely not NP-hard (as in the general case), testing whether $P_{\frac{1}{2}}(A, b) = P_I$ is an NP-hard problem by itself. Let us briefly comment on the complexity of this problem for fractional matching and stable set polytopes. Of course, the $\{0, \frac{1}{2}\}$ -closure of the fractional matching polytope as considered in Example 2.1 is always integral (and its linear description is TDI).

The more interesting case is that of the fractional stable set polytope and the complexity of recognizing (strong) t -perfection. Recall that in this case, the membership problem for the $\{0, \frac{1}{2}\}$ -closure, as given in Example 2.2, can be solved in polynomial time [114, 122]. This implies that there is a coNP certificate for t -perfection: It suffices to exhibit a fractional vertex x^* of the $\{0, \frac{1}{2}\}$ -closure of the fractional stable set polytope along with a corresponding basis. Then one can verify in polynomial time that x^* is in the $\{0, \frac{1}{2}\}$ -closure (using a membership oracle) and that it is indeed a vertex (by simply checking the basis). This observation can be found in Chapter 9 of [125] and readily generalizes in the following way.

Proposition 2.40. *Given $A \in \mathbb{Z}^{m \times n}$ and $b \in \mathbb{Z}^m$ with $P = \{x \in \mathbb{R}^n: Ax \leq b\}$, deciding whether $P_{\frac{1}{2}}(A, b) = P_I$ is in coNP when the membership problem for $P_{\frac{1}{2}}(A, b)$ is in P.*

Whether recognizing t -perfection is in NP or in P is not known (see Chapter 9 of [125]). However, some classes of t -perfect graphs are known to be polynomial-time recognizable, including claw-free t -perfect graphs [46] and bad- K_4 -free graphs [115]. A bad- K_4 is a non- t -perfect subdivision of K_4 . For example, the graph in Figure 2.1 is bad- K_4 -free.

Interestingly, these two classes of graphs are also strongly t -perfect [47, 175]. Recall that strong t -perfection implies t -perfection (cf. Proposition 2.6); whether the converse implication is also true in general is not known (see [177]).

Both fractional matching and stable set polytopes are set packing polytopes. We stress that our hardness results of Section 2.2 do not directly apply here as the polytopes arising in our reduction are not of set packing type. Is recognizing integrality of the $\{0, \frac{1}{2}\}$ -closure of a set packing polyhedron also NP-hard? We leave this as an open question. While the membership problem for the $\{0, \frac{1}{2}\}$ -closure is known to be strongly NP-hard even in the set packing case [152], there are some positive results for *approximate* linear optimization over the Gomory-Chvátal and $\{0, \frac{1}{2}\}$ -closures of set packing polyhedra: It follows from a result of [154] that the optimization problem over the Gomory-Chvátal closure of a set packing polyhedron P admits a PTAS. This means that one can optimize any given linear objective function over P' in polynomial time, up to an arbitrary fixed precision. (In fact, the PTAS of [154] applies to a more general class of packing problems; related approximation results for set *covering* problems were obtained in [18, 103, 154].) For approximating the $\{0, \frac{1}{2}\}$ -closure of set packing polyhedra, there is a significantly less involved PTAS [42].

Our results of Sections 2.3.3 and 2.3.4 leave open many interesting questions, too. First, as we saw in Section 2.3.3, the $\{0, \frac{1}{2}\}$ -closure of a set packing polyhedron is integral if the associated binary clutter is ideal, and the linear description of the $\{0, \frac{1}{2}\}$ -closure is even TDI if the clutter has the MFMC property (Proposition 2.18 and Theorem 2.21). It would be interesting to explore to which extent these statements generalize beyond the set packing case.

Second, for the two special families of set packing polyhedra that served as our recurring examples throughout this chapter, the associated clutters can be tested for idealness in polynomial time (see Corollary 2.39). It is therefore natural to wonder whether this is true more generally.

Question 2.41. *What is the computational complexity of recognizing 0/1 matrices A for which $\mathcal{F}_+(A, \mathbf{1})$ is ideal?*

Note that testing (not necessarily binary) clutters for being ideal (or having the MFMC property) is coNP-complete when the members of the clutter are explicitly given as part of the input [76]. In contrast, the input to the recognition problem of Question 2.41 only consists of A (from which one can easily write down the system of congruences defining $\mathcal{F}_+(A, \mathbf{1})$).

Here, we stress that for the special case of A being an edge-node incidence matrix M_G , we were only able to conclude that ideal clutters of the form $\mathcal{F}_+(M_G, \mathbf{1})$ can be recognized efficiently because of three key reasons. First, idealness is in this case equivalent to the minor $\mathcal{F}(M_G, \mathbf{1})$ having the MFMC property (see Corollary 2.35); this follows from, second, Seymour's characterization of Q_6 as the unique minimal binary clutter without the MFMC property [183] stated in Proposition 2.32; and, third, a Q_6 minor can be detected in polynomial time using the algorithm of Truemper [193] (see Proposition 2.38). It is unclear whether any of these three ingredients may help answer Question 2.41 for more general 0/1 matrices, not just edge-node incidence matrices. (Recall from Section 2.3.4 that the node-edge case is trivial.)

In particular, our proof of the first of the three ingredients given in Section 2.3.4 crucially relies on knowing when clutters of T -cuts and odd cycles are ideal, as characterized in [86, 126] (see Propositions 2.23 and 2.27). One may ask whether there is a forbidden minor characterization of ideal binary clutters, more generally, analogous to Seymour's result for the MFMC property. This is, in fact, a long-standing open question. Seymour [183] conjectured that every non-ideal binary clutter must have one of exactly three minors, two of which are \mathcal{O}_5 , the clutter of odd cycles of odd- K_5 , and its blocker $b(\mathcal{O}_5)$. This conjecture is known to be true for special classes of binary clutters: For example, both results of [86, 126] mentioned above can be regarded as special cases (see also [60]). More general results were obtained in [2, 61]. Despite this progress, Seymour's conjecture remains open. We refer to [1, 60, 61] for more background.

The other central result used in Section 2.3.4 is the characterization of edge-node incidence matrices with the EJ property given in [114]. Recall that the EJ property actually pertains to the Gomory-Chvátal closure of a polyhedron. Yet for linear systems with a 2-regular constraint matrix A , the EJ property of A implies integrality of the $\{0, \frac{1}{2}\}$ -closure, as explained in Section 2.3.4. All known classes of matrices with the EJ property are 2-regular (see [74]), including the edge-node incidence matrices of [114] and integral *binet matrices* [5]. So a natural question is whether Theorem 2.25 may extend to all 2-regular 0/1 matrices:

Question 2.42. *For a given 2-regular 0/1 matrix A , does A have the EJ property if $\mathcal{F}_+(A, \mathbf{1})$ is ideal?*

Note that integral 2-regular matrices can only have entries in $\{0, \pm 1, \pm 2\}$. If we allow entries equal to ± 2 , then the implication conjectured in Question 2.42 is false: The matrix

$$A_4 = \begin{pmatrix} 2 & 1 & 1 & 0 \\ 2 & 1 & 0 & 1 \\ 2 & 0 & 1 & 1 \end{pmatrix} \quad (2.33)$$

is 2-regular but does not have the EJ property [74]. The associated clutter $\mathcal{F}_+(A_4, \mathbf{1})$ is Q_6 except that its ground set has a seventh element (corresponding to the first column of A_4) that is never used in any member. Recall from Section 2.3.4 that Q_6 is the clutter of odd cycles of odd- K_4 and is therefore ideal by Proposition 2.27.

In fact, the matrix A_4 in (2.33) is conjectured to be one of two forbidden matrix minors for 2-regular matrices with the EJ property [113] (see [74]). Here, minors are matrices obtained by a sequence of the operations (i) to (iii) defined in Section 2.3.4 as well as a pivoting operation that refines operation (iv). Additionally, one is allowed to divide by 2 a row whose entries are in $\{0, \pm 2\}$. Each of these operations preserves both the EJ property and 2-regularity [74]. For example, we saw that M_{K_4} is the minimal edge-node incidence matrix without the EJ property [114] (see Proposition 2.28). However, M_{K_4} is not minimal among *all* 2-regular matrices without the EJ property, as it has an A_4 minor.

Finally, let us comment on the complexity of the binary clutter problem for clutters that satisfy either of the two integrality properties considered in Sections 2.3.3 and 2.3.4. In [193], Truemper did not only show how to test for the MFMC property efficiently (see Proposition 2.38), he also gave a polynomial-time algorithm for the binary clutter

problem (or, rather, its matroid port counterpart) for clutters with the MFMC property. This generalizes the well-known facts that minimum-weight s - t -paths and s - t -cuts in graphs with nonnegative edge weights can be found in polynomial time. For instance, the result of [193] was used in [150] to describe broader classes of polynomially solvable cases of the separation problem for the $\{0, \frac{1}{2}\}$ -closure, as mentioned in Section 2.1.

In the case of the binary clutters associated with fractional matching and stable set polytopes, however, Truemper's result in [193] does not have any interesting implications for separating $\{0, \frac{1}{2}\}$ -cuts efficiently. Indeed, by Theorems 2.33 and 2.36, we know that the associated clutters only have the MFMC property if the underlying graph is bipartite. Yet both the fractional matching polytope and the fractional stable set polytope of a bipartite graph are already integral, so the separation problem for the $\{0, \frac{1}{2}\}$ -closure is trivial in this case. It should be noted, however, that the relevant binary clutter for separating $\{0, \frac{1}{2}\}$ -cuts for the fractional stable set polytope of a graph G with edge-node incidence matrix M_G is, in fact, not $\mathcal{F}_+(M_G, \mathbf{1})$ but its minor $\mathcal{F}(M_G, \mathbf{1})$ (see Section 2.3.1).

To the best of our knowledge, it is not known whether idealness would already suffice to render the binary clutter problem easy. Positive answers are known for special classes of ideal clutters such as those associated with T -cuts and T -joins (their blockers) or odd cycles [86, 114, 122, 161]. Some further evidence may be provided by the following observation that was made, e.g., in [122]: Take a binary clutter \mathcal{F} given as a system of congruences, and suppose that one has an algorithm that solves the binary clutter problem for \mathcal{F} and all nonnegative weights. Such an algorithm can be used as a separation oracle for the associated set covering polyhedron $Q(\mathcal{F})$. Using this oracle and the ellipsoid method, one can therefore optimize any (nonnegative) linear objective function over $Q(\mathcal{F})$ in oracle-polynomial time [124, 125]. Now suppose that \mathcal{F} is ideal. Then linear optimization over $Q(\mathcal{F})$ is, in fact, the same as linear optimization over the incidence vectors of the members of $b(\mathcal{F})$ by Proposition 2.20. This implies that, given an oracle for the binary clutter problem over an ideal binary clutter, the binary clutter problem over its blocker can be solved in oracle-polynomial time (provided a congruence system for the blocker; see Proposition 2.19 for a special case).

This provides an interesting contrast to the situation for *non-ideal* clutters. To illustrate this, consider the clutter of odd cycles in $(G, E(G))$ for a given graph G . Let us denote this clutter by \mathcal{F} . By definition, every cover of \mathcal{F} must contain at least one edge from each odd cycle in G . This means that deleting all edges of a cover leaves a bipartite subgraph. In other words, the complements of the covers of \mathcal{F} are the edge sets of bipartite subgraphs of G . Each such edge set is contained in a cut, namely the cut whose shores are the two node classes of the bipartition. Hence, each cover of \mathcal{F} contains a cut in G , which implies that the blocker of \mathcal{F} consists of all complements of cuts in G (see, e.g., [111, 122, 126]). For given nonnegative edge weights, finding a minimum-weight member of the blocker $b(\mathcal{F})$ therefore amounts to finding a maximum-weight cut in G . This problem is the well-known *max-cut problem*, which is strongly NP-hard even for uniform edge weights [107], while the binary clutter problem for \mathcal{F} can always be solved in polynomial time [114, 122], as pointed out in [108] (see also [150]).

Note that this reduction is one of the ways to prove strong NP-completeness of the binary clutter problem, as mentioned in Section 2.3.1 (see [49, 111, 122, 126]). Combined

with the above observation about the complexity of the binary clutter problem for *ideal* clutters, it also implies that the max-cut problem can be solved in polynomial time on graphs G for which $(G, E(G))$ has no odd- K_5 minor (such graphs are called *weakly bipartite*), as observed in [122].

Chapter 3

Limitations of the Hyperplane Separation Bound for the Extension Complexity of Polytopes

The results of this chapter appear in [44]. The presentation of the material below is similar to the paper.

3.1 Introduction

In the previous chapter, we discussed integral polyhedra whose linear descriptions can be obtained from relaxations by means of certain cutting planes. Yet for linear programming over a polyhedron $P \subseteq \mathbb{R}^n$, integral or not, one does not need to have a complete linear description of P available. It suffices to know the description of a polyhedron $Q \subseteq \mathbb{R}^m$ and an affine map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that $\pi(Q) = P$, since for any linear objective function $c \in \mathbb{R}^n$, the optimal values of the LPs $\max \{c^\top x: x \in P\}$ and $\max \{c^\top \pi(y): y \in Q\}$ are the same. Assuming that P is pointed, we may express the latter LP using no more constraints and variables than Q has facets (see [134]): First, equations may be removed by eliminating variables. Second, if the resulting full-dimensional polyhedron is not pointed, we may replace it with its projection onto the orthogonal complement of the lineality space. This means that the minimum number of facets of any polyhedron Q that affinely projects onto P is a measure of how small an LP with feasible region P can possibly be (where we ignore the encoding lengths of the coefficients). This minimum number is called the *extension complexity* of P , denoted by $\text{xc}(P)$. An *extension* of P is a polyhedron Q together with an affine map π such that $\pi(Q) = P$. A linear description of such a polyhedron Q is called an *extended formulation* for P , whose *size* is defined as the number of its inequalities. In other words, $\text{xc}(P)$ is the minimum size of any extended formulation for P .

The extension complexity of a polyhedron may be significantly smaller than the number of its facets. For instance, consider again the *spanning tree polytope* of a connected graph $G = (V, E)$ that was introduced in Chapter 1. In this chapter, we will denote it by $P_{\text{st}}(G)$. Recall that $P_{\text{st}}(G)$ is defined as the convex hull of the incidence vectors of the spanning trees in G , i.e.,

$$P_{\text{st}}(G) = \text{conv} \{ \chi(T) \in \{0, 1\}^E : T \subseteq E \text{ is a spanning tree in } G \}. \quad (3.1)$$

As shown by Edmonds [82], $P_{\text{st}}(G)$ is completely determined by the following linear

system:

$$\begin{aligned} x(E) &= |V| - 1 \\ x(E(U)) &\leq |U| - 1 \quad \text{for all } \emptyset \neq U \subseteq V \\ x &\geq \mathbf{0} \end{aligned} \tag{3.2}$$

In general, the number of facets of $P_{\text{st}}(G)$ is exponential in $|V|$. Nonetheless, as mentioned in Chapter 1, there are extended formulations of size $O(|V||E|)$ due to Wong [204] and Martin [153] (see also [6, 58, 205]). Special classes of graphs admit even smaller extended formulations: For instance, Williams [200] gave a formulation of size $O(|V|)$ for planar graphs. More generally, if G is embedded in a surface of fixed genus, the extension complexity of $P_{\text{st}}(G)$ was shown to be at most $O(|V|^{3/2})$ in [102]. This upper bound was recently generalized to proper minor-closed classes of graphs [8]. However, if $G = K_n$, the best known upper bound on $\text{xc}(P_{\text{st}}(K_n))$ is the one obtained by counting the inequalities in Martin or Wong’s formulations [153, 204]; namely, $O(n^3)$.

Lower bounding the extension complexity, on the other hand, seems even more difficult: How does one prove that *no* extended formulation of a certain size exists? Rothvoß [166], using a counting argument, showed that there are 0/1 polytopes with exponential extension complexity. Yet this result does not imply any lower bounds for concrete polytopes. In a seminal paper from 1991, Yannakakis [205] proved that TSP polytopes and matching polytopes of complete graphs do not admit extended formulations of subexponential size that satisfy certain symmetry assumptions. Since then, lower bounds on sizes of *symmetric* extended formulations have been studied for other families of polytopes as well [137]. However, in search of unconditional lower bounds on the extension complexity (without symmetry assumptions), additional tools are required.

In his paper [205], Yannakakis also provided an algebraic handle on the extension complexity. His beautiful result uses the concept of a *slack matrix*. For the spanning tree polytope as described by the constraints in (3.2), such a slack matrix has one row for each inequality in (3.2) and one column for each spanning tree T in G . The corresponding entry is the slack of the vertex $\chi(T)$ in the respective inequality. In an analogous fashion, one can associate a slack matrix with the set of vertices of any polytope P and a given linear description of it. Yannakakis’ theorem in [205] now states that the minimum number r for which this slack matrix can be written as the product of two nonnegative matrices with r as the intermediate dimension is equal to $\text{xc}(P)$.

This result laid the ground for many advances in recent years, most notably the exponential lower bounds for the extension complexity of TSP polytopes, cut and correlation polytopes, stable set polytopes [101, 136], and (perfect) matching polytopes [167]. The results of [101, 136] rely on a close connection to the communication complexity of 0/1 matrices (see also [95]). This connection was already explored by Yannakakis [205] and gives rise to a lower bound on the extension complexity of a polytope known as the *rectangle covering lower bound* (see [100]). Being a *combinatorial* lower bound, it only depends on the zero/nonzero pattern of a slack matrix, which encodes the vertex-facet incidence structure of the polytope. It is known that the rectangle covering lower bound, when applied to the slack matrix of a polytope P , is at least the dimension of P [100]. For the spanning tree polytope of K_n , we thus obtain

an immediate lower bound of $\Omega(n^2)$ on its extension complexity. The question whether this bound can be improved, e.g., to match the upper bound of $O(n^3)$, is open (see [8, 141, 199]). Khoshkhah and Theis [141] showed that the rectangle covering lower bound is at most $O(n^2 \log n)$. They asked whether using non-combinatorial techniques instead may lead to stronger lower bounds on the extension complexity of $P_{\text{st}}(K_n)$. This question provides the motivation for our work in this chapter.

One candidate for such a non-combinatorial lower bound is the *hyperplane separation bound*. It was proposed by Fiorini [99] and applied by Rothvoß [167] in his proof of the exponential lower bound for the matching polytope. Like the rectangle covering lower bound, it is derived from Yannakakis' algebraic characterization of the extension complexity in terms of slack matrices. We show that for the slack matrix of $P_{\text{st}}(K_n)$ obtained from the linear description in (3.2), the hyperplane separation technique fails to produce a lower bound stronger than $\Omega(n^2)$. In this sense, the trivial dimension bound is already at least as strong. Our proof in Section 3.3.1 relies on a dual interpretation of the method, which will be explained in Section 3.2.

The limitations of the hyperplane separation method can be observed in another family of well-understood polytopes as well, which is the subject of Section 3.3.2. Recall that a *zonotope* in \mathbb{R}^n is the Minkowski sum of a finite number of line segments, i.e., sets of the form $[x, y] := \text{conv}(\{x, y\})$ for some $x, y \in \mathbb{R}^n$. For a given graph $G = (V, E)$ on $V = [n]$, a zonotope is called a *graphic zonotope* of G if it is the Minkowski sum of $|E|$ line segments in the directions $\{e_j - e_i\}_{ij \in E}$ (see [162]). Every graphic zonotope of G is the affine linear image of the hypercube $[0, 1]^E$ and, hence, its extension complexity is at most $2|E| \leq n(n-1)$. In fact, no smaller extended formulation is known to date, not even for *completion time polytopes*, a well-known subclass of graphic zonotopes of K_n . They have been described by Wolsey [202] (who also first observed the fact that they are zonotopes; see the remark in [135]) and Queyranne [163]. For the simplest of all completion time polytopes, however, the extension complexity is known: Goemans [116] gave an asymptotically minimal extended formulation of size $\Theta(n \log n)$ for the n th *permutahedron*, which is defined as $\text{conv}\{(\sigma(1), \dots, \sigma(n)) : \sigma \in \mathfrak{S}_n\}$, where \mathfrak{S}_n denotes the symmetric group on $[n]$. The lower bound in [116] is established via a purely combinatorial argument. Since any two graphic zonotopes of K_n are combinatorially equivalent (see Section 3.3.2), $\Omega(n \log n)$ is therefore best possible for any combinatorial lower bound on the extension complexity of graphic zonotopes of K_n . In Section 3.3.2, we show that the hyperplane separation bound is at most a constant when applied to our linear description of graphic zonotopes (which generalizes the canonical linear description of completion time polytopes in [163, 202]) and the resulting slack matrix.

At the same time, we stress that our negative results do not rule out the possibility of obtaining meaningful bounds for different slack matrices. For instance, one may rescale the inequalities describing a given polytope or add redundant linear inequalities to the description. Section 3.4 studies the effect of these operations on the hyperplane separation bound. In particular, the hyperplane separation bound is not invariant under scaling the rows and columns of a given slack matrix. This is a property that is shared with the norm-based lower bounds of similar flavour introduced by Fawzi and Parrilo [96, 97]. Which scalings of rows and columns produce the strongest bounds is left as an open question in [96]. We address this issue in Section 3.4 and provide a partial answer:

If one normalizes the rows in such a way that the maximum entry in every row equals one, and proceeds analogously with the columns, the hyperplane separation bound will not decrease. Our analysis in Section 3.4 also shows that carefully adding redundant rows or columns can only increase the bound.

3.2 Preliminaries

Before we introduce the hyperplane separation bound and present some of its fundamental properties in Section 3.2.2, we first define slack matrices of polytopes more formally and briefly review some of the other techniques for lower bounding the extension complexity. Note that, even though the results below are stated in terms of polytopes only, they are valid for all polyhedra, bounded or not (see [57]).

3.2.1 Slack Matrices and Nonnegative Factorizations

Let $S \in \mathbb{R}_{\geq 0}^{m \times n}$ be a nonnegative matrix. The *nonnegative rank* of S , denoted by $\text{rk}_+(S)$, is defined as the smallest number $r \in \mathbb{N}$ such that $S = UV$ for two nonnegative matrices $U \in \mathbb{R}_{\geq 0}^{m \times r}$ and $V \in \mathbb{R}_{\geq 0}^{r \times n}$. Such a factorization of S is called a *nonnegative factorization*. The following equivalent characterization of the nonnegative rank can be found, e.g., in [55, 121] (see also [57]). We provide a brief proof.

Proposition 3.1 (Corollary 2.2 in [55], see [121]). *Let S be a nonnegative matrix. Then $\text{rk}_+(S)$ is the minimum $r \in \mathbb{N}$ such that S can be written as the sum of r nonnegative matrices of rank one.*

Proof. Let $S = UV$ for two nonnegative matrices U and V . Let us denote the columns of U by $u^{(1)}, \dots, u^{(r)}$ and the rows of V by $v^{(1)}, \dots, v^{(r)}$. Then $UV = \sum_{i=1}^r u^{(i)}(v^{(i)})^\top$ and each of the summands is easily seen to have rank (at most) one. Conversely, every nonnegative rank-one matrix must be of the form uv^\top for some nonzero column vectors u and v . \square

Now let $P \subseteq \mathbb{R}^n$ be a polytope given by $P = \text{conv}(X) = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ for some finite set $X = \{x^{(1)}, \dots, x^{(k)}\} \subseteq \mathbb{R}^n$ and $A \in \mathbb{R}^{m_A \times n}$, $b \in \mathbb{R}^{m_A}$, $B \in \mathbb{R}^{m_B \times n}$, and $d \in \mathbb{R}^{m_B}$. We further assume that each inequality in $Bx \leq d$ defines a nonempty face of P . The $m_B \times k$ matrix whose j th column equals $d - Bx^{(j)}$ is a *slack matrix* of P . If X is the set of vertices of P , we refer to the corresponding slack matrix as the slack matrix of P with respect to the linear description above. In particular, slack matrices of polytopes are nonnegative by definition. Yannakakis [205] showed a striking relationship between their nonnegative rank and the extension complexity of the associated polytopes.

Proposition 3.2 ([205]). *Let S be a slack matrix of a polytope P . Then $\text{xc}(P) = \text{rk}_+(S)$.*

Yannakakis' result is a powerful tool, especially for lower bounding the extension complexity of polytopes, since one may work with the nonnegative rank instead. However, computing this quantity for a given nonnegative matrix S is a difficult problem by itself: Deciding whether $\text{rk}_+(S)$ equals the (linear) rank of S , denoted by $\text{rk}(S)$, is NP-hard [195]. Note that we always have $\text{rk}_+(S) \geq \text{rk}(S)$ (see [55]). If S is the slack matrix of a

polytope P of dimension $d \geq 1$, then $\text{rk}(S) = d + 1$ (see Theorem 14 in [119]). Thus, by Proposition 3.2, the dimension of a polytope is a first, simple lower bound on its extension complexity, as mentioned in Section 3.1.

A more refined bound considers the support of $S = (s_{ij})$, which – analogously to the support of a vector – is defined as $\text{supp}(S) = \{(i, j) : s_{ij} \neq 0\}$. By Proposition 3.1, the nonnegative rank of S is the minimum number of nonnegative rank-one matrices whose sum is S . Observe that the support of each summand in such an additive decomposition is a *rectangle*, i.e., a set of the form $I \times J$ for some nonempty subsets of rows I and columns J . All rectangles together must cover $\text{supp}(S)$. So the minimum number of rectangles needed to cover $\text{supp}(S)$ is a lower bound on $\text{rk}_+(S)$, the so-called *rectangle covering lower bound* that was already mentioned in Section 3.1; see [100] for a detailed study of this bound (see also [121, 205]).

The rectangle covering lower bound has been successfully employed, e.g., in [101, 136] to prove exponential lower bounds on the extension complexity of correlation polytopes, which appear as faces or linear images of cut, stable set, and TSP polytopes. For matching polytopes, however, this lower bound performs rather poorly: The support of their slack matrices may be covered using only a number of rectangles that is polynomial in the number of nodes of the graph, as observed in [205] (see also [100, 167]). Rothvoß' breakthrough result of [167], showing that the extension complexity of the matching polytope of K_n is $2^{\Omega(n)}$, therefore relies on a different bounding technique that was first proposed by Fiorini [99], the so-called *hyperplane separation bound*. This is the technique that this chapter is concerned with.

3.2.2 The Hyperplane Separation Bound

To formally state the hyperplane separation bound, we need the following notation. For two matrices $A = (a_{ij}), B = (b_{ij}) \in \mathbb{R}^{m \times n}$, we let $\|A\|_\infty := \max_{i,j} |a_{ij}|$ and denote by $\langle A, B \rangle$ the *Frobenius inner product* of A and B , i.e., $\langle A, B \rangle := \sum_{i=1}^m \sum_{j=1}^n a_{ij} b_{ij}$. Throughout this chapter, we will also use the same notation $\langle a, b \rangle$ for the inner product $a^\top b$ of two vectors $a, b \in \mathbb{R}^n$.

Proposition 3.3 ([99], see [167]). *Let $S \in \mathbb{R}_{\geq 0}^{m \times n}$ not identically zero, and let $\mathcal{R}_{m,n}$ denote the set of rank-one matrices in $\{0, 1\}^{m \times n}$. We further let*

$$\text{hsb}(S) := \sup \left\{ \frac{\langle S, X \rangle}{\|S\|_\infty \rho(X)} : X \in \mathbb{R}^{m \times n} \right\}, \quad (3.3)$$

where $\rho(X) := \max \{\langle R, X \rangle : R \in \mathcal{R}_{m,n}\}$ for every $X \in \mathbb{R}^{m \times n}$. Then $\text{rk}_+(S) \geq \text{hsb}(S)$.

By a slight abuse of terminology, we call the quantity $\text{hsb}(S)$ itself the *hyperplane separation bound* (of S). After normalizing X such that $\rho(X) = 1$ in the definition of $\text{hsb}(S)$, we may rewrite (3.3) as follows:

$$\begin{aligned} \|S\|_\infty \text{hsb}(S) &= \sup \{ \langle S, X \rangle : X \in \mathbb{R}^{m \times n}, \rho(X) = 1 \} \\ &= \sup \{ \langle S, X \rangle : X \in \mathbb{R}^{m \times n}, \rho(X) \leq 1 \} \\ &= \max \{ \langle S, X \rangle : X \in \mathbb{R}_{\geq 0}^{m \times n}, \langle X, R \rangle \leq 1 \ \forall R \in \mathcal{R}_{m,n} \}. \end{aligned} \quad (3.4)$$

In the last step, we used the fact that the supremum of $\langle S, \cdot \rangle$ is finite: Any $X \in \mathbb{R}^{m \times n}$ with $\rho(X) \leq 1$ satisfies $\langle R, X \rangle \leq 1$ for all R with singleton support, that is, every entry of X is at most one. As S is nonnegative, the sum of its entries is an upper bound on $\langle S, X \rangle$.

Note that (3.4) is an LP. From strong LP duality, we obtain the following dual characterization of the hyperplane separation bound. This already appears in [181], even though it is stated in a slightly differently form (in terms of the Minkowski gauge function of the set $\text{conv}(\mathcal{R}_{m,n})$, to be precise).

Proposition 3.4 (see [181]). *Let S and $\mathcal{R}_{m,n} =: \mathcal{R}$ be defined as in Proposition 3.3. Then*

$$\text{hsb}(S) = \min \left\{ \|S\|_\infty^{-1} \sum_{R \in \mathcal{R}} y_R : y \in \mathbb{R}_{\geq 0}^{\mathcal{R}}, \sum_{R \in \mathcal{R}} y_R R = S \right\}. \quad (3.5)$$

With this dual characterization of the hyperplane separation bound, we may also give a simple proof of Proposition 3.3.

Proof of Proposition 3.3. We start by showing the following useful fact, a slightly stronger version of which is used in [167].

Claim ([167]). $\text{conv}(\mathcal{R}_{m,n}) \supseteq \{A \in [0, 1]^{m \times n} : \text{rk}(A) = 1\}$.

Proof of Claim. Let $A \in [0, 1]^{m \times n}$ be of rank one. Then $A = vw^\top$ for nonnegative vectors $v \in \mathbb{R}^m, w \in \mathbb{R}^n$. Possibly after rescaling, both v and w may be assumed to be $[0, 1]$ -valued. Then $v \in [0, 1]^m$ can be written as a convex combination $v = \sum_{i=1}^p \lambda_i v^{(i)}$ for $v^{(i)} \in \{0, 1\}^m, \lambda_i \geq 0$, and $\sum_{i=1}^p \lambda_i = 1$. Similarly, $w = \sum_{j=1}^q \mu_j w^{(j)}$ for $w^{(j)} \in \{0, 1\}^n, \mu_j \geq 0$, and $\sum_{j=1}^q \mu_j = 1$. It follows that

$$A = vw^\top = \sum_{i=1}^p \sum_{j=1}^q \lambda_i \mu_j \cdot \underbrace{v^{(i)}(w^{(j)})^\top}_{\in \mathcal{R}_{m,n}} \quad \text{and} \quad \sum_{i=1}^p \sum_{j=1}^q \lambda_i \mu_j = 1.$$

Hence, A is a convex combination of matrices from $\mathcal{R}_{m,n}$ as desired. \diamond

Using the claim, we may prove that $\text{hsb}(S) \leq \text{rk}_+(S)$ as follows. Since the hyperplane separation bound is invariant under scaling S (see Section 3.2.2), we may assume without loss of generality that $\|S\|_\infty = 1$. Let $r := \text{rk}_+(S)$. Then, by Proposition 3.1, S is the sum of r nonnegative rank-one matrices $A_1, \dots, A_r \in \mathbb{R}_{\geq 0}^{m \times n}$. For all $k \in [r]$, we even have that $A_k \in [0, 1]^{m \times n}$ since $\|A_k\|_\infty \leq \|S\|_\infty = 1$. Writing $\mathcal{R} := \mathcal{R}_{m,n}$ for short, the claim above implies that $A_k = \sum_{R \in \mathcal{R}} y_R^k R$ for some coefficients $y_R^k \geq 0$ with $\sum_{R \in \mathcal{R}} y_R^k = 1$, for all $k \in [r]$. This means that the vector $y \in \mathbb{R}^{\mathcal{R}}$ defined by $y_R = \sum_{k=1}^r y_R^k$ for all $R \in \mathcal{R}$ is a feasible solution of the LP in (3.5). Hence, $\text{hsb}(S) \leq \sum_{R \in \mathcal{R}} y_R = r$ by Proposition 3.4. \square

With the primal and dual characterizations of the hyperplane separation bound stated in (3.4) and (3.5), it is easy to determine the exact value of $\text{hsb}(S)$ for small matrices S .

Example 3.5. Consider the matrix

$$S = \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$$

and let

$$X = \frac{1}{4} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

It is easily verified that $\langle R, X \rangle \leq 1$ for all 0/1 matrices of rank one in $\mathcal{R}_{2,2}$. Hence, $\text{hsb}(S) \geq \frac{1}{2} \langle S, X \rangle = \frac{3}{2}$ by (3.4). On the other hand, we can express S as

$$S = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} + \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

where all summands are in $\mathcal{R}_{2,2}$. It follows from (3.5) that $\text{hsb}(S) \leq \frac{3}{2}$ and, thus, $\text{hsb}(S) = \frac{3}{2}$. \diamond

Note that the feasible region of the dual LP in (3.5) corresponds to a particular type of nonnegative factorization of S , namely the decomposition of S into the weighted sum of 0/1 matrices of rank one. In particular, if all weights are 0 or 1, such a decomposition is equivalent to a factorization of S with 0/1 factors by the proof of Proposition 3.1. We therefore obtain the following corollary to Proposition 3.4, which will be the key ingredient of our proofs in Section 3.3.

Corollary 3.6. *Let S be a nonnegative matrix. If S has a nonnegative factorization of rank r with both factors being 0/1 matrices, then $\text{hsb}(S) \leq r \|S\|_{\infty}^{-1}$.*

Before we proceed, let us remark that, by its definition, the hyperplane separation bound $\text{hsb}(S)$ is invariant under multiplying S by positive scalars, under transposition, and under permutations of rows and columns of S , respectively. It further satisfies the following two useful properties on submatrices, both of which are immediate consequences of Proposition 3.4.

Lemma 3.7. *Let $S = \begin{pmatrix} A \\ B \end{pmatrix}$ for nonnegative matrices A and B . Then*

$$(i) \quad \|S\|_{\infty} \text{hsb}(S) \geq \|A\|_{\infty} \text{hsb}(A),$$

$$(ii) \quad \text{hsb}(S) \leq \text{hsb}(A) + \text{hsb}(B). \text{ If } A = B, \text{ then } \text{hsb}(S) = \text{hsb}(A).$$

Proof. Observe that a matrix $R \in \{0, 1\}^{m \times n}$ of rank one is easily turned into a 0/1 matrix of rank (at most) one in a different dimension $m' \times n$, either by padding R with zero rows (if $m' > m$) or by deleting rows (if $m' < m$). The second operation, when applied to each matrix R in the support of an optimal solution of (3.5) for S , shows (i). To prove (ii), we apply the first operation to optimal solutions for A and B , respectively, and obtain

$$\|S\|_{\infty} \text{hsb}(S) \leq \|A\|_{\infty} \text{hsb}(A) + \|B\|_{\infty} \text{hsb}(B).$$

Since $\|S\|_{\infty} = \max\{\|A\|_{\infty}, \|B\|_{\infty}\}$, the statement follows. \square

3.3 Limitations of the Hyperplane Separation Bound

In the previous section, we saw that the hyperplane separation bound is small for nonnegative matrices that admit 0/1 factorizations of low rank. Here, we will apply this observation to the canonical slack matrices of two families of polytopes, spanning tree polytopes and graphic zonotopes. In both cases, we will see that the best known upper bounds on their nonnegative rank (see Section 4.1) can, in fact, be achieved via a factorization into 0/1 matrices.

3.3.1 Spanning Tree Polytopes

Let $G = (V, E)$ be a connected graph. Recall from Section 3.1 that the spanning tree polytope of G , denoted by $P_{\text{st}}(G) \subseteq \mathbb{R}^E$, is given by

$$x(E) = |V| - 1 \quad (3.6)$$

$$x(E(U)) \leq |U| - 1 \quad \text{for all } \emptyset \neq U \subseteq V \quad (3.7)$$

$$x \geq \mathbf{0} \quad (3.8)$$

Theorem 3.8. *Let $G = (V, E)$ be a connected graph and let S_G denote the slack matrix of $P_{\text{st}}(G)$ with respect to the description (3.6)–(3.8). Then $\text{hsb}(S_G) = O(|E|)$.*

Proof. Since there are $|E|$ many nonnegativity constraints in (3.8), it suffices to consider the row submatrix of S_G restricted to inequalities (3.7) only, which we will denote by S_G again. The bound for the entire slack matrix then follows from Lemma 3.7(ii).

Let us index the rows of S_G by the nonempty subsets of V and the columns by the spanning trees in G . The entry in row $U \subseteq V$ and column T equals $c(U, T) - 1$, where $c(U, T)$ denotes the number of connected components of the subgraph $(U, T \cap E(U))$. First, observe that

$$\|S_G\|_\infty \geq \frac{1}{2}|V| - 1. \quad (3.9)$$

Indeed, if T is a spanning tree in G and $U \subseteq V$ a stable set in T , then $c(U, T) = |U|$. Because T is a bipartite graph, both node classes in a bipartition are stable sets in T . At least one of them must be of size $\frac{1}{2}|V|$.

Based on the extended formulation given in [153], Conforti, Cornuéjols, and Zambelli [56] showed that S_G admits a nonnegative factorization of rank $O(|V||E|)$ where both factors are 0/1 matrices. For the sake of clarity, we include a proof. Even though our proof is slightly different than the one of [56], it is also inspired by [153].

Claim ([56]). *There is a nonnegative factorization of S_G of rank $2|V||E|$ where both factors are 0/1 matrices.*

Proof of Claim. For a spanning tree T in G , let $\tau(T)$ be the set of all triples of nodes $(i, j, k) \in V^3$ such that j is the neighbour of i on the unique i - k -path in T . From each nonempty subset $U \subseteq V$, we choose an arbitrary representative $k(U) \in U$. For every triple $(i, j, k) \in V^3$ where $ij \in E$, define the set

$$R(i, j, k) := \{U \subseteq V : i \in U, j \notin U, k = k(U)\} \times \{T \text{ spanning tree} : (i, j, k) \in \tau(T)\}.$$

For each of these $2|V||E|$ triples (i, j, k) , there is a unique 0/1 matrix indexed in the same way as S_G whose support is $R(i, j, k)$. We claim that these matrices, which clearly are of rank at most one, add up to S_G . To see this, let $\emptyset \neq U \subseteq V$ and T be a spanning tree in G , and let $c = c(U, T)$ for short. It suffices to show that

$$|\{(i, j, k) \in V^3 : ij \in E, (U, T) \in R(i, j, k)\}| = c - 1.$$

If $c = 1$, there is nothing to prove. Now let $c \geq 2$, and let F_1, \dots, F_c be the connected components of the subgraph $(U, T \cap E(U))$. Without loss of generality, we may assume that $k(U) \in V(F_c)$. For $l \in [c - 1]$, we say that a path in T *connects* $k(U)$ and F_l if one endpoint of the path is $k(U)$, the other one is a node in $V(F_l)$, and no internal node of the path belongs to $V(F_l)$. For all $l \in [c - 1]$, there exists a unique path in T connecting $k(U)$ and F_l . Let i_l be its endpoint in $V(F_l)$, and let j_l be the neighbour of i_l on the path. Then $j_l \notin U$, and $(U, T) \in R(i_l, j_l, k(U))$ for all $l \in [c - 1]$.

Conversely, if $(U, T) \in R(i, j, k)$ for some $(i, j, k) \in V^3$ with $ij \in E$, then $k = k(U)$ and $i \notin V(F_c)$, say, $i \in V(F_1)$. Since $j \notin U$, the path connecting i and $k(U)$ in T cannot visit any other node in $V(F_1)$. Hence, it connects $k(U)$ and F_1 and we conclude that $(i, j) = (i_1, j_1)$.

This shows that the sets $R(i, j, k)$ induce a decomposition of S_G into $2|V||E|$ summands which are 0/1 matrices. The claim follows from Proposition 3.1. \diamond

Combining the claim with (3.9), it follows from Corollary 3.6 that

$$\text{hsb}(S_G) \leq \frac{2|V||E|}{|V|/2 - 1} = O(|E|).$$

This concludes the proof of the theorem. \square

3.3.2 Graphic Zonotopes

Recall from Section 3.1 that a graphic zonotope of a graph $G = (V, E)$ with $V = [n]$ is the Minkowski sum of a finite number of line segments, each of which is parallel to the difference of two standard basis vectors $e_j - e_i$ for some $ij \in E$. Let $A = (a_{ij}) \in \mathbb{R}_{\geq 0}^{n \times n}$ be symmetric and nonnegative. With the matrix A , we associate a zonotope $Z(A) \subseteq \mathbb{R}^n$ as follows:

$$Z(A) := \sum_{1 \leq j \leq n} a_{jj} e_j + \sum_{1 \leq i < j \leq n} a_{ij} [e_i, e_j]. \quad (3.10)$$

Up to translations, the graphic zonotopes of graphs on n nodes are exactly those of the form (3.10) for some symmetric and nonnegative $n \times n$ matrix A (where $a_{ij} > 0$ if and only if $ij \in E$).

Graphic zonotopes are generalized permutahedra: For $n \in \mathbb{N}$, the n th permutahedron is the convex hull of the vectors $(\sigma(1), \dots, \sigma(n)) \in \mathbb{Z}^n$ for all permutations $\sigma \in \mathfrak{S}_n$. It is a well-known fact that permutahedra are zonotopes (see, e.g., [206]); in fact, the n th permutahedron equals $Z(A)$ for the $n \times n$ all-one matrix A (see Example 7.15 in [206]). A minimal linear description of the n th permutahedron is the following (see, e.g., [104, Section 2.2]):

$$\left\{ x \in \mathbb{R}^n : x([n]) = \binom{n+1}{2}, x(S) \geq \binom{|S|+1}{2} \text{ for all } S \subseteq [n] \right\}$$

Let us now generalize this description to all graphic zonotopes $Z(A)$. For a symmetric nonnegative matrix $A \in \mathbb{R}_{\geq 0}^{n \times n}$, we define the set function $g_A: 2^{[n]} \rightarrow \mathbb{R}$ by

$$[n] \supseteq S \mapsto g_A(S) := \sum_{\substack{i,j \in S: \\ i \leq j}} a_{ij}.$$

We will first argue that g_A is supermodular. Recall that a set function $g: 2^{[n]} \rightarrow \mathbb{R}$ is *supermodular* (see [104]) if

$$g(S \cup T) + g(S \cap T) \geq g(S) + g(T) \quad \text{for all } S, T \subseteq [n]. \quad (3.11)$$

Observe that g_A is supermodular if and only if the set function \bar{g}_A defined by $\bar{g}_A(S) = 2g_A(S) - \sum_{i \in S} a_{ii}$ for $S \subseteq [n]$ is supermodular. Since A is symmetric by hypothesis, $\bar{g}_A(S)$ is precisely the sum of all entries of the square submatrix of A obtained after deleting all rows and columns whose index is not in S . Supermodularity of \bar{g}_A follows from the fact that all entries of A are nonnegative.

The *supermodular base polytope* (see, e.g., [104]) of a supermodular function $g: 2^{[n]} \rightarrow \mathbb{R}$ with $g(\emptyset) = 0$ is defined as

$$B(g) = \{x \in \mathbb{R}^n : x([n]) = g([n]), x(S) \geq g(S) \text{ for all } S \subseteq [n]\}. \quad (3.12)$$

Using standard arguments (see [81]), one can show that for a symmetric nonnegative matrix A , the zonotope $Z(A)$ is the supermodular base polytope of g_A . For the sake of completeness, we include a short proof of this fact.

Lemma 3.9. *Let $A \in \mathbb{R}_{\geq 0}^{n \times n}$ be symmetric. Then $Z(A) = B(g_A)$.*

Proof. It suffices to show that, for every linear objective function $w \in \mathbb{R}^n$, the minima of w over $Z(A)$ and $B(g_A)$ coincide. After a permutation of the coefficients of w , we may assume that $w_1 \geq \dots \geq w_n$. The greedy rule (see [81]) then implies that a minimizer x^* over $B(g_A)$ is given by

$$x_j^* := g_A([j]) - g_A([j-1]) = \sum_{i=1}^j a_{ij}, \quad j \in [n].$$

Minimizing w over the zonotope $Z(A)$ can be done over each summand in the Minkowski sum in (3.10) individually. For $1 \leq i < j \leq n$, it is easy to see that the minimum of w on the line segment $[a_{ij}e_j, a_{ij}e_i]$ is attained in the first endpoint since $w_i \geq w_j$. Hence, the minimum of w over $Z(A)$, too, is attained in x^* . \square

From the proof of Lemma 3.9, we conclude that the vertices of $Z(A)$ are in correspondence with the permutations in \mathfrak{S}_n via the map

$$\mathfrak{S}_n \ni \sigma \quad \mapsto \quad x^\sigma \in \mathbb{R}^n; \quad x_j^\sigma = \sum_{\substack{i \in [n]: \\ \sigma(i) \leq \sigma(j)}} a_{ij}, \quad j \in [n]. \quad (3.13)$$

Now define a matrix M_A with one row for every nonempty proper subset of $[n]$ and one column for every permutation in \mathfrak{S}_n as follows: If x^σ denotes the vertex of $Z(A)$

induced by $\sigma \in \mathfrak{S}_n$ via (3.13), the entry of M_A in row $S \subsetneq [n]$, $S \neq \emptyset$, and column σ equals

$$x^\sigma(S) - g_A(S) = \sum_{\substack{i \in [n], j \in S: \\ \sigma(i) \leq \sigma(j)}} a_{ij} - \sum_{\substack{i, j \in S: \\ \sigma(i) \leq \sigma(j)}} a_{ij} = \sum_{\substack{i \notin S, j \in S: \\ \sigma(i) \leq \sigma(j)}} a_{ij}, \quad (3.14)$$

using symmetry of A in the first equation. Thus, M_A is precisely the slack matrix of $Z(A)$ with respect to the linear description (3.12), possibly with repeated columns.

Before we state the main result of this section, observe that the support of M_A is independent of the actual entries of A if A is strictly positive. In this case, g_A is strictly supermodular, that is, the inequality in (3.11) is strict for all subsets $S, T \subseteq [n]$ neither of which is contained in the other. It follows that all inequalities in (3.12) for $\emptyset \neq S \subsetneq [n]$ define facets of $Z(A)$ [163]. Further, the map (3.13) is a bijection if A is strictly positive, which implies that $Z(A)$ and the n th permutahedron (and, more generally, $Z(A')$ for any other symmetric and strictly positive $n \times n$ matrix A') are combinatorially equivalent.

Theorem 3.10. *Let $A \in \mathbb{R}_{\geq 0}^{n \times n}$ be symmetric, and let M_A be the slack matrix of $Z(A)$ with respect to (3.12). Then $\overline{\text{hsb}}(M_A) \leq 4$.*

Proof. For every pair $i, j \in [n]$, $i \neq j$, let

$$R(i, j) := \{S \subseteq [n] : i \notin S, j \in S\} \times \{\sigma \in \mathfrak{S}_n : \sigma(i) \leq \sigma(j)\}$$

and let $\widehat{R}(i, j)$ denote the unique 0/1 matrix indexed like M_A whose support equals $R(i, j)$. Note that $\widehat{R}(i, j)$ has rank one and, by (3.14),

$$\sum_{i \neq j} a_{ij} \widehat{R}(i, j) = M_A.$$

Since the expression in (3.14) is less than or equal to $\sum_{i \notin S, j \in S} a_{ij}$ with equality if $\sigma([n] \setminus S) = [n - |S|]$, we have that

$$\|M_A\|_\infty = \max_{S \subseteq [n]} \sum_{i \notin S, j \in S} a_{ij}.$$

This quantity is, in fact, the maximum weight of a cut in the graph underlying the graphic zonotope $Z(A)$ where the (nonnegative) edge weights are given by A . Since there is always a cut whose weight is at least half the total weight of the edges (see, e.g., Theorem 5.3 in [201]), we conclude that

$$\|M_A\|_\infty \geq \frac{1}{2} \sum_{i < j} a_{ij} = \frac{1}{4} \sum_{i \neq j} a_{ij}.$$

The theorem follows from an application of Corollary 3.6. \square

In particular, Theorem 3.10 applies to completion time polytopes. Let us briefly discuss how they fit into the framework of graphic zonotopes. Consider n jobs with processing times $p = (p_1, \dots, p_n) \in \mathbb{R}_{> 0}^n$ to be scheduled on a single machine. Every permutation $\sigma \in \mathfrak{S}_n$ defines a feasible schedule without idle time where job j is

completed at time $C_j^\sigma := \sum_{i: \sigma(i) \leq \sigma(j)} p_i$ for $j \in [n]$. The *completion time polytope* associated with p is denoted by $P_{\text{ct}}(p)$ and is defined as

$$P_{\text{ct}}(p) := \text{conv} \{(C_1^\sigma, \dots, C_n^\sigma) \in \mathbb{R}^n : \sigma \in \mathfrak{S}_n\}.$$

Now let $A \in \mathbb{R}_{>0}^{n \times n}$ be a symmetric positive rank-one matrix. Then A can be written as $A = pp^\top$ for some $p \in \mathbb{R}_{>0}^n$. This can be seen as follows. Since the matrix A has rank one, it must be of the form $A = uv^\top$ for some $u, v \in \mathbb{R}_{>0}^n$. Note that $u_1 v = v_1 u$ by symmetry of A . As $u_1, v_1 > 0$, we may define $p := (\sqrt{v_1/u_1})u$ and obtain $A = pp^\top$ as desired. This means that $Z(A)$ is the image of $P_{\text{ct}}(p)$ under the linear transformation $(x_1, \dots, x_n) \mapsto (p_1 x_1, \dots, p_n x_n)$. Up to this transformation, the inequalities in (3.12) coincide with the canonical linear description of $P_{\text{ct}}(p)$ in [163, 202]. When $p_j = 1$ for all $j \in [n]$, the completion time polytope $P_{\text{ct}}(p)$ is the n th permutahedron and, as we saw earlier, is equal to $Z(A)$ with the $n \times n$ all-one matrix for A .

This special case also shows that our analysis in the proof of Theorem 3.10 is asymptotically tight. Namely, if A is the $n \times n$ all-one matrix, then

$$\|M_A\|_\infty = \left\lfloor \frac{n}{2} \right\rfloor \cdot \left\lceil \frac{n}{2} \right\rceil = \begin{cases} \frac{1}{4}(n-1)(n+1) & \text{if } n \text{ odd} \\ \frac{1}{4}n^2 & \text{if } n \text{ even} \end{cases}$$

and therefore

$$\|M_A\|_\infty^{-1} \sum_{i \neq j} a_{ij} = \begin{cases} 4 \frac{n}{n+1} & \text{if } n \text{ odd} \\ 4 \frac{n-1}{n} & \text{if } n \text{ even} \end{cases}$$

which tends to 4 as $n \rightarrow \infty$.

Both Theorem 3.10 and the analogous result for spanning tree polytopes stated in Theorem 3.8 crucially rely on the fact that the associated slack matrices admit 0/1 factorizations of low rank. The nonnegative rank itself, of course, does not depend on which slack matrix we choose for a given polytope. Recall that this is a consequence of Proposition 3.2. The situation for the hyperplane separation bound, however, is fundamentally different, as we will see next.

3.4 Diagonal Scalings and Redundancy

Before we address the issue of choosing a slack matrix more rigorously, let us highlight the difference between the nonnegative rank and the hyperplane separation bound with two examples.

Example 3.11. Consider the standard hypercube $C_n = [0, 1]^n$ and let S_n denote its slack matrix with respect to the (minimal) description $C_n = \{x \in \mathbb{R}^n : \mathbf{0} \leq x \leq \mathbf{1}\}$. The inequality $\mathbf{1}^\top x \geq 0$ is valid for C_n (defining the vertex $\mathbf{0}$). Adding this inequality to the minimal description of C_n adds one row to S_n , which equals the sum of the rows corresponding to the n facets defined by $x \geq \mathbf{0}$. Let S'_n denote the slack matrix with this additional row. Then we have $\|S'_n\|_\infty = n$ and $\|S_n\|_\infty \text{hsb}(S_n) = \|S'_n\|_\infty \text{hsb}(S'_n)$. Thus, $\text{hsb}(S'_n) = \frac{1}{n} \text{hsb}(S_n)$. \diamond

Not even slack matrices with respect to *minimal* linear descriptions behave identically under the hyperplane separation bound:

Example 3.12. The n -simplex spanned by the standard basis vectors in \mathbb{R}^n and the origin is the set of all $x \in \mathbb{R}^n$ satisfying

$$\begin{aligned} x_1 + \cdots + x_n &\leq 1 \\ x_i &\geq 0 \quad \text{for all } i \in [n-1] \\ \lambda x_n &\geq 0 \end{aligned}$$

for any $\lambda > 0$. Each inequality is facet-defining. Modulo permutations of rows and columns, the associated slack matrix $S_{n,\lambda}$ is obtained from the $(n+1) \times (n+1)$ identity by multiplying the first row by λ . Clearly, $n+1 \geq \text{rk}_+(S_{n,\lambda}) \geq \text{rk}(S_{n,\lambda}) = n+1$ (see Section 3.2), so we must have equality everywhere.

Now suppose that $\lambda \geq 1$. Then $\|S_{n,\lambda}\|_\infty = \lambda$ and $\text{hsb}(S_{n,\lambda}) \leq \frac{n}{\lambda} + 1$, since $S_{n,\lambda}$ is the sum of $n+1$ matrices with a single nonzero entry each on the diagonal. In fact, one can even show that $\text{hsb}(S_{n,\lambda}) = \frac{n}{\lambda} + 1$ by noting that the $(n+1) \times (n+1)$ matrix X with 1 on the diagonal and -1 elsewhere is a solution to the LP (3.4) with $\langle S_{n,\lambda}, X \rangle = \lambda + n$. Thus, $\text{hsb}(S_{n,\lambda})$ may be arbitrarily bad compared to $\text{rk}_+(S_{n,\lambda})$. \diamond

Let us now generalize these observations. We first study the effect that scaling rows or columns of a nonnegative matrix S has on $\text{hsb}(S)$. A *positive diagonal scaling* of S is a matrix S' which can be written as $S' = D_1 S D_2$ where D_1 and D_2 are positive diagonal matrices, i.e., diagonal matrices with positive diagonal elements. Note that $\text{rk}_+(S') = \text{rk}_+(S)$ (see Lemma 2.8 in [55]), while in Example 3.12 we have seen that the hyperplane separation bound may indeed change (see also [96, 97]). The following theorem is our main ingredient in this section.

Theorem 3.13. *Let $S \in \mathbb{R}_{\geq 0}^{m \times n}$ not identically zero, and let $D \in \mathbb{R}_{\geq 0}^{m \times m}$ be a positive diagonal matrix. Then*

$$\text{hsb}(DS) \leq \text{hsb}(S) \frac{\|D\|_\infty \|S\|_\infty}{\|DS\|_\infty}.$$

Proof. Let $X \in \mathbb{R}^{m \times n}$ be a feasible solution of the LP in (3.4) that is optimal for $S' := DS$. We will denote the i th rows of S and X by $s^{(i)}$ and $x^{(i)}$, respectively, and the i th diagonal element of D by $d_i > 0$. First, observe that $\langle s^{(i)}, x^{(i)} \rangle \geq 0$ for all $i \in [m]$: Indeed, if $\langle s^{(i)}, x^{(i)} \rangle < 0$ for some i , let S'_{-i} and X_{-i} be the matrices obtained from S' and X by deleting the i th row. Then

$$\begin{aligned} \|S'_{-i}\|_\infty \text{hsb}(S'_{-i}) &\geq \langle S'_{-i}, X_{-i} \rangle \\ &= \langle S', X \rangle - d_i \langle s^{(i)}, x^{(i)} \rangle \\ &> \langle S', X \rangle = \|S'\|_\infty \text{hsb}(S'), \end{aligned}$$

contradicting Lemma 3.7(i). We conclude that

$$\begin{aligned} \langle S', X \rangle &= \sum_{i=1}^m d_i \langle s^{(i)}, x^{(i)} \rangle \leq \sum_{i=1}^m \|D\|_\infty \langle s^{(i)}, x^{(i)} \rangle = \|D\|_\infty \langle S, X \rangle \\ &\leq \|D\|_\infty \|S\|_\infty \text{hsb}(S). \end{aligned}$$

□

Since the hyperplane separation bound is invariant under transposition (cf. Section 3.2), Theorem 3.13 immediately generalizes to positive diagonal scalings.

Corollary 3.14. *Let $S \in \mathbb{R}_{\geq 0}^{m \times n}$ not identically zero, and let $D_1 \in \mathbb{R}_{\geq 0}^{m \times m}$, $D_2 \in \mathbb{R}_{\geq 0}^{n \times n}$ be positive diagonal matrices. Then*

$$\frac{\|S\|_\infty}{\|D_1^{-1}\|_\infty \|D_1 S D_2\|_\infty \|D_2^{-1}\|_\infty} \text{hsb}(S) \leq \text{hsb}(D_1 S D_2) \leq \text{hsb}(S) \frac{\|D_1\|_\infty \|S\|_\infty \|D_2\|_\infty}{\|D_1 S D_2\|_\infty}.$$

Proof. Follows from Theorem 3.13, using the fact that $S = D_1^{-1}(D_1 S D_2)D_2^{-1}$. \square

Consider a positive diagonal scaling S' of S whose nonzero rows and columns are normalized with respect to $\|\cdot\|_\infty$. We say that S' is both *row-* and *column-normalized*. If $S' = D_1 S D_2$ with D_1 and D_2 chosen in such a way that $D_1 S$ is row-normalized or $S D_2$ is column-normalized, then $\|D_1^{-1}\|_\infty \|D_2^{-1}\|_\infty = \|S\|_\infty$ and Corollary 3.14 implies that $\text{hsb}(S') \geq \text{hsb}(S)$. In other words, if one first normalizes the rows of S and then normalizes the columns of the resulting matrix (or vice versa), $\text{hsb}(S)$ will only increase. In general, though, not every row- and column-normalized diagonal scaling S' results from a pair of diagonal matrices that satisfy this additional requirement:

Example 3.15 (Example 3.5 continued). For the matrix

$$S = \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix},$$

we computed $\text{hsb}(S) = \frac{3}{2}$ in Example 3.5. If we normalize the rows (or the columns) of S , we obtain the row- and column-normalized rescaling $\frac{1}{2}S$, whose hyperplane separation bound is the same as that of S (see Section 3.2.2). Now consider the positive diagonal matrix

$$D = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 1 \end{pmatrix}.$$

If we rescale both rows and columns of S simultaneously using D , the resulting matrix is

$$S' = D S D = \begin{pmatrix} \frac{1}{4} & 1 \\ 1 & 1 \end{pmatrix}.$$

Like $\frac{1}{2}S$, the matrix S' is row- and column-normalized. However, neither is $D S$ row-normalized nor is $S D$ column-normalized.

In fact, one can show that $\text{hsb}(S') = \frac{7}{4} > \frac{3}{2} = \text{hsb}(S)$ by exhibiting a primal-dual pair of solutions to the LPs (3.4) and (3.5) (as we did in Examples 3.5 and 3.12). Indeed, the matrix

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

is a primal feasible solution of (3.4) with $\langle S', X \rangle = \frac{7}{4}$. A corresponding dual solution is given by the following decomposition of S' :

$$S' = \frac{1}{4} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} + \frac{3}{4} \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix} + \frac{3}{4} \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}.$$

\diamond

Rescaling rows and columns is not the only operation that has an effect on the hyperplane separation bound of a nonnegative matrix S . Suppose that we add to S a row (column) which is a nonnegative linear combination of rows (columns) of S and is therefore redundant (cf. Example 3.11). This operation, too, leaves the nonnegative rank of S unchanged (see Lemmas 2.6 and 2.7 in [55]). The next result bounds the gain on $\text{hsb}(S)$.

Theorem 3.16. *Let $S \in \mathbb{R}_{\geq 0}^{m \times n}$ and $S' := \begin{pmatrix} S \\ w^\top S \end{pmatrix}$ for a vector $w \in \mathbb{R}_{\geq 0}^m$ such that $\|w^\top S\|_\infty \leq \|S\|_\infty$. Then we have*

$$\text{hsb}(S) \leq \text{hsb}(S') \leq \text{hsb}(S) \max\{1, \|w\|_1\}.$$

Proof. The first inequality follows from Lemma 3.7(i). In order to show the second inequality, suppose first that $\|w\|_1 \leq 1$. Let $X \in \mathbb{R}^{m \times n}$, $x \in \mathbb{R}^m$ such that $\begin{pmatrix} X \\ x \end{pmatrix} \in \mathbb{R}^{(m+1) \times n}$ is an optimal solution of the LP in (3.4) for S' . Adding $w_i x$ to the i th row of X for every $i \in [m]$, we obtain a matrix $X' \in \mathbb{R}^{m \times n}$ satisfying $\langle S, X' \rangle = \langle S, X \rangle + \langle w^\top S, x \rangle$. It remains to show that X' is a feasible solution of the LP in (3.4) for S . To this end, let $R \in \mathcal{R}_{m,n}$, and let $I \subseteq [m]$ denote the set of indices of rows that R is supported in. Note that all rows in I are equal to some $r \in \{0, 1\}^n$. Hence,

$$\langle X', R \rangle = \langle X, R \rangle + \langle x, r \rangle \sum_{i \in I} w_i.$$

If $\langle x, r \rangle \leq 0$, then $\langle X', R \rangle \leq \langle X, R \rangle \leq 1$ because X is feasible. Otherwise, $\langle X', R \rangle \leq \langle X, R \rangle + \langle x, r \rangle \leq 1$ since $\sum_{i \in I} w_i \leq \|w\|_1 \leq 1$ and $\begin{pmatrix} R \\ r \end{pmatrix} \in \mathcal{R}_{m+1,n}$.

Now suppose that $\|w\|_1 > 1$. Let $\bar{w} := \|w\|_1^{-1} w$ and consider the matrix

$$S'' := \begin{pmatrix} S \\ \bar{w}^\top S \end{pmatrix}.$$

Since $\|\bar{w}\|_1 = 1$, the first case of the proof implies that $\text{hsb}(S'') = \text{hsb}(S)$. At the same time, S' is a diagonal scaling of S'' where the maximum diagonal element equals $\max\{1, \|w\|_1\} = \|w\|_1$. The statement follows from Theorem 3.13. \square

We conclude this section with some remarks on Theorem 3.16. First, there is no loss of generality in requiring that $\|w^\top S\|_\infty \leq \|S\|_\infty$ above: If $\|w^\top S\|_\infty > \|S\|_\infty$, one may “normalize” the redundant row $w^\top S$ by replacing w with $\|S\|_\infty \|w^\top S\|_\infty^{-1} w$. Corollary 3.14 guarantees that this will not decrease $\text{hsb}(S')$. (In fact, the matrix S'_n from Example 3.11 requires such a normalization for Theorem 3.16 to apply; this can be achieved by dividing the additional row by n .) Secondly, we note that the theorem applies to any w with $\|w\|_1 \leq 1$ since $\|w^\top S\|_\infty \leq \|w\|_1 \|S\|_\infty$. In this case, we obtain $\text{hsb}(S') = \text{hsb}(S)$. In other words, adding a convex combination of rows of S has no effect on $\text{hsb}(S)$. Finally, the statement of Theorem 3.16 easily extends to the case of adding multiple rows and, by transposition, columns.

3.5 Further Notes and Open Questions

For both families of polytopes studied in this chapter and their canonical slack matrices, we have shown that the hyperplane separation technique is unable to improve on the

currently best known lower bounds on their extension complexity. Unlike the nonnegative rank, the hyperplane separation bound depends on the particular choice of slack matrix. By making a more careful choice, it is possible that the technique does indeed yield more meaningful bounds than the ones in Section 3.3. In particular, Section 3.4 proposes two potential avenues: normalizing rows and columns, and introducing redundancy by adding nonnegative linear combinations of rows or columns. Both operations will only strengthen the hyperplane separation bound (under the assumptions discussed in Section 3.4) while preserving the nonnegative rank.

How much can one gain by these operations when applying them to the specific slack matrices in Section 3.3? While we leave this as an open question, we briefly discuss the effect of normalization. For the spanning tree polytope $P_{\text{st}}(G)$, normalizing the rows of its canonical slack matrix S_G produces a matrix S'_G whose columns are normalized as well. To see this, consider the row submatrix of S_G consisting of all rows for subsets of exactly two nodes. This is a 0/1 matrix (see the proof of Theorem 3.8). Unless G is a tree itself, each column of this submatrix has a one entry: For each spanning tree T in G , there is some edge that is not contained in T . This means that normalizing the rows of S_G will not affect these one entries, and therefore S'_G is column-normalized. Now Corollary 3.14 and the remark thereafter imply that $\text{hsb}(S_G) \leq \text{hsb}(S'_G) \leq \|S_G\|_\infty \text{hsb}(S_G) = O(|V||E|)$. This *a priori* analysis at least suggests that the hyperplane separation method may indeed achieve a lower bound that is closer to the known upper bound of $O(|V||E|)$.

In the case of the slack matrix M_A of a graphic zonotope $Z(A)$, the potential benefit of normalization seems to depend heavily on A . In fact, as a byproduct of the proof of Theorem 3.10, we have that the maximum in row $S \subseteq [n]$ of M_A equals the weight of the cut induced by S in the graph underlying $Z(A)$. Therefore, the gain on $\text{hsb}(M_A)$ that one can expect from normalizing the rows of M_A according to Corollary 3.14 (which actually produces a matrix M'_A that is column-normalized as well) is at most the ratio of the maximum and the minimum weight of a cut. While this ratio can grow arbitrarily large, it depends on A . For instance, if A is the $n \times n$ all-one matrix (that is, if $Z(A)$ is the n th permutahedron), normalizing M_A does not help much: Since every cut in K_n has at least $n - 1$ and at most $\lfloor n/2 \rfloor \lceil n/2 \rceil$ edges, we obtain $\text{hsb}(M'_A) = O(n)$. However, recall that $\text{rk}_+(M'_A) = \text{rk}_+(M_A) = \Theta(n \log n)$.

Another question which is left open by our analysis in Section 3.4 is a refined version of Corollary 3.14: Which positive diagonal scalings yield the strongest bounds among those that are both row- and column-normalized? Note that Corollary 3.14 does not necessarily imply an improvement in the hyperplane separation bound for every such row- and column-normalized positive diagonal scaling. For instance, consider again the matrices S and D from Example 3.15. Multiplying S on both sides with the positive diagonal matrix D , the lower bound of Corollary 3.14 evaluates to $\frac{1}{2} \text{hsb}(S)$. However, we saw that rescaling with D actually strictly increases the hyperplane separation bound of S .

Finally, how significant is the effect of adding redundant rows or columns (cf. Theorem 3.16) compared to the possible gain achieved by the best diagonal scalings?

Chapter 4

Circuits under Projections of Polyhedra

The results of this chapter are joint work with Steffen Borgwardt and appear in [28]. Sections 4.3 to 4.5 and parts of Sections 4.1 and 4.2 largely coincide with our paper.

4.1 Introduction

In the previous chapter, we saw that spanning tree polytopes and completion time polytopes – despite their number of facets being exponential in the dimension – admit extended formulations of polynomial size. Such *compact* extended formulations have been discovered time and again for polyhedra in combinatorial optimization. Prominent examples are independence polytopes of regular matroids [7] (which generalize spanning tree polytopes), parity polytopes [51, 205], and subtour elimination polytopes for the TSP [204] (which are relaxations of TSP polytopes), to name but a few. Linear descriptions of all of these polytopes in the original space generally have an exponential number of inequalities. In some cases, such as for dominants of cut polytopes [56] or fixed-shape partition polytopes from clustering [16, 26, 29, 43, 131], compact extended formulations have been shown to exist while a complete description in the original space is not even known. We refer the reader to the surveys [56, 134, 203] for an overview (see also Chapter 4 of [57]).

Interestingly, also the $\{0, \frac{1}{2}\}$ -closure of the fractional stable set polytope from Chapter 2 admits a compact extended formulation: The separation problem for the odd-cycle inequalities can be expressed as a compact LP that follows the separation algorithm of [114, 122], as observed in [205] (see also [56]). This means that one can solve the stable set problem in t -perfect graphs by solving a single LP of polynomial size, without relying on the equivalence of optimization and separation via the ellipsoid method [124].

Probably the most famous algorithm for solving LPs is the Simplex method [70]. It finds an optimal solution by tracing an augmenting path along edges of the polyhedron that is the feasible region of the LP. The Simplex method belongs to a more general family of LP algorithms called *circuit augmentation schemes* [72]. These algorithms augment a solution along a more general set of directions, the *circuits* of the feasible region [120], until the optimum is found. We postpone formal definitions to Section 4.2. For now, let us think of the circuits of a polyhedron $P = \{x: Ax = b, Bx \leq d\}$ as “potential” edge directions in the sense that each circuit appears as an edge direction of $\{x: Ax = b', Bx \leq d'\}$ for some choice of b' and d' [98]. In particular, the set of circuits includes the actual edge directions of P .

For polyhedra in combinatorial optimization, circuit directions often admit a nice combinatorial characterization (see [133]). The circuits of network flow polyhedra, for

example, correspond to directed cycles in the underlying directed graph (see, e.g., [94]). Many network flow algorithms that iteratively push flow along cycles can therefore be regarded as special circuit augmentation schemes (see, e.g., [72, 94, 132]). This connection, in turn, has been a source of inspiration for generalizations of network flow algorithms to augmentation schemes for general LPs or even IPs (see, e.g., [72, 156, 179] and the references therein). This suggests that understanding the circuits of polyhedra in combinatorial optimization may help find or analyze combinatorial algorithms. We saw above that for some of these polyhedra, however, a complete linear description is not even available and one has to resort to an extended formulation. In those cases, it would be desirable if one could describe the set of circuits of the original polyhedron in terms of the circuits of its extension. Recall that the original polyhedron is the image of its extension under an affine projection (which we may assume to be linear, possibly after a translation). It is a well-known fact that all edge directions of the original polyhedron are images of edge directions of the extension under the same projection map. Is this true, more generally, for the set of circuits? In other words, are the circuits of the original polyhedron projections of circuits of the extension? This question provides the motivation for the work in this chapter.

Note that the close relationship between the edge directions of a polyhedron and of an extension is what the shadow vertex pivot rule [25] for the Simplex method implicitly relies on. This pivot rule constructs a Simplex path by following edges of 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 [68, 184, 197], in the study of diameters of polyhedra [48, 67] (see also Chapter 5), and in recent work on strong bounds for the performance of the Simplex method on 0/1 polytopes [20]. Can one exploit similar ideas to design efficient circuit augmentation schemes? Again, this requires an understanding of how circuits of polyhedra behave under taking projections.

As we will see, their behaviour is fundamentally different from that of the edge directions. For the sake of brevity, we say that a polyhedron P *inherits* its circuits from an extension Q with $\pi(Q) = P$ for some linear map π , if all circuits of P are images of circuits of Q under π . We first demonstrate in Section 4.3 that polyhedra do not necessarily inherit their circuits from extensions: We show how to construct counterexamples with a minimal number of facets, vertices, and extreme rays in every dimension greater than 2, with corresponding extensions just one dimension higher. Our construction extends to a relevant family of polyhedra in combinatorial optimization: fixed-shape partition polytopes. One can even find counterexamples with an exponential gap in the number of unique directions between the subset of circuits that are inherited and the entire set of circuits.

In Section 4.4, we then consider the following three natural questions:

- (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 Q ?
- (Q3) For which polyhedra Q does every polyhedron P that is a linear projection of Q inherit its circuits from Q ?

It is not difficult to give two sufficient conditions for membership in the classes stated in questions (Q1) and (Q2), respectively: As we will see in Section 4.2, injective linear maps π define linear isomorphisms between Q and $\pi(Q)$ and thus satisfy the property that (Q1) asks for. Further, polyhedra in which all circuits are edge directions clearly belong to the class of polyhedra for (Q2). What is more interesting is that these properties are, in fact, both sufficient and *necessary*, thus completely resolving questions (Q1) and (Q2). We prove this in Section 4.4 and also give a partial answer to question (Q3), showing that in dimension 4 or higher, no polyhedron with a non-degenerate vertex has the property that (Q3) asks for.

Our results imply that, whenever a polyhedron P *does* inherit all circuits from another polyhedron Q (with a non-degenerate vertex) under some linear 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 a linear isomorphism between P and Q or because P has no circuits other than the 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 will conclude this chapter with a family of examples of such nontrivial combinations that guarantee inheritance.

4.2 Preliminaries

We begin by formally introducing the set of circuits of polyhedra. As circuits have been widely studied in various contexts in polyhedral theory and optimization (see, e.g., [9, 19, 22, 33, 34, 35, 36, 69, 72, 73, 94, 105, 109, 120, 127, 133, 158, 165] and the references therein), there are several ways to characterize them. We present some of these characterizations here and another one later, in Section 5.2. Easily accessible introductions to the topic of circuits and their fundamental properties are also given in [98, 132]. We follow these two references for all basic definitions and concepts.

4.2.1 Circuits of Polyhedra

Definition 4.1 ([120, 187]). *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a polyhedron. A nonzero vector $g \in \mathbb{R}^n$ is a circuit of P if*

- (i) $g \in \ker(A)$ and
- (ii) $\text{supp}(Bg)$ is inclusion-minimal in the collection $\{\text{supp}(By) : y \in \ker(A), y \neq \mathbf{0}\}$.

Here, $\ker(A)$ denotes the kernel of A . For polyhedra given by linear systems in standard form $Ax = b, x \geq \mathbf{0}$, the circuits are precisely the support-minimal vectors in $\ker(A) \setminus \{\mathbf{0}\}$. These vectors are also referred to as *elementary vectors* in the literature [165].

Geometrically speaking, the circuits of a polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ are those directions in \mathbb{R}^n that are (i) parallel to the affine subspace defined by $Ax = b$ (which, given a minimal description, is the affine hull of P), and (ii) parallel to an inclusion-maximal subset of hyperplanes defined by $Bx \leq d$ (see [132, 133]).

To avoid unnecessary technicalities, we will assume all polyhedra in this chapter to be pointed. In this case, we have the following well-known equivalent characterization

of the set of circuits [34, 133]. For the sake of completeness, we include a proof, which follows the proof of Lemma 1 in [34].

Proposition 4.2 ([34, 133]). *The circuits of a pointed polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ are precisely 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 $\begin{pmatrix} A \\ B' \end{pmatrix}$ is $n - 1$.*

Proof. Let B' be a row submatrix of B and let g be a nonzero vector in $\ker \begin{pmatrix} A \\ B' \end{pmatrix}$. Any vector $y \in \ker(A) \setminus \{\mathbf{0}\}$ with $\text{supp}(By) \subseteq \text{supp}(Bg)$ must also be in $\ker \begin{pmatrix} A \\ B' \end{pmatrix}$. So if the rank of $\begin{pmatrix} A \\ B' \end{pmatrix}$ is $n - 1$, i.e., its kernel is one-dimensional, then y is a multiple of g . This implies that $\text{supp}(By) = \text{supp}(Bg)$ and therefore g is a circuit of P .

To prove the converse implication, we may assume that B' is the maximal row submatrix of B such that $g \in \ker(B')$. Note that the rank of $\begin{pmatrix} A \\ B' \end{pmatrix}$ is at most $n - 1$ since $\begin{pmatrix} A \\ B \end{pmatrix}$ has rank n and $g \neq \mathbf{0}$. Now suppose that the rank of $\begin{pmatrix} A \\ B' \end{pmatrix}$ is strictly less than $n - 1$. Then we may add rows of B to $\begin{pmatrix} A \\ B' \end{pmatrix}$ to obtain a matrix of rank $n - 1$. Its kernel is generated by some $y \in \ker(A) \setminus \{\mathbf{0}\}$ with $\text{supp}(By) \subsetneq \text{supp}(Bg)$. Thus, g is not a circuit of P . \square

Since any edge of P is defined by $\dim(P) - 1$ linearly independent inequalities from $Bx \leq d$, it follows from Proposition 4.2 that all edge directions of P are also circuits. In fact, the set of circuits of P can be shown to consist precisely of all edge directions of polyhedra $\{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ as the right-hand sides b and d vary [98] (see also [120, 187]).

Note that the set of circuits of P is infinite: By definition, every nonzero multiple of a circuit is a circuit again. However, up to rescaling, each circuit of P is uniquely defined by the support of its product with B (see [120]). This suggests a finite representation of the set of circuits, e.g., by normalizing them to have co-prime integer components (assuming rational A and B). The set of all such finitely many representatives is usually denoted by $\mathcal{C}(A, B)$ in the literature (see [34, 35, 72, 73, 98, 133]). In this thesis, we will use the same notation for the *entire* set of circuits of P , as it will be convenient to not assume any kind of normalization of the circuits and rather view the set $\mathcal{C}(A, B)$ as a finite union of one-dimensional linear subspaces of \mathbb{R}^n . (We shall elaborate on this point of view in Section 5.2.)

It is important to note that the set of circuits of a polyhedron P depends on the linear description of P (although not on the right-hand sides): Adding redundant inequalities to a linear system may enlarge the set of circuits. However, if P is given by a *minimal* description $Ax = b, Bx \leq d$, then the set of its circuits is indeed independent of the description (see, e.g., [132]). This means that we may in this case write $\mathcal{C}(P)$ for $\mathcal{C}(A, B)$. Throughout this chapter, we will assume that all polyhedra are given by minimal descriptions, unless stated otherwise.

4.2.2 Circuits under Affine Equivalence of Polyhedra

We make the assumption of minimality for another key reason: Among all polyhedra $P \subseteq \mathbb{R}^n$ and $Q \subseteq \mathbb{R}^m$ such that $P = \pi(Q)$ for some linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$, we would like to characterize those triples (P, Q, π) for which all circuits of P are inherited from Q under π , i.e., $\mathcal{C}(P) \subseteq \pi(\mathcal{C}(Q))$. If we allow Q to be given by *any*, not necessarily minimal

linear description, then, by adding redundant inequalities, we may blow up the set of circuits of Q until it contains a preimage for each circuit of P . Likewise, by introducing suitable redundancy to the description of P , we could make any given nonzero vector a circuit of P and destroy the desired inheritance property. So the minimality assumption is crucial for obtaining meaningful statements about the inheritance of circuits under projections. We can thus speak of inheritance as a purely geometric property of (P, Q, π) .

This intuition is further supported by the observation that inheritance is preserved under affine equivalence of polyhedra. Recall that two polyhedra $P \subseteq \mathbb{R}^n$ and $Q \subseteq \mathbb{R}^m$ are *affinely (linearly) isomorphic* (or *equivalent*) if there exists an affine (linear) map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that $\pi(Q) = P$ and, for all $x \in P$, there is 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.

Lemma 4.3. *Let $Q \subseteq \mathbb{R}^m$ be a pointed polyhedron and let $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ be an affine map such that Q and $\pi(Q)$ are affinely isomorphic. Then $\mathcal{C}(\pi(Q)) = \pi(\mathcal{C}(Q)) - \pi(\mathbf{0})$.*

To prove this, we need the following well-known fact about projections of polyhedra (see, e.g., [100, Proposition 2.1] or [206, Lemma 7.10] for the polytopal case).

Remark 4.4 (see [100, 206]). *Let $P \subseteq \mathbb{R}^n, Q \subseteq \mathbb{R}^m$ be polyhedra such that $P = \pi(Q)$ for an affine map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$, and let F be a face of P . Then the set $\pi^{-1}(F) := \{y \in Q: \pi(y) \in F\}$ is a face of Q . Indeed, if $F = \{x \in P: a^\top x = \beta\}$ for an inequality $a^\top x \leq \beta$ that is valid for P , then $\pi^{-1}(F)$ is given by $\pi^{-1}(F) = \{y \in Q: a^\top \pi(y) = \beta\}$.*

We are now ready to give a proof of Lemma 4.3.

Proof of Lemma 4.3. We start by observing that translating a polyhedron does not change its set of circuits, as translations only affect the right-hand sides of linear descriptions. We may therefore assume that π is a *linear* isomorphism between Q and $P := \pi(Q)$.

Let $P = \{x \in \mathbb{R}^n: A_P x = b_P, B_P x \leq d_P\}$ and $Q = \{x \in \mathbb{R}^m: A_Q x = b_Q, B_Q x \leq d_Q\}$ be minimal linear descriptions of P and Q , respectively, and define the polyhedron

$$\tilde{Q} := \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^m: x = \pi(y), A_Q x = b_Q, B_Q x \leq d_Q\}. \quad (4.1)$$

Now consider the two linear projection maps $\pi_x: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $\pi_y: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ which send (x, y) to x and y , respectively. Clearly, $\pi_x(\tilde{Q}) = P$ and $\pi_y(\tilde{Q}) = Q$. Moreover, both projections actually define linear isomorphisms between \tilde{Q} and P or Q , respectively: For each $y \in Q$, there clearly is a unique x such that $(x, y) \in \tilde{Q}$; and for each $x \in P$, there is a unique y for which $\pi(y) = x$ and therefore $(x, y) \in \tilde{Q}$, since Q and $\pi(Q) = P$ are linearly isomorphic by hypothesis.

We may therefore characterize the circuits of \tilde{Q} in two different ways, namely in terms of those of Q and of P . First, note that the linear description of \tilde{Q} given in (4.1) is minimal: It is obtained from the minimal description of Q by adding equality constraints only. Further, the equality constraint matrix is of the form

$$\begin{pmatrix} \mathbf{I}_n & * \\ \mathbf{0} & A_Q \end{pmatrix}.$$

and therefore has rank $n + \text{rk}(A_Q) = n + m - \dim(Q) = n + m - \dim(\tilde{Q})$. It follows directly from the definition that the set of circuits of \tilde{Q} is given by

$$\mathcal{C}(\tilde{Q}) = \{(\pi(f), f) : f \in \mathcal{C}(Q)\}. \quad (4.2)$$

Next, recall that each inequality in $B_P x \leq d_P$ defines a facet of P . Since $\pi_x(\tilde{Q}) = P$, each inequality in $B_P \pi_x(x, y) \leq d_P$ therefore defines a proper face of \tilde{Q} by Remark 4.4, which must be of at least the same dimension. As \tilde{Q} and P are linearly isomorphic, each such face is, in fact, a facet. It follows that also the system $x = \pi(y), A_Q y = b_Q, B_P x \leq d_P$ is a minimal description of \tilde{Q} by noting that $B_P \pi_x(x, y) = B_P x$. Using this description to characterize the circuits, we conclude that a nonzero vector (g, f) is a circuit of \tilde{Q} if and only if $g = \pi(f), f \in \ker(A_Q) \setminus \{\mathbf{0}\}$, and $B_P g$ is support-minimal among all nonzero vectors (g, f) with this property. Since the affine hulls of linearly isomorphic polyhedra are also linearly isomorphic, π must be a bijection between $\ker(A_P)$ and $\ker(A_Q)$. Here, we used that $A_P x = b_P$ and $A_Q y = b_Q$ define the affine hulls of P and Q , respectively, by virtue of P and Q being given by minimal descriptions. This means that $g = \pi(f)$ for some $f \in \ker(A_Q)$ if and only if $g \in \ker(A_P)$. Hence,

$$\mathcal{C}(P) = \pi_x(\mathcal{C}(\tilde{Q})) = \pi(\mathcal{C}(Q))$$

where the second identity follows from (4.2). \square

As noted in the proof of Lemma 4.3, translations are special affine isomorphisms and leave the set of circuits invariant. Thus, we may restrict ourselves to *linear* projection maps in the following.

Lemma 4.3 will be one of the key ingredients of our proofs later in this chapter. It further provides a simple sufficient condition for the inheritance of circuits: If $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$ is an injective linear map, then it is clear that any polyhedron $Q \subseteq \mathbb{R}^m$ and its image $\pi(Q)$ are linearly isomorphic and therefore $\mathcal{C}(\pi(Q)) = \pi(\mathcal{C}(Q))$ by Lemma 4.3.

Lemma 4.3 also has the following consequence, which we believe to be of independent interest in the study of circuits: 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. More precisely, we obtain the following corollary to Lemma 4.3.

Corollary 4.5. *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a pointed polyhedron given by a minimal linear description with $B \in \mathbb{R}^{m \times n}$. Define the affine map $\sigma: \mathbb{R}^n \rightarrow \mathbb{R}^m$ by $x \mapsto d - Bx$. Then $\sigma(P)$ is a polyhedron with a minimal description in standard form 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 [100]). 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. Applying Lemma 4.3, we obtain $\mathcal{C}(\sigma(P)) = \sigma(\mathcal{C}(P)) - d = B \cdot \mathcal{C}(P)$. \square

We point out that Corollary 4.5 contrasts with the behaviour 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 [34] that this conversion may introduce exponentially many new circuits. Corollary 4.5 suggests that this behaviour is a consequence of splitting the variables and not of introducing slack variables, which is what applying the slack map σ defined in Corollary 4.5 implicitly does as well. More precisely, $\sigma(P)$ is the projection of

$$\tilde{P} := \{(x, s) \in \mathbb{R}^n \times \mathbb{R}^m : Ax = b, Bx + s = d, s \geq \mathbf{0}\}$$

onto the slack variables s . By the same argument as in the proof of Lemma 4.3, \tilde{P} and P (and, hence, \tilde{P} and $\sigma(P)$) are affinely isomorphic. Characterizing $\sigma(P)$ via \tilde{P} adds the benefit that one can derive an explicit standard form representation of $\sigma(P)$ from the description of \tilde{P} , using a projection technique of Balas and Pulleyblank [12] (see also [57, Theorem 3.46]): For a basis $\{(u^{(1)}, v^{(1)}), \dots, (u^{(l)}, v^{(l)})\}$ of $\ker((B^\top | A^\top))$, we have that

$$\sigma(P) = \left\{ s \in \mathbb{R}^m : s \geq \mathbf{0}, (u^{(i)})^\top s = (u^{(i)})^\top d \text{ for all } i \in [l] \right\}.$$

In some sense, Corollary 4.5 may be regarded as an extension of the definition of circuits as elementary vectors [165]: The above result shows that the circuits of *any* polyhedron (not just the ones in standard form) are the elementary vectors of some linear subspace.

We next resolve another easy case when circuits are inherited under linear projections.

4.2.3 Inheritance of Edge Directions

Every 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. For the sake of clarity, we provide a short proof.

Lemma 4.6. *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. Let e be an edge of P . Since P is pointed by hypothesis, e has a vertex v . Then, by Remark 4.4, 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' . It suffices to show that $\pi^{-1}(e)$ has an edge that contains both v' and some point w' with $\pi(w') \neq v$. Indeed, for this point, $v' - w'$ is an edge direction of Q with $\pi(v' - w') = v - \pi(w') \neq \mathbf{0}$ and $\pi(w') \in e$.

So suppose for the sake of contradiction that no such edge exists. Then all edges of $\pi^{-1}(e)$ that are incident with v' must be contained in $\pi^{-1}(v)$. Hence, $\pi^{-1}(v)$ contains the feasible cone of $\pi^{-1}(e)$ at v' , and therefore $\pi^{-1}(v) \supseteq \pi^{-1}(e)$, a contradiction. \square

Lemma 4.6 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 . For example, each circuit of a simplex is an edge direction since any $d - 1$ facets of a d -simplex define an edge. Likewise, each circuit of a simplicial cone is an edge direction, too. Circuits and edge directions therefore also coincide for Cartesian products of simplices (or simplicial cones), as we will see next.

Given two polyhedra P_1 and P_2 , the nonempty faces of their product $P_1 \times P_2$ are exactly the sets of the form $F_1 \times F_2$ where F_i is a nonempty face of P_i for $i \in \{1, 2\}$ (see [206]). In particular, all edge directions of $P_1 \times P_2$ are of the form $(g^{(1)}, \mathbf{0})$ or $(\mathbf{0}, g^{(2)})$ for an edge direction $g^{(i)}$ of P_i . The next result shows that this is true, more generally, for the set of circuits of $P_1 \times P_2$. This was shown in Lemma 3.9 of [39] for polyhedra P_i in canonical form. We restate and prove the result in all generality here.

Proposition 4.7 ([39]). *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 $\ker(A^{(1)}) \times \ker(A^{(2)})$. Without loss of generality, 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

As hypercubes are products of simplices (namely, line segments), also their circuits are exactly the edge directions by Proposition 4.7. Beyond hypercubes and simplices, though, polyhedra with this property are rare – in general, the set of circuits may be much larger than the set of edge directions. Interesting classes of polyhedra that do have the property that circuits and edge directions coincide are Birkhoff polytopes, as we will see in Section 4.3.3, and fractional matching polytopes [73, 169] (see Chapter 2).

Also note that in dimension 2, the sets of circuits and edge directions clearly are the same; this follows directly from the fact that facets of two-dimensional polyhedra are edges. The next result is a direct consequence of this observation.

Corollary 4.8. *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 $\mathcal{C}(P) \subseteq \pi(\mathcal{C}(Q))$.*

Proof. Suppose that $\dim(Q) \leq 3$. Then either $\dim(P) = 3$ and therefore also $\dim(Q) = 3$, or $\dim(P) \leq 2$. In the first case, P and Q must be linearly isomorphic, so $\mathcal{C}(P) = \pi(\mathcal{C}(Q))$ by Lemma 4.3. In the second case, each circuit of P is an edge direction, as we saw above, so Lemma 4.6 implies that $\mathcal{C}(P) \subseteq \pi(\mathcal{C}(Q))$. \square

Corollary 4.8 suggests that in order to find a counterexample for the inheritance of circuits, i.e., a triple (P, Q, π) such that $\pi(Q) = P$ but $\mathcal{C}(P) \not\subseteq \pi(\mathcal{C}(Q))$, one has to look for a polyhedron P in dimension 3 or higher. In fact, such counterexamples exist in all dimensions 3 and higher, as we will see next.

4.3 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 behaviour of edge directions stated in Lemma 4.6. We begin by constructing a family of provably minimal counterexamples.

4.3.1 A Family of Minimal Counterexamples

The essential building block for our constructions is a carefully chosen family of linear projections. For all $m, n \in \mathbb{N}$ with $m > n \geq 3$, we define a matrix

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

and let $\pi_{n,m}: \mathbb{R}^m \rightarrow \mathbb{R}^n$ 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 4.9. *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 pointed cone, R_n is a pointed cone with vertex $\mathbf{0}$ again, spanned by the first $n+1$ column vectors of the matrix $\Pi_{n,m}$. 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 4.1). 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 $2e_n$, 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 $e_1 - e_2$, 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 multiples of edge directions of R_n only. \square

The construction in Lemma 4.9 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 $C_m := [0, 1]^m$ here.




Figure 4.1: The cone $R_3 = \pi_{3,4}(\mathbb{R}_{\geq 0}^4)$ from Lemma 4.9, 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 4.10. *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. Clearly, $\mathbf{0} \in \pi(C_m) \subseteq \mathbb{R}_{\geq 0}^n$. Thus, $\mathbf{0}$ is a vertex of $\pi(C_m)$. Observe that $\pi(\mathbb{R}_{\geq 0}^m)$ is the feasible 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 C_m and $\mathbb{R}_{\geq 0}^m$ are full-dimensional, it follows that $\mathcal{C}(\pi(C_m)) \supseteq \mathcal{C}(\pi(\mathbb{R}_{\geq 0}^m))$. We further have that $\mathcal{C}(C_m) = \mathcal{C}(\mathbb{R}_{\geq 0}^m)$, so the first part of the statement follows from Lemma 4.9.

Note that $\pi(C_m)$ is the linear projection of a hypercube and, as such, a zonotope (see Figure 4.2). Equivalently, $\pi(C_m)$ can be written as the Minkowski sum of $n + 1$ line segments $[\mathbf{0}, \pi_{n,m}(e_i)]$ for every $i \in [n + 1]$. Since every edge of a zonotope is a translate of one of the line segments from which it is generated, the second part of the statement follows. \square




Figure 4.2: The zonotope $\pi_{3,4}(C_4)$ from Lemma 4.10. The two facets that yield the circuit direction e_3 are highlighted.

From the proof of Lemma 4.10, we can extract 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 4.11. *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 4.10 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 4.7. 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 4.8, the minimal dimension of any polyhedron which does not inherit its circuits from an extension is 3. Indeed, the lowest-dimensional counterexamples given in Lemmas 4.9 and 4.10 are 3-dimensional. The family of cones given in Lemma 4.9 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 4.6). 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 4.10 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 $S_m := \{x \in \mathbb{R}^m : x \geq \mathbf{0}, \mathbf{1}^\top x \leq 1\}$ for $m \in \mathbb{N}$.

Lemma 4.12. *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 $P_n = \{x \in R_n : \mathbf{1}^\top x \leq 2\}$ since the entries of each nonzero column of $\Pi_{n,m}$ sum to 2 (see Figure 4.1). In particular, this implies that the inequality $\mathbf{1}^\top x \leq 2$ 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 4.1. 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 4.9 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 4.9 and 4.12, we obtain the following theorem.

Theorem 4.13. *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 4.9 and 4.12 have the same number of vertices and extreme rays as their projections. Indeed, such a canonical “simplex extension” exists for every pointed polyhedron (see, e.g., [100]) and will be the starting point for proving Theorem 4.20 in Section 4.4.

Next, we will see that not only do there exist counterexamples that fail the inheritance of circuits, but the difference in the number of unique circuits (up to scaling) that are inherited and those that are not may be exponentially large in the dimension.

4.3.2 Counterexamples by Counting Circuits

In this section, we will be “counting” circuits. Recall from Section 4.2.1 that we work with the (infinite) set of circuits in this chapter, not assuming any normalization scheme. For the purpose of counting, however, we will speak of *unique circuit directions* of a polyhedron, by which we mean any finite set of representatives such that no two are multiples of one another. Our main observation is that there exist polyhedra that have many unique circuit directions and, at the same time, are projections of polyhedra with very few unique circuit directions. So merely by counting unique circuit directions, one can prove that not all circuits can be inherited. More specifically, our goal is to prove the following theorem.

Theorem 4.14. *For all $n \geq 3$, there exist 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 $O(n^2)$.*

The proof of Theorem 4.14 relies on the interplay between the basic solutions of a polytope and the circuits in its homogenization. Let us recall the relevant concepts first.

A *basic solution* of a linear system $Ax = b, Bx \leq d$ in variables $x \in \mathbb{R}^n$ is a solution of $Ax = b, B'x = d'$ where $(B' | d')$ is a row submatrix of $(B | d)$ such that $\begin{pmatrix} A \\ B' \end{pmatrix}$ has full column rank n . Note that basic solutions need not be feasible. The *feasible* basic solutions are precisely the vertices of the polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$. By a slight abuse of terminology, we will call a basic solution of $Ax = b, Bx \leq d$ also a basic solution of P , and denote the set of all such basic solutions by $\mathcal{B}(P)$. With this notation, $\mathcal{B}(P) \cap P$ is the set of vertices $\mathcal{V}(P)$ of P .

Notice the similarity between the definition of $\mathcal{B}(P)$ and the characterization of the set of circuits $\mathcal{C}(P)$ for pointed polyhedra P given in Proposition 4.2. 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 [206], we define the *homogenization* of a polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ as

$$\text{hom}(P) := \{(t, x) \in \mathbb{R} \times \mathbb{R}^n : t \geq 0, Ax - bt = \mathbf{0}, Bx - dt \leq \mathbf{0}\}. \quad (4.3)$$

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 4.15. *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a pointed polyhedron given by a minimal description. Up to rescaling, the circuits of $\text{hom}(P)$ with respect to the system (4.3) 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 $(\gamma, g) \neq (0, \mathbf{0})$ such that $Ag - b\gamma = \mathbf{0}$. Suppose first that $\gamma \neq 0$. Possibly after scaling, we may assume that $\gamma = 1$. By Proposition 4.2, $(1, g)$ is a circuit of $\text{hom}(P)$ if and only if there is row submatrix $(-d' | B')$ of $(-d | B)$ such that $B'g = d'$ and

$$\left(\begin{array}{c|c} -b & A \\ \hline -d' & B' \end{array} \right) \quad (4.4)$$

has rank n . We claim that the rank of this matrix is the same as that of $\begin{pmatrix} A \\ B' \end{pmatrix}$, which implies that $(1, g)$ is a circuit of $\text{hom}(P)$ if and only if $g \in \mathcal{B}(P)$. Indeed, there is nothing to prove if $b = \mathbf{0}$ and $d' = \mathbf{0}$ since deleting a zero column from a matrix does not change its rank. So suppose that not both b and d' are zero. As $Ag = b$ and $B'g = d'$ by hypothesis, it follows that $g \neq \mathbf{0}$. Hence, the first column of the matrix (4.4) is in the column span of $\begin{pmatrix} A \\ B' \end{pmatrix}$, which means that $\begin{pmatrix} A \\ B' \end{pmatrix}$ has the same rank as (4.4).

Now suppose that $\gamma = 0$ (and thus $g \neq \mathbf{0}$). Again, we know from Proposition 4.2 that for $(0, g)$ to be a circuit of $\text{hom}(P)$, there must exist a row submatrix $(-d' | B')$ of $(-d | B)$ such that $B'g = \mathbf{0}$ and the rank of the following matrix is n :

$$\left(\begin{array}{c|c} -b & A \\ \hline -d' & B' \\ \hline -1 & \mathbf{0}^\top \end{array} \right) \quad (4.5)$$

Note here that even if the submatrix given to us by applying Proposition 4.2 does not include the bottom row corresponding to the constraint $t \geq 0$, we may add this row without increasing the rank to $n + 1$, since $g \neq \mathbf{0}$ and $\text{hom}(P)$ is pointed. Now the rank of the matrix (4.5) is n if and only if the rank of $\begin{pmatrix} A \\ B' \end{pmatrix}$ is $n - 1$. By Proposition 4.2, this is equivalent with $g \in \mathcal{C}(A, B)$. \square

Using Lemma 4.15, the basic solutions of polyhedra may also be characterized as follows.

Corollary 4.16. *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a pointed polyhedron. Then $\mathcal{B}(P)$ consists of all vectors $g \in \mathbb{R}^n$ such that $Ag = b$ and $\text{supp}(Bg - d)$ is minimal in the collection $\{\text{supp}(By - d) : y \in \mathbb{R}^n, Ay = b\}$.*

Proof. Let $\text{hom}(P)$ be defined as in (4.3). By Lemma 4.15(ii), the basic solutions of P are those vectors $g \in \mathbb{R}^n$ for which $(1, g)$ is a circuit of $\text{hom}(P)$, that is, for which $Ag - b = \mathbf{0}$ and $Bg - d$ is support-minimal in $\{By - d : y \in \mathbb{R}^n, Ay = b\}$ by definition. \square

The crucial observation for proving Theorem 4.14 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 linear 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 4.15. 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 (4.3) (possibly except $t \geq 0$) defines a facet of $\text{hom}(P)$ and all vectors (γ, g) satisfying the condition in Lemma 4.15(ii) are indeed circuits.

To prove Theorem 4.14, it therefore suffices to exhibit a family of polytopes with polynomially many (in the dimension) vertices but exponentially many basic solutions (with respect to 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*

$$Q_n := \left\{ x \in \mathbb{R}^n : x^\top y \leq 1 \text{ for all } y \in \{-1, 1\}^n \right\},$$

suitably cropped by intersecting it with a hypercube, satisfies all the desired properties. This will complete the proof of Theorem 4.14.

Lemma 4.17. *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 standard basis 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, no vertex 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

A careful analysis of the counterexamples in Section 4.3.1 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 4.3.1. 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 4.9 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 e_3 by Lemma 4.15.

4.3.3 Fixed-Shape Partition Polytopes

All constructions seen so far may seem specifically designed so as to fail the inheritance of circuits. However, there do not only exist pathological counterexamples. We conclude this section with an example from combinatorial optimization that exhibits this undesirable behaviour despite a number of favourable properties.

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 $\sum_{i=1}^k \kappa_i = n$. Popular clustering objectives like least-squares assignments can be computed by linear optimization over so-called *fixed-shape partition polytopes* (see [16, 26, 29, 30, 43, 106, 130, 131]). These are defined as follows. Given a data set $X = \{x^{(1)}, \dots, x^{(n)}\} \subseteq \mathbb{R}^d$, a number of clusters $k \in \mathbb{N}$, and prescribed cluster sizes $\kappa := (\kappa_1, \dots, \kappa_k) \in \mathbb{N}^k$, the associated *fixed-shape partition polytope*, denoted by $P(X, k, \kappa)$, is the convex hull of all feasible *clustering vectors* $(c^{(1)}, \dots, c^{(k)}) \in (\mathbb{R}^d)^k$, where $c^{(i)} = \sum_{x \in C_i} x$ is the sum of all items assigned to cluster C_i .

An explicit inequality description of $P(X, k, \kappa)$ is not known. However, the following system in variables $y \in \mathbb{R}^{k \times n}$ is an extended formulation for $P(X, k, \kappa)$ (see [27]):

$$\begin{aligned} \sum_{j=1}^n y_{ij} &= \kappa_i & \text{for all } i \in [k] \\ \sum_{i=1}^k y_{ij} &= 1 & \text{for all } j \in [n] \\ y &\geq \mathbf{0} \end{aligned} \tag{4.6}$$

The corresponding projection map is the linear map $\pi_X: y \mapsto (c^{(1)}, \dots, c^{(k)})$ defined by $c^{(i)} = \sum_{j=1}^n y_{ij} \cdot x^{(j)}$ for all $i \in [k]$. The polytope described by the linear system (4.6) belongs to a widely studied class of polytopes called *transportation polytopes* (see, e.g., [38, 71, 146]). Here, the k clusters represent the suppliers with supplies given by κ , and the n items are the customers with demands 1 each. The vectors $y \in \mathbb{R}^{k \times n}$ satisfying (4.6) describe feasible commodity flows from the suppliers to the customers. For given $n = |X|$, k , and κ , let us denote the transportation polytope (4.6) by $T(n, k, \kappa)$. The

special case $k = n$ and $\kappa_i = 1$ for all $i \in [n]$ yields the well-known n th Birkhoff polytope [14], which is the convex hull of all $n \times n$ permutation matrices, and is also known as the perfect matching polytope of the complete bipartite graph $K_{n,n}$ (see [206]).

Also for general k and κ , the structure of the polytopes $T(n, k, \kappa)$ is well-understood: They are 0/1 polytopes whose constraint matrix in (4.6) is totally unimodular [146]. Edges of $T(n, k, \kappa)$ correspond to cyclical exchanges of items between clusters [26, 27, 32, 146], where a subset of clusters are ordered in a cycle, and one item from each cluster is transferred to the next along the cycle. Moreover, the circuits of $T(n, k, \kappa)$ may be characterized in exactly the same way, showing that they coincide with the edge directions [33].

Is there a similar characterization of the circuits of the fixed-shape partition polytopes $P(X, k, \kappa)$? Interestingly, the edges of $P(X, k, \kappa)$, too, correspond to cyclical exchanges of items between clusters, for any choice of parameters X, k, κ [29, 106, 131]. Despite this promising fact, somewhat surprisingly, new circuits may appear in the projection onto $P(X, k, \kappa)$, even for small d, n , and k .

Lemma 4.18. *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 4.12, where we may assume without loss of generality that $x^{(n)} = \mathbf{0}$ and the remaining $n - 1$ vertices are labelled in arbitrary order. (Note that 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

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

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 4.3 and 4.12. \square

Despite the negative statement of Lemma 4.18, we stress that there do exist classes of fixed-shape partition polytopes in which all circuits are inherited from the transportation-type extensions, even though this happens for one of the two trivial reasons stated in Lemmas 4.3 and 4.6. For instance, $T(n, k, \kappa)$ is a fixed-shape partition polytope itself for any n and k , using the standard basis vectors in \mathbb{R}^n as item locations, as observed

in [33]. Similarly, suppose that we augment a given data set $X \subseteq \mathbb{R}^d$ of size $|X| = n$ with the standard basis 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 the proof of Lemma 4.3. In particular, the resulting polytope is affinely isomorphic to $T(n, k, \kappa)$.

4.4 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 P , the extension polyhedron Q , and the projection map π from Q to P – play for the inheritance of circuits. Our discussion is driven by three natural questions, first stated in Section 4.1:

- (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 Q ?
- (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 4.4.1 and of the polyhedra for question (Q2) in Section 4.4.2. 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 4.4.3. In Section 4.4.4, we explain why our characterizations are best possible. We do so by exhibiting combinations of polyhedra and maps that lead to an inheritance of circuits, but where neither of them exhibits the aforementioned properties.

4.4.1 Inheritance Based on the Projection Map

In Lemma 4.3, 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 4.19. *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 4.13 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 4.19 will be that, in any fixed dimension, the particular projection used

to obtain Theorem 4.13 can be exchanged for any other one after a suitable linear transformation of the domain space.

Proof of Theorem 4.19. By Lemma 4.3, 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 4.13, 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 4.3, we conclude that

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

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

4.4.2 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 4.20. *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 4.20:

Theorem 4.21. *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 [10, 11] 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 4.22 ([10, 11]). *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 $i \in [p]$:*

$$\begin{aligned} \lambda &\geq \mathbf{0} \\ \sum_{i=1}^p \lambda_i &= 1 \\ A^{(i)}x^{(i)} &= b^{(i)}\lambda_i \quad \text{for all } i \in [p] \\ B^{(i)}x^{(i)} &\leq d^{(i)}\lambda_i \quad \text{for all } i \in [p] \end{aligned} \tag{4.7}$$

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 4.22.

Lemma 4.23. *Let $P \subseteq \mathbb{R}^n$ be a pointed polyhedron and let \mathcal{P} and $Q_{\mathcal{P}}$ be defined as in Proposition 4.22. Up to rescaling, the circuits of $Q_{\mathcal{P}}$ with respect to the system (4.7) 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. First, note that all polyhedra in the collection \mathcal{P} must be pointed since $\text{conv}(\bigcup \mathcal{P}) = P$ is pointed. 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 4.7.

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 Corollary 4.16. In order to prove the claim, observe that the vector $(\mathbf{e}_1 - \mathbf{e}_2, \frac{1}{f_1}g^{(1)}, -\frac{1}{f_2}g^{(2)}, \mathbf{0}, \dots, \mathbf{0}) \in \mathbb{R}^p \times (\mathbb{R}^n)^p$ is a circuit of $Q_{\mathcal{P}}$, too. Then $\text{supp}(f) = \text{supp}(\mathbf{e}_1 - \mathbf{e}_2)$ by definition of the set of circuits. 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 4.21. The ‘‘only if’’ part immediately follows from Lemma 4.6. 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 a multiple 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 by hypothesis, we have that $\dim(F_{\{u, v\}}) \geq 2$. 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

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

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 4.22. We claim that under the projection given in Proposition 4.22, none of the circuits of $Q_{\mathcal{P}}$ maps to a multiple of g . By Lemma 4.23, the circuits of $Q_{\mathcal{P}}$ are either sent 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). To see that case (ii) is not possible either, recall that any two distinct members of \mathcal{P} can be strictly separated by a hyperplane whose normal vector is orthogonal to g , which means that g is parallel to this hyperplane. We conclude that g is not inherited from $Q_{\mathcal{P}}$.

Now suppose that P is unbounded. Then $P = \bar{P} + \text{rec}(P)$ where $\bar{P} := \text{conv}(\mathcal{V}(P))$ 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 \bar{P} is a polytope, the first part of the proof yields an extension $Q_{\bar{P}}$ of \bar{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 $Q_{\bar{P}} \times \mathbb{R}_{\geq 0}^q$. It is an extension of P , where the corresponding projection first maps $Q_{\bar{P}} \times \mathbb{R}_{\geq 0}^q$ to $\bar{P} \times \text{rec}(P)$ and then applies the map $(x, y) \mapsto x + y$. By Proposition 4.7, g is not inherited from $Q_{\bar{P}} \times \mathbb{R}_{\geq 0}^q$ 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 a minimal linear description if we follow Proposition 4.22. 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 4.12 is, in fact, a special case of the more general extension $Q_{\mathcal{P}}$ used in the proof of Theorem 4.20: For a polytope P and the decomposition $\mathcal{P} := \{\{v\} : v \in \mathcal{V}(P)\}$, the polyhedron $Q_{\mathcal{P}}$ and the simplex $S_{|\mathcal{V}(P)|-1}$ are affinely equivalent. Proposition 4.22 also generalizes another result from Section 4.3: Let P be a polytope and consider $\mathcal{P} := \{\{\mathbf{0}\}, \{1\} \times P\}$; then $Q_{\mathcal{P}}$ as defined in Proposition 4.22 equals $\{(t, x) \in \text{hom}(P) : t \leq 1\}$.

4.4.3 (No) Inheritance Based on the Extension Polyhedron

In Section 4.3.1, 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 4.24. *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 4.9, 4.10 and 4.12 and identify a common strategy: All polyhedra that we projected from in Section 4.3.1 have a non-degenerate vertex at the origin $\mathbf{0}$, and their feasible 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 4.3.1 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 (see Section 4.2.1), 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 4.3.1. The remainder of this section is dedicated to the proof details.

Proof of Theorem 4.24. After an affine transformation, we may assume that Q is full-dimensional, $\mathbf{0}$ is a non-degenerate vertex of Q , and the feasible cone of Q at $\mathbf{0}$ equals $\mathbb{R}_{\geq 0}^m$. For all $\alpha \in \mathbb{N} \setminus \{1\}$, we define the matrix

$$\Pi_\alpha := \left(\begin{array}{cccc|c} \alpha & 1 & 0 & 0 & \mathbf{0} \\ 0 & 0 & \alpha & 1 & \\ 0 & \alpha - 1 & 0 & \alpha - 1 & \\ \hline & \mathbf{0} & & & \alpha \mathbf{I}_{m-4} \end{array} \right) \in \mathbb{R}^{(m-1) \times m}$$

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 4.3.

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 4.9. 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 Section 4.2.1 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 feasible cone of $\pi_\alpha(Q)$ at the vertex $\mathbf{0}$. \square

4.4.4 Inheritance in Nontrivial Instances

Theorems 4.19, 4.20 and 4.24 imply that, beyond the trivial cases that we saw in Section 4.2, 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 4.19 and 4.20. In this sense, Theorems 4.19 and 4.20 are the best possible statements.

Lemma 4.25. *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, and not all circuits of P are edge directions.*

Proof. We again modify the construction from Lemma 4.12. For $n \geq 3$ and $m \geq n + 3$, we define the matrix

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

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 4.3). 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 4.12. For the case $n = 3$, we refer to

Figure 4.3. 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 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 4.9. 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, $e_3 \in \mathcal{C}(P'_n)$ is not an edge direction of P'_n , and S_m clearly has a non-degenerate vertex. \square




Figure 4.3: The polytope P'_3 from the proof of Lemma 4.25.

4.5 Further Notes and Open Questions

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 4.4).

For fixed-shape partition polytopes, we demonstrated in Section 4.3.3 that their circuits cannot be characterized via projecting from their transportation-type extensions. Still, a characterization of the circuits of fixed-shape partition polytopes would help compute robustness measures for clusterings [29], and would lead to improved methods for gradual transitions between so-called separable clusterings [40].

Chapter 5

Hirsch Counterexamples and the Circuit Diameter Conjecture

This chapter is based on joint work with Alexander E. Black and Steffen Borgwardt [21]. The results of Section 5.3 are presented in largely the same way in our paper.

5.1 Introduction

The vertices and bounded edges of a polyhedron P naturally define a connected graph, simply called the *graph* (or *1-skeleton*) of P . The diameter of this graph, i.e., the maximum length of a shortest path between any two vertices, is referred to as the *combinatorial diameter* of P . The famous *Hirsch conjecture*, first posed by Warren B. Hirsch in 1957 (see [70]), claimed an upper bound of $f - d$ on the combinatorial diameter of any d -dimensional polyhedron with f facets. It was disproved for unbounded polyhedra by Klee and Walkup [145] in 1967, who found a 4-dimensional unbounded counterexample with 8 facets whose combinatorial diameter is strictly greater than $8 - 4 = 4$. The Hirsch conjecture for polytopes, however, was only disproved much later, in 2012, by Santos [170]. He showed that there is a polytope in dimension 43 with 86 facets and diameter at least 44. Today, the arguably most important open question in the field is the *polynomial Hirsch conjecture*, which asks whether the combinatorial diameter is bounded by a polynomial in f and d (see [170]). Note that the existence of a strongly polynomial pivot rule for the Simplex method would require this conjecture to be true (see [144, Section 3]).

The currently best known upper bound on the combinatorial diameter is due to Sukegawa [188] and is of the form $(f - d)^{\log O(d)}$ for a d -dimensional polyhedron with f facets, improving the first subexponential bound of Kalai and Kleitman [138] (see also [192], and [189] for a recent refined asymptotic analysis). For polytopes, a linear bound in fixed dimension d was given by Larman [147] and Barnette [17]. Further note that the Hirsch conjecture is true for 0/1 polytopes, as proved by Naddef [159] (see also [178]), and for certain other special classes of polyhedra (see, e.g., [13, 37]). The *polynomial Hirsch conjecture* has been proved for polyhedra with totally unimodular constraint matrices [78] and, more generally, with integral constraint matrices for which the absolute values of all subdeterminants are polynomially bounded [24, 48, 67, 91]. We refer the reader to the surveys [143, 144, 171, 207] for a detailed account of combinatorial diameter bounds and the history of the Hirsch conjecture.

To gain a better understanding of diameters of polyhedra, the notion of the combinatorial diameter has been relaxed and generalized in a number of different ways (see, e.g.,

[65, 75, 93, 142, 143, 171] and the references therein). This chapter is concerned with one such natural relaxation that was introduced by Borgwardt, Finhold, and Hemmecke [35]. Paths in the graph of a polyhedron P “walk” from vertex to vertex along edges of P (which is why we will also refer to them as *edge walks*). Let us relax this notion of a walk and allow the next point to be any point on the boundary of P (not necessarily a vertex) as long as it can be reached from the previous point by moving maximally along a *circuit* of P . We refer to these walks as *circuit walks* and formally define them as follows. Recall from Section 4.2.1 that $\mathcal{C}(A, B)$ denotes the set of circuits of a polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$.

Definition 5.1. *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a polyhedron. Given two vertices u and v of P , a circuit walk from u to v on P is a finite sequence of points $y^{(0)}, y^{(1)}, \dots, y^{(k)} \in P$ such that $y^{(0)} = u, y^{(k)} = v$, and for all $i \in [k]$, we have that*

- (i) $y^{(i)} - y^{(i-1)} \in \mathcal{C}(A, B)$
- (ii) $y^{(i)} + \varepsilon(y^{(i)} - y^{(i-1)}) \notin P$ for all $\varepsilon > 0$.

Condition (ii) ensures that each step i is of maximal length. Given a circuit walk $y^{(0)}, y^{(1)}, \dots, y^{(k)}$, the number k of steps is called its *length*. Following [35], we define the *circuit diameter* of P as the maximum length of a shortest circuit walk between any pair of vertices of P . Note that this distance is not necessarily symmetric: The number of steps required to walk from a vertex u to another vertex v may not be the same as from v to u . This is because, unlike edge walks, circuit walks are not necessarily reversible (see Example 1 in [35], or Figure 5.2 in Section 5.2). Another technical issue is the dependence of the set of circuits of P on its linear description. Therefore, also the circuit diameter may depend on the particular description. However, as explained in Section 4.2.1, this issue can be resolved by restricting to *minimal* linear descriptions. In this case, speaking of the circuit diameter of a polyhedron as a geometric property is perfectly well-defined (see also [132]).

The circuit analogue of the Hirsch conjecture, the so-called *circuit diameter conjecture* [35], asks whether the circuit diameter of all d -dimensional polyhedra with f facets is at most $f - d$. This conjecture is open, as of writing this thesis. Interestingly, if one considers sequences of points that satisfy all properties of a circuit walk possibly except for condition (ii) in Definition 5.1, then one can always find such a sequence with endpoints u and v whose length is at most $f - d$, for any two vertices u and v of a d -dimensional polyhedron with f facets [36] (see also Section 5.2).

For circuit diameters, however, which require condition (ii), the best known upper bounds depend on the linear descriptions of polyhedra: It is known that circuit diameters of rational polyhedra are weakly polynomially bounded in the encoding length of their linear description [73]. Moreover, results of [69, 94] imply strongly polynomial bounds (i.e., depending only on the number of variables) for polyhedra in standard form with coefficients of polynomially bounded encoding length. We note that the bounds of [69, 73, 94] were obtained in the context of circuit augmentation schemes, where the circuit diameter is studied as a proxy for the number of augmentations, much in the same way that the combinatorial diameter lower bounds the number of iterations of the Simplex method. Moreover, specific bounds on circuit diameters have been obtained for certain families of polyhedra in combinatorial optimization [35, 38, 133].

As the set of circuits of a polyhedron includes all edge directions (see Section 4.2.1), edge walks clearly are special circuit walks. This means that the circuit diameter of any polyhedron is a lower bound on its combinatorial diameter. Thus, any counterexample to the circuit diameter conjecture would also be a counterexample to its combinatorial relative, the Hirsch conjecture. So a natural question, raised by Borgwardt, Finhold, and Hemmecke [35], is the following:

“It is an immediate interesting open question whether the counterexamples to the Hirsch conjecture [145, 170] give rise to counterexamples to [the circuit diameter conjecture] [...].” [35, page 2]

Stephen and Yusun [185] showed that the Klee-Walkup counterexample [145] to the *unbounded* Hirsch conjecture does not carry over to the circuit setting. In fact, the circuit diameter conjecture is true for all polyhedra with $d = 4$ and $f = 8$ [39].

In this chapter, we study the original *bounded* Hirsch counterexample of Santos [170] and the subsequent improvements by Matschke, Santos, and Weibel [155]. They are based on a class of polytopes called *spindles*. These are polytopes with two distinguished vertices u and v (the *apices*) such that each facet contains exactly one of u and v . Put differently, a spindle is the intersection of two translated pointed polyhedral cones emanating from the apices u and v such that each apex is in the interior of the opposite cone. A simple example of a spindle is a hypercube (with any pair of antipodal vertices as apices).

Santos discovered that in order to construct a Hirsch counterexample, it suffices to find a spindle for which the combinatorial distance between its apices (which he calls the *length* of the spindle) is strictly greater than the dimension. A spindle may have many facets, so even with a length exceeding the dimension, the two apices may not be at a distance that immediately violates the Hirsch bound of $f - d$. However, Santos showed in [170] that in this case, one can lift the spindle to one dimension higher to obtain a new spindle with only one more facet while the length is guaranteed to increase by at least one. So this construction leaves $f - d$ invariant while increasing the length. Iterating the lifting procedure for long enough eventually yields a spindle whose length does indeed violate the Hirsch bound. (A more detailed account of this construction is given in Section 5.3.2.)

In [170], Santos gave a highly degenerate 5-dimensional spindle (called S_5^{48} in the following) with 48 facets and length 6. He then concluded via his iterative construction that there is a spindle with 86 facets in dimension 43 and length at least 44, which is greater than $86 - 43$ and therefore violates the Hirsch conjecture. In a follow-up to Santos' work, Matschke, Santos, and Weibel [155] found two spindles, which we denote by S_5^{28} and S_5^{25} , also of dimension 5 and length 6 but with fewer facets (28 and 25, respectively). These lead to counterexamples in lower dimensions 23 and 20, respectively. Note that this 20-dimensional Hirsch counterexample is the lowest-dimensional one known to date. The authors of [155] also carried out the steps of Santos' construction for both S_5^{28} and S_5^{25} explicitly, obtaining linear descriptions of corresponding (high-dimensional) Hirsch counterexamples.

In Section 5.3, we consider Santos' original spindle S_5^{48} from [170] as well as the two smaller ones S_5^{28} and S_5^{25} from [155] and prove that their circuit length is at most 5.

Here, the *circuit length* of a spindle is the circuit analogue of its (combinatorial) length and denotes the maximum length of a shortest circuit walk from one apex of the spindle to the other one. Our result implies that the key property of the three spindles – having a length greater than the dimension – is no longer true in the circuit setting.

The main technical ingredient of our circuit length bounds in Section 5.3.1 is a simple sufficient condition for short circuit walks to exist on polygons, which we explain in Section 5.2. We apply this to certain 2-faces of the three spindles. Interestingly, our proof exhibits that for all three spindles, the apices can be connected by circuit walks of length at most 5 with no more than two non-edge steps. With the same arguments, in Section 5.3.2, we are also able to verify computationally that the circuit length of the two explicit non-Hirsch spindles constructed from S_5^{28} and S_5^{25} as given in [155] satisfies the Hirsch bound of $f - d$. In other words, the two vertices whose combinatorial distance makes each of these two spindles a Hirsch counterexample by Santos' arguments are no longer a distance apart that violates the Hirsch bound when using circuit walks.

We stress that our results only concern the circuit *length* and do not imply that the circuit *diameter* is small. We address the relationship between these two notions in the context of the circuit diameter conjecture in Section 5.4. It must further be noted that, in contrast to the combinatorial setting, circuit diameters and circuit lengths are *not* preserved under combinatorial equivalence. In particular, there can be two realizations of the same polyhedron with different circuit diameters (see, e.g., [186] or Example 2 in [35]). While the proofs of our circuit length bounds do depend on the geometry of the particular realizations of the three spindles S_5^{48} , S_5^{28} , and S_5^{25} in Section 5.3.1, our analysis extends to all realizations satisfying mild conditions. For example, these conditions are still satisfied after slight perturbations. It will remain open whether *all* realizations of these polytopes satisfy these conditions, or, more generally, whether all realizations have a circuit length bounded by the dimension. (Note here that circuit diameters and circuit lengths are preserved under *affine* equivalence; this is a consequence of Lemma 4.3.)

5.2 Preliminaries

We saw in Section 4.2.1 that the set of circuits of a pointed polyhedron admits many equivalent characterizations. In this chapter, we shall be working with yet another characterization, which is closely related to oriented matroids (see [9, 19, 160, 206]) and will be explained next. Our arguments for bounding the circuit length of the spindles in Section 5.3 are rather directly developed from this characterization, as we will see in Section 5.2.3.

5.2.1 Circuits Revisited

We first need some terminology from [33, 206]. A *hyperplane arrangement* in \mathbb{R}^n is a finite set of hyperplanes $\{H_1, \dots, H_m\}$ where each H_i is of the form $H_i = \{x \in \mathbb{R}^n : (b^{(i)})^\top x = 0\}$ for some nonzero vector $b^{(i)} \in \mathbb{R}^n$. If the vectors $b^{(i)}$ are the rows of a matrix $B \in \mathbb{R}^{m \times n}$, we denote this hyperplane arrangement by $\mathcal{H}(B)$. To simplify notation, we write $(Bx)_i$ for $(b^{(i)})^\top x$ in this case. Note that $\mathcal{H}(B)$ partitions the entire space \mathbb{R}^n into polyhedral cones with pairwise disjoint relative interiors. Each such cone

is the set of all $x \in \mathbb{R}^n$ satisfying

$$\begin{aligned} (Bx)_i &= 0 & \text{for all } i \in I^0 \\ (Bx)_i &\geq 0 & \text{for all } i \in I^+ \\ (Bx)_i &\leq 0 & \text{for all } i \in I^- \end{aligned} \tag{5.1}$$

for subsets $I^0, I^+, I^- \subseteq [m]$ that partition $[m]$. We call the cones of the form (5.1) the *faces* of $\mathcal{H}(B)$. Note that the set of faces of $\mathcal{H}(B)$ is finite and is closed under taking faces (in the ordinary polyhedral sense).

The hyperplane arrangements that are relevant for our purposes are those for which $B \in \mathbb{R}^{m \times n}$ is the constraint matrix of a full-dimensional pointed polyhedron $P = \{x \in \mathbb{R}^n : Bx \leq d\}$. In this case, following [33], we call $\mathcal{H}(B)$ the *elementary arrangement* of P . The union of the one-dimensional faces (the *extreme rays*) of the elementary arrangement of P , excluding the origin $\mathbf{0}$, is precisely the set of circuits of P , as observed in [33]. This can be seen, for example, using Proposition 4.2.

Even though any polyhedron given by a minimal linear description $\{x \in \mathbb{R}^{n'} : Ax = b, Bx \leq d\}$ is affinely equivalent to a full-dimensional polyhedron of the above form in dimension $n = n' - \text{rk}(A)$ (see [33]), it will be more convenient for our exposition below to extend the terminology of [33] to polyhedra given by general linear systems, possibly with equality constraints. So let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a pointed polyhedron, and let $\mathcal{H}(B)$ be defined as above. The intersection of each face of $\mathcal{H}(B)$ with $\ker(A)$ is a polyhedral cone again. By a slight abuse of terminology, we call these cones the *faces* of the *restricted elementary arrangement* of P (restricted to $\ker(A)$). Again, it follows from Proposition 4.2 that $\mathcal{C}(A, B)$, the set of circuits of P , is the union of all one-dimensional faces of the restricted elementary arrangement of P , where $\mathbf{0}$ is excluded. As with the elementary arrangement, the set of all faces of the *restricted elementary arrangement* is easily seen to be closed under taking faces. The minimal face of the restricted elementary arrangement of a pointed polyhedron is $\{\mathbf{0}\}$.

With this characterization of the set of circuits, the connection to oriented matroids that was mentioned above can be made precise. (Incidentally, this connection also justifies the use of the term “circuit”.) To simplify the explanation, we assume a full-dimensional pointed polyhedron $P = \{x \in \mathbb{R}^n : Bx \leq d\}$ where $B \in \mathbb{R}^{m \times n}$. Let F be a face of its elementary arrangement $\mathcal{H}(B)$. Note that for any two vectors x and y in the relative interior of F , the signs of $(Bx)_i$ and $(By)_i$ agree for all $i \in [m]$. So each face of the elementary arrangement is uniquely identified by a sign vector in $\{0, +, -\}^m$. With respect to a partial order on sign vectors that prefers 0 over any of $+, -$, the minimal sign vectors from $\{0, +, -\}^m$ form the *circuits* of an oriented matroid (see Lecture 6 of [206] or [9, 19, 160] for details). These correspond precisely to the extreme rays of $\mathcal{H}(B)$ which, in turn, are generated by the circuits of P , as we saw above.

Before we proceed, let us briefly digress and explain how the above characterization of the circuits leads to the well-known concept of conformal sums. This concept is at the core of a fundamental property of the set of circuits, which may be regarded as additional motivation for why the Hirsch bound of $f - d$ is of interest for bounding circuit diameters.

5.2.2 Conformal Sums of Circuits

Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a pointed polyhedron with an $m \times n$ matrix B . Two vectors $x, y \in \mathbb{R}^n$ are *sign-compatible with respect to B* if $(Bx)_i(By)_i \geq 0$ for all $i \in [m]$. With this definition, we may state the so-called *sign-compatible representation property* (or *conformal sum property*) [120] (see also [9, 165]) of the set of circuits of P .

Proposition 5.2 ([120]). *Let $x \in \ker(A) \setminus \{0\}$. Then x can be expressed as $x = \sum_{i=1}^k g^{(i)}$ where $k \leq n$ and, for all $i \in [k]$,*

(i) $g^{(i)} \in \mathcal{C}(A, B)$ and

(ii) $g^{(i)}$ and x are sign-compatible with respect to B and $\text{supp}(Bg^{(i)}) \subseteq \text{supp}(Bx)$.

Such a decomposition into the sum of sign-compatible circuits is called a *conformal sum of circuits*. Note that the set of all vectors in $\ker(A)$ (not necessarily circuits) that satisfy condition (ii) is a polyhedral cone, as observed in the proof of Lemma 2 in [36]. Stated in a slightly different form, this observation is as follows.

Proposition 5.3 (see [36]). *For $x \in \ker(A)$, let F be the set of all vectors $g \in \ker(A)$ such that g and x are sign-compatible with respect to B and $\text{supp}(Bg) \subseteq \text{supp}(Bx)$. Then F is the minimal face of the restricted elementary arrangement of P containing x .*

Proof. By definition, a vector $g \in \ker(A)$ is in F if and only if, for all $i \in [m]$, we have that

$$(Bg)_i \begin{cases} = 0 & \text{if } (Bx)_i = 0 \\ \geq 0 & \text{if } (Bx)_i > 0 \\ \leq 0 & \text{if } (Bx)_i < 0 \end{cases}$$

Hence, F is the intersection of $\ker(A)$ with the face (5.1) of the elementary arrangement of P where $I^0 = [m] \setminus \text{supp}(Bx)$ and an index $i \in \text{supp}(Bx)$ is in I^+ or I^- depending on the sign of $(Bx)_i$. Clearly, $x \in F$. So any face of the restricted elementary arrangement of P that contains x must be a face of F . However, $(Bx)_i \neq 0$ for all $i \in I^+ \cup I^-$, which implies that no proper face of F can contain x . Therefore, F is the minimal face of the restricted elementary arrangement of P with this property. \square

Note that all vectors in the face F as defined in Proposition 5.3 are not only sign-compatible with x but also pairwise sign-compatible. This means that the elementary arrangement $\mathcal{H}(B)$ partitions $\ker(A)$ into cones of pairwise sign-compatible vectors. In particular, we obtain the following corollary.

Corollary 5.4. *For every face F of the restricted elementary arrangement of P , all vectors in F are pairwise sign-compatible.*

The sign-compatible representation property (Proposition 5.2) now readily follows from Proposition 5.3 and standard polyhedral theory:

Proof of Proposition 5.2. Let F be the minimal face of the restricted elementary arrangement of P such that $x \in F$. Then F is a polyhedral cone of dimension at most n that is generated by circuits of P . It follows from Carathéodory's theorem that (possibly

after scaling) x can be expressed as a sum of $k \leq n$ circuits $g^{(1)}, \dots, g^{(k)} \in \mathcal{C}(A, B) \cap F$. Since $g^{(i)} \in F$ for all $i \in [k]$, all of them satisfy condition (ii) in Proposition 5.2 by Proposition 5.3. \square

Applied to differences of vertices of a polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$, the sign-compatible representation property has an interesting consequence for finding circuit walks on P . Namely, take two distinct vertices u and v of P . Then $v - u \in \ker(A) \setminus \{\mathbf{0}\}$. This means that $v - u$ can be expressed as a conformal sum of $k \leq n$ circuits $g^{(1)}, \dots, g^{(k)}$ of P by Proposition 5.2. Now define a sequence of points $y^{(0)}, \dots, y^{(k)}$ by $y^{(0)} = u$ and $y^{(j)} = y^{(j-1)} + g^{(j)}$ for $j \in [k]$. In particular, the last point in this sequence is $y^{(k)} = v$. An important observation is that each point $y^{(j)}$ in this sequence is in P , no matter how the circuits were ordered. This is implicitly used in the proof of Lemma 2 in [36]. For the sake of clarity, we give an explicit proof of this fact.

Lemma 5.5. *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a pointed polyhedron. Given $u, v \in P$, let $v - u = \sum_{i=1}^k g^{(i)}$ be a conformal sum of circuits in $\mathcal{C}(A, B)$. Then $u + \sum_{i=1}^j g^{(i)} \in P$ for all $j \in [k]$.*

Proof. First, note that all points $u + \sum_{i=1}^j g^{(i)}$ for $j \in [k]$ are in the zonotope Z given by

$$Z := u + \sum_{i=1}^k [0, 1]g^{(i)}.$$

It therefore suffices to show that $Z \subseteq P$.

Suppose for the sake of contradiction that there is some point $z \in Z \setminus P$. Since $u \in P$ and $g^{(i)} \in \ker(A)$ for all $i \in [k]$ by the definition of $\mathcal{C}(A, B)$, we must have $(Bz)_l > d_l$ for some $l \in [m]$. By Proposition 5.3, the circuits $g^{(i)}$ from the decomposition of $v - u$ must be in the minimal face of the restricted elementary arrangement of P that contains $v - u$. Let us denote this face by F . Now consider the vectors $z - u$ and $v - z$. Since $z \in Z$, we have that $z - u = \sum_{i=1}^k \alpha_i g^{(i)}$ for some $\alpha_i \in [0, 1], i \in [k]$. Thus, $v - z = (v - u) - (z - u) = \sum_{i=1}^k (1 - \alpha_i)g^{(i)}$. As F is a cone and $\alpha_i \in [0, 1]$ for all $i \in [k]$, this implies that both $z - u$ and $v - z$ are in F . Hence, by Corollary 5.4, they must be sign-compatible with respect to B . In particular, this means that $(Bz - Bu)_l \cdot (Bv - Bz)_l \geq 0$. However, $(Bu)_l \leq d_l$ and $(Bv)_l \leq d_l$ (by the hypothesis that $u, v \in P$) whereas $(Bz)_l > d_l$. So $(Bz - Bu)_l > 0$ and $(Bv - Bz)_l < 0$, a contradiction. Thus, $Z \subseteq P$ as desired. \square

Comparing the properties of the sequence $y^{(0)}, \dots, y^{(k)}$ with the definition of a circuit walk (Definition 5.1), we therefore almost obtained a circuit walk from u to v : All points are feasible by Lemma 5.5 and the difference of any two consecutive points is a circuit by construction. The only requirement that may not be met, however, is requirement (ii) in Definition 5.1, which states that each circuit step in the sequence must be of maximal length. Sequences that satisfy all of the properties of circuit walks except possibly the maximality requirement (ii) are referred to as *feasible circuit walks* (as opposed to *maximal* ones) in [36].

We note that the length of such a feasible circuit walk always satisfies the Hirsch bound of $f - d$. Indeed, it follows from a careful analysis of our proof of Proposition 5.2

that for every pair of vertices u, v of a d -dimensional polyhedron P with f facets, the difference vector $v - u$ can be expressed as a conormal sum of at most d' circuits of P , where d' is the dimension of the minimal face of the restricted elementary arrangement containing $v - u$. Equivalently, d' is the dimension of the minimal face of P containing both u and v . Using an argument from [159] (see also [178, 206]), one can show that $d' \leq f - d$. Namely, if u and v do not share a facet, then $f \geq 2d$ and therefore $d' = d \leq f - d$. Otherwise, any facet that contains both u and v can have at most $f - 1$ facets itself, so $d' \leq (f - 1) - (d - 1) = f - d$ by induction on d .

The above discussion suggests the following question: Among all conormal circuit decompositions of differences of two given vertices, does there always exist one that may be turned into an actual circuit walk (with steps of maximal length)? Unfortunately, no, as the following example shows (see also Lemma 1 of [36]).

Example 5.6. Consider the polygon P in Figure 5.1 and the two vertices u and v . The shaded region R is the intersection of the minimal faces of the (restricted) elementary arrangement of P that contain $v - u$ and $u - v$, respectively, when translated over to u and v , respectively. For every conormal sum of circuits that yields $v - u$, the zonotope as defined in the proof of Lemma 5.5 must be contained in R . However, the only boundary points of P contained in R are u and v , and $v - u$ is not a circuit (as is easily checked). So starting from either of the vertices u or v , any maximal circuit step must therefore leave R . \diamond




Figure 5.1

So to find short circuit walks of *maximal* step length, other tools are required.

5.2.3 Short Circuit Walks on Polygons

The main tool that we will leverage in this chapter is based on the following simple but useful fact.

Remark 5.7. *The recession cone of every pointed polyhedron is a face of its restricted elementary arrangement.*

This has the following interesting consequence for circuit walks on polytopes. If the deletion of some inequalities from its linear description makes a polytope P unbounded, then there is always a single circuit step (of maximal length) from any point in P to

some point on a facet defined by one of the deleted inequalities. More specifically, we prove the following statement, which is the key technical lemma of this chapter.

Lemma 5.8. *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d, B'x \leq d'\}$ be a polytope such that $Bx \leq d$ contains no implicit equations. Suppose that the polyhedron P' given by $P' := \{x \in \mathbb{R}^n : Ax = b, B'x \leq d'\}$ is unbounded. Then, for all $y^{(0)} \in P$, there is a point $y^{(1)}$ on a facet of P defined by an inequality from $Bx \leq d$ such that $y^{(1)}$ can be reached from $y^{(0)}$ in at most one circuit step.*

Proof. First, note that the restricted elementary arrangement of P is a refinement of that of P' since P' is obtained from the linear description of P by omitting inequalities. By Remark 5.7, the recession cone of P' is a face of the restricted elementary arrangement of P' . As such, it is a union of faces of the restricted elementary arrangement of P . Since P' is unbounded by hypothesis, the recession cone of P' is of dimension at least one and therefore contains a circuit g of P .

Now let $y^{(0)} \in P$. Then also $y^{(0)} \in P'$ since $P \subseteq P'$. It follows that the ray given by $y^{(0)} + \mu g$ for all $\mu \geq 0$ is contained in P' and follows a circuit direction. As the affine hulls of P and P' are the same by hypothesis, this ray is also contained in the affine hull of P . Since P is bounded but P' is not, the ray must therefore intersect a facet of P defined by one of the inequalities from $Bx \leq d$ in a point $y^{(1)}$. \square

Specializing Lemma 5.8 to polygons, we obtain a simple sufficient condition for the existence of short circuit walks that terminate at one of a specific set of target vertices.

Lemma 5.9. *Let $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$ be a polygon such that $Bx \leq d$ contains no implicit equations, and let \mathcal{V} be a subset of its vertices. Further, let $C(P, \mathcal{V}) := \{x \in \mathbb{R}^n : Ax = b, B'x \leq d'\}$, where the system $B'x \leq d'$ consists of all inequalities from $Bx \leq d$ that are not tight at any vertex $v \in \mathcal{V}$. If $C(P, \mathcal{V})$ is unbounded, then for all vertices $y^{(0)} \notin \mathcal{V}$ of P there is a circuit walk of length at most 2 from $y^{(0)}$ to some $v \in \mathcal{V}$ on P .*

Proof. First, note that any vertex $v \in \mathcal{V}$ can be reached via a single circuit step starting from any point on an edge of P incident with v . In particular, this is true for the two adjacent vertices of v . Thus, if every vertex of P is adjacent to a vertex in \mathcal{V} , then we are done, so suppose there exists a vertex $y^{(0)}$ of P that is not adjacent to any vertex in \mathcal{V} . By Lemma 5.8, there is a point $y^{(1)}$ on an edge of P such that $y^{(1)} - y^{(0)}$ is a circuit of P and the edge that contains $y^{(1)}$ is incident with some vertex $v \in \mathcal{V}$. So v can be reached from $y^{(1)}$ in at most one circuit step. \square

In particular, Lemma 5.9 guarantees that for every vertex v of a polygon P , there is always a circuit walk of length at most 2 from any other vertex to v if the omission of all inequalities that are tight at v would make P unbounded. See Figure 5.2 for a schematic picture. While the statement and proof of Lemma 5.9 are phrased in terms of a vertex $y^{(0)}$, the statement remains true for any point in P .

Lemma 5.9 will play a central role in our discussion of Santos' counterexamples to the Hirsch conjecture in the next section.




Figure 5.2: A polygon P with a vertex v and the corresponding polyhedron $C(P, \{v\})$ as defined in Lemma 5.9. The recession cone of $C(P, \{v\})$ is generated by the two dashed edge directions. Starting at the vertex $y^{(0)}$, we can follow any of them to a point $y^{(1)}$ on one of the two edges of P incident with v . This gives a circuit walk on P from $y^{(0)}$ to v of length 2.

5.3 Bounded Hirsch Counterexamples

In this section, we study the three 5-dimensional spindles S_5^{48} , S_5^{28} , and S_5^{25} that provide the basis of Santos' construction for obtaining bounded Hirsch counterexamples. The key property of these spindles is that their length of 6 strictly exceeds their dimension 5. In contrast, we prove in Section 5.3.1 that their circuit length is at most 5.

Our arguments rely on a careful analysis of certain 2-faces and, as we will see, extend to all realizations of the spindles that satisfy mild assumptions, including slight perturbations. As an additional benefit, the arguments also enable us in Section 5.3.2 to verify directly that neither of the two explicit (high-dimensional) Hirsch counterexamples from [155], which are specific instances of Santos' construction, has a circuit length exceeding the dimension.

5.3.1 The Circuit Length of the 5-Dimensional Spindles

For each of the three spindles S_5^{48} , S_5^{28} , and S_5^{25} , we will show the following: Within 3 edge steps from one apex, one can reach a 2-face that contains the other apex and satisfies the condition of Lemma 5.9 when deleting all inequalities that are tight for the other apex. We first explain the arguments in detail for the 5-dimensional spindle S_5^{48} with 48 facets and length 6 that Santos' original counterexample in [170] is constructed from. Following [170, Theorem 3.1], it is given by the minimal description

$S_5^{48} = \{x \in \mathbb{R}^5 : A^+x \leq \mathbf{1}, A^-x \leq \mathbf{1}\}$ where A^+ and A^- are the matrices

$$A^+ = \begin{pmatrix} 1 & \pm 18 & 0 & 0 & 0 \\ 1 & 0 & \pm 18 & 0 & 0 \\ 1 & 0 & 0 & \pm 45 & 0 \\ 1 & 0 & 0 & 0 & \pm 45 \\ 1 & \pm 15 & 15 & 0 & 0 \\ 1 & \pm 15 & -15 & 0 & 0 \\ 1 & 0 & 0 & \pm 30 & 30 \\ 1 & 0 & 0 & \pm 30 & -30 \\ 1 & 0 & \pm 10 & 40 & 0 \\ 1 & 0 & \pm 10 & -40 & 0 \\ 1 & \pm 10 & 0 & 0 & 40 \\ 1 & \pm 10 & 0 & 0 & -40 \end{pmatrix} \begin{matrix} 1^\pm \\ 2^\pm \\ 3^\pm \\ 4^\pm \\ 5^\pm \\ 6^\pm \\ 7^\pm \\ 8^\pm \\ 9^\pm \\ 10^\pm \\ 11^\pm \\ 12^\pm \end{matrix}, \quad A^- = \begin{pmatrix} -1 & 0 & 0 & 0 & \pm 18 \\ -1 & 0 & 0 & \pm 18 & 0 \\ -1 & \pm 45 & 0 & 0 & 0 \\ -1 & 0 & \pm 45 & 0 & 0 \\ -1 & 0 & 0 & 15 & \pm 15 \\ -1 & 0 & 0 & -15 & \pm 15 \\ -1 & \pm 30 & 30 & 0 & 0 \\ -1 & \pm 30 & -30 & 0 & 0 \\ -1 & 40 & 0 & \pm 10 & 0 \\ -1 & -40 & 0 & \pm 10 & 0 \\ -1 & 0 & 40 & 0 & \pm 10 \\ -1 & 0 & -40 & 0 & \pm 10 \end{pmatrix} \begin{matrix} 13^\pm \\ 14^\pm \\ 15^\pm \\ 16^\pm \\ 17^\pm \\ 18^\pm \\ 19^\pm \\ 20^\pm \\ 21^\pm \\ 22^\pm \\ 23^\pm \\ 24^\pm \end{matrix}$$

with 24 rows each, labelled 1^\pm to 24^\pm . We will also use these labels for the corresponding inequalities and the facets that they define. The two apices of S_5^{48} are $v^+ = (1, 0, 0, 0, 0)$ and $v^- = (-1, 0, 0, 0, 0)$. The spindle S_5^{48} can equivalently be written as $S_5^{48} = (C^+ + v^+) \cap (C^- + v^-)$ for the two cones $C^+ = \{x \in \mathbb{R}^5 : A^+x \leq \mathbf{0}\}$ and $C^- = \{x \in \mathbb{R}^5 : A^-x \leq \mathbf{0}\}$.

Among the properties of S_5^{48} listed in [170] are the following two symmetries. They are direct consequences of the symmetry in the coefficients of A^\pm .

Proposition 5.10 ([170]). *The following linear transformations of \mathbb{R}^5 leave S_5^{48} invariant:*

(i) $(x_1, x_2, x_3, x_4, x_5) \mapsto (-x_1, x_5, x_4, x_2, x_3),$

(ii) $(x_1, x_2, x_3, x_4, x_5) \mapsto (x_1, x_3, x_2, x_5, x_4).$

We note that the transformation given in Proposition 5.10(i) switches the roles of A^+ and A^- and of v^+ and v^- , while (ii) is a permutation of coordinates that preserves the matrices A^\pm up to reordering rows. With these observations we are now able to bound the circuit length of S_5^{48} .

Theorem 5.11. *The circuit length of S_5^{48} is at most 5.*

Proof. By Proposition 5.10(i), S_5^{48} is symmetric under a linear transformation that sends v^+ to v^- and vice versa. Hence, from any circuit walk from v^+ to v^- on S_5^{48} , one can immediately obtain a circuit walk from v^- to v^+ by applying the same linear transformation. It therefore suffices to find a circuit walk of length at most 5 in one of the directions, say from v^+ to v^- , to prove the statement.

Let F be the 2-face of S_5^{48} defined by inequalities 15^+ , 19^+ , and 21^+ . Figure 5.3 shows the graph of F , along with shortest paths from v^+ to four vertices of F that are at distance 3 from v^+ in the graph of S_5^{48} (graph computations were done using Polymake [110]). For this face F , we now claim that the polyhedron $C(F, \{v^-\})$ as defined in Lemma 5.9 is unbounded. It then follows from Lemma 5.9 that from either of the four vertices that are at distance 3 from v^+ , there is a circuit walk on F to the

other apex v^- of length at most 2. Since circuits of F are also circuits of S_5^{48} , this means that there is a circuit walk on S_5^{48} of length at most 5 from v^+ to v^- .

To verify that $C(F, \{v^-\})$ is unbounded, it suffices to show the following property:

Claim (*). *The polyhedron Q given by inequalities 13^\pm to 24^\pm together with 4^\pm is bounded.*

Before we give a proof, let us first argue why the claim implies unboundedness of $C(F, \{v^-\})$. The inequalities 13^\pm to 24^\pm and 4^\pm that determine Q must also be facet-defining for Q since each of them defines a facet of S_5^{48} and $S_5^{48} \subseteq Q$. Let F_Q be the 2-face of Q defined by 15^+ , 19^+ , and 21^+ . Since $S_5^{48} \subseteq Q$, we also have that $F \subseteq F_Q$. Hence, v^- and its two adjacent vertices $4^+13^+17^+23^+$ and $4^-13^-17^-23^-$ in Figure 5.3 are also vertices of F_Q . Moreover, any edge of F_Q that does not contain v^- can only be defined by 4^+ or 4^- because v^- is in all other facets of Q . In fact, both 4^+ and 4^- define edges of F_Q . To see this, note that the (proper) face of F_Q defined by 4^+ contains an edge of F (namely, the edge between vertices $4^+13^+17^+23^+$ and 4^+11^+ in Figure 5.3) and, hence, must be an edge of F_Q itself. A similar argument shows that 4^- defines an edge of F_Q , too, which is incident with the vertex $4^-13^-17^-23^-$. These two edges of F_Q must be distinct as the two vertices $4^+13^+17^+23^+$ and $4^-13^-17^-23^-$ would otherwise be adjacent on F_Q and therefore also on F ; a contradiction (see Figure 5.3).

Now suppose that Q is bounded. Then F_Q is bounded, which means that the two edges of F_Q defined by 4^+ and 4^- must intersect in a vertex w of F_Q . Further observe that for each vertex of F in Figure 5.3, at most one of the two inequalities 4^\pm is tight. This means that w cannot be a vertex of F . Thus, F_Q has exactly four vertices, namely v^- , $4^+13^+17^+23^+$, $4^-13^-17^-23^-$, and w . The first three are also vertices of F and therefore satisfy all inequalities that define edges of F . Since F_Q is bounded and $F \subseteq F_Q$, any inequality that defines an edge of F but not of F_Q must cut off a vertex of F_Q , and that can only be w . Since each edge of $C(F, \{v^-\})$ is either defined by 4^\pm or by an inequality that defines an edge of F but not of F_Q , the recession cone of $C(F, \{v^-\})$ is therefore precisely the 2-dimensional feasible cone of F_Q at w . This means that $C(F, \{v^-\})$ is unbounded as desired.

It remains to prove Claim (*). By definition of Q , the recession cone of Q consists of all vectors $x \in C^-$ that further satisfy the two facet-defining inequalities 4^\pm for the cone C^+ (with right-hand side 0). Since the two rows 4^\pm of A^+ add up to $(2, 0, 0, 0, 0)$, all vectors x in the recession cone of Q must satisfy $x_1 \leq 0$. Further, the sum of all rows 13^\pm to 24^\pm of A^- equals $(-24, 0, 0, 0, 0)$. This implies that $(-1, 0, 0, 0, 0)$ is in the interior of the polar cone of C^- . For the linear objective function $(-1, 0, 0, 0, 0)^\top x$, the origin $\mathbf{0}$ is therefore the unique maximizer over C^- . Hence, any nonzero vector $x \in C^-$ must satisfy $x_1 > 0$. The recession cone of Q therefore only contains $\mathbf{0}$, which means that Q must be bounded. \square

Our proof of Theorem 5.11 exhibits that it is possible to verify that the circuit length of S_5^{48} is strictly less than the combinatorial length through the geometry of its 2-faces: From each of the vertices $7^+11^+15^+19^+21^+$ and $8^+12^+15^+19^+21^+$ in Figure 5.3 (which are at distance 3 from v^+), there is an edge walk of length 3 to v^- that stays on the 2-face F . In contrast, two circuit steps suffice to reach v^- as shown above.




Figure 5.3: The subgraph of S_5^{48} induced by all vertices of the 2-face F defined by inequalities 15^+ , 19^+ , and 21^+ , and by vertices on a shortest path from v^+ to a vertex of F . Vertex labels indicate which inequalities are tight. For the vertices of F , we omitted the facets containing F from their labels. The four highlighted vertices are at distance 3 from v^+ .

We further note that Claim (*) is the only geometric property of S_5^{48} used in the proof of Theorem 5.11. The remainder of the argument does not depend on the particular realization. Indeed, consider a different realization \tilde{S}_5^{48} whose facets (and facet-defining inequalities) are labelled 1^\pm to 24^\pm again such that facets of \tilde{S}_5^{48} and S_5^{48} with the same label are combinatorially equivalent. Suppose that \tilde{S}_5^{48} satisfies Claim (*). As an immediate consequence of the proof of Theorem 5.11, there is a circuit walk of length at most 5 from one of the apices of \tilde{S}_5^{48} to the other one. However, since \tilde{S}_5^{48} might not be symmetric under the linear transformation in Proposition 5.10(i), this does not directly imply the existence of a short walk in the converse direction as well. In order to extend our bound on the circuit length to \tilde{S}_5^{48} , we require validity of a “symmetric” version of Claim (*). More precisely, we obtain the following corollary.

Corollary 5.12. *Let \tilde{S}_5^{48} be a realization of S_5^{48} with facets and facet-defining inequalities labelled 1^\pm to 24^\pm such that each facet is combinatorially equivalent to the facet of S_5^{48} with the same label. Suppose that \tilde{S}_5^{48} satisfies both of the following properties:*

- (i) *The polyhedron given by inequalities 13^\pm to 24^\pm together with one of the pairs 3^\pm or 4^\pm is bounded.*
- (ii) *The polyhedron given by inequalities 1^\pm to 12^\pm together with one of the pairs 15^\pm or 16^\pm is bounded.*

Then the circuit length of \tilde{S}_5^{48} is at most 5.

In particular, Corollary 5.12 applies to all realizations resulting from mild perturbations of S_5^{48} .

Proof. Under the permutation of coordinates given in Proposition 5.10(ii) (which is an involution and leaves S_5^{48} invariant) the pairs of inequalities 3^\pm and 4^\pm , and 15^\pm and 16^\pm are in correspondence. This means that S_5^{48} has a 2-face G that is linearly

isomorphic to the 2-face F from the proof of Theorem 5.11 and satisfies Claim (*) with 3^\pm instead of 4^\pm . Since any realization of S_5^{48} has two faces that are combinatorially equivalent with F and G , respectively, we conclude from the proof of Theorem 5.11 that there is a circuit walk of length at most 5 from the apex not contained in F or G to the other one on any realization with property (i).

For the converse direction, recall that S_5^{48} is also invariant under the linear transformation given in Proposition 5.10(i). For the pairs of facets defined by inequalities 3^\pm and 4^\pm , the corresponding facets in the image are defined by 15^\pm and 16^\pm . Thus, if we assume property (ii), the existence of a circuit walk of length at most 5 in the converse direction follows directly from the proof of Theorem 5.11 again. \square

For the spindle S_5^{48} as realized in [170, Theorem 3.1], we are even able to determine its exact circuit length.

Theorem 5.13. *The circuit length of S_5^{48} is 2.*

Proof. Consider again the face F of S_5^{48} defined by inequalities 15^+ , 19^+ , and 21^+ from the proof of Theorem 5.11. Let $y = (0, \frac{1}{45}, \frac{1}{90}, \frac{1}{90}, \frac{7}{360}) = \frac{5}{8} \cdot 7^+ 11^+ 15^+ 19^+ 21^+ + \frac{3}{8} \cdot 4^+ 11^+ 15^+ 19^+ 21^+$ (where we use the vertex labelling of Figure 5.3 to refer to the corresponding coordinate vectors). By this construction, y is a point in the relative interior of the edge of F defined by 11^+ . Moreover, each of the two vectors $360(y - v^\pm) = (\pm 360, 8, 4, 4, 7)$ is a circuit: $(360, 8, 4, 4, 7)$ is parallel to the four facets 12^- , 15^+ , 19^+ , and 21^+ ; and $(-360, 8, 4, 4, 7)$ is parallel to 11^+ , 15^- , 20^- , and 22^- . In both cases, it can be verified by a direct computation that the corresponding rows of A^\pm are linearly independent. Since the edge of F defined by 11^+ neither contains v^+ nor v^- , the point y can be reached from either of v^\pm via a circuit step of maximal step length each. We conclude that both the sequence v^+, y, v^- and its reverse are circuit walks of length 2.

Further, the vector $(v^+ - v^-)/2 = (1, 0, 0, 0, 0)$ is not parallel to any facet of S_5^{48} and thus cannot be a circuit. Hence, the circuit length of S_5^{48} is 2. \square

Santos' original example constructed from S_5^{48} is not the lowest-dimensional bounded Hirsch counterexample known to date. In [155], Matschke, Santos, and Weibel gave two smaller counterexamples, both of which are constructed from 5-dimensional spindles of length 6 with 28 and 25 facets, respectively. The first one, S_5^{28} , from [155, Corollary 2.9] is given by the minimal description $S_5^{28} = \{x \in \mathbb{R}^5 : A^+ x \leq \mathbf{1}, A^- x \leq \mathbf{1}\}$ for the matrices

$$A^+ = \begin{pmatrix} 1 & \pm 18 & 0 & 0 & 0 \\ 1 & 0 & 0 & \pm 30 & 0 \\ 1 & 0 & 0 & 0 & \pm 30 \\ 1 & 0 & 5 & 0 & \pm 25 \\ 1 & 0 & -5 & 0 & \pm 25 \\ 1 & 0 & 0 & 18 & \pm 18 \\ 1 & 0 & 0 & -18 & \pm 18 \end{pmatrix} \begin{matrix} 1^\pm \\ 2^\pm \\ 3^\pm \\ 4^\pm \\ 5^\pm \\ 6^\pm \\ 7^\pm \end{matrix}, \quad A^- = \begin{pmatrix} -1 & 0 & 0 & \pm 18 & 0 \\ -1 & 0 & \pm 30 & 0 & 0 \\ -1 & \pm 30 & 0 & 0 & 0 \\ -1 & 25 & 0 & 0 & \pm 5 \\ -1 & -25 & 0 & 0 & \pm 5 \\ -1 & 18 & \pm 18 & 0 & 0 \\ -1 & -18 & \pm 18 & 0 & 0 \end{pmatrix} \begin{matrix} 8^\pm \\ 9^\pm \\ 10^\pm \\ 11^\pm \\ 12^\pm \\ 13^\pm \\ 14^\pm \end{matrix}$$

with 14 rows each, labelled 1^\pm to 14^\pm . Again, the two apices are $v^+ = (1, 0, 0, 0, 0)$ and $v^- = (-1, 0, 0, 0, 0)$, and we can write S_5^{28} in the form $(C^+ + v^+) \cap (C^- + v^-)$

for the two cones $C^+ = \{x \in \mathbb{R}^5: A^+x \leq \mathbf{0}\}$ and $C^- = \{x \in \mathbb{R}^5: A^-x \leq \mathbf{0}\}$. Our arguments for bounding the circuit length of S_5^{48} easily carry over to S_5^{28} by analyzing the 2-faces of S_5^{28} . The following result is the analogous statement to Theorem 5.11 and Corollary 5.12.

Corollary 5.14. *The circuit length of S_5^{28} is at most 5. The same bound holds for all realizations of S_5^{28} with facets and facet-defining inequalities labelled 1^\pm to 14^\pm such that each facet is combinatorially equivalent to the facet of S_5^{28} with the same label, and such that*

- (i) *the polyhedron given by inequalities 8^\pm to 14^\pm together with one of the pairs 2^\pm or 3^\pm is bounded, and*
- (ii) *the polyhedron given by inequalities 1^\pm to 7^\pm together with one of the pairs 9^\pm or 10^\pm is bounded.*

Proof. The proof strategy is identical to the proofs of Theorem 5.11 and Corollary 5.12. We only give the necessary modifications here.

Each apex of S_5^{28} is contained in a 2-face with a vertex at distance 3 from the other apex in the graph of S_5^{28} (see the graphs in Figure 5.4). The pairs of facet-defining inequalities 2^\pm , 3^\pm , 9^\pm , and 10^\pm now take the role that 3^\pm , 4^\pm , 15^\pm , and 16^\pm took for S_5^{48} (cf. Corollary 5.12). Hence, properties (i) and (ii) are the analogues of Claim (*) from the proof of Theorem 5.11 (and properties (i) and (ii) in Corollary 5.12) and are therefore sufficient conditions for the existence of circuit walks of length at most 5 between the apices of any realization of S_5^{28} .

It remains to show that S_5^{28} itself satisfies properties (i) and (ii). Note that S_5^{28} is not symmetric under the transformations in Proposition 5.10: To switch the roles of the apices while leaving S_5^{28} invariant, rows 4^\pm and 5^\pm would have to correspond with 11^\pm and 12^\pm , which is impossible to achieve by permutating coordinates and flipping signs. Similarly, no permutation of coordinates (except for the identity) preserves A^\pm . This means that we cannot use the same argument as in the proof of Corollary 5.12 to reduce all four pairs of inequalities 2^\pm , 3^\pm , 9^\pm , and 10^\pm to just one. However, summing over all rows of A^+ yields $(14, 0, 0, 0, 0)$, which implies that this vector is in the interior of the polar cone of C^+ . Similarly, the sum of all rows of A^- , which is the vector $(-14, 0, 0, 0, 0)$, is in the interior of the polar cone of C^- . For each of the pairs 2^\pm , 3^\pm , 9^\pm , and 10^\pm , the two corresponding rows of A^+ or A^- add up to $(\pm 2, 0, 0, 0, 0)$, respectively. So by the same argument as in the proof of Claim (*) in the proof of Theorem 5.11, it then follows that, in fact, all four polyhedra described in (i) and (ii) are bounded. \square

As for S_5^{48} , we are able to establish a circuit length of exactly 2 for the particular realization of S_5^{28} in [155, Corollary 2.9].

Theorem 5.15. *The circuit length of S_5^{28} is 2.*

Proof. A direct computation shows that the five inequalities 3^+ , 6^+ , 10^+ , 11^+ , and 13^+ define a vertex $y = (0, \frac{1}{30}, \frac{2}{90}, \frac{2}{90}, \frac{1}{30})$ of S_5^{28} (one of the highlighted vertices in Figure 5.4a). By a direct computation, one can verify that both difference vectors




Figure 5.4: Four subgraphs of S_5^{28} induced by all vertices of a 2-face F and by vertices on a shortest path from the apex not contained in F to a vertex of F . The face F is defined by inequalities (a) $10^+, 11^+, 13^+$, or (b) $8^+, 9^+, 13^+$, or (c) $1^+, 2^+, 6^+$, or (d) $3^+, 4^+, 6^+$, respectively. Vertex labels indicate which inequalities are tight. For the vertices of F , we omitted the facets containing F from their labels. In each of the graphs (a) to (d), the two highlighted vertices are at distance 3 from both v^+ and v^- .

$90(y - v^\pm) = (\pm 90, 3, 2, 2, 3)$ are circuits: $(90, 3, 2, 2, 3)$ is parallel to facets 3^- , 7^- , 10^+ , and 13^+ ; and $(-90, 3, 2, 2, 3)$ is parallel to 3^+ , 6^+ , 10^- , and 14^- . Hence, v^+, y, v^- is a reversible circuit walk of length 2. Since the vector $(v^+ - v^-)/2 = (1, 0, 0, 0, 0)$ is not parallel to any facet of S_5^{28} , it cannot be a circuit of S_5^{28} . It follows that the circuit length of S_5^{28} is 2. \square

Our framework for bounding the circuit lengths of S_5^{48} and S_5^{28} also applies to the remaining spindle from [155, Theorem 2.14], which has 25 facets and is given by the minimal description $S_5^{25} = \{x \in \mathbb{R}^5 : A^+x \leq \mathbf{1}, A^-x \leq \mathbf{1}\}$ where

$$A^+ = \begin{pmatrix} 1 & 0 & 0 & 0 & 32 \\ 1 & 0 & 0 & 0 & -32 \\ 1 & 0 & 0 & 21 & -7 \\ 1 & 0 & 0 & -21 & -7 \\ 1 & 0 & 0 & 20 & -4 \\ 1 & 0 & 0 & -20 & -4 \\ 1 & 0 & 0 & 16 & -15 \\ 1 & 0 & 0 & -16 & -15 \\ 1 & \frac{3}{50} & -\frac{1}{25} & 0 & -30 \\ 1 & -\frac{3}{50} & -\frac{1}{25} & 0 & 30 \\ 1 & \frac{3}{1000} & \frac{7}{1000} & 0 & -\frac{159}{5} \\ 1 & -\frac{3}{1000} & \frac{7}{1000} & 0 & \frac{159}{5} \end{pmatrix}, \quad A^- = \begin{pmatrix} -1 & 60 & 0 & 0 & 0 \\ -1 & -55 & 0 & 0 & 0 \\ -1 & 0 & 76 & 0 & 0 \\ -1 & 0 & -33 & 0 & 0 \\ -1 & 44 & 34 & 0 & 0 \\ -1 & 8 & -30 & 0 & 0 \\ -1 & -34 & 36 & 0 & 0 \\ -1 & -2 & -32 & 0 & 0 \\ -1 & -20 & 0 & \frac{1}{5} & -\frac{1}{5} \\ -1 & \frac{2999}{50} & 0 & -\frac{3}{25} & -\frac{1}{5} \\ -1 & \frac{299999}{5000} & 0 & 0 & \frac{1}{100} \\ -1 & -\frac{549}{10} & 0 & \frac{1}{5000} & \frac{1}{800} \\ -1 & -54 & 0 & \frac{1}{500} & -\frac{1}{80} \end{pmatrix}.$$

The two apices of S_5^{25} again are $v^+ = (1, 0, 0, 0, 0)$ and $v^- = (-1, 0, 0, 0, 0)$. For the two cones $C^+ = \{x \in \mathbb{R}^5 : A^+x \leq \mathbf{0}\}$ and $C^- = \{x \in \mathbb{R}^5 : A^-x \leq \mathbf{0}\}$, we can write S_5^{25} as $S_5^{25} = (C^+ + v^+) \cap (C^- + v^-)$.

Corollary 5.16. *The circuit length of S_5^{25} is at most 5. The same bound holds for all realizations of S_5^{25} with facets and facet-defining inequalities labelled 1 to 25 such that each facet is combinatorially equivalent to the facet of S_5^{25} with the same label, and such that*

- (i) *the polyhedron given by inequalities 13 to 25 together with one of the pairs 1, 2 or 3, 4 is bounded, and*
- (ii) *the polyhedron given by inequalities 1 to 12 together with one of the pairs 13, 14 or 15, 16 is bounded.*

Proof. Examples of the relevant 2-faces and their graphs are given in Figure 5.5. Note that for each such face, the two critical facet-defining inequalities are one of the pairs 1, 2, or 3, 4, or 13, 14, or 15, 16 in (i) and (ii). The proof of the second part of the statement is therefore analogous to the proofs for S_5^{48} and S_5^{28} above (Theorem 5.11 and Corollaries 5.12 and 5.14).

It remains to prove that S_5^{25} satisfies (i) and (ii). We again follow the proof strategy for Claim (*) in the proof of Theorem 5.11. We show that for each of the pairs 1, 2

and 13, 14, there is a nonnegative linear combination of the two corresponding rows of A^+ or A^- whose negative is in the interior of the polar cone of C^- or C^+ , respectively. It then follows that the respective polyhedra in (i) and (ii) are bounded, since their recession cones only consist of $\mathbf{0}$.

To see this, first observe that we can write the vector $(-1, 0, 0, 0, 0)$ as a nonnegative linear combination of rows 13 and 14 of A^- . (In fact, the same holds true for rows 15 and 16.) Further, consider a linear combination of the rows of A^+ that assigns coefficient $\frac{21}{8}$ for row 1, $\frac{40}{7}$ for rows 11 and 12, and coefficient 1 for all other rows. This linear combination yields a positive multiple of $(1, 0, 0, 0, 0)$. Since all coefficients are positive, we conclude that $(1, 0, 0, 0, 0)$ is in the interior of C^+ .

For rows 1 and 2 of A^+ , we proceed analogously. Their sum equals $(2, 0, 0, 0, 0)$. To show that $(-2, 0, 0, 0, 0)$ is in the interior of the polar cone of C^- , consider the linear combination of rows 22, 24, and 25 of A^- with coefficients 1, 380, and 22, respectively. The resulting vector is $(-403, -\beta, 0, 0, 0)$ where $\beta = 549 \cdot 38 + 54 \cdot 22 - \frac{2999}{50} > 0$. Further, the linear combination of rows 16 and 17 of A^- with coefficients $\frac{34}{33 \cdot 44}\beta$ and $\frac{1}{44}\beta$, respectively, yields the vector $(-\frac{67}{33 \cdot 44}\beta, \beta, 0, 0, 0)$. Adding both vectors, we thus obtain a positive multiple of $(-1, 0, 0, 0, 0)$ from a linear combination of rows 16, 17, 22, 24, and 25 of A^- where all coefficients are positive since $\beta > 0$. This means that the vector $(-1, 0, 0, 0, 0)$ is in the interior of the cone generated by rows 16, 17, 22, 24, and 25 of A^- . Note that this cone is contained in the polar cone of C^- , which is generated by all rows of A^- , and it is full-dimensional since the five rows 16, 17, 22, 24, and 25 are linearly independent. Hence, $(-1, 0, 0, 0, 0)$ must also be in the interior of the polar cone of C^- . \square

In contrast to the statements of Theorems 5.13 and 5.15 for the other two spindles, the circuit length of S_5^{25} as given in [155, Theorem 2.14] is at least 3. This can be verified computationally by a brute-force enumeration of all points $y^{(1)}$ on the boundary of S_5^{25} that can be reached from v^- via a single circuit step (S_5^{25} has 17454 circuits). For no such point $y^{(1)}$, the vector $v^+ - y^{(1)}$ is a circuit direction.

We conclude this section with some remarks on our proofs. For all three spindles S_5^{48} , S_5^{28} , and S_5^{25} , the faces given in Figures 5.3 to 5.5 are not the only 2-faces that satisfy the prerequisite of Lemma 5.9. In fact, we enumerated all 2-faces using Polymake [110] and found that for each of the three spindles there are 32 such 2-faces that contain one of the apices (16 for each apex). Each of them is combinatorially equivalent to one of the examples given in Figures 5.3 to 5.5. Moreover, for any such 2-face, the two facet-defining inequalities that are relevant for verifying the boundedness condition in Lemma 5.9 are one of the pairs given in Corollaries 5.12, 5.14 and 5.16.

Finally, our bounds on the circuit length of S_5^{48} , S_5^{28} , and S_5^{25} are robust under mild perturbations as long as they retain the properties in Corollaries 5.12, 5.14 and 5.16, respectively. However, we do not know whether, in fact, *all* realizations of the three spindles satisfy these properties. We leave this as an open question.




Figure 5.5: Four subgraphs of S_5^{25} induced by all vertices of a 2-face F and by vertices on a shortest path from the apex not contained in F to a vertex of F . The face F is defined by inequalities (a) 15, 17, 21, or (b) 13, 18, 22, or (c) 2, 8, 9, or (d) 3, 5, 9, respectively. Vertex labels indicate which inequalities are tight. For the vertices of F , we omitted the facets containing F from their labels. In each of the graphs (a) to (d), the highlighted vertices are at distance 3 from the apex not contained in F .

5.3.2 The Circuit Length of the 20- and 23-Dimensional Hirsch Counterexamples Satisfies the Hirsch Bound

Santos' original disproof of the bounded Hirsch conjecture in [170] crucially relies on finding a degenerate spindle whose (combinatorial) length is greater than its dimension. We have shown that in the circuit setting, neither Santos' original spindle S_5^{48} nor any of the subsequent improvements S_5^{28} and S_5^{25} from [155] meet this requirement: All three spindles (and slight perturbations thereof) have circuit length at most 5. This suggests that applying Santos' construction from [170] to either of them might not yield a counterexample to the circuit diameter conjecture. In this section, we provide further evidence for this.

For the two smaller of the three spindles, the steps of Santos' construction have been explicitly carried out by Matschke, Santos, and Weibel [155], resulting in inequality descriptions of two explicit Hirsch counterexamples. Using our arguments developed in Section 5.3.1, we may even verify that the circuit length of these two explicitly given spindles is indeed at most their dimension. To see how our techniques also apply here, we first explain Santos' construction in more detail. As the original construction in [170] is stated in terms of *prismatoids*, the polar duals of spindles, we briefly repeat it in the language of spindles here.

Let $S_d^f \subseteq \mathbb{R}^d$ be a d -dimensional spindle with f facets and length l where $f > 2d$ and $l > d$. We denote the apices of S_d^f by u and v . Since $f > 2d$, at least one of the apices, say u , is degenerate. Now choose an arbitrary facet F of S_d^f that contains the other apex v and perform the following *wedge operation*: Let H^+ and H^- be two (non-parallel) hyperplanes in \mathbb{R}^{d+1} such that each of them intersects the interior of $S_d^f \times \mathbb{R} \subseteq \mathbb{R}^{d+1}$ and $H^+ \cap H^- \supseteq F \times \{0\}$. For the two polyhedra W^\pm given by $W^\pm = (S_d^f \times \mathbb{R}) \cap H^\pm$, we then define $W_F(S_d^f) = \text{conv}(W^+ \cup W^-)$. Note that by construction, W^\pm are affinely equivalent embeddings of the spindle S_d^f into the two hyperplanes H^\pm . Hence, $W_F(S_d^f)$ is a polytope again. We call $W_F(S_d^f)$ a *wedge (on S_d^f) over the facet F* . See Figure 5.6 for an illustration of the wedge operation. Note that this operation can increase the circuit diameter by at most one [39].

The wedge $W_F(S_d^f)$ has $f + 1$ facets in dimension $d + 1$ and is almost a spindle: Each facet either contains the vertex $(v, 0)$ or the edge between u^+ and u^- , where u^\pm denotes the apex of W^\pm distinct from $(v, 0)$. To get a spindle from $W_F(S_d^f)$, we carefully perturb the facets of $W_F(S_d^f)$ that contain the edge between u^\pm so as to make an interior point of this edge become a vertex (the new apex; see Figure 5.6c). If the perturbation is done appropriately as described in [170], the resulting spindle S_{d+1}^{f+1} has length at least $l + 1$. In fact, by the proof of Theorem 2.6 in [170], carefully perturbing a *single* facet suffices to increase the length as desired.

If this wedge-plus-perturbation operation is iteratively applied $f - 2d$ times to S_d^f , we obtain an $(f - d)$ -dimensional spindle S_{f-d}^{2f-2d} with $2f - 2d$ facets and length at least $l + f - 2d$. So if $l > d$, then the length of S_{f-d}^{2f-2d} exceeds $f - d$, which means that the spindle S_{f-d}^{2f-2d} violates the Hirsch conjecture.




Figure 5.6: (a) The initial spindle S_d^f is degenerate (the orange line at u indicates an “extra” facet incident with u). (b) By wedging over a facet F that contains the other apex v , we obtain the wedge $W_F(S_d^f)$ with two facets W^\pm that are affinely equivalent with S_d^f . The adjacent vertices u^\pm now correspond to the apex u of S_d^f . (c) By perturbing the orange facet of the wedge, we get a spindle S_{d+1}^{f+1} with apices \tilde{u} and \tilde{v} .

In [155], Matschke, Santos, and Weibel explicitly built and computationally checked two Hirsch counterexamples resulting from S_5^{28} and S_5^{25} via Santos’ construction described above. The resulting spindles are of length 24 and 21 in dimension 23 and 20, respectively (see also [3] for a recent formal verification). The authors remark that carrying out the steps of the construction in such a way that the length indeed increases as desired was computationally feasible only for the two smaller spindles S_5^{28} and S_5^{25} and not for S_5^{48} (see also Santos’ remark in Section 1 of his paper [170]). For those two spindles, we verified computationally that our proof technique from Section 5.3.1 for bounding their circuit length also transfers to the explicit counterexamples themselves obtained by Matschke, Santos, and Weibel.

Using the inequality descriptions and vertex adjacencies provided in [155, 198], we found that slight perturbations of the 2-faces in Figures 5.4 and 5.5 still appear as 2-faces (with the same combinatorics) after the final wedge-plus-perturbation step. Furthermore, our computations show that on the final spindle, the length of a shortest edge walk from each apex to those 2-faces increases by exactly the number of times we

wedge over a facet that contains the apex. For instance, the 20-dimensional explicit counterexample from [155, Theorem 1.3] based on S_5^{25} has a 2-face that is a perturbed equivalent of the 2-face in Figure 5.5c. The original face of S_5^{25} could be reached within 3 edge steps from one apex, and the other apex was a vertex of the face already. Now, in dimension 20, the equivalent face can be reached within $3 + 8 = 11$ edge steps from one apex, and its vertex that was the other apex in dimension 5 now is at distance 7 from the new apex. The perturbations applied by Matschke, Santos, and Weibel are small enough for the properties (i) and (ii) in Corollary 5.16 to still hold. By our arguments in Section 5.3.1, this means that there is a circuit walk of length at most 2 on the perturbed 2-face that connects the two edge walks to and from the face to give a circuit walk of total length at most $11 + 2 + 7 = 20$. Also for the other 2-faces of S_5^{25} in Figure 5.5 and those of S_5^{28} in Figure 5.4, we verified computationally that the length 5 circuit walks via those faces can be extended in a completely analogous way to obtain circuit walks of the desired length in higher dimension. As an immediate consequence of these observations, we obtain the following corollaries.

Corollary 5.17. *The circuit length of the 20-dimensional spindle with 40 facets given in [155, Theorem 1.3] is at most 20.*

Corollary 5.18. *The circuit length of the 23-dimensional spindle with 46 facets given in [155, 198] is at most 23.*

We stress that these two explicit Hirsch counterexamples result from a particular sequence of wedge-plus-perturbation operations applied to S_5^{25} and S_5^{28} , respectively. However, the steps of Santos' construction are not uniquely determined: The choice of the facet to wedge over is arbitrary (as long as it contains the right apex), and so is the choice of the facet that is perturbed. Different choices may lead to different counterexamples. Nonetheless, our arguments from Section 5.3.1 allow us to make the following observation: Regardless of how the steps of Santos' construction applied to S_5^{28} or S_5^{25} are executed, the 2-faces that our circuit length bounds for the 5-dimensional spindles crucially relied on will be preserved up to slight changes.

To see this, consider the first wedge on S_5^{25} (or S_5^{28}) over an arbitrary facet (both apices are degenerate). Let us denote the two facets that are affinely equivalent with S_5^{25} by W^\pm with apices v and u^\pm , as in the sketch in Figure 5.6b. Thus, all 2-faces of S_5^{25} in Figure 5.5 also appear as 2-faces of W^\pm (up to an affine transformation). If we now perturb a facet according to Santos' construction (one of the facets that contains u^\pm), then one of the vertices u^\pm , say u^+ , must be cut off in order to get a spindle again (cf. Figure 5.6c). Note that the only degenerate vertices of S_5^{25} , and therefore of W^+ , are the apices (this can be verified computationally, e.g., using Polymake [110]). So by a slight perturbation of the chosen facet, a 2-face of W^+ that contains u^+ will either be slightly perturbed without changing the combinatorics, or it will become a 2-face where combinatorially the only change is that u^+ is replaced with two new, adjacent vertices (the edge between them must then be defined by the perturbed facet). Moreover, 2-faces of W^+ that do not contain u^+ are unaffected (up to slight perturbations) since the facet that we perturb contains u^+ . In either case, Lemma 5.9 guarantees that on the resulting 2-face, two circuit steps still suffice to reach a vertex that corresponds to the apex of S_5^{25} contained in the corresponding face of S_5^{25} .

The above observation also applies to S_5^{28} by noting that all vertices of S_5^{28} other than the apices are non-degenerate. However, this is not true for the vertices of S_5^{48} . Therefore, we cannot directly conclude that any wedge-plus-perturbation operation according to Santos' construction will preserve the 2-face in Figure 5.3 or its symmetric equivalents in the proofs of Theorem 5.11 and Corollary 5.12.

We conclude this chapter by providing some more context on the role that spindles and their circuit length may play for resolving the circuit diameter conjecture.

5.4 Further Notes on Circuit Lengths and Circuit Diameters

In the previous section, we proved that the circuit length of the spindles used by Matschke, Santos, and Weibel [155, 170] to build counterexamples to the Hirsch conjecture is at most their dimension. Our results suggest the following question:

Question 5.19. *Is the circuit length of every d -dimensional spindle at most d ?*

An answer to Question 5.19 is not known, not even for d -dimensional spindles with exactly $2d$ facets (see Conjecture 3.8 of [39]). These so-called *Dantzig figures* [145] are intersections of two d -dimensional simplicial cones. Recall that all Hirsch counterexamples obtained via Santos' construction from [170], in particular the ones from [155] that we analyzed in Section 5.3.2, are indeed Dantzig figures – and the fact that they *are* Hirsch counterexamples is precisely due to their large length. In fact, this is no coincidence: In the same paper [145] that contains their counterexample to the unbounded Hirsch conjecture, Klee and Walkup proved that a number of seemingly more specialized variants of the Hirsch conjecture all are equivalent to the original statement of the conjecture. More precisely, they consider the so-called *d -step conjecture*, which is the Hirsch conjecture specialized to d -dimensional polyhedra with $2d$ facets. The other equivalent conjecture from [145] asks whether the combinatorial *length* (not the *diameter*) of d -dimensional Dantzig figures is bounded by d .

Each of these two variants of the Hirsch conjecture naturally has a circuit counterpart, as considered by Borgwardt, Stephen, and Yusun [39]. They showed that the circuit diameter conjecture is equivalent to its d -step variant. Whether this is also true for the circuit version of the Dantzig figure variant of the Hirsch conjecture is not known [39]. More generally, it is open whether answering Question 5.19 in the positive is sufficient to prove the circuit diameter conjecture.

As for Question 5.19 itself, note that any spindle whose circuit length exceeds the dimension must be at least 5-dimensional, since spindles up to dimension $d \leq 4$ are known to have (combinatorial, and thus also circuit) length at most d [172]. In fact, by leveraging Lemma 5.8 once again, we may give a partial answer to Question 5.19 for a large class of spindles in any dimension. Recall that a spindle with apices u and v is of the form $(C + u) \cap (-D + v)$ for two pointed cones C and D . If one of these cones is contained in the other one, there is always a short circuit walk from the apex of the wider cone to that of the narrower one:

Theorem 5.20. *Let P be a d -dimensional spindle with apices u and v , given by $P = (C + u) \cap (-D + v)$ for two pointed cones C and D . If $D \subseteq C$, then there is a circuit walk of length at most d from u to v on P .*

Theorem 5.20 easily follows from an inductive application of Lemma 5.8. For the sake of the following proof, we extend the definition of a circuit walk given in Section 5.1 to arbitrary feasible starting points (not just vertices), as in [39].

Proof of Theorem 5.20. We will prove the following stronger statement.

Claim. *Let F be a face of P with $v \in F$. For any starting point $y^{(0)} \in F$, there is a circuit walk of length at most $\dim(F)$ from $y^{(0)}$ to v on F .*

This clearly implies the statement of the theorem for the choice $F = P$ and $y^{(0)} = u$. We prove the claim by induction on the dimension of F . If $F = \{v\}$, there is nothing to prove. So suppose that $\dim(F) \geq 1$, and let $C = \{x: Ax \leq \mathbf{0}\}$ and $D = \{x: Bx \leq \mathbf{0}\}$ be minimal linear descriptions of the cones C and D for two matrices A and B . Then F is obtained from the linear description of $P = \{x: A(x - u) \leq \mathbf{0}, B(x - v) \geq \mathbf{0}\}$ by changing some of the inequalities in $B(x - v) \geq \mathbf{0}$ to equations. Let B' be the maximal row submatrix of B such that $B'(x - v) = \mathbf{0}$ for all $x \in F$. We now define the polyhedron $F_v := \{x: A(x - u) \leq \mathbf{0}, B'(x - v) = \mathbf{0}\}$. Note that F_v is obtained from F by omitting all inequalities that are tight for v , since all those are facet-defining inequalities for the cone $-D + v$.

We claim that F_v is unbounded. Then, by Lemma 5.8, we know that for any $y^{(0)} \in F$, there is a facet $G \ni v$ of F and a point $y^{(1)} \in G$ such that $y^{(1)}$ can be reached from $y^{(0)}$ in at most one circuit step, using a circuit of F . By the induction hypothesis, there is a circuit walk of length at most $\dim(G) = \dim(F) - 1$ from $y^{(1)}$ to v on G . Together with the first step from $y^{(0)}$ to $y^{(1)}$, this yields a circuit walk of length at most $\dim(F)$ from $y^{(0)}$ to v on F as desired. Note that all steps follow circuit directions of F since circuits of G are also circuits of F .

It remains to show that F_v is indeed unbounded. Note that the feasible cone of F at v is given by $\{x \in -D: B'x = \mathbf{0}\}$. Since $D \subseteq C$ by hypothesis, it follows that

$$\{x \in D: B'x = \mathbf{0}\} \subseteq \{x \in C: B'x = \mathbf{0}\} = \{x: Ax \leq \mathbf{0}, B'x = \mathbf{0}\} = \text{rec}(F_v).$$

In other words, the recession cone of F_v contains the negative of the feasible cone of F at v , which is at least one-dimensional by the hypothesis that $\dim(F) \geq 1$. This concludes the proof of the claim. \square

A refined version of Theorem 5.20 plays a central role for bounding the lengths of *monotone* circuit walks in [21]. These are circuit walks for which each step is strictly increasing with respect to some given linear objective function. Note that edge walks traced by the Simplex method must be monotone. Motivated by this, a stronger version of the Hirsch conjecture asks whether for all linear objective functions c over a d -dimensional polyhedron with f facets, there is a monotone edge walk of length at most $f - d$ from any vertex to a vertex maximizer of c . This so-called *monotone Hirsch conjecture* is false: Todd [191] found a 4-dimensional polytope with 8 facets and a linear objective function such that 5 monotone steps are required to reach the

unique maximizer from a particular starting vertex. Interestingly, the Todd polytope is a spindle and those two vertices, the starting vertex and the maximizer, are its apices. One may refine Theorem 5.20 to prove that, in this case, there is a monotone *circuit* walk of length at most 4. With extra work, one can even show that short monotone circuit walks exist for all vertices and across all linear objective functions, proving that Todd's counterexample is not a counterexample to the monotone version of the circuit diameter conjecture. This is elaborated in [21].

Finally, let us remark that resolving the circuit diameter conjecture may shed some light on why the Hirsch conjecture is false [36, 39]. As explained in Section 5.1, every edge walk is a circuit walk. The only way in which circuit walks truly generalize edge walks is the larger set of directions that they are composed of. All other properties of edge walks (feasibility, maximality of step lengths) are maintained. This means that, if the circuit diameter conjecture were true, then it would be the more restrictive directions that force longer walks in the combinatorial setting. On the other hand, if not even the circuit diameter satisfied the Hirsch bound, the reason for this would be the maximality of steps: Recall from Section 5.2.2 that, by virtue of the sign-compatible representation property of circuits, one may always find *feasible* circuit walks that satisfy the Hirsch bound of $f - d$, between any pair of vertices of a given polyhedron, when non-maximal circuit steps are allowed.

Appendix A

Polyhedral Theory

Here, we briefly summarize some relevant results from polyhedral theory. An extensive treatment can be found, e.g., in [174, 206].

Polyhedra. A *polyhedron* in \mathbb{R}^n is a set of the form $P = \{x \in \mathbb{R}^n : Bx \leq d\}$ for some matrix $B \in \mathbb{R}^{m_B \times n}$ and some vector $d \in \mathbb{R}^{m_B}$. The linear system $Bx \leq d$ is a *linear description* of P . Inequalities from $Bx \leq d$ that are *tight* (i.e., satisfied at equality) for all $x \in P$ are called *implicit equations*. We may therefore represent any polyhedron $P \subseteq \mathbb{R}^n$ by a system of the form $Ax = b, Bx \leq d$ for $A \in \mathbb{R}^{m_A \times n}, B \in \mathbb{R}^{m_B \times n}$ and $b \in \mathbb{R}^{m_A}, d \in \mathbb{R}^{m_B}$, where $Bx \leq d$ contains no implicit equations. If no constraint in $Ax = b, Bx \leq d$ can be removed from the system without changing its set of solutions, we call this system a *minimal* (or *irredundant*) linear description of P . Every polyhedron has a minimal linear description.

The *affine hull* of a polyhedron $P \subseteq \mathbb{R}^n$, denoted by $\text{aff}(P)$, is the intersection of all affine subspaces of \mathbb{R}^n containing P . We call its dimension the *dimension* of P , denoted by $\dim(P)$ (where $\dim(\emptyset) = -1$ by convention). If P is given by a minimal linear description $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$, then $\text{aff}(P) = \{x \in \mathbb{R}^n : Ax = b\}$ and $\dim(P) = n - \text{rk}(A)$.

The *relative interior* of a polyhedron P is the interior of P with respect to its affine hull $\text{aff}(P)$.

Faces. Let $P \subseteq \mathbb{R}^n$ be a polyhedron. A linear inequality $a^\top x \leq \beta$ with $a \in \mathbb{R}^n, \beta \in \mathbb{R}$ is *valid* for P if P is contained in the halfspace $\{x \in \mathbb{R}^n : a^\top x \leq \beta\}$. The intersection of P with the corresponding hyperplane $\{x \in \mathbb{R}^n : a^\top x = \beta\}$ is a *face* of P . Each face of P is a polyhedron again. If P is pointed, then the minimal faces are those of dimension zero; they are called *vertices*. One-dimensional faces are called *edges* and the faces of dimension $\dim(P) - 1$ are the *facets* of P . The polyhedron P itself is a face, too; all other faces are called *proper*. Every proper face is the intersection of all facets containing it. k -dimensional faces are often simply called *k-faces*. Two vertices u and v are *adjacent* if $\text{conv}\{u, v\}$ is an edge.

Every vertex of a d -dimensional polyhedron is contained in at least d facets. A vertex is called *non-degenerate* if it is contained in the minimum number of d facets, and *degenerate* otherwise.

Recession Cone and Lineality Space. A *cone* is a set $C \subseteq \mathbb{R}^n$ such that $\lambda x + \mu y \in C$ for all $x, y \in C$ and all $\lambda, \mu \in \mathbb{R}_{\geq 0}$. Such a cone C is *generated* by some set $R \subseteq \mathbb{R}^n$ if all vectors in C can be expressed as finite linear combinations of vectors in R with

nonnegative coefficients (so-called *nonnegative linear combinations*). If R is a finite set, C is said to be *finitely generated*. A cone that is generated by linearly independent vectors is called *simplicial*.

Cones of the form $C = \{x \in \mathbb{R}^n : Ax \leq \mathbf{0}\}$ for some matrix A with n columns are called *polyhedral cones*. All cones encountered in this thesis are polyhedral. A *ray* in \mathbb{R}^n is a set of the form $\{x + \mu g : \mu \in \mathbb{R}_{\geq 0}\}$ for some nonzero vector $g \in \mathbb{R}^n$ (the *direction* of the ray) and $x \in \mathbb{R}^n$. The *recession cone* of a polyhedron $P \subseteq \mathbb{R}^n$, denoted by $\text{rec}(P)$, is the cone consisting of all directions of rays contained in P , i.e.,

$$\text{rec}(P) := \{g \in \mathbb{R}^n : x + \mu g \in P \text{ for all } x \in P \text{ and all } \mu \in \mathbb{R}_{\geq 0}\}.$$

The *lineality space* of P , denoted by $\text{ls}(P)$, is the linear subspace of \mathbb{R}^n consisting of all directions of lines contained in P , i.e.,

$$\text{ls}(P) := \{g \in \mathbb{R}^n : x + \mu g \in P \text{ for all } x \in P \text{ and all } \mu \in \mathbb{R}\}.$$

If the polyhedron P is given by $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$, then $\text{rec}(P) = \{x : Ax = \mathbf{0}, Bx \leq \mathbf{0}\}$ and $\text{ls}(P)$ is the kernel of the matrix $\begin{pmatrix} A \\ B \end{pmatrix}$. In particular, $\text{rec}(P)$ is a polyhedral cone. P is said to be *pointed* if it contains no line, i.e., if $\text{ls}(P) = \{\mathbf{0}\}$. Equivalently, P is pointed if and only if $\begin{pmatrix} A \\ B \end{pmatrix}$ has full column rank. Clearly, if P is pointed, so is $\text{rec}(P)$.

Pointed polyhedral cones have a unique vertex at the origin $\mathbf{0}$. Their one-dimensional faces are called *extreme rays*. For a pointed polyhedron P , the extreme rays of $\text{rec}(P)$ are also referred to as the extreme rays of P .

Given a vertex v of a pointed polyhedron $P = \{x \in \mathbb{R}^n : Ax = b, Bx \leq d\}$, let $B'x \leq d'$ be the subsystem of $Bx \leq d$ consisting of all inequalities that are tight for v . Then the cone $\{x \in \mathbb{R}^n : Ax = \mathbf{0}, B'x \leq \mathbf{0}\}$ is called the *feasible cone of P at v* . Its extreme rays are generated by the directions of all edges of P that are incident with v . All feasible cones of P are of the same dimension, namely $\dim(P)$.

Polytopes. A *polytope* in \mathbb{R}^n is the convex hull of a finite set of vectors in \mathbb{R}^n . If these vectors are in $\{0, 1\}^n$, the polytope is a *0/1 polytope*. A special case is the *standard hypercube* (or simply *0/1 cube*) given by $\text{conv}(\{0, 1\}^n) = [0, 1]^n$. Two-dimensional polytopes are called *polygons*. The convex hull of affinely independent vectors is a *simplex*. A simplex with $n + 1$ vertices is sometimes called an n -simplex. Hypercubes, simplices, and polygons are *simple* polytopes, i.e., all their vertices are non-degenerate.

The *Minkowski sum* of two sets $X, Y \subseteq \mathbb{R}^n$ is defined as $X + Y := \{x + y : x \in X, y \in Y\}$. The theorem of Minkowski-Weyl states that every polyhedron P is the Minkowski sum of a polytope and a finitely generated cone. In particular, any pointed polyhedron P can be expressed as $P = \text{conv}(\mathcal{V}(P)) + \text{rec}(P)$, where $\mathcal{V}(P)$ denotes the set of vertices of P . This implies that polytopes are exactly the bounded polyhedra.

Two polytopes P and Q are *combinatorially equivalent* if there is a bijection between their vertices that preserves the vertex sets of facets. So we may collect polytopes into equivalence classes (*combinatorial types*), which are sometimes referred to as *combinatorial polytopes*. A *realization* of a combinatorial polytope is a particular embedding with concrete vertex coordinates in some space. In this thesis, a realization of a (geometric, not combinatorial) polytope P simply refers to any other polytope Q which is combinatorially equivalent to P .

Polarity. Given a set $X \subseteq \mathbb{R}^n$, the *polar* (or *polar dual*) of X is the set $X^\circ := \{y \in \mathbb{R}^n : x^\top y \leq 1 \text{ for all } x \in X\}$. If P is a polytope which contains the origin $\mathbf{0}$ in its interior, then P° is a polytope again and $(P^\circ)^\circ = P$. Moreover, suppose that P is given by $P = \{x \in \mathbb{R}^n : Bx \leq d\}$, then P° is the convex hull of all row vectors of B . So the vertices and facets of P are in bijection with the facets and vertices of P° . For example, the polar of an n -simplex is an n -simplex again.

Polarity for cones plays a central role in LP duality. Given a cone $C \subseteq \mathbb{R}^n$, its polar (the *polar cone* of C) is of the form $C^\circ = \{y \in \mathbb{R}^n : x^\top y \leq 0 \text{ for all } x \in C\}$. If C is a polyhedral cone of the form $C = \{x \in \mathbb{R}^n : Ax \leq \mathbf{0}\}$, then the polar cone of C is the cone generated by the rows of A . For a polyhedron P and any given linear objective function c , a vertex v of P is a maximizer of c over P if and only if c is in the polar cone of the feasible cone of P at v (also called the *normal cone of P at v*). Moreover, v is the unique maximizer if c is in the relative interior of the normal cone.

Integral Polyhedra. A polyhedron $P \subseteq \mathbb{R}^n$ is said to be *integral* if $P = \text{conv}(P \cap \mathbb{Z}^n)$. Equivalently, P is integral if and only if all nonempty faces of P contain integral points. In particular, integral *polytopes* are precisely those for which all vertices are integral. Another sufficient and necessary condition for a polyhedron P to be integral, which is used in Chapter 2, is that $\max\{c^\top x : x \in P\} \in \mathbb{Z}$ for all integral vectors c for which the maximum is finite.

References

- [1] A. Abdi. *Ideal clutters*. PhD thesis, University of Waterloo, 2018. (Cited on pages [22](#), [23](#), [27](#), [34](#), [39](#), [51](#).)
- [2] A. Abdi and B. Guenin. The minimally non-ideal binary clutters with a triangle. *Combinatorica*, 39(4):719–752, 2019. (Cited on page [51](#).)
- [3] X. Allamigeon, Q. Canu, and P.-Y. Strub. A formal disproof of Hirsch conjecture. In *Proceedings of the 12th ACM SIGPLAN International Conference on Certified Programs and Proofs. CPP 2023*, page 17–29, 2023. (Cited on page [115](#).)
- [4] G. Appa and B. Kotnyek. Rational and integral k -regular matrices. *Discrete Mathematics*, 275(1-3):1–15, 2004. (Cited on page [37](#).)
- [5] G. Appa, B. Kotnyek, K. Papalamprou, and L. Pitsoulis. Optimization with binet matrices. *Operations Research Letters*, 35(3):345–352, 2007. (Cited on page [51](#).)
- [6] M. Aprile. Extended formulations for matroid polytopes through randomized protocols. *Operations Research Letters*, 50(2):145–149, 2022. (Cited on page [56](#).)
- [7] M. Aprile and S. Fiorini. Regular matroids have polynomial extension complexity. *Mathematics of Operations Research*, 47(1):540–559, 2022. (Cited on page [71](#).)
- [8] M. Aprile, S. Fiorini, T. Huynh, G. Joret, and D. R. Wood. Smaller extended formulations for spanning tree polytopes in minor-closed classes and beyond. *arXiv preprint*, arXiv:2106.11945, 2021. (Cited on pages [7](#), [56](#), [57](#).)
- [9] A. Bachem and W. Kern. *Linear Programming Duality: An Introduction to Oriented Matroids*. Universitext. Springer, 1984. (Cited on pages [73](#), [98](#), [99](#), [100](#).)
- [10] E. Balas. Disjunctive programming and a hierarchy of relaxations for discrete optimization problems. *SIAM Journal on Algebraic Discrete Methods*, 6(3):466–486, 1985. (Cited on page [88](#).)
- [11] E. Balas. Disjunctive programming: Properties of the convex hull of feasible points. *Discrete Applied Mathematics*, 89(1):3–44, 1998. (Cited on page [88](#).)
- [12] E. Balas and W. Pulleyblank. The perfectly matchable subgraph polytope of a bipartite graph. *Networks*, 13(4):495–516, 1983. (Cited on page [77](#).)
- [13] M. L. Balinski. The Hirsch conjecture for dual transportation polyhedra. *Mathematics of Operations Research*, 9(4):629–633, 1984. (Cited on page [95](#).)
- [14] M. L. Balinski and A. Russakoff. On the assignment polytope. *SIAM Review*, 16(4):516–525, 1974. (Cited on page [86](#).)

-
- [15] F. Barahona. Reducing matching to polynomial size linear programming. *SIAM Journal on Optimization*, 3(4):688–695, 1993. (Cited on page 6.)
- [16] 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. (Cited on pages 71, 85.)
- [17] D. Barnette. An upper bound for the diameter of a polytope. *Discrete Mathematics*, 10(1):9–13, 1974. (Cited on page 95.)
- [18] D. Bienstock and M. Zuckerberg. Subset algebra lift operators for 0-1 integer programming. *SIAM Journal on Optimization*, 15(1):63–95, 2004. (Cited on page 50.)
- [19] A. Björner, M. L. Vergnas, B. Sturmfels, N. White, and G. Ziegler. *Oriented Matroids*. Cambridge University Press, 1999. (Cited on pages 73, 98, 99.)
- [20] A. E. Black, J. A. De Loera, S. Kafer, and L. Sanità. On the Simplex method for 0/1 polytopes. *arXiv preprint*, arXiv:2111.14050, 2021. (Cited on page 72.)
- [21] A. E. Black, S. Borgwardt, and M. Brugger. On the circuit diameter conjecture for counterexamples to the Hirsch conjecture. *arXiv preprint*, arXiv:2302.03977, 2023. (Cited on pages 10, 95, 118, 119.)
- [22] R. G. Bland. New finite pivoting rules for the Simplex method. *Mathematics of Operations Research*, 2(2):103–107, 1977. (Cited on page 73.)
- [23] B. Bollobás. *Graph Theory: An Introductory Course*, volume 63 of *Graduate Texts in Mathematics*. Springer, 2012. (Cited on page 27.)
- [24] N. Bonifas, M. Di Summa, F. Eisenbrand, N. Hähnle, and M. Niemeier. On subdeterminants and the diameter of polyhedra. *Discrete & Computational Geometry*, 52(1):102–115, 2014. (Cited on page 95.)
- [25] K. H. Borgwardt. *The Simplex Method: A Probabilistic Analysis*, volume 1 of *Algorithms and Combinatorics*. Springer Berlin, Heidelberg, 1987. (Cited on page 72.)
- [26] S. Borgwardt. *A combinatorial optimization approach to constrained clustering*. PhD thesis, Technische Universität München, 2010. (Cited on pages 71, 85, 86.)
- [27] S. Borgwardt. On the diameter of partition polytopes and vertex-disjoint cycle cover. *Mathematical Programming, Series A*, 141(1):1–20, 2013. (Cited on pages 85, 86.)
- [28] S. Borgwardt and M. Brugger. Circuits in extended formulations. *arXiv preprint*, arXiv:2208.05467, 2022. (Cited on pages 10, 71.)
- [29] S. Borgwardt and F. Happach. Good clusterings have large volume. *Operations Research*, 67(1):215–231, 2019. (Cited on pages 71, 85, 86, 93.)

-
- [30] S. Borgwardt and S. Onn. Efficient solutions for weight-balanced partitioning problems. *Discrete Optimization*, 21(C):71–84, 2016. (Cited on page 85.)
- [31] S. Borgwardt and C. Viss. An implementation of steepest-descent augmentation for linear programs. *Operations Research Letters*, 48(3):323–328, 2020. (Cited on page 8.)
- [32] S. Borgwardt and C. Viss. Constructing clustering transformations. *SIAM Journal on Discrete Mathematics*, 35(1):152–178, 2021. (Cited on page 86.)
- [33] S. Borgwardt and C. Viss. Circuit walks in integral polyhedra. *Discrete Optimization*, 44(1):100566, 2022. (Cited on pages 73, 86, 87, 98, 99.)
- [34] S. Borgwardt and C. Viss. A polyhedral model for enumeration and optimization over the set of circuits. *Discrete Applied Mathematics*, 308:68–83, 2022. (Cited on pages 8, 73, 74, 77.)
- [35] S. Borgwardt, E. Finhold, and R. Hemmecke. On the circuit diameter of dual transportation polyhedra. *SIAM Journal on Discrete Mathematics*, 29(1):113–121, 2015. (Cited on pages 9, 73, 74, 96, 97, 98.)
- [36] S. Borgwardt, J. A. De Loera, and E. Finhold. Edges vs circuits: A hierarchy of diameters in polyhedra. *Advances in Geometry*, 16(4):511–530, 2016. (Cited on pages 73, 96, 100, 101, 102, 119.)
- [37] S. Borgwardt, J. A. De Loera, and E. Finhold. The diameters of network-flow polytopes satisfy the Hirsch conjecture. *Mathematical Programming*, 171:283–309, 2018. (Cited on page 95.)
- [38] 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. (Cited on pages 85, 96.)
- [39] S. Borgwardt, T. Stephen, and T. Yusun. On the circuit diameter conjecture. *Discrete & Computational Geometry*, 60(3):558–587, 2018. (Cited on pages 78, 97, 114, 117, 118, 119.)
- [40] 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. (Cited on page 93.)
- [41] S. Boyd and W. R. Pulleyblank. Facet generating techniques. In W. Cook, L. Lovász, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, pages 33–55. Springer, 2009. (Cited on pages 4, 15.)
- [42] L. Brandl and A. S. Schulz. Primal separation and approximation for the $\{0, 1/2\}$ -closure. *arXiv preprint*, arXiv:2311.03909, 2023. (Cited on page 50.)
- [43] A. Brieden and P. Gritzmann. On optimal weighted balanced clusterings: Gravity bodies and power diagrams. *SIAM Journal on Discrete Mathematics*, 26:415–434, 2012. (Cited on pages 71, 85.)

-
- [44] M. Brugger. Limitations of the hyperplane separation technique for bounding the extension complexity of polytopes. *Operations Research Letters*, 49(6):896–901, 2021. (Cited on pages 10, 55.)
- [45] M. Brugger and A. S. Schulz. On the complexity of recognizing integrality and total dual integrality of the $\{0, 1/2\}$ -closure. *Operations Research Letters*, 50(1):16–19, 2022. (Cited on pages 10, 13.)
- [46] H. Bruhn and O. Schaudt. Claw-free t -perfect graphs can be recognized in polynomial time. *SIAM Journal on Discrete Mathematics*, 30(2):832–855, 2016. (Cited on page 49.)
- [47] H. Bruhn and M. Stein. t -perfection is always strong for claw-free graphs. *SIAM Journal on Discrete Mathematics*, 24(3):770–781, 2010. (Cited on page 50.)
- [48] T. Brunsch and H. Röglin. Finding short paths on polytopes by the shadow vertex algorithm. In F. Fomin, R. Freivalds, M. Kwiatkowska, and D. Peleg, editors, *Automata, Languages, and Programming. ICALP 2013*, volume 7965 of *Lecture Notes in Computer Science*, pages 279–290. Springer, 2013. (Cited on pages 72, 95.)
- [49] A. Caprara and M. Fischetti. $\{0, 1/2\}$ -Chvátal-Gomory cuts. *Mathematical Programming*, 74:221–235, 1996. (Cited on pages 5, 6, 9, 13, 15, 16, 22, 24, 25, 26, 27, 28, 52.)
- [50] A. Caprara, M. Fischetti, and A. N. Letchford. On the separation of maximally violated mod- k cuts. *Mathematical Programming*, 87:37–56, 2000. (Cited on page 16.)
- [51] R. D. Carr and G. Konjevod. Polyhedral combinatorics. In H. J. Greenberg, editor, *Tutorials on Emerging Methodologies and Applications in Operations Research*, pages 2–1–2–46. Springer, 2005. (Cited on page 71.)
- [52] K. Chandrasekaran, L. A. Végh, and S. S. Vempala. The cutting plane method is polynomial for perfect matchings. *Mathematics of Operations Research*, 41(1):23–48, 2016. (Cited on page 6.)
- [53] V. Chvátal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics*, 4(4):305–337, 1973. (Cited on pages 3, 4, 13, 14.)
- [54] V. Chvátal. On certain polytopes associated with graphs. *Journal of Combinatorial Theory, Series B*, 18(2):138–154, 1975. (Cited on pages 7, 14.)
- [55] J. E. Cohen and U. G. Rothblum. Nonnegative ranks, decompositions, and factorizations of nonnegative matrices. *Linear Algebra and its Applications*, 190:149–168, 1993. (Cited on pages 58, 67, 69.)
- [56] M. Conforti, G. Cornuéjols, and G. Zambelli. Extended formulations in combinatorial optimization. *Annals of Operations Research*, 204(1):97–143, 2013. (Cited on pages 62, 71.)

-
- [57] M. Conforti, G. Cornuéjols, and G. Zambelli. *Integer Programming*, volume 271 of *Graduate Texts in Mathematics*. Springer, 2014. (Cited on pages [10](#), [13](#), [58](#), [71](#), [77](#).)
- [58] M. Conforti, V. Kaibel, M. Walter, and S. Weltge. Subgraph polytopes and independence polytopes of count matroids. *Operations Research Letters*, 43(5):457–460, 2015. (Cited on page [56](#).)
- [59] W. Cook. Operations that preserve total dual integrality. *Operations Research Letters*, 2(1):31–35, 1983. (Cited on page [34](#).)
- [60] G. Cornuéjols. *Combinatorial Optimization: Packing and Covering*. SIAM, Philadelphia, 2001. (Cited on pages [16](#), [17](#), [22](#), [26](#), [34](#), [36](#), [38](#), [44](#), [45](#), [47](#), [49](#), [51](#).)
- [61] G. Cornuéjols and B. Guenin. Ideal binary clutters, connectivity, and a conjecture of Seymour. *SIAM Journal on Discrete Mathematics*, 15(3):329–352, 2002. (Cited on page [51](#).)
- [62] G. Cornuéjols and Y. Li. When the Gomory-Chvátal closure coincides with the integer hull. *Operations Research Letters*, 46(2):251–256, 2018. (Cited on pages [5](#), [18](#).)
- [63] G. Cornuéjols and B. Novick. Ideal 0,1 matrices. *Journal of Combinatorial Theory, Series B*, 60(1):145–157, 1994. (Cited on page [32](#).)
- [64] G. Cornuéjols, D. Lee, and Y. Li. On the rational polytopes with Chvátal rank 1. *Mathematical Programming*, 179:21–46, 2020. (Cited on pages [5](#), [18](#).)
- [65] F. Criado and F. Santos. Topological prisms and small simplicial spheres of large diameter. *Experimental Mathematics*, 31(2):461–473, 2022. (Cited on page [96](#).)
- [66] W. H. Cunningham and A. B. Marsh. A primal algorithm for optimum matching. In M. L. Balinski and A. J. Hoffman, editors, *Polyhedral Combinatorics*, pages 50–72. Springer, 1978. (Cited on pages [15](#), [44](#).)
- [67] D. Dadush and N. Hähnle. On the shadow simplex method for curved polyhedra. *Discrete & Computational Geometry*, 56:882–909, 2016. (Cited on pages [72](#), [95](#).)
- [68] D. Dadush and S. Huiberts. A friendly smoothed analysis of the simplex method. *SIAM Journal on Computing*, 49(5):STOC18–449, 2020. (Cited on page [72](#).)
- [69] D. Dadush, Z. K. Koh, B. Natura, and L. A. Végh. On circuit diameter bounds via circuit imbalances. In *Integer Programming and Combinatorial Optimization*, pages 140–153, Cham, 2022. Springer International Publishing. (Cited on pages [73](#), [96](#).)
- [70] G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, 1963. (Cited on pages [8](#), [9](#), [71](#), [95](#).)

-
- [71] J. A. De Loera and E. D. Kim. *Combinatorics and geometry of transportation polytopes: An update*, volume 625 of *Contemporary Mathematics*. American Mathematical Society, 2014. (Cited on page 85.)
- [72] 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. (Cited on pages 8, 71, 72, 73, 74.)
- [73] J. A. De Loera, S. Kafer, and L. Sanità. Pivot rules for circuit-augmentation algorithms in linear optimization. *SIAM Journal on Optimization*, 32(3):2156–2179, 2022. (Cited on pages 8, 73, 74, 78, 96.)
- [74] A. Del Pia, A. Musitelli, and G. Zambelli. On matrices with the Edmonds-Johnson property arising from bidirected graphs. *Journal of Combinatorial Theory, Series B*, 130:49–91, 2018. (Cited on pages 39, 51, 132.)
- [75] A. Deza, T. Terlaky, and Y. Zinchenko. A continuous d -step conjecture for polytopes. *Discrete & Computational Geometry*, 41(2):318–327, 2009. (Cited on page 96.)
- [76] G. Ding, L. Feng, and W. Zang. The complexity of recognizing linear systems with certain integrality properties. *Mathematical Programming*, 114(2):321–334, 2008. (Cited on page 50.)
- [77] E. A. Dinic. Algorithm for solution of a problem of maximum flow in a network with power estimation. *Soviet Mathematics – Doklady*, 11:1277–1280, 1970. (Cited on page 7.)
- [78] M. Dyer and A. Frieze. Random walks, totally unimodular matrices, and a randomised dual simplex algorithm. *Mathematical Programming*, 64(1-3):1–16, 1994. (Cited on page 95.)
- [79] J. Edmonds. Maximum matching and a polyhedron with 0, 1-vertices. *Journal of research of the National Bureau of Standards B*, 69:125–130, 1965. (Cited on pages 4, 5, 14, 28, 35.)
- [80] J. Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965. (Cited on page 5.)
- [81] J. Edmonds. Submodular functions, matroids, and certain polyhedra. In R. Guy, H. Hanani, N. Sauer, and J. Schönheim, editors, *Combinatorial Structures and Their Applications*, pages 69–87. Gordon and Breach, New York, 1970. (Cited on page 64.)
- [82] J. Edmonds. Matroids and the greedy algorithm. *Mathematical Programming*, 1(1):127–136, 1971. (Cited on pages 6, 7, 55.)
- [83] J. Edmonds and D. R. Fulkerson. Bottleneck extrema. *Journal of Combinatorial Theory*, 8(3):299–306, 1970. (Cited on page 23.)

-
- [84] J. Edmonds and R. Giles. A min-max relation for submodular functions on graphs. *Annals of Discrete Mathematics*, 1:185–204, 1977. (Cited on page 21.)
- [85] J. Edmonds and E. L. Johnson. Matching: A well-solved class of integer linear programs. In *Combinatorial Structures and their Applications*, pages 89–92. Gordon and Breach, New York, 1970. (Cited on page 37.)
- [86] J. Edmonds and E. L. Johnson. Matching, Euler tours and the Chinese postman. *Mathematical Programming*, 5:88–124, 1973. (Cited on pages 35, 37, 51, 52.)
- [87] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM*, 19(2):248–264, 1972. (Cited on page 7.)
- [88] F. Eisenbrand. On the membership problem for the elementary closure of a polyhedron. *Combinatorica*, 19(2):297–300, 1999. (Cited on pages 15, 25.)
- [89] F. Eisenbrand. *Gomory-Chvátal cutting planes and the elementary closure of polyhedra*. PhD thesis, Universität des Saarlandes, 2000. (Cited on page 31.)
- [90] F. Eisenbrand and A. S. Schulz. Bounds on the Chvátal rank of polytopes in the 0/1-cube. *Combinatorica*, 23:245–261, 2003. (Cited on page 4.)
- [91] F. Eisenbrand and S. Vempala. Geometric random edge. *Mathematical Programming*, 164:325–339, 2017. (Cited on page 95.)
- [92] F. Eisenbrand, G. Rinaldi, and P. Ventura. Primal separation for 0/1 polytopes. *Mathematical Programming*, 95:475–491, 2003. (Cited on page 6.)
- [93] F. Eisenbrand, N. Hähnle, A. Razborov, and T. Rothvoß. Diameter of polyhedra: Limits of abstraction. *Mathematics of Operations Research*, 35(4):786–794, 2010. (Cited on page 96.)
- [94] F. Ekbatalani, B. Natura, and L. A. Végh. Circuit imbalance measures and linear programming. *Surveys in Combinatorics*, pages 64–114, 2022. (Cited on pages 72, 73, 96.)
- [95] Y. Faenza, S. Fiorini, R. Grappe, and H. R. Tiwary. Extended formulations, non-negative factorizations, and randomized communication protocols. *Mathematical Programming*, 153:75–94, 2015. (Cited on page 56.)
- [96] H. Fawzi and P. A. Parrilo. Lower bounds on nonnegative rank via nonnegative nuclear norms. *Mathematical Programming*, 153:41–66, 2015. (Cited on pages 57, 67.)
- [97] H. Fawzi and P. A. Parrilo. Self-scaled bounds for atomic cone ranks: applications to nonnegative rank and cp-rank. *Mathematical Programming*, 158:417–465, 2016. (Cited on pages 57, 67.)
- [98] E. Finhold. *Circuit diameters and their application to transportation problems*. PhD thesis, Technische Universität München, 2015. (Cited on pages 8, 71, 73, 74.)

-
- [99] S. Fiorini. Personal communication in [167], 2013. (Cited on pages 7, 57, 59.)
- [100] 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. (Cited on pages 7, 56, 59, 75, 76, 82.)
- [101] 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):17, 2015. (Cited on pages 6, 56, 59.)
- [102] S. Fiorini, T. Huynh, G. Joret, and K. Pashkovich. Smaller extended formulations for the spanning tree polytope of bounded-genus graphs. *Discrete & Computational Geometry*, 57:757–761, 2017. (Cited on page 56.)
- [103] S. Fiorini, T. Huynh, and S. Weltge. Strengthening convex relaxations of 0/1-sets using boolean formulas. *Mathematical Programming*, 190(1-2):467–482, 2021. (Cited on page 50.)
- [104] S. Fujishige. *Submodular Functions and Optimization*, volume 58 of *Annals of Discrete Mathematics*. Elsevier, 2nd edition, 2005. (Cited on pages 63, 64.)
- [105] K. Fukuda and T. Terlaky. Criss-cross methods: A fresh view on pivot algorithms. *Mathematical Programming*, 79:369–395, 1997. (Cited on page 73.)
- [106] K. Fukuda, S. Onn, and V. Rosta. An adaptive algorithm for vector partitioning. *Journal of Global Optimization*, 25:305–319, 2003. (Cited on pages 85, 86.)
- [107] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979. (Cited on pages 30, 52.)
- [108] G. Gastou and E. L. Johnson. Binary group and Chinese postman polyhedra. *Mathematical Programming*, 34:1–33, 1986. (Cited on page 52.)
- [109] J. B. Gauthier, J. Desrosiers, and M. Lübbecke. Decomposition theorems for linear programs. *Operations Research Letters*, 42(8):553–557, 2014. (Cited on page 73.)
- [110] E. Gawrilow and M. Joswig. `polymake`: A framework for analyzing convex polytopes. In *Polytopes – Combinatorics and Computation*, volume 29 of *DMV Seminar*, pages 43–73. Birkhäuser, Basel, 2000. (Cited on pages 105, 112, 116.)
- [111] J. F. Geelen and B. Guenin. Packing odd circuits in Eulerian graphs. *Journal of Combinatorial Theory, Series B*, 86(2):280–295, 2002. (Cited on pages 16, 25, 36, 52.)
- [112] A. M. H. Gerards. A min-max relation for stable sets in graphs with no odd- K_4 . *Journal of Combinatorial Theory, Series B*, 47(3):330–348, 1989. (Cited on page 46.)
- [113] A. M. H. Gerards. Personal communication in [74], 2007. (Cited on page 51.)

-
- [114] A. M. H. Gerards and A. Schrijver. Matrices with the Edmonds-Johnson property. *Combinatorica*, 6(4):365–379, 1986. (Cited on pages 5, 14, 15, 16, 17, 29, 36, 37, 39, 40, 49, 51, 52, 71.)
- [115] A. M. H. Gerards and F. B. Shepherd. The graphs with all subgraphs t -perfect. *SIAM Journal on Discrete Mathematics*, 11(4):524–545, 1998. (Cited on pages 36, 49.)
- [116] M. X. Goemans. Smallest compact formulation for the permutahedron. *Mathematical Programming*, 153(1):5–11, 2015. (Cited on page 57.)
- [117] A. V. Goldberg and R. E. Tarjan. Finding minimum-cost circulations by canceling negative cycles. *Journal of the ACM*, 36(4):873–886, 1989. (Cited on page 7.)
- [118] R. E. Gomory. An algorithm for integer solutions to linear programs. In R. L. Graves and P. Wolfe, editors, *Recent Advances in Mathematical Programming*, pages 269–302. McGraw-Hill, 1963. (Cited on pages 3, 13.)
- [119] J. Gouveia, R. Grappe, V. Kaibel, K. Pashkovich, R. Z. Robinson, and R. R. Thomas. Which nonnegative matrices are slack matrices? *Linear Algebra and its Applications*, 439(10):2921–2933, 2013. (Cited on page 59.)
- [120] J. E. Graver. On the foundations of linear and integer programming I. *Mathematical Programming*, 9:207–226, 1975. (Cited on pages 8, 71, 73, 74, 100.)
- [121] D. A. Gregory and N. J. Pullman. Semiring rank: Boolean rank and nonnegative rank factorizations. *Journal of Combinatorics, Information and System Sciences*, 8(3):223–233, 1983. (Cited on pages 58, 59.)
- [122] M. Grötschel and W. R. Pulleyblank. Weakly bipartite graphs and the max-cut problem. *Operations Research Letters*, 1(1):23–27, 1981. (Cited on pages 15, 16, 25, 29, 49, 52, 53, 71.)
- [123] M. Grötschel and K. Truemper. Decomposition and optimization over cycles in binary matroids. *Journal of Combinatorial Theory, Series B*, 46(3):306–337, 1989. (Cited on page 16.)
- [124] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1:169–197, 1981. (Cited on pages 5, 15, 52, 71.)
- [125] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer, 1988. (Cited on pages 25, 49, 52.)
- [126] B. Guenin. A characterization of weakly bipartite graphs. *Journal of Combinatorial Theory, Series B*, 83(1):112–168, 2001. (Cited on pages 16, 17, 25, 36, 37, 39, 51, 52.)
- [127] R. Hemmecke, S. Onn, and R. Weismantel. A polynomial oracle-time algorithm for convex integer minimization. *Mathematical Programming*, 126:97–117, 2011. (Cited on pages 8, 73.)

- [128] D. S. Hochbaum, N. Megiddo, J. Naor, and A. Tamir. Tight bounds and 2-approximation algorithms for integer programs with two variables per inequality. *Mathematical Programming*, 62(1-3):69–83, 1993. (Cited on page 37.)
- [129] A. J. Hoffman and J. B. Kruskal. Integral boundary points of convex polyhedra. In H. W. Kuhn and A. W. Tucker, editors, *Linear Inequalities and Related Systems*, pages 223–246. Princeton University Press, 1956. (Cited on page 48.)
- [130] 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. (Cited on page 85.)
- [131] 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. (Cited on pages 71, 85, 86.)
- [132] S. Kafer. *Polyhedral diameters and applications to optimization*. PhD thesis, University of Waterloo, 2022. (Cited on pages 8, 72, 73, 74, 96.)
- [133] S. Kafer, K. Pashkovich, and L. Sanità. On the circuit diameter of some combinatorial polytopes. *SIAM Journal on Discrete Mathematics*, 33(1):1–25, 2019. (Cited on pages 71, 73, 74, 96.)
- [134] V. Kaibel. Extended formulations in combinatorial optimization. *Optima*, 85, 2011. (Cited on pages 55, 71.)
- [135] V. Kaibel and K. Pashkovich. Constructing extended formulations from reflection relations. In M. Jünger and G. Reinelt, editors, *Facets of Combinatorial Optimization*, pages 77–100. Springer, 2013. (Cited on page 57.)
- [136] V. Kaibel and S. Weltge. A short proof that the extension complexity of the correlation polytope grows exponentially. *Discrete & Computational Geometry*, 53(2):397–401, 2015. (Cited on pages 6, 56, 59.)
- [137] V. Kaibel, K. Pashkovich, and D. O. Theis. Symmetry matters for sizes of extended formulations. *SIAM Journal on Discrete Mathematics*, 26(3):1361–1382, 2012. (Cited on page 56.)
- [138] G. Kalai and D. J. Kleitman. A quasi-polynomial bound for the diameter of graphs of polyhedra. *Bulletin of the American Mathematical Society*, 26(2):315–316, 1992. (Cited on page 95.)
- [139] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller, J. W. Thatcher, and J. D. Bohlinger, editors, *Complexity of Computer Computations*, pages 85–103. Springer, 1972. (Cited on pages 2, 18.)
- [140] R. M. Karp and C. H. Papadimitriou. On linear characterizations of combinatorial optimization problems. *SIAM Journal on Computing*, 11(4):620–632, 1982. (Cited on page 4.)

-
- [141] K. Khoshkhan and D. O. Theis. On the combinatorial lower bound for the extension complexity of the spanning tree polytope. *Operations Research Letters*, 46(3):352–355, 2018. (Cited on pages 7, 57.)
- [142] E. D. Kim. Polyhedral graph abstractions and an approach to the Linear Hirsch Conjecture. *Mathematical Programming*, 143(1):357–370, 2014. (Cited on page 96.)
- [143] E. D. Kim and F. Santos. An update on the Hirsch conjecture. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 112(2):73–98, 2010. (Cited on pages 95, 96.)
- [144] V. Klee and P. Kleinschmidt. The d -step conjecture and its relatives. *Mathematics of Operations Research*, 12(4):718–755, 1987. (Cited on page 95.)
- [145] V. Klee and D. Walkup. The d -step conjecture for polyhedra of dimension $d < 6$. *Acta Mathematica*, 133:53–78, 1967. (Cited on pages 9, 95, 97, 117.)
- [146] 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. (Cited on pages 85, 86.)
- [147] D. G. Larman. Paths on polytopes. *Proceedings of the London Mathematical Society*, 3(1):161–178, 1970. (Cited on page 95.)
- [148] A. Lehman. A solution of the Shannon switching game. *Journal of the Society for Industrial and Applied Mathematics*, 12(4):687–725, 1964. (Cited on page 23.)
- [149] A. Lehman. On the width-length inequality. *Mathematical Programming*, 16(1):245–259, 1979. (Cited on page 33.)
- [150] A. N. Letchford. Binary clutter inequalities for integer programs. *Mathematical Programming*, 98:201–221, 2003. (Cited on pages 16, 24, 25, 26, 32, 33, 52.)
- [151] A. N. Letchford, G. Reinelt, and D. O. Theis. Odd minimum cut sets and b -matchings revisited. *SIAM Journal on Discrete Mathematics*, 22(4):1480–1487, 2008. (Cited on page 28.)
- [152] A. N. Letchford, S. Pokutta, and A. S. Schulz. On the membership problem for the $\{0, 1/2\}$ -closure. *Operations Research Letters*, 39(5):301–304, 2011. (Cited on pages 15, 18, 22, 25, 50.)
- [153] R. K. Martin. Using separation algorithms to generate mixed integer model reformulations. *Operations Research Letters*, 10(3):119–128, 1991. (Cited on pages 6, 56, 62.)
- [154] M. Mastrolilli. High degree sum of squares proofs, Bienstock–Zuckerberg hierarchy, and Chvátal–Gomory cuts. *SIAM Journal on Optimization*, 30(1):798–822, 2020. (Cited on page 50.)

-
- [155] B. Matschke, F. Santos, and C. Weibel. The width of five-dimensional prisms. In *Proceedings of the London Mathematical Society*, volume 110, pages 647–672, 2015. (Cited on pages [9](#), [97](#), [98](#), [104](#), [108](#), [109](#), [111](#), [112](#), [114](#), [115](#), [116](#), [117](#).)
- [156] S. T. McCormick and A. Shioura. Minimum ratio canceling is oracle polynomial for linear programming, but not strongly polynomial, even for networks. *Operations Research Letters*, 27(5):199–207, 2000. (Cited on page [72](#).)
- [157] K. Menger. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10:96–115, 1927. (Cited on page [17](#).)
- [158] S. Müller and G. Regensburger. Elementary vectors and conformal sums in polyhedral geometry and their relevance for metabolic pathway analysis. *Frontiers in Genetics*, 7:90, 2016. (Cited on page [73](#).)
- [159] D. Naddef. The Hirsch conjecture is true for $(0, 1)$ -polytopes. *Mathematical Programming*, 45(1-3):109–110, 1989. (Cited on pages [95](#), [102](#).)
- [160] J. G. Oxley. *Matroid Theory*. Oxford Graduate Texts in Mathematics. Oxford University Press, 2006. (Cited on pages [22](#), [26](#), [98](#), [99](#).)
- [161] M. W. Padberg and M. R. Rao. Odd minimum cut-sets and b -matchings. *Mathematics of Operations Research*, 7(1):67–80, 1982. (Cited on pages [5](#), [15](#), [16](#), [27](#), [28](#), [52](#).)
- [162] A. Postnikov, V. Reiner, and L. Williams. Faces of generalized permutohedra. *Documenta Mathematica*, 13:207–273, 2008. (Cited on page [57](#).)
- [163] M. Queyranne. Structure of a simple scheduling polyhedron. *Mathematical Programming*, 58(1-3):263–285, 1993. (Cited on pages [57](#), [65](#), [66](#).)
- [164] R. Rizzi. A simple minimum T -cut algorithm. *Discrete Applied Mathematics*, 129(2-3):539–544, 2003. (Cited on page [28](#).)
- [165] 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. (Cited on pages [73](#), [77](#), [100](#).)
- [166] T. Rothvoß. Some 0/1 polytopes need exponential size extended formulations. *Mathematical Programming*, 142(1-2):255–268, 2013. (Cited on page [56](#).)
- [167] T. Rothvoß. The matching polytope has exponential extension complexity. *Journal of the ACM*, 64(6):41, 2017. (Cited on pages [6](#), [7](#), [56](#), [57](#), [59](#), [60](#), [132](#).)
- [168] T. Rothvoß and L. Sanità. 0/1 polytopes with quadratic Chvátal rank. *Operations Research*, 65(1):212–220, 2017. (Cited on page [4](#).)
- [169] L. Sanità. 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, 2018. (Cited on pages [8](#), [78](#).)

-
- [170] F. Santos. A counterexample to the Hirsch conjecture. *Annals of Mathematics*, 176(1):383–412, 2012. (Cited on pages 9, 95, 97, 104, 105, 108, 114, 115, 117.)
- [171] F. Santos. Recent progress on the combinatorial diameter of polytopes and simplicial complexes. *TOP*, 21:426–460, 2013. (Cited on pages 95, 96.)
- [172] F. Santos, T. Stephen, and H. Thomas. Embedding a pair of graphs in a surface, and the width of 4-dimensional prisms. *Discrete & Computational Geometry*, 47:569–576, 2012. (Cited on page 117.)
- [173] A. Schrijver. On cutting planes. *Annals of Discrete Mathematics*, 9:291–296, 1980. (Cited on page 3.)
- [174] A. Schrijver. *Theory of Linear and Integer Programming*. Wiley, 1998. (Cited on pages 10, 13, 15, 34, 37, 46, 48, 121.)
- [175] A. Schrijver. Strong t -perfection of bad- K_4 -free graphs. *SIAM Journal on Discrete Mathematics*, 15(3):403–415, 2002. (Cited on pages 46, 50.)
- [176] A. Schrijver. A short proof of Guenin’s characterization of weakly bipartite graphs. *Journal of Combinatorial Theory, Series B*, 85(2):255–260, 2002. (Cited on page 39.)
- [177] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003. (Cited on pages 6, 7, 10, 16, 17, 22, 39, 40, 46, 50.)
- [178] A. S. Schulz. On the relative complexity of 15 problems related to 0/1-integer programming. In W. J. Cook, L. Lovász, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, pages 399–428. Springer, 2009. (Cited on pages 6, 7, 25, 26, 95, 102.)
- [179] A. S. Schulz and R. Weismantel. The complexity of generic primal algorithms for solving general integer programs. *Mathematics of Operations Research*, 27(4):681–692, 2002. (Cited on page 72.)
- [180] A. S. Schulz, R. Weismantel, and G. M. Ziegler. 0/1-Integer programming: Optimization and augmentation are equivalent. In P. Spirakis, editor, *Algorithms – ESA ’95. ESA 1995*, volume 979 of *Lecture Notes in Computer Science*, pages 473–483. Springer, 1995. (Cited on page 7.)
- [181] J. Seif. Bounding techniques for extension complexity. Master’s thesis, Université Libre de Bruxelles, 2017. (Cited on page 60.)
- [182] P. D. Seymour. The forbidden minors of binary clutters. *Journal of the London Mathematical Society*, s2-12(3):356–360, 1976. (Cited on page 23.)
- [183] P. D. Seymour. The matroids with the max-flow min-cut property. *Journal of Combinatorial Theory, Series B*, 23(2-3):189–222, 1977. (Cited on pages 17, 22, 32, 36, 39, 44, 45, 49, 50, 51.)

- [184] 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. (Cited on page 72.)
- [185] T. Stephen and T. Yusun. The circuit diameter of the Klee-Walkup polyhedron. *Electronic Notes in Discrete Mathematics, The Eight European Conference on Combinatorics, Graph Theory and Applications, EuroComb 2015*, 49:505–512, 2015. (Cited on pages 9, 97.)
- [186] T. Stephen and T. Yusun. Circuit diameter and Klee-Walkup constructions. *arXiv preprint*, arXiv:1503.05252, 2015. (Cited on page 98.)
- [187] B. Sturmfels and R. R. Thomas. Variation of cost functions in integer programming. *Mathematical Programming*, 77(2):357–387, 1997. (Cited on pages 73, 74.)
- [188] N. Sukegawa. Improving bounds on the diameter of a polyhedron in high dimensions. *Discrete Mathematics*, 340(9):2134–2142, 2017. (Cited on page 95.)
- [189] N. Sukegawa. An asymptotically improved upper bound on the diameter of polyhedra. *Discrete & Computational Geometry*, 62:690–699, 2019. (Cited on page 95.)
- [190] T. Terlaky and S. Zhang. Pivot rules for linear programming: A survey on recent theoretical developments. *Annals of Operations Research*, 46:203–233, 1993. (Cited on page 8.)
- [191] M. J. Todd. The monotonic bounded Hirsch conjecture is false for dimension at least 4. *Mathematics of Operations Research*, 5(4):599–601, 1980. (Cited on page 118.)
- [192] M. J. Todd. An improved Kalai–Kleitman bound for the diameter of a polyhedron. *SIAM Journal on Discrete Mathematics*, 28(4):1944–1947, 2014. (Cited on page 95.)
- [193] K. Truemper. Max-flow min-cut matroids: Polynomial testing and polynomial algorithms for maximum flow and shortest routes. *Mathematics of Operations Research*, 12(1):72–96, 1987. (Cited on pages 16, 25, 31, 49, 50, 51, 52.)
- [194] F. Tseng and K. Truemper. A decomposition of the matroids with the max-flow min-cut property. *Discrete Applied Mathematics*, 15(2-3):329–364, 1986. (Cited on page 49.)
- [195] S. A. Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2010. (Cited on page 58.)
- [196] P. Ventura and F. Eisenbrand. A compact linear program for testing optimality of perfect matchings. *Operations Research Letters*, 31(6):429–434, 2003. (Cited on page 6.)

-
- [197] 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. (Cited on page 72.)
- [198] C. Weibel. Counter-examples to the Hirsch conjecture. <https://sites.google.com/site/christophweibel/research/hirsch-conjecture>. (Cited on pages 115, 116.)
- [199] S. Weltge. *Sizes of linear descriptions in combinatorial optimization*. PhD thesis, Otto-von-Guericke-Universität Magdeburg, 2016. (Cited on pages 7, 57.)
- [200] J. C. Williams. A linear-size zero-one programming model for the minimum spanning tree problem in planar graphs. *Networks*, 39(1):53–60, 2002. (Cited on page 56.)
- [201] D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011. (Cited on page 65.)
- [202] L. A. Wolsey. Mixed integer programming formulations for production planning and scheduling problems. Invited talk at the 12th International Symposium on Mathematical Programming, MIT, Cambridge, 1985. (Cited on pages 57, 66.)
- [203] L. A. Wolsey. Using extended formulations in practice. *Optima*, 85, 2011. (Cited on page 71.)
- [204] 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. (Cited on pages 6, 56, 71.)
- [205] M. Yannakakis. Expressing combinatorial optimization problems by linear programs. *Journal of Computer and System Sciences*, 43(3):441–466, 1991. (Cited on pages 7, 56, 58, 59, 71.)
- [206] G. M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer, New York, NY, 2012. (Cited on pages 63, 75, 78, 83, 86, 98, 99, 102, 121.)
- [207] G. M. Ziegler. Who solved the Hirsch conjecture. *Documenta Mathematica Extra Volume: Optimization Stories*, pages 75–85, 2012. (Cited on page 95.)