Technische Universität München
Fakultät für Mathematik
Lehrstuhl für Angewandte Geometrie und Diskrete Mathematik

Circuit diameters

and their application to transportation problems

Elisabeth Finhold

Vollständiger Abdruck der von der Fakultät für Mathematik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender:		Univ.-Prof. Dr. Dr. Jürgen Richter-Gebert
Prüfer der Dissertation:	1.	Univ.-Prof. Dr. Raymond Hemmecke
	2.	Prof. Dr. Jesús A. De Loera University of California, Davis (USA)
	3.	Univ.-Prof. Dr. Rüdiger Schultz Universität Duisburg-Essen

Die Dissertation wurde am 20.11.2014 bei der Technischen Universität München eingereicht und durch die Fakultät für Mathematik am 19.02.2015 angenommen.

Circuit diameters

and their application to transportation problems

ELISABETH FINHOLD

Zusammenfassung

In dieser Arbeit führen wir den Zirkuit-Durchmesser (circuit diameter) ein, eine Verallgemeinerung des kombinatorischen Durchmessers von Polyedern. Wir untersuchen, wie viele Zirkuit-Schritte wir brauchen um je zwei Ecken eines Polyeders mit einem Zirkuit-Weg zu verbinden. Diese Wege verlaufen entlang von Zirkuit-Richtungen, das heißt, parallel zu potentiellen Kanten des Polyeders. Indem wir fordern, dass die Zirkuit-Wege zusätzliche Kriterien erfüllen, ergeben sich diverse Kategorien von Zirkuit-Durchmessern. Wir setzen all diese Klassen zueinander in Relation, wodurch eine Hierarchie von Durchmessern entsteht. Diese Vielzahl miteinander verwandter Konzepte ermöglicht es uns, neue Erkenntnisse über den kombinatorischen Durchmesser zu gewinnen. Beispielsweise bilden schwächere Zirkuit-Durchmesser untere Schranken an den kombinatorischen Durchmesser.

Wie man all diese Zirkuit-Durchmesser-Konzepte konkret ausnutzen kann, zeigen wir am Beispiel von Transportpolytopen und dualen Netzwerkfluss-Polyedern. Wir beweisen obere und untere Schranken an die verschiedenen Durchmesser, wobei wir besonderes Augenmerk auf die Schranke legen, die in Verbindung mit der berühmten Hirsch-Vermutung steht.

Abstract

The present thesis introduces the circuit diameter of polyhedra as a generalization to the combinatorial diameter. It tells us how many circuit steps we need to connect any two vertices of a polyhedron with a circuit walk. Such a walk goes along circuit directions, that is, every step is parallel to a potential edge of the polyhedron. Several notions of circuit diameters arise from putting further restrictions on these walks. We relate all these categories in a comprehensive hierarchy. Therein, the generalized diameters reveal some interesting lower bounds for the combinatorial diameter and thus their study might help for a better understanding of the latter one.

By investigating the hierarchy for transportation polytopes and dual network flow polyhedra, we demonstrate on two concrete examples how one can exploit the availability of diameters of different strength. In this we focus on proving upper and lower bounds on the several circuit diameters, where the bound induced by the famous Hirsch conjecture is of special interest.

Acknowledgments

First and foremost, my thanks go to my adviser Prof. Raymond Hemmecke for his guidance and advice in the past years, for putting trust in me, supporting and encouraging me, and for giving me the opportunity to work on the idea of the circuit diameter that turned out to be a great topic for my thesis.

I want to thank my co-authors Steffen Borgwardt, Jesús A. De Loera, Raymond Hemmecke, and Jake Miller for the enjoyable collaboration. Thank you for all the significant results we achieved and I may use in this thesis. Special thanks go to Prof. De Loera for inviting me to UC Davis this year. I experienced a great time and I took great benefit from it, as it gave rise to important parts of my thesis. Further, I would like to thank Steffen for all his valuable hints and comments.

I very much appreciate the pleasant environment at Technische Universität München at the department of Prof. Gritzmann and thanks to the TopMath program. I am certainly grateful that I got the chance to do research at this early point. The people involved in the TopMath always were a great help - Agnieszka Baumgärtel, Prof. Brokate, Carl-Friedrich Kreiner, and Denise Lichtig, only to name a few. I acknowledge the support from the graduate program TopMath of the Elite Network of Bavaria and the TopMath Graduate Center of TUM Graduate School at Technische Universität München.

But above all, my deepest thanks go to my family and close friends for their steady support and trust, for all the words of cheer and their never-ending patience and understanding – DANKE –

Contents

Introduction

An optimization problem over a polyhedron can be solved via a simple augmentation scheme: As long as the current solution is not optimal, improve it! The essential ingredient for such an algorithm is a *test set*, a set of vectors that contains an augmenting direction for every non-optimal solution. Such a set is also called an *optimality certificate*, as we have found the optimum if we can no longer improve our current solution along any vector in the test set.

In this thesis we consider a particularly interesting test set, the *circuits* [14, 15, 28]. These vectors are associated with the polyhedron that defines the feasible region of a linear program. They allow a particularly nice, intuitive interpretation: We will see that for a polyhedron of the form

$$P = \{\, \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \,\}$$

the circuits are given by all potential edge directions that appear as the right-hand sides $\mathbf{b}$ and $\mathbf{d}$ vary. Consequently, all these polyhedra, that result from translating the defining hyperplanes, admit the same set of circuits.

Using the circuits as a test set creates a path from an initial vertex $\mathbf{v}^0$ to an optimal vertex $\mathbf{v}^*$ of the associated optimization problem, illustrated in Figure 1. Observe that every augmenting direction is parallel to an edge and thus a circuit. We call a sequence of (not necessarily monotonic) *circuit steps* that connects two vertices of a polyhedron a *circuit walk*.

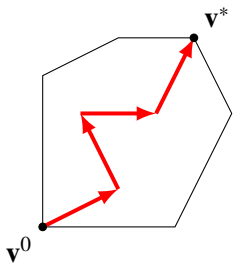

Figure 1: A circuit walk to the optimal vertex.

An immediate question is the following: Starting at any vertex of the polyhedron, how many augmentation steps do we need to reach an optimal vertex? It is not hard to see that in the example depicted in Figure 1 we could have used far less steps, see Figure 2.

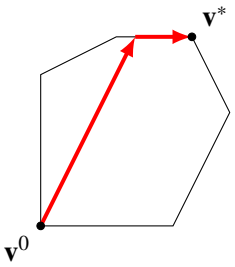

Figure 2: A shorter circuit walk to the optimal vertex.

In [10] it was recently shown that the number of augmentation steps is bounded by the number of circuits if we always use a steepest-descent circuit direction and apply a step of maximum length. This thesis is concerned with the question of how many circuit steps we need to an optimal vertex of the polyhedron if we assume a 'perfect' selection rule. In this we even disregard monotonicity of the paths with respect to a certain objective function. In an attempt to answer this question, we introduce and study different types of *circuit diameters* of polyhedra.

Among these categories we will also recover the combinatorial diameter. This is a classical field of research in the theory of linear programming as it is of interest in the context of a best-case performance of the Simplex algorithm. The *Simplex method* is a well-known special case of augmentation algorithms along circuits: Every augmenting step goes along an actual edge of the polyhedron, from one vertex to a better one. We call such a sequence of steps an *edge walk*. The *combinatorial diameter* is the maximum number of edges we need to connect any two vertices of the polyhedron, and hence it is a lower bound on the number of steps the Simplex algorithm might take. It is still open whether there is a polynomial upper bound on the combinatorial diameter and thus whether there is hope for a pivot rule for the Simplex algorithm that takes only polynomially many steps. The connection to the Simplex algorithm becomes even more direct when investigating the *monotone* diameter. Here we consider monotone edge-walks, visiting vertices of non-decreasing objective values with respect to a certain linear functional.

One of the most famous problems associated with the combinatorial diameter is the *Hirsch conjecture*. In 1957, Warren M. Hirsch suggested an upper bound of $f - d$ on the combinatorial diameter any d-dimensional polyhedron with f facets [8]. Only ten years later, Klee and Walkup showed that the conjecture does not hold for unbounded polyhedra of dimension ≥ 4 [21]. This led to the bounded version of the Hirsch conjecture, claiming an upper bound of $f - d$ on the combinatorial diameter of *polytopes*. It was a long-standing open question, until Santos recently came up with the first counterexample in dimension 43 with 86 facets and diameter at least 44 [27]. Nevertheless, in the past decades validity of the Hirsch conjecture was proved for certain classes of polyhedra, such as $(0, 1)$-polytopes [24], dual transportation polyhedra [1] and special combinations of the dimension d and the number of facets f (see [20] for a survey). However, despite great efforts the Hirsch conjecture remains open for most classes of polyhedra; even the more general polynomial Hirsch conjecture, asking only for a polynomial diameter bound in d and

f. Yet, the latter one is not disproved either!

Being confronted with these long-standing open question in the context of the combinatorial diameter, we have to look out for new approaches. Introducing the circuit diameter as a generalization to the combinatorial diameter might help for a better understanding of this difficult field.

The present thesis is concerned with the very foundations of the circuit diameter, based on our definitions and results in [3] and [5]. Furthermore, it contains a demonstration of these concepts by considering transportation polytopes [4] and dual network flow polyhedra [5, 6].

Let us now have a closer look at the general ideas behind the different types of circuit diameters. To define these notions, we extend the very specific edge walks related to the combinatorial diameter to different kinds of circuit walks between the vertices of the polyhedron (as already seen in Figure 1 and Figure 2 in the context of circuit augmentation algorithms). While the former only use the *actual edges* of the polyhedron and thus never leave its boundary, the latter ones go along *potential edge directions*, the circuits. In particular, a circuit walk may enter the interior of the polyhedron, which possibly yields fewer steps. Figure 3 illustrates four different types of circuit walks in a two-dimensional polytope that arise from imposing or relaxing certain restrictions on the circuit steps as stated below the respective images.

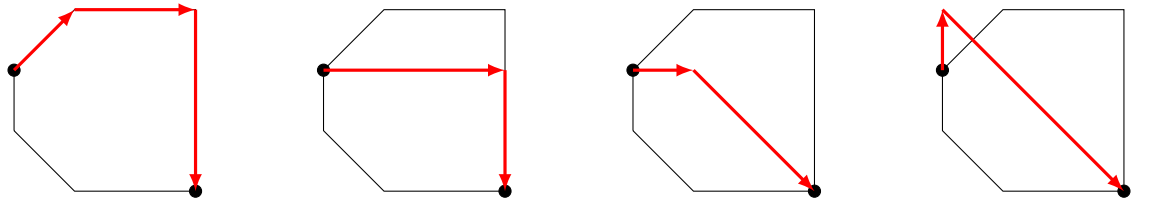

edge walk | circuit walk (maximum length) | weak circuit walk (arbitrary length) | soft circuit walk (possibly infeasible)

Figure 3: Four different types of circuit walks.

The first picture shows a traditional edge walk, but for the other examples we liberate ourselves from the stringent conditions of edge walks by allowing to leave the boundary of the polyhedron.

In the second picture we see a circuit walk that uses *maximum length steps*, that is, we always go along the circuit directions as far as possible, potentially also through the interior of the polyhedron. Our original definition of the circuit diameter in [5] is based on this kind of walks. The third and fourth picture depict even more general circuit walks, which we introduced in [3]. On *weak circuit walks* we skip maximality of the steps but we stay inside the polyhedron. In contrast, we can even leave the polyhedron along circuit directions on a *soft circuit walk*. These four main types of circuit walks are complemented by some more specific ones that arise from forbidding repeated and backwards steps, or asking for sign-compatible walks [3].

Now, every such type of circuit walk induces a circuit distance: The *circuit distance* from one vertex to another one is the minimum number of steps of a corresponding circuit walk, and the *circuit diameter* of a polyhedron is the maximum circuit distance among its vertices. Note that

each notion of circuit diameter can be seen as an indication for a best-case performance of circuit augmentation algorithms with corresponding properties.

We are particularly interested in how all these different circuit diameters are related. We answer this question in a comprehensive hierarchy of circuit distance concepts. The diameters associated with the four types of circuit walks in Figure 3 form the 'key chain' of this hierarchy: Note that the circuit walks become less restrictive from the left to the right and thus the corresponding distances trivially bound each other. In particular, the combinatorial diameter is an upper bound on all these circuit diameters, as an edge walk is just a special kind of circuit walk. On the other hand, the less restrictive notions provide some interesting lower bounds on the combinatorial diameter. What makes this so useful is the fact that the weaker categories are typically much easier to compute. Also, we often find similarities in the approaches for tackling the combinatorial and the 'original' circuit diameter. For these reasons, the availability of circuit diameter concepts of varying strength is an extremely valuable tool.

We demonstrate and exploit this by discussing the hierarchy for two families of polyhedra: *Transportation polytopes* and *dual network flow polyhedra*. These well-known problems are frequently used in combinatorial optimization to model a variety of problems.

Transportation problems [12, 13, 22, 29] describe the distribution of some product from M supply locations to N demand locations under minimum costs. Even though they are among the most fundamental problems in mathematical programming, it is still open whether the Hirsch conjecture holds for transportation polytopes; currently the best upper bounds are linear.

In contrast, the Hirsch conjecture is known to be true for *dual transportation polyhedra* [1]. This motivates the investigation of the other notions of circuit diameters for such polyhedra. In addition, we aim at generalizing the results to the case of *dual network flow polyhedra*, which are associated to minimum-cost flows through a network.

When studying these two families of polyhedra we focus on proving upper and lower bounds on the main categories of circuit diameters. For both types of polyhedra we easily obtain upper bounds on the weak and on the soft circuit diameter, which are actually tight. On the other hand, when investigating the two stronger diameter categories we often obtain only rough upper bounds unless we restrict ourselves to special cases. Typically, the combinatorial diameter is even more involved than the 'original' circuit diameter. However, our studies of the circuit diameter serve as a basis for addressing the combinatorial diameter, we often use similar approaches when proving upper or lower bounds on these notions of diameter.

Structure of the thesis and main results

In Chapter 1 we introduce circuits as the most fundamental notion of this thesis. We summarize all definitions and results we need for the forthcoming discussions and illustrate them on some elementary examples.

Chapter 2 is the core part of the thesis. Here we formally introduce the different categories of

circuit walks, circuit distances and circuit diameters, as already initiated in the introduction. We then relate them to each other, resulting in our hierarchy. We also provide and discuss general upper bounds, in particular we show that the weaker categories of circuit diameter satisfy the bound induced by the Hirsch conjecture. Finally, we study the circuit diameter concepts in dimension two where we can refine our previous results. This is joint work with Steffen Borgwardt, Jesús A. De Loera and Raymond Hemmecke [3, 5].

In the last two independent chapters we investigate the 'key chain' of the hierarchy for two particular classes of polyhedra. In this we aim at upper and lower bounds on these diameters, where the bound induced by the Hirsch conjecture is of particular interest.

In Chapter 3 we consider the transportation polytopes. We prove that the Hirsch conjecture holds for $2\times N$ and $3\times N$ transportation polytopes. For the $2\times N$ case can even prove the monotone Hirsch conjecture with a slightly stronger bound. We obtain similar results for the circuit diameter, but the discussion itself reveals some specifics of this diameter category. For the general $M\times N$ transportation polytopes we show that the weak and the soft circuit diameter satisfy the Hirsch bound and that this bound is indeed tight. This is joint work with Steffen Borgwardt, Jesús A. De Loera and Jake Miller [4].

In Chapter 4 we then turn to dual network flow polyhedra. We prove linear bounds in the special case of dual transportation polyhedra. In the general setup of dual network flow polyhedra, we show that the weak and soft circuit diameter remain linear, while we prove quadratic upper bounds for the stronger categories. This is joint work with Steffen Borgwardt and Raymond Hemmecke [5, 6].

The thesis concludes with an outlook concerning future work and open questions on the circuit diameter. In particular, we emphasize a potential extension of circuit diameters to an integral setting of lattice points in a polyhedron.

Chapter 1

Circuits

We now start with a thorough introduction of the circuits, which are the most fundamental notion of this thesis: All chapters hereafter are based on this set of vectors. The circuits are associated with integral matrices $A \in \mathbb{Z}^{m_A \times n}$ and $B \in \mathbb{Z}^{m_B \times n}$ that define polyhedra of the form

$$P = \{ \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \}.$$

They only depend on A and B, but are independent of the rational right-hand sides $\mathbf{b} \in \mathbb{Q}^{m_A}$, $\mathbf{d} \in \mathbb{Q}^{m_B}$; we even allow $\mathbf{b} \in \mathbb{Q}_\infty^{m_A}$, $\mathbf{d} \in \mathbb{Q}_\infty^{m_B}$, where $\mathbb{Q}_\infty := \mathbb{Q} \cup \{\pm\infty\}$. In fact, we will see that the circuits provide an optimality certificate (or test set) for a whole family of optimization problems with the right-hand sides $\mathbf{b}$ and $\mathbf{d}$ and the (linear) objective function as parameters.

Circuits as optimality certificates were introduced by Jack E. Graver in his seminal paper [14]. Most of the definitions and results we provide in the following were already stated therein. However, for completeness we prove all fundamental results we need throughout the thesis in this section. This also serves as an introduction to the topic and we can introduce our notation thereby.

We begin with a basic definition of circuits based on the support of vectors, followed by some fundamental observations on this set. We then turn to three equivalent descriptions of the circuits as we will make use of these characterizations in the forthcoming chapters. In this, the *representation property* immediately implies that circuits are indeed optimality certificates for linear programs. We close the section by providing some examples, including a remark on how a certain representation of the polyhedron affects the associated set of circuits.

Before we start, we make two assumptions on the matrices A and B. In the later chapters we want our polyhedra to have vertices. A vertex $\mathbf{v}$ of $P = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}, B\mathbf{x} \leq \mathbf{d}\}$ is determined by $n - \operatorname{rank}(A)$ linearly independent inequalities that are tight. Thus we assume $\operatorname{rank}\binom{A}{B} = n$ throughout the thesis, as otherwise the polyhedron cannot have vertices at all. We further assume that $\operatorname{rank}\binom{A}{B} = \operatorname{rank}(A) + \operatorname{rank}(B)$, which can easily be obtained by removing those rows (inequalities) from B that are linear combinations of the rows (equations) of A. This defines the same polyhedron, but we will see that it does not change the associated set of circuits.

1.1 Definition and fundamental properties

The definition of a circuit is based on the *support* of a vector: It is given by the indices of the non-zero components of $\mathbf{v} \in \mathbb{R}^n$ as $\operatorname{supp}(\mathbf{v}) = \{\, i \in \{1, \dots, n\} : \; v_i \neq 0 \,\}$. Note that for $\mathbf{y} \in \ker(A)$, $\operatorname{supp}(B\mathbf{y}) = \emptyset$ implies $\mathbf{y} = \mathbf{0}$ as we assumed $\operatorname{rank}\binom{A}{B} = n$. We say that a vector $\mathbf{v}$ is (set-wise) *support-minimal in a set* $S \subseteq \mathbb{R}^n$ if $\mathbf{v} \in S$ and there is no $\mathbf{w} \in S$ with $\operatorname{supp}(\mathbf{w}) \subsetneq \operatorname{supp}(\mathbf{v})$. We now can define the circuits of matrices A and B.

Definition 1.1 (See Def. 3.8 in [14])**.** *The* circuits $C(A, B)$ *associated with matrices A and B are those non-zero vectors $\mathbf{g} \in \ker(A) = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{0}\}$ for which $B\mathbf{g}$ is support-minimal in the set $\{\, B\mathbf{y} : \mathbf{y} \in \ker(A) \backslash \{\mathbf{0}\} \,\}$, where $\mathbf{g}$ is normalized to (coprime) integer components.*

There are a few remarks to make about this definition. First note that it is indeed independent of $\mathbf{b}$ and $\mathbf{d}$. That is why the set of circuits was called the *universal test set* in [14]. The term 'circuit' for such vectors was first used in [28]. The set is always symmetric (that is, if $\mathbf{g} \in C(A, B)$ then $-\mathbf{g} \in C(A, B)$) and finite. The latter follows immediately from the next lemma, as vectors $B\mathbf{y} \in \mathbb{R}^{m_B}$ admit only finitely many supports.

Lemma 1.2 (See Prop. 3.9 in [14])**.** *Let $\mathbf{g} \in C(A, B)$ and $\mathbf{v} \in \ker(A)$ with $\operatorname{supp}(B\mathbf{g}) = \operatorname{supp}(B\mathbf{v})$. Then $\mathbf{v} = \alpha\mathbf{g}$ for some $\alpha \in \mathbb{R}$.*

Proof. As $\mathbf{g} \in \ker(A)\backslash\{\mathbf{0}\}$, there is some $i \in \{1, \dots, m_B\}$ such that $(B\mathbf{g})_i \neq 0$ and hence $(B\mathbf{v})_i \neq 0$. Let $\mathbf{w} \coloneqq \mathbf{g} - \left[(B\mathbf{g})_i/(B\mathbf{v})_i\right] \cdot \mathbf{v}$. Then we have $\mathbf{w} \in \ker(A)$ and further $\operatorname{supp}(B\mathbf{w}) \subsetneq \operatorname{supp}(B\mathbf{g})$ as $(B\mathbf{w})_i = (B\mathbf{g})_i - \left[(B\mathbf{g})_i/(B\mathbf{v})_i\right] \cdot (B\mathbf{v})_i = 0$. Hence we have $\mathbf{w} = \mathbf{0}$ by support-minimality of $\mathbf{g}$ and thus $\mathbf{v} = \left[(B\mathbf{v})_i/(B\mathbf{g})_i\right] \cdot \mathbf{g}$. □

We want to highlight two special settings in which circuits frequently appear. Often one simply imposes upper or lower bounds on the variable $\mathbf{x}$ instead of having a whole system of inequalities $B\mathbf{x} \leq \mathbf{d}$. In other words, we have $B = \pm I_n$, the identity matrix in $\mathbb{R}^{n\times n}$, and thus $\operatorname{supp}(B\mathbf{x}) = \operatorname{supp}(\mathbf{x})$. Then our definition of circuits is closely related to the *elementary vectors* of $\ker(A)$ as introduced in [26]. These are *all* support-minimal elements in $\ker(A)\backslash\{\mathbf{0}\}$. Observe that there are infinitely many such elements as the elementary vectors can be scaled arbitrarily: The elementary vectors are lines through the origin, with the origin excluded. However, Lemma 1.2 tells us that when restricting to those vectors whose components have greatest common divisor one, we recover the circuits $C(A, I_n)$, which we simply denote $C(A)$. In particular, every vector of minimal-support in $\ker(A)\backslash\{\mathbf{0}\}$ has two representatives in $C(A)$, some circuit and its negative (compare Lemma 1.2). Further, note that our definition of circuits is consistent with the term 'circuits' used in matroid theory: The support of a support-minimal vector in $\ker(A)\backslash\{\mathbf{0}\}$ describes an inclusion-minimal set of linearly dependent columns of A.

Another special case is that of having only inequality-constraints, that is, $A = O$, a matrix with all-zero entries. We then write $C_{\leq}(B) \coloneqq C(O, B)$. This special case will be particularly useful for investigating two-dimensional examples.

Finally, we want to point out the following: The set of circuits associated with the matrices A and B does not change when we modify B by

1. swapping rows of B,
2. multiplying a row of B by a non-zero scalar,
3. adding linear combinations of rows of A and B to another row of B
4. adding a row to the matrix B that is a linear combination of other rows of A and B, or
5. removing a row from B that is a linear combination of other rows of A and B,

as long as all entries remain integral. Even though these operations might change the support of some $B\mathbf{x}$, support-minimality is preserved. This follows immediately from the linear dependencies among the rows and from $A\mathbf{g} = \mathbf{0}$ for all circuits $\mathbf{g}$. In particular this implies that the polyhedra

$$\{\,\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d}\,\},\ \{\,\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \geq \mathbf{d}\,\} \text{ and } \{\,\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ \mathbf{l} \leq B\mathbf{x} \leq \mathbf{u}\,\}$$

all admit the same sets of circuits since it holds that $C(A, B) = C(A, -B) = C\left(A, \binom{B}{-B}\right)$. Therefore, we may without loss of generality add lower bounds on $B\mathbf{x}$ or replace the upper bounds by lower bounds.

1.2 Circuits as optimality certificates

We now aim at proving the test set property for circuits. Therefore, we have to consider sign-compatible vectors: Two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ are *sign-compatible with respect to the matrix B* defining the polyhedron $P = \{\,\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d}\,\}$, if $B\mathbf{v}$ and $B\mathbf{w}$ belong to the same orthant of $\mathbb{R}^{m_B}$, that is, their i-th components $(B\mathbf{v})_i$ and $(B\mathbf{w})_i$ satisfy $(B\mathbf{v})_i \cdot (B\mathbf{w})_i \geq 0$ for all $i = 1, \ldots, m_B$. Note that for $B = I_n$ this definition becomes much simpler. We first need the following

Lemma 1.3 (Cf. Lemma 3.11 in [14])**.** *If $\mathbf{v} \neq \mathbf{0}$ is not support-minimal in $\{\,B\mathbf{y} : \mathbf{y} \in \ker(A)\backslash\{\mathbf{0}\}\,\}$, then there is an element $\mathbf{w} \in \ker(A)\backslash\{0\}$ such that $\mathbf{w}$ and $\mathbf{v}$ are sign-compatible with respect to B and $\operatorname{supp}(B\mathbf{w}) \subsetneq \operatorname{supp}(B\mathbf{v})$.*

Proof. As $B\mathbf{v}$ is not support-minimal in $\{\,B\mathbf{y} : \mathbf{y} \in \ker(A)\backslash\{\mathbf{0}\}\,\}$, there is $\mathbf{y} \in \ker(A)\backslash\{0\}$ such that $\operatorname{supp}(B\mathbf{y}) \subsetneq \operatorname{supp}(B\mathbf{v})$. We can assume without loss of generality that there is some $i \in \{1, \ldots, m_B\}$ such that $(B\mathbf{v})_i(B\mathbf{y})_i > 0$ (otherwise we replace $\mathbf{y}$ by $-\mathbf{y}$). Let $\mathbf{w} := \mathbf{v} - \alpha\mathbf{y}$, where $\alpha := \min\{\,(B\mathbf{v})_i/(B\mathbf{y})_i : (B\mathbf{v})_i(B\mathbf{y})_i > 0\,\} \geq 0$. Let $(B\mathbf{v})_i = 0$. Then $(B\mathbf{w})_i = -\alpha(B\mathbf{y})_i = 0$ as $\operatorname{supp}(B\mathbf{y}) \subseteq \operatorname{supp}(B\mathbf{v})$ and hence $\operatorname{supp}(B\mathbf{w}) \subseteq \operatorname{supp}(B\mathbf{v})$. It remains to show sign-compatibility, that is, $0 \leq (B\mathbf{w})_i(B\mathbf{v})_i = (B\mathbf{v})_i^2 - \alpha(B\mathbf{y})_i(B\mathbf{v})_i$. For $(B\mathbf{y})_i(B\mathbf{v})_i \leq 0$ non-negativity of this term is obvious. Otherwise we have (by definition of α) that $(B\mathbf{v})_i^2 - \alpha(B\mathbf{y})_i(B\mathbf{v})_i \geq (B\mathbf{v})_i^2 - \left[(B\mathbf{v})_i/(B\mathbf{y})_i\right](B\mathbf{y})_i(B\mathbf{v})_i = 0$. □

With this we now are able to state a crucial property of circuits.

Proposition 1.4 (Representation property, see Theorem 3.12 in [14])**.** *Every element $\mathbf{v} \in \ker(A)$ can be written as a positive linear sign-compatible sum of elements in $C(A, B)$, that is, we have a representation $\mathbf{v} = \sum_{i=1}^{k} \alpha_i \mathbf{g}^i$, such that for all $i = 1, \ldots, k$*

(i) $\mathbf{g}^i \in C(A,B)$,
(ii) $\mathbf{g}^i$ *and* $\mathbf{v}$ *are sign-compatible and* $\operatorname{supp}(B\mathbf{g}^i) \subseteq \operatorname{supp}(B\mathbf{v})$ *, and*
(iii) $\alpha_i \in \mathbb{R}_+ := \{\alpha \in \mathbb{R} : \alpha \geq 0\}$.

Proof. Let $\mathbf{v} \in \ker(A)$. If $\mathbf{v} = \mathbf{0}$, there is nothing to show. Else, if $B\mathbf{v}$ is support-minimal in the set $\{ B\mathbf{y} : \mathbf{y} \in \ker(A)\backslash\{\mathbf{0}\} \}$, by Lemma 1.2 there is $\mathbf{g} \in C(A,B)$ such that $\mathbf{v} = \alpha\mathbf{g}$ for some $\alpha \in \mathbb{R}$ and we can choose $\alpha > 0$ as the set of circuits is symmetric. So assume $\mathbf{v}$ is not support-minimal in $\{ B\mathbf{y} : \mathbf{y} \in \ker(A)\backslash\{\mathbf{0}\} \}$. Then by Lemma 1.3 there is some $\mathbf{w} \in \ker(A)\backslash\{0\}$ such that $\mathbf{w}$ and $\mathbf{v}$ are sign-compatible with respect to B and $\operatorname{supp}(B\mathbf{w}) \subsetneq \operatorname{supp}(B\mathbf{v})$. Let $\mathbf{w}' := \mathbf{v} - \alpha\mathbf{w}$ for $\alpha := \min\{ (B\mathbf{v})_i/(B\mathbf{w})_i : (B\mathbf{w})_i \neq 0\} \geq 0$. Then $\operatorname{supp}(B\mathbf{w}') \subsetneq \operatorname{supp}(B\mathbf{v})$, as clearly $\operatorname{supp}(B(\mathbf{v} - \alpha\mathbf{w})) \subseteq \operatorname{supp}(B\mathbf{v})$, and $B(\mathbf{v} - \alpha\mathbf{w})$ has at least one extra zero-entry by definition of α. Further $\mathbf{w}'$ and $\mathbf{v}$ are sign-compatible as $(B\mathbf{v})_i(B\mathbf{w}')_i = (B\mathbf{v})_i^2 - \alpha(B\mathbf{v})_i(B\mathbf{w})_i \geq (B\mathbf{v})_i^2 - [(B\mathbf{v})_i/(B\mathbf{w})_i]\,(B\mathbf{v})_i(B\mathbf{w})_i = 0$. Hence we found a representation $\mathbf{v} = \alpha\mathbf{w} + \mathbf{w}'$ for vectors $\alpha\mathbf{w}$ and $\mathbf{w}'$ sign-compatible with $\mathbf{v}$ and $\operatorname{supp}(B(\alpha\mathbf{w})), \operatorname{supp}(B\mathbf{w}') \subsetneq \operatorname{supp}(B\mathbf{v})$. The claim follows inductively. □

Note that the circuits $\mathbf{g}^i$ are also pairwise sign-compatible with respect to B. The following corollary now gives us an alternative definition for circuits.

Corollary 1.5 (Circuit definition via representation property)**.** *The set of circuits* $C(A,B)$ *is (up to scaling) the unique inclusion-minimal set admitting the representation property posed in Proposition 1.4.*

Proof. By Proposition 1.4, $C(A,B)$ has this representation property. On the other hand, (ii) and (iii) from Proposition 1.4 together with Lemma 1.2 imply that such a set must contain a representative $\alpha\mathbf{g}$, $\alpha \in \mathbb{R}_+$ of every circuit $\mathbf{g} \in C(A,B)$. □

Now it is not hard to show that the circuits are indeed a test set for maximizing (or equivalently minimizing) linear objectives, as already proven in [14].

Corollary 1.6 (See Cor. 3.13 in [14])**.** *Let* $\mathbf{x}^0$ *be a non-optimal solution of*

$$\max\{ \mathbf{c}^\intercal\mathbf{x} : \ A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d},\ \mathbf{x} \in \mathbb{R}^n \} \ .$$

Then there is an augmenting direction $\mathbf{g} \in C(A,B)$ *and some* $\alpha \in \mathbb{R}_+$ *such that* $\mathbf{x}^0 + \alpha\mathbf{g} \in P = \{ \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \}$ *and* $\mathbf{c}^\intercal(\mathbf{x}^0 + \alpha\mathbf{g}) > \mathbf{c}^\intercal(\mathbf{x}^0)$.

Proof. Let $\mathbf{x}^* \in P$ be a better solution of the linear program, that is, $\mathbf{c}^\intercal\mathbf{x}^* > \mathbf{c}^\intercal\mathbf{x}^0$. Then $\mathbf{x}^* - \mathbf{x}^0 \in \ker(A)$ and by Proposition 1.4 it can be written as $\mathbf{x}^* - \mathbf{x}^0 = \sum_{i=1}^k \alpha_i\mathbf{g}^i$, satisfying (i),(ii),(iii). Observe that $B\mathbf{x}^0 \leq \mathbf{d}$ and $B(\mathbf{x}^0 + \sum_{i=1}^k \alpha_i\mathbf{g}^i) = B\mathbf{x}^* \leq \mathbf{d}$ together with sign-compatibility of $B\mathbf{g}^i$ and $B(\mathbf{x}^* - \mathbf{x}^0)$ imply that $B(\mathbf{x}^0 + \alpha_i\mathbf{g}^i) \leq \mathbf{d}$ for $i = 1, \ldots, k$. Thus we have $\mathbf{x}^0 + \alpha_i\mathbf{g}^i \in P$ for all i. We further know that $0 < \mathbf{c}^\intercal(\mathbf{x}^* - \mathbf{x}^0) = \mathbf{c}^\intercal(\sum_{i=1}^k \alpha_i\mathbf{g}^i) = \sum_{i=1}^k \alpha_i\mathbf{c}^\intercal\mathbf{g}^i$ and hence there is some i such that $\mathbf{c}^\intercal\mathbf{g}^i > 0$. Therefore, we found $\mathbf{x}^0 + \alpha_i\mathbf{g}^i$ satisfying the claim. □

This proves the test set property of $C(A, B)$ for linear objectives for any choice of $\mathbf{b}$ and $\mathbf{d}$. It is not hard to show that the circuits are in fact the unique inclusion-minimal set with this property (up to scaling). Note that this result can easily be extended to the minimization of separable convex functions $f\colon \mathbb{R}^n \to \mathbb{R}$, i.e. f is convex and can be written as $f(\mathbf{x}) = \sum_{i=1}^{n} f_i(x_i)$ for convex functions $f_1, \dots, f_n\colon \mathbb{R} \to \mathbb{R}$.

Observe that the circuits can also be used to check whether the linear program is bounded: This is the case if and only if there is no $\mathbf{g} \in C(A, B)$ with $B\mathbf{g} \leq \mathbf{0}$ and $\mathbf{c}^\intercal \mathbf{g} > \mathbf{0}$ (if we maximize); see for example Prop. 3.15 in [14].

Finally, we want to point out that one has to be careful when using circuits in an augmentation procedure as an infinite zig-zagging can occur, see Section 4 in [15]. An example for a strategy that ensures termination is the steepest-descent approach from [10].

1.3 Further characterizations of circuits

The definition of circuits via the representation property is directly connected to a third characterization that is based on minimal generating sets of cones: Denote by B_i the i-th row of B and consider the hyperplanes $\{\, \mathbf{x} \in \mathbb{R}^n : \ B_i\mathbf{x} = \mathbf{0} \,\}$. These hyperplanes partition $\mathbb{R}^n$ into 2^{m_B} rational polyhedral cones of the form $C_\sim = \{\, \mathbf{x} \in \mathbb{R}^n : B\mathbf{x} \sim \mathbf{0} \,\}$, where $\sim \in \{\leq, \geq\}^{m_B}$. All elements in such a cone $C_\sim$ are pairwise sign-compatible. Let $H_\sim$ the unique minimal generating set of $C_\sim \cap \ker(A)$ over $\mathbb{R}$, where the components of each generator are scaled to integers with greatest common divisor one. Then we get the following characterization of the circuits.

Proposition 1.7. *The circuits of A and B are given by*

$$C(A, B) = \bigcup_{\sim \in \{\leq, \geq\}^{m_B}} H_\sim \ .$$

Proof. Let $\mathbf{g} \in C(A, B)$. Then $\mathbf{g} \in C_\sim$ for some $\sim \in \{\leq, \geq\}^{m_B}$ and $\mathbf{g}$ can be written as $\mathbf{g} = \sum_{i=1}^{k} \beta_i \mathbf{h}^i$ for $\beta_i \in \mathbb{R}_+$, $\mathbf{h}^i \in H_\sim$. As all coefficients are positive we must have $\operatorname{supp}(B\mathbf{h}^i) \subseteq \operatorname{supp}(B\mathbf{g})$ for all $i = 1, \dots, k$. As $\mathbf{h}^i \neq \mathbf{0}$, Lemma 1.2 implies that $\mathbf{h}^i = \alpha_i \mathbf{g}$ for $\alpha_i \in \mathbb{R}$, $i = 1, \dots, k$. Thus $\mathbf{g} = \mathbf{h}^i$ for all i as the generators are normalized and hence $\mathbf{g} \in H_\sim$.

Now let $\mathbf{h} \in H_\sim$ for some $\sim \in \{\leq, \geq\}^{m_B}$. Then $\mathbf{h} \in \ker(A)$. By Proposition 1.4 there is a representation $\mathbf{h} = \sum_{i=1}^{k} \alpha_i \mathbf{g}^i$ with (i),(ii),(iii), in particular $\mathbf{g}^i \in C_\sim$ by (ii). Thus $\mathbf{h} = \mathbf{g}^i$ for $i = 1, \dots, k$ as $\mathbf{h}$ belongs to the minimal normalized generating set of $C_\sim \cap \ker(A)$, in particular $\mathbf{h} \in C(A, B)$. □

Our last description of the set of circuits is of a different nature: We show that the circuits coincide with all potential edge directions of the polyhedron $P = \{\, \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \,\}$ that can appear when we fix the matrices A and B, but let the right-hand sides $\mathbf{b}$ and $\mathbf{d}$ vary. In [28] this was proven for the special case $B = I_n$ and [14] states that all these edge directions are contained in $C(A, B)$. We now prove that these two sets are indeed the same. Note that an edge

of the polyhedron $P = \{\, \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \,\}$ is determined by $n - \operatorname{rank}(A) - 1$ linearly independent inequalities that are tight.

Theorem 1.8. *The set of circuits $C(A, B)$ coincides with the edge directions of polyhedra*

$$P = \{\, \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \,\}$$

that can appear for any choice of rational $\mathbf{b}$ *and* $\mathbf{d}$*.*

Proof. Let e be an edge of $P = \{\, \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \,\}$. Let $r := \operatorname{rank}(A)$. Then e is determined by $n - r - 1$ linearly independent inequalities of the system $B\mathbf{x} \leq \mathbf{d}$ that are tight for all elements of e, that is, there is a submatrix $B^e \in \mathbb{R}^{l \times n}$ of B with $\operatorname{rank}(B^e) = n - r - 1$ and a corresponding subvector $\mathbf{d}^e \in \mathbb{R}^l$ of $\mathbf{d}$ such that $\binom{A}{B^e}\mathbf{v} = \binom{\mathbf{b}}{\mathbf{d}^e}$ for all $\mathbf{v} \in e$. Observe that $\operatorname{rank}\binom{A}{B^e} = n - 1$. Let $\mathbf{g}$ be the direction of e (that is, $\mathbf{g} = \mathbf{w} - \mathbf{v}$ for two distinct vectors $\mathbf{v}, \mathbf{w} \in e$). Then $\mathbf{g} \in \ker\binom{A}{B^e}$ and this kernel is one-dimensional. Every element $\mathbf{y} \in \ker(A)\backslash\{\mathbf{0}\}$ with $\operatorname{supp}(B\mathbf{y}) \subseteq \operatorname{supp}(B\mathbf{g})$ must satisfy $B^e\mathbf{y} = 0$. Thus $\mathbf{y} \in \ker\binom{A}{B^e}$, and hence $\mathbf{y} = \lambda\mathbf{g}$ for some $\lambda \in \mathbb{R}$. Therefore, $B\mathbf{g}$ is support-minimal in $\{\, B\mathbf{y} : \mathbf{y} \in \ker(A)\backslash\{\mathbf{0}\} \,\}$.

Now let $\mathbf{g}$ be a circuit. We set $\mathbf{b} := \mathbf{0}$ and define $\mathbf{d}$ as the component-wise maximum of $B\mathbf{g}$ and $\mathbf{0}$. We claim that $P = \{\, \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \,\}$ has an edge of direction $\mathbf{g}$. Clearly $\mathbf{0}, \mathbf{g} \in P$. It remains to show that these two vectors are contained in a common one-dimensional face of P. Let $B^0 \in \mathbb{R}^{l \times n}$ be the submatrix of B that satisfies $B^0\mathbf{g} = \mathbf{0}$, where l is chosen maximal (i.e. $(B\mathbf{g})_i \neq 0$ for all rows i of B not contained in B^0). Observe that the support of $B\mathbf{g}$ corresponds to the rows of B that are not in B^0. Assume $\operatorname{rank}(B^0) < n - r - 1$. Then we can add rows of B to B^0 to obtain a matrix $\bar{B}$ of rank $n - r - 1$. But then $\mathbf{g} \notin \ker\binom{A}{\bar{B}} = \{\lambda\mathbf{x} : \lambda \in \mathbb{R}\}$ for some $\mathbf{x} \neq \mathbf{0}$. Thus we have $\operatorname{supp}(B\mathbf{x}) \subsetneq \operatorname{supp}(B\mathbf{g})$, a contradiction to $\mathbf{g}$ being a circuit. □

We will frequently use this quite intuitive definition when studying *circuit diameters* in the forthcoming chapters.

1.4 Examples

We now give two examples to illustrate the prior definitions and characterizations of circuits, one for each of the special cases $A = O$ and $B = I_n$.

Example 1.9. *Let* $B = \begin{pmatrix} -1 & 0 \\ 2 & -1 \end{pmatrix}$*. The circuits are given by*

$$C_{\leq}(B) = \left\{ \pm\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ 2 \end{pmatrix} \right\}.$$

To see this, we compute the vectors $B\mathbf{g}$ *for all circuits* $\mathbf{g}$*:*

$$B \cdot \pm\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \pm\begin{pmatrix} 0 \\ -1 \end{pmatrix}, \quad B \cdot \pm\begin{pmatrix} 1 \\ 2 \end{pmatrix} = \pm\begin{pmatrix} -1 \\ 0 \end{pmatrix}.$$

Obviously, both products $\pm\binom{0}{-1}$ *and* $\pm\binom{-1}{0}$ *are support-minimal in* $\{ B\mathbf{y} : \mathbf{y} \in \ker(A)\backslash\{\mathbf{0}\} = \mathbb{R}^2\backslash\{\mathbf{0}\} \}$ *and there can be not other vectors of minimal support. Thus the corresponding factors* $\binom{0}{1}$ *and* $\binom{1}{2}$ *constitute the set of circuits as they are integral vectors and their components have greatest common divisor one.*

Now consider the polyhedra

$$P^1 = \left\{ \mathbf{x} \in \mathbb{R}^2 : \ B\mathbf{x} \leq \mathbf{d} \right\}, \quad \textit{where } \mathbf{d} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

and

$$P^2 = \left\{ \mathbf{x} \in \mathbb{R}^2 : \ \mathbf{l} \leq B\mathbf{x} \leq \mathbf{u} \right\}, \quad \textit{where } \mathbf{l} = \begin{pmatrix} -1 \\ -2 \end{pmatrix}, \ \mathbf{u} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

We know that both polyhedra admit the same set of circuits. The polyhedra look as follows:

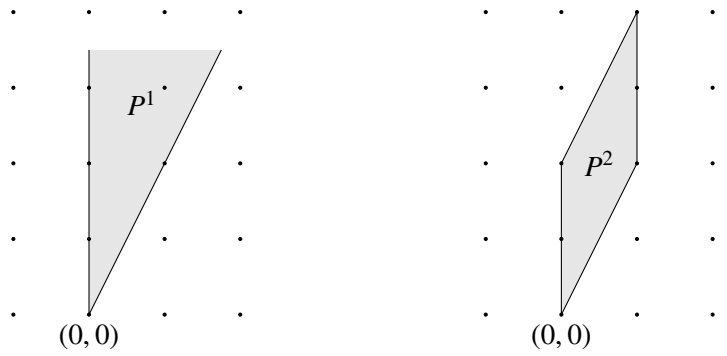

Figure 1.1: Two polyhedra defined by the same matrix B.

P^1 *is an unbounded cone, while* P^2 *is a polytope. We see that the circuits* $\left\{\pm\binom{0}{1}, \ \pm\binom{1}{2}\right\}$ *are indeed the directions of the edges of the polyhedra, as stated in Theorem 1.8.*

Example 1.10. *Let* $A = \begin{pmatrix} -1 & 0 & 1 & 0 \\ 2 & -1 & 0 & 1 \end{pmatrix}$, $B = I_4$. *Then*

$$C(A) = \left\{ \pm \begin{pmatrix} 0 \\ 1 \\ 0 \\ 1 \end{pmatrix}, \pm \begin{pmatrix} 1 \\ 2 \\ 1 \\ 0 \end{pmatrix}, \pm \begin{pmatrix} 1 \\ 0 \\ 1 \\ -2 \end{pmatrix} \right\}.$$

Clearly, all these elements lie in the kernel of A and it is not hard to check that there are no non-zero elements in that kernel of strictly smaller support.

To illustrate the representation property, consider $(-2, 3, -2, 7)^\intercal \in \ker(A)$. *We can express this*

vector as a sign-compatible non-negative linear combination of circuits:

$$\begin{pmatrix} -2 \\ 3 \\ -2 \\ 7 \end{pmatrix} = 3 \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \\ 1 \end{pmatrix} + 2 \cdot \begin{pmatrix} -1 \\ 0 \\ -1 \\ 2 \end{pmatrix}$$

Circuits for equivalent descriptions of polyhedra

We want to point out that the particular representation of a polyhedron affects the set of circuits in the following sense: The polyhedra

$$P = \{ \mathbf{x} \in \mathbb{R}^n : B\mathbf{x} \leq \mathbf{d} \}$$

and

$$P' = \left\{ \begin{pmatrix} \mathbf{x} \\ \mathbf{x}' \end{pmatrix} \in \mathbb{R}^{n+m_B} : \begin{pmatrix} B & I_{m_B} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{x}' \end{pmatrix} = \mathbf{d},\ \mathbf{x}' \geq \mathbf{0} \right\}$$

are equivalent, but this is not necessarily true for the associated circuits $C_{\leq}(B)$ and $C(A)$, where $A = \begin{pmatrix} B & I_{m_B} \end{pmatrix}$ as Example 1.11 shows. More precisely, we will see that $C(A)$ can contain additional elements, as the circuits associated with A allow for more elements of minimal support.

Example 1.11. *Observe that the matrix B in Example 1.9 and the matrix A in Example 1.10 are indeed related via $A = \begin{pmatrix} B & I_2 \end{pmatrix}$. But $C_{\leq}(B)$ contains four vectors while $C(A)$ consists of six elements. Hence the sets of circuits are not equivalent.*

However, one can deduce the set $C_{\leq}(B)$ from the circuits of $A = \begin{pmatrix} B & I_{m_B} \end{pmatrix}$.

Lemma 1.12. *Let $B \in \mathbb{Z}^{m_B \times n}$, $A := \begin{pmatrix} B & I_{m_B} \end{pmatrix}$ and $\mathbf{g} \in \mathbb{R}^n$. Then $\mathbf{g} \in C_{\leq}(B)$ if and only if there is some $\begin{pmatrix} \mathbf{g} \\ \mathbf{g}' \end{pmatrix} \in C(A)$ such that $\mathbf{g}'$ is support-minimal in the set $\left\{ \mathbf{y}' : \begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \in C(A) \text{ for some } \mathbf{y} \in \mathbb{R}^n \right\}$.*

Proof. Let $\mathbf{g} \in C_{\leq}(B)$. Then $B\mathbf{g}$ is support-minimal $\{B\mathbf{y} : \mathbf{y} \in \mathbb{R}^n\} = \{-B\mathbf{y} : \mathbf{y} \in \mathbb{R}^n\}$ and $\begin{pmatrix} \mathbf{g} \\ -B\mathbf{g} \end{pmatrix} \in \ker(A)$. Let $\mathbf{g}' := -B\mathbf{g}$. Then $\begin{pmatrix} \mathbf{g} \\ \mathbf{g}' \end{pmatrix} \in \ker(A)$ and $\mathbf{g}'$ is support-minimal in $\{\mathbf{y}' : \begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \in \ker(A)\backslash\{\mathbf{0}\}$ for some $\mathbf{y} \in \mathbb{R}^n\}$. But then also $\begin{pmatrix} \mathbf{g} \\ \mathbf{g}' \end{pmatrix}$ is support-minimal in $\ker(A)\backslash\{\mathbf{0}\}$ and hence $\begin{pmatrix} \mathbf{g} \\ \mathbf{g}' \end{pmatrix} \in C(A)$. Further, as $\mathbf{g}'$ is support-minimal in $\{\mathbf{y}' : \begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \in \ker(A)\backslash\{\mathbf{0}\}$ for some $\mathbf{y} \in \mathbb{R}^n\}$, it is also support-minimal in the subset $\{\mathbf{y}' : \begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \in C(A)\}$.

Now let $\begin{pmatrix} \mathbf{g} \\ \mathbf{g}' \end{pmatrix} \in C(A)$ such that $\mathbf{g}'$ is support-minimal in the set $\left\{ \mathbf{y}' : \begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \in C(A) \text{ for some } \mathbf{y} \in \mathbb{R}^n \right\}$. For $\mathbf{g}$ being a circuit it is enough to show that $\mathbf{g}' = -B\mathbf{g}$ is support-minimal in $\{-B\mathbf{y} : \mathbf{y} \in \mathbb{R}^n\backslash\{\mathbf{0}\}\}$. Assume that $\mathbf{g}'$ is not support-minimal in this set. Then there is some $\mathbf{y} \in \mathbb{R}^n\backslash\{\mathbf{0}\}$ such that $\operatorname{supp}(\mathbf{y}') \subsetneq \operatorname{supp}(\mathbf{g}')$ for $\mathbf{y}' := -B\mathbf{y}$. By case-assumption, $\begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \notin C(A)$, but $\begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \in \ker(A)\backslash\{\mathbf{0}\}$ by definition of $\mathbf{y}'$. Hence there is $\begin{pmatrix} \mathbf{x} \\ \mathbf{x}' \end{pmatrix} \in C(A)$ such that $\operatorname{supp}(\mathbf{x}) \subseteq \operatorname{supp}(\mathbf{y})$ and $\operatorname{supp}(\mathbf{x}') \subseteq \operatorname{supp}(\mathbf{y}') \subsetneq \operatorname{supp}(\mathbf{g}')$, that is, $\mathbf{g}'$ is not support-minimal in $\left\{ \mathbf{y}' : \begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \in C(A) \text{ for some } \mathbf{y} \in \mathbb{R}^n \right\}$, a contradiction. □

Example 1.11 cont'd. *To deduce $C_{\leq}(B)$ from $C(A)$, we have to look for support-minimality in the 'slack', i.e. within the vectors formed by the third and forth components of the elements in $C(A)$. These are*

$$\left\{ \mathbf{y}' : \begin{pmatrix} \mathbf{y} \\ \mathbf{y}' \end{pmatrix} \in C(A) \text{ for some } \mathbf{y} \in \mathbb{R}^2 \right\} = \left\{ \pm \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \pm \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \pm \begin{pmatrix} 1 \\ -2 \end{pmatrix} \right\}.$$

The first two elements are support-minimal in this set. Hence the vectors formed by the first and second components of the corresponding vectors in $C(A)$, $\pm \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and $\pm \begin{pmatrix} 1 \\ 2 \end{pmatrix}$, constitute the set $C_{\leq}(B)$ as already shown in Example 1.9.

Outlook

In the next Chapter 2 we will introduce a hierarchy of circuit distances. These distances are based on circuit walks, which connect two vertices of a polyhedron by a number of steps along circuit directions. Proposition 1.4 tells us that there always is such a walk, even a sign-compatible one. Further, we know that we can reach any vertex of the polyhedron from any point within the polyhedron along circuits by Corollary 1.6. All these results will be important for the upcoming investigations. The characterization of circuits as edge directions from Theorem 1.8 will be stressed in Chapter 3 and 4, where we study the circuit distances of transportation polytopes and dual network flow polyhedra, respectively.

Chapter 2

A hierarchy of circuit distances

We now come to the key part of the present thesis: In this chapter we introduce the different categories of circuit diameters. Most of these classes are less restrictive than the combinatorial diameter: Instead of only going along the edges of the polyhedron we liberate ourselves from these restrictions. We try to go from one vertex of the polyhedron to another one more efficiently by possibly going through the interior of the polyhedron, along linear combinations of circuits. We are even willing to leave the feasible region.

As already outlined in the introduction, the purpose of this approach is two-fold: On the one hand we have the intimate connection of circuit distances to a best-case performance of augmentation algorithms along circuit directions. In particular, this could serve as an indication on how to use circuits in such algorithms. Recall that augmentation algorithms are a generalization of the Simplex method.

On the other hand, we hope that the circuit diameters as a generalization to the classical combinatorial diameter might shed some light on this difficult field of research and help us for a better understanding. In particular, the combinatorial diameter is bounded below by diameters that have much weaker properties and therefore may be much easier to bound or to compute. We will exploit this in Chapter 3 for transportation polytopes and in Chapter 4 for dual network flow polyhedra.

This chapter is based on joint work with Steffen Borgwardt, Jesús A. De Loera and Raymond Hemmecke, published in [3] and [5]. It is structured the following way: In Section 2.1 we begin with a formal introduction of circuit walks, circuit distances, and circuit diameters, along with an overview of their characteristics. In Section 2.2 we then present the core result, the hierarchy of circuit distances, followed by the proofs of the relations claimed therein. In Section 2.3 we have a look at general upper bounds for circuit diameters. In particular we will see that many of the diameter categories satisfy the bound induced by the Hirsch conjecture. We conclude the chapter by investigating the hierarchy in dimension two in Section 2.4. In this special case many of the circuit diameter concepts coincide such that the hierarchy ‘collapses’ into only few distinct categories.

2.1 Definitions and preliminaries

In this section we first define a very general notion of *distance* based on circuits. We then introduce additional properties for these walks, which leads to a variety of circuit distances. After explaining these categories, we summarize some basic characteristics of circuit walks and circuit distances.

2.1.1 Circuit walks

Let $P = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}, B\mathbf{x} \leq \mathbf{d}\}$ be a polyhedron and let $C(A, B)$ be the set of circuits associated with the matrices A and B. For a pair of vertices $\mathbf{v}^{(1)}$, $\mathbf{v}^{(2)}$ of P, we call a sequence $\mathbf{v}^{(1)} = \mathbf{y}^{(0)}, \ldots, \mathbf{y}^{(k)} = \mathbf{v}^{(2)}$ a *circuit walk of length k* if for all $i = 0, \ldots, k-1$ we have $\mathbf{y}^{(i+1)} - \mathbf{y}^{(i)} = \alpha_i \mathbf{g}^i$ for some circuit $\mathbf{g}^i \in C(A, B)$ and some $\alpha_i > 0$. Note that because we are allowing the α_i to be arbitrary real non-negative numbers, circuit walks can be infinite, but we restrict our attention to those that are finite. We are particularly interested in circuit walks that satisfy certain combinations of additional properties:

(e) If $\mathbf{y}^{(i)}$ and $\mathbf{y}^{(i+1)}$ are neighboring vertices of the polyhedron (vertices connected by an edge) for all $i = 0, \ldots, k-1$, we call the walk an *edge walk*. This is the term that corresponds to the classical combinatorial diameter.

(m) If the extension multipliers α_i are maximal, i.e. if $\mathbf{y}^{(i)} + \alpha \mathbf{g}^i$ is infeasible for all $\alpha > \alpha_i$, we say that the walk is of *maximum extension length* or simply *maximal*. Otherwise, we say that the extension is of *arbitrary length*.

(f) If $\mathbf{y}^{(i)} \in P$ for all $i = 0, \ldots, k-1$, then we say the circuit walk is *feasible*.

(r) If no circuit is repeated, then we say the walk is *non-repetitive*.

(b) If no pair of circuits $\mathbf{g}^i$, $-\mathbf{g}^i$ is used, we say the walk is *non-backwards*.

(s) If all the circuits are pairwise sign-compatible and are all sign-compatible with the vector $\mathbf{v}^{(2)} - \mathbf{v}^{(1)}$, we say the circuit walk is *sign-compatible*. Recall that two vectors $\mathbf{x}$ and $\mathbf{y}$ are sign-compatible (with respect to the matrix B), if $(B\mathbf{x})_i \cdot (B\mathbf{y})_i \geq 0$ for all $i = 1, \ldots, m_B$.

2.1.2 Categories of circuit distances and circuit diameters

We now define the *circuit distance* from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ as the minimum length of a circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$. A circuit walk that realizes the circuit distance is called a *shortest* or *optimal* walk. The *circuit diameter* of a polyhedron P is the maximum circuit distance between any two vertices of P. Considering different types of circuit walks yields different categories of circuit distances and diameters.

The main goal of this chapter is to relate these distances to each other. We now introduce a notation for this discussion: $\mathcal{CD}$ refers to the circuit distance defined on circuit walks without any further restrictions. We call it the *soft circuit distance*. When we consider circuit distances with

respect to circuit walks that satisfy some of the properties stated in Section 2.1.1, we denote this by small subscript letters as used in the above list. For example, the classical *combinatorial diameter* is defined on circuit walks that always go along an edge to the next vertex. These are maximal steps that always remain feasible, so we write $C\mathcal{D}_{efm}$ for the combinatorial distance. By dropping the edge walk property, we obtain $C\mathcal{D}_{fm}$. This is the *circuit distance* we originally introduced in [5]. We call the category $C\mathcal{D}_f$, defined on feasible circuit walks with steps of arbitrary length, the *weak circuit distance*. These four types of circuit distances will form the 'key chain' in the hierarchy of distances we introduce in Section 2.2.

Note that circuit distances on which we impose more restrictions trivially bound less restrictive ones from above. We denote this with a weak-inequality sign. For example, the 'key chain' looks as follows:

$$C\mathcal{D}_{efm} \;\geq\; C\mathcal{D}_{fm} \;\geq\; C\mathcal{D}_f \;\geq\; C\mathcal{D}\,.$$

If we have '$\geq$' and there is a polyhedron with a pair vertices such the respective circuits distances actually differ, we use '$>$'. For example, we will see that $C\mathcal{D}_{fm} > C\mathcal{D}_f$. Note that this notation is transitive: Clearly $C\mathcal{D}_{fm} \geq C\mathcal{D}_f \geq C\mathcal{D}$ implies $C\mathcal{D}_{fm} \geq C\mathcal{D}$ and $C\mathcal{D}_{fm} > C\mathcal{D}_f \geq C\mathcal{D}$ implies $C\mathcal{D}_{fm} > C\mathcal{D}$. We sometimes want to refer to several combinations at the same time. We then use brackets to emphasize this, for example $C\mathcal{D}_{fm(b)} \geq C\mathcal{D}_{f(b)}$ means that both statements hold, $C\mathcal{D}_{fmb} \geq C\mathcal{D}_{fb}$ and $C\mathcal{D}_{fm} \geq C\mathcal{D}_f$.

Note that we may equally use this '$C\mathcal{D}$-notation' to refer to the circuit *diameter* of a polyhedron as the maximum distance over all pairs of vertices.

Theoretically, there are many different categories of circuit distances that arise from combining the properties above. However, we will not consider every such combination in our hierarchy in Section 2.2 for two reasons: Some of these combinations are not well-defined and there are combinations that actually refer to the same concepts.

For the latter, note that certain combinations of properties already imply other properties. For example, every edge walk is feasible and maximal, that is, $C\mathcal{D}_e = C\mathcal{D}_{efm}$, the classes coincide. Thus one can use both notations equally. However, in this chapter we stick to $C\mathcal{D}_{efm}$ to keep all the properties in sight. Similarly, maximal circuit walks are always feasible and thus $C\mathcal{D}_m = C\mathcal{D}_{fm}$. In Lemma 2.14 we will see that same holds for sign-compatible circuit walks, that is, $C\mathcal{D}_s = C\mathcal{D}_{fs}$. Note that we even have $C\mathcal{D}_{e(b)(r)} = C\mathcal{D}_{efm(b)(r)}$ and $C\mathcal{D}_{m(b)(r)} = C\mathcal{D}_{fm(b)(r)}$.

Circuit distance categories can also coincide in the sense that the *optimal* circuit walks satisfy additional properties. For example, every optimal circuit walk without any restrictions is non-backwards and non-repetitive as we could combine opposite or repeated steps to one single circuit step, yielding a shorter circuit walk. Thus we have $C\mathcal{D} = C\mathcal{D}_b = C\mathcal{D}_r$. Again by Lemma 2.14 this is also true for sign-compatible walks, so we have $C\mathcal{D}_{fs} = C\mathcal{D}_{fsbr} = C\mathcal{D}_s = C\mathcal{D}_{sbr}$.

Next we come to combinations of properties that are not well-defined. This happens if we require sign-compatibility for edge walks or for feasible maximal walks.

Lemma 2.1. *In dimension $n = 2$, there is a polytope with a pair of vertices such that there is no*

feasible maximal sign-compatible circuit walk from one vertex to the other one. In particular there is no sign-compatible edge walk.

Proof. Consider the polytope

$$P = \left\{ \mathbf{x} \in \mathbb{R}^2 : \mathbf{l} \leq B\mathbf{x} \leq \mathbf{u} \right\}$$

defined by

$$B = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & -1 \\ 1 & -2 \end{pmatrix}, \quad \mathbf{l} = \begin{pmatrix} 0 \\ 0 \\ -1 \\ -3 \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} \infty \\ 6 \\ 4 \\ \infty \end{pmatrix}.$$

All possible edge directions $\mathbf{g}$ of P (the circuits) are given by

$$\pm\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ -1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \pm\begin{pmatrix} 2 \\ 1 \end{pmatrix},$$

and the corresponding vectors $B\mathbf{g}$ are

$$\pm\begin{pmatrix} 0 \\ 1 \\ -1 \\ -2 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ 0 \\ 2 \\ 3 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ 2 \\ 0 \\ -1 \end{pmatrix}, \pm\begin{pmatrix} 2 \\ 3 \\ 1 \\ 0 \end{pmatrix}.$$

We want to perform a circuit walk from $\mathbf{v}^{(1)} = (2, -2)^\intercal$ to $\mathbf{v}^{(2)} = (1, 2)^\intercal$. We have $B \cdot \left(\mathbf{v}^{(2)} - \mathbf{v}^{(1)}\right) = (-1, 3, -5, -6)^\intercal$. The only sign-compatible circuits are $(0, 1)^\intercal$ (as $B \cdot (0, 1)^\intercal = (0, 1, -1, -2)^\intercal$) and $(-1, 1)^\intercal$ (as $B \cdot (-1, 1)^\intercal = (-1, 0, 2, 3)^\intercal$). But choosing the directions $(0, 1)^\intercal$ and $(-1, 1)^\intercal$, respectively, as a first feasible maximal circuit step at $\mathbf{v}^{(1)}$ we arrive at $(0, 0)^\intercal$ and $(2, 2\frac{1}{2})^\intercal$, respectively. From neither of these points we can go to $\mathbf{v}^{(2)}$ using only the circuits $(0, 1)^\intercal$ and $(-1, 1)^\intercal$, which are the only ones that are sign-compatible with $\mathbf{v}^{(2)} - \mathbf{v}^{(1)}$.

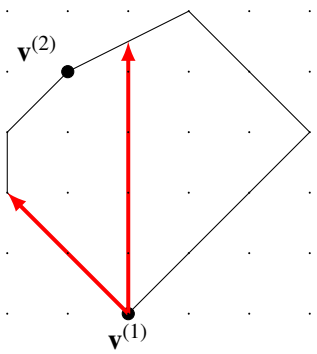

Figure 2.1: All feasible maximal circuit steps at $\mathbf{v}^{(1)}$ that are sign-compatible with $\mathbf{v}^{(2)} - \mathbf{v}^{(1)}$.

□

In contrast, we already know from our discussion in the previous chapter that any two vertices of a polyhedron $P = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}, B\mathbf{x} \leq \mathbf{d}\}$ are connected by a feasible sign-compatible circuit walk if we allow arbitrary length steps, and thus $C\mathcal{D}_{fs}$ is well-defined (see Proposition 1.4). In Lemma 2.18 we will further prove that there are good upper bounds on the length of such walks.

2.1.3 Fundamental observations

First of all, let us point out that there are classes of polyhedra for which all notions of circuit diameters coincide: In simplices any two vertices are connected by an edge and thus all circuit diameters are just one. Further, for any n-dimensional zonotope, all circuit diameters equal n; the n-dimensional cube is a particularly simple special case. Recall that a zonotope is point-symmetric with respect to its center of gravity. Vertices that correspond to each other with respect to this point symmetry are connected by an edge walk of length exactly n. Using any set of circuits and no restrictions on the walk we cannot do any better, as the circuits here correspond to the actual, existing edge directions.

We continue with some elementary characteristics of circuit walks and circuit distances.

Reversible circuit walks

Given a circuit walk from a vertex $\mathbf{v}^{(1)}$ to a vertex $\mathbf{v}^{(2)}$ we get a circuit walk from $\mathbf{v}^{(2)}$ to $\mathbf{v}^{(1)}$ by simply reversing the steps. This obviously preserves the properties feasible **(f)**, non-repetitive **(r)**, non-backwards **(b)** and sign-compatible **(s)**, and it turns an edge walk into another edge walk **(e)**. Thus, circuit walks of all categories except $C\mathcal{D}_{fm(b)(r)}$ are reversible. In particular the corresponding distances are *symmetric*, that is, a certain circuit distance from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ coincides with the respective circuit distance from $\mathbf{v}^{(2)}$ to $\mathbf{v}^{(1)}$. But if we do not have an edge walk, maximality **(m)** is not necessarily maintained and thus circuit walks of type $C\mathcal{D}_{fm(b)(r)}$ are not always reversible, neither is the distance symmetric, as illustrated in the following example.

Example 2.2. *Consider the polyhedron*

$$P = \{\mathbf{x} \in \mathbb{R}^2 : \mathbf{l} \leq B\mathbf{x} \leq \mathbf{u}\}$$

given by

$$B = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad \mathbf{l} = \begin{pmatrix} 0 \\ -\infty \\ 0 \\ -1 \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} \infty \\ 2 \\ 4 \\ 6 \end{pmatrix}.$$

P is a two-dimensional polytope with six vertices whose circuits are given by

$$C_{\leq}(B) = \left\{ \pm\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ -1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\},$$

as these are obviously all possible edge directions (the polytope is depicted in Figure 2.2). Let us have a look at circuit walks between $\mathbf{v}^{(1)}$ *and* $\mathbf{v}^{(2)}$.

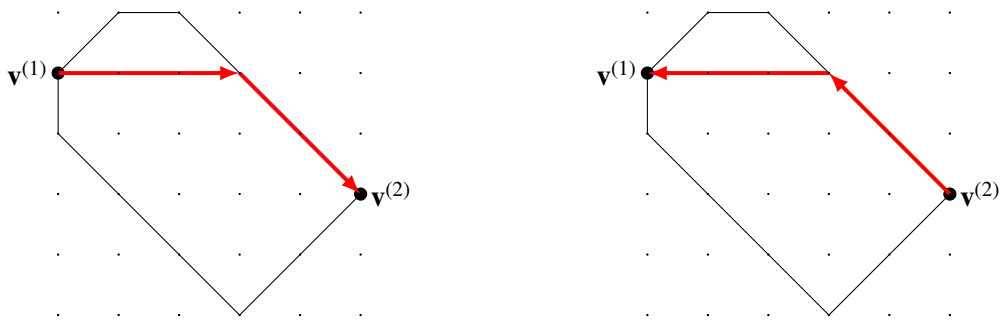

Figure 2.2: An optimal circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ and the reversed walk.

The walk from $\mathbf{v}^{(1)}$ *to* $\mathbf{v}^{(2)}$ *is maximal, but the reversed walk is not. We can even show that there is no feasible maximal circuit walk of length two from* $\mathbf{v}^{(2)}$ *to* $\mathbf{v}^{(1)}$ *and thus the circuit distances* $\mathcal{CD}_{fm}$ *do not coincide. To see this, we illustrate all maximal feasible circuit steps that could be applied at* $\mathbf{v}^{(2)}$.

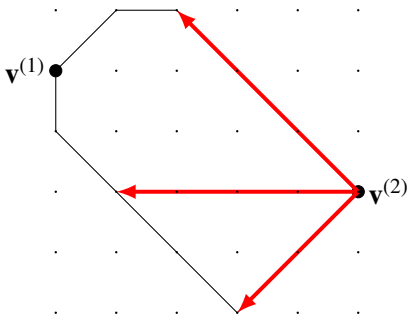

Figure 2.3: Feasible maximal circuit steps at $\mathbf{v}^{(2)}$.

No matter which circuit direction we choose for a first maximal step starting at $\mathbf{v}^{(2)}$*, we cannot go to* $\mathbf{v}^{(1)}$ *with only one more step.*

Commutative circuit walks

Given a circuit walk from a vertex $\mathbf{v}^{(1)}$ to a vertex $\mathbf{v}^{(2)}$ we can construct alternative circuit walks from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ by rearranging the circuit steps. We say that a circuit walk is *commutative* if it does not matter in which order we apply the steps; all circuit walks obtained by rearranging fall into the same category as the original one.

While circuit walks in general have to be regarded as ordered sequences of vectors, commutative walks can be interpreted simply as linear combinations of the form $\mathbf{v}^{(2)} - \mathbf{v}^{(1)} = \sum_{i=1}^{k} \alpha_i \mathbf{g}^i$ for circuits $\mathbf{g}^i$. In particular, the non-restricted walks of type $\mathcal{CD}$ are commutative and thus, in a way, the soft circuit distance $\mathcal{CD}$ is just a linear algebra bound on the diameter given in terms of the support (or number of elements) of a linear combination of circuits.

It is not hard to see that almost all other types of circuit walks are not commutative as in general rearranging does not preserve feasibility. However, this is the case for sign-compatible circuit walks: It follows immediately from the representation $\mathbf{v}^{(2)} = \mathbf{v}^{(1)} + \sum_{i=1}^{k} \alpha_i \mathbf{g}^i$ for circuits $\mathbf{g}^i$ that are sign-compatible with $\mathbf{v}^{(2)} - \mathbf{v}^{(1)}$ and thus satisfy $\operatorname{supp}(B\mathbf{g}^i) \subseteq \operatorname{supp}(B(\mathbf{v}^{(2)} - \mathbf{v}^{(1)}))$.

This significant observation shows that sign-compatibly is a natural and extremely useful property for circuit walks. For example, in [18] sign-compatible walks play a crucial role in showing that there is a selection strategy such that only polynomially many circuit augmentation steps are needed to reach an optimal solution (recall that this is still unresolved for the Simplex method). However, it is still open how to implement this greedy-type augmentation oracle in polynomial time.

Perturbing the right-hand side

Recall that it suffices to consider generic polyhedra to bound the combinatorial diameter, as by a perturbation any polyhedron can be turned into a generic one whose diameter is at least as big as that of the original polyhedron. It is not clear whether the same is true for the other notions circuit diameters.

In the following example we will see that a perturbation of the right-hand side vector might affect the circuit diameter in a way that is hard to predict. In particular, it may change the circuit diameter while keeping the combinatorial structure of the polyhedron. Note that such a perturbation does not change the set of circuits.

Example 2.3. *Consider the polyhedron*

$$\tilde{P} = \{\, \mathbf{x} \in \mathbb{R}^2 : \mathbf{l} \leq B\mathbf{x} \leq \tilde{\mathbf{u}} \,\}$$

given by

$$B = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad \mathbf{l} = \begin{pmatrix} 0 \\ -\infty \\ 0 \\ -1 \end{pmatrix}, \quad \tilde{\mathbf{u}} = \begin{pmatrix} \infty \\ 2 \\ 4 \\ 4 \end{pmatrix},$$

and the polyhedron P as defined in Example 2.2. As P and $\tilde{P}$ are both defined by the matrix B, these polyhedra possess the same sets of circuits. Further they have identical combinatorial structure:

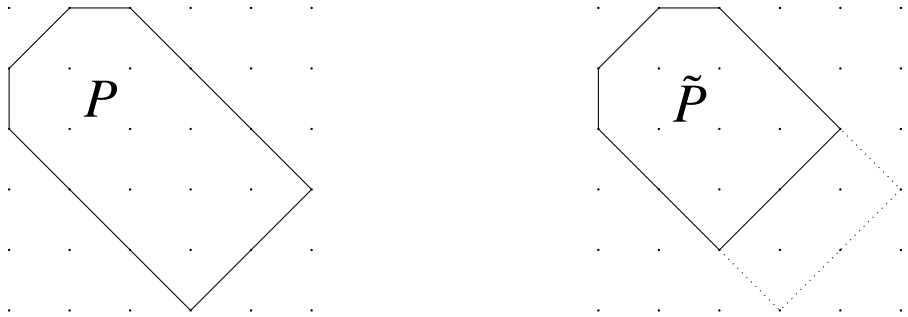

Figure 2.4: Two polyhedra with same combinatorial structure.

As already seen in Example 2.2, the circuit diameter CD_{fm} of P equals three. In contrast, it can easily be checked that in $\tilde{P}$ any two vertices can be connected by a feasible maximal circuit walk of length at most two.

2.2 The circuit distance hierarchy

We now turn to the main result of this chapter - our hierarchy of circuit distances. Observe that the results on the circuit distances readily transfer to statements about the circuit diameters of polyhedra. Thus, we can equally view it as a hierarchy of circuit diameters.

Theorem 2.4. *The circuit distances satisfy a hierarchy as depicted in Figure 2.5.*

This hierarchy states the relations between certain categories of circuit distances that arise from combining the properties introduced in the previous section. We want to point out that these categories are indeed a viable choice: We are able to prove that almost all circuit distances in the hierarchy are distinct (as we have '>') and our discussion in Section 2.1.2 showed that we included all relevant classes. Note that for sake of having a clear layout the lower left and lower right parts refer to the same categories CD_{fbr} and CD_{fs}. Distance notions that share properties are grouped in boxes.

The remaining section is concerned with the proof of Theorem 2.4. We have to show that the distances indeed satisfy the claimed relations. For most of the 'weak' inequalities this follows immediately as we are just imposing additional constraints. In order to prove the 'strict' inequalities, we exhibit polytopes with pairs of vertices $v^{(1)}$, $v^{(2)}$ for which the lengths of optimal walks with the respective properties indeed differ. The numbers at the '<' -symbols refer to the lemmas that show these strict inequalities.

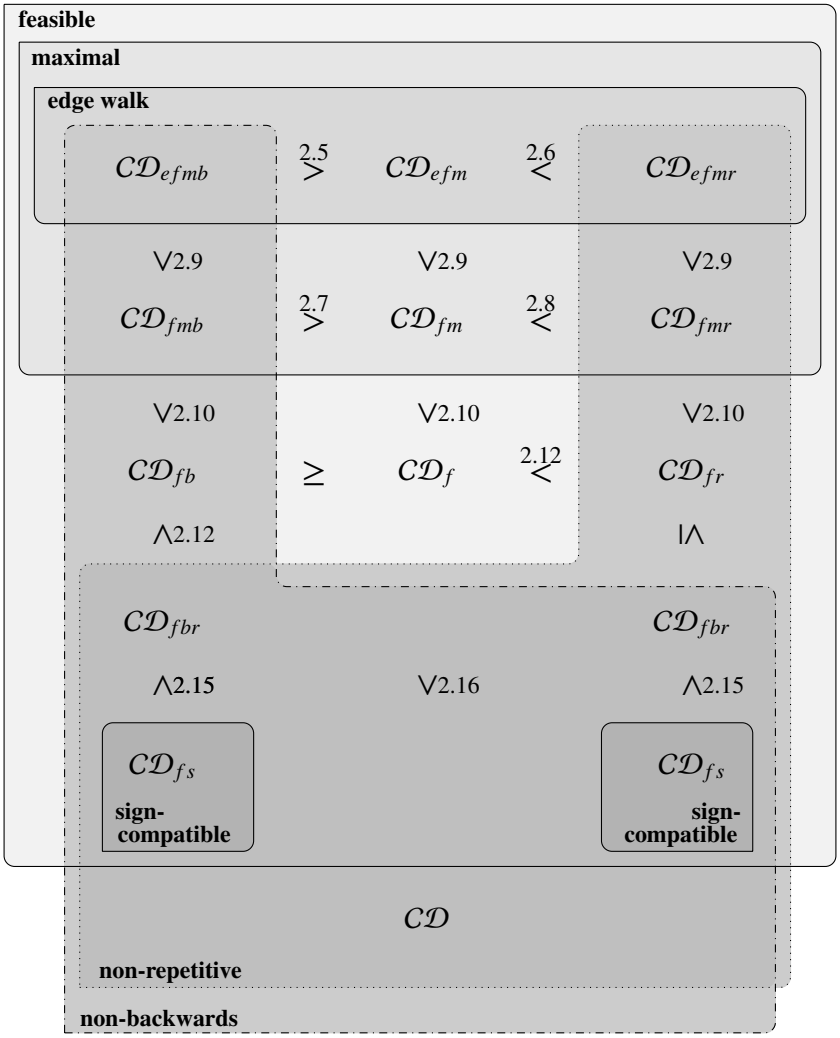

Figure 2.5: A hierarchy of circuit distances.

The top horizontal layer of the table contains the distances defined on edge walks. An edge walk always is both feasible and maximal, so there are only combinations that contain all of these properties at the same time. We distinguish between $\mathcal{CD}_{efm}$ and $\mathcal{CD}_{efmb}, \mathcal{CD}_{efmr}$. Recall that $\mathcal{CD}_{efms}$ is not included in the hierarchy as it is not necessarily well-defined, see Lemma 2.1. Same holds for $\mathcal{CD}_{fms}$. As we just impose additional constraints we immediately get $\mathcal{CD}_{efmb} \geq \mathcal{CD}_{efm}$ and $\mathcal{CD}_{efmr} \geq \mathcal{CD}_{efm}$. The following two Lemmas show that allowing the opposite or the repeated use of an edge direction can indeed yield a shorter edge walk and thus we have strict inequalities.

Lemma 2.5 ($\mathcal{CD}_{emfb} > \mathcal{CD}_{efm}$)**.** *For $n = 2$, there is a polytope with a pair of vertices for which the unique optimal edge walk is backwards. Hence the distances $\mathcal{CD}_{efmb}$ and $\mathcal{CD}_{efm}$ differ in this case.*

Proof. In the polytope below, the unique non-backwards edge walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ has length four, while there is an edge walk of length three that uses edges in opposite directions.

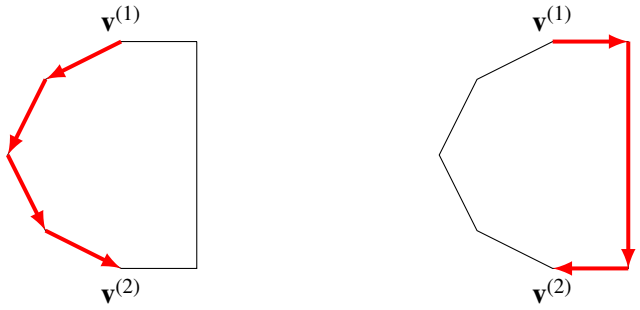

Figure 2.6: An optimal non-backwards edge walk and a backwards edge walk.

□

Lemma 2.6 ($\mathcal{CD}_{efmr} > \mathcal{CD}_{efm}$)**.** *For $n = 3$, there is a polytope with a pair of vertices for which the unique optimal edge walk is repetitive. Hence the distances $\mathcal{CD}_{efmr}$ and $\mathcal{CD}_{efm}$ differ in this case.*

Proof. We construct a polytope with the claimed property by cutting off vertices of a three-dimensional cube as illustrated in the following figures:

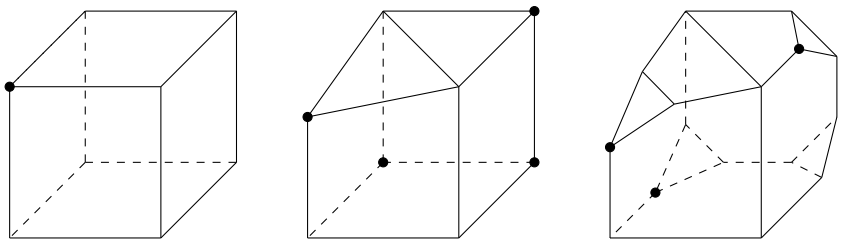

Figure 2.7: Constructing the polytope by cutting off vertices (marked with dots).

We obtain the polytope below, in which there is a repetitive edge walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ of length four. It is easy to check that any other edge walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ has length at least five.

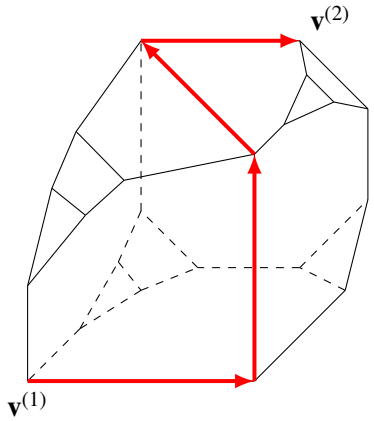

Figure 2.8: Unique optimal edge walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$.

□

We now relax the constraint **(e)** and allow circuit walks through the interior of the polyhedron. For feasible maximal circuit walks we again distinguish between $\mathcal{CD}_{fm}$ and $\mathcal{CD}_{fmb}$, $\mathcal{CD}_{fmr}$. In Lemmas 2.7 and 2.8 we prove that these concepts do not coincide. We first exhibit a polytope in which going backwards yields shorter circuit walks, that is, $\mathcal{CD}_{fm} < \mathcal{CD}_{fmb}$.

Lemma 2.7 ($\mathcal{CD}_{fmb} > \mathcal{CD}_{fm}$)**.** *For $n = 2$, there is a polytope with a pair of vertices for which every optimal feasible maximal circuit walk is backwards. Hence the distances $\mathcal{CD}_{fmb}$ and $\mathcal{CD}_{fm}$ differ in this case.*

Proof. We consider the polytope on 11 vertices depicted in Figure 2.9; the lower subfigure is a zoomed-in view on the right part of the polygon. The edge directions are given by

$$\begin{pmatrix}1\\1\end{pmatrix}, \begin{pmatrix}1\\0\end{pmatrix}, \begin{pmatrix}-1\\1\end{pmatrix}, \begin{pmatrix}3\\-10\end{pmatrix}, \begin{pmatrix}2\\-10\end{pmatrix}, \begin{pmatrix}1\\-10\end{pmatrix}, \begin{pmatrix}1\\10\end{pmatrix}, \begin{pmatrix}2\\10\end{pmatrix}, \begin{pmatrix}3\\10\end{pmatrix}.$$

There is a feasible maximal circuit walk of length three from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(5)}$ that is backwards.

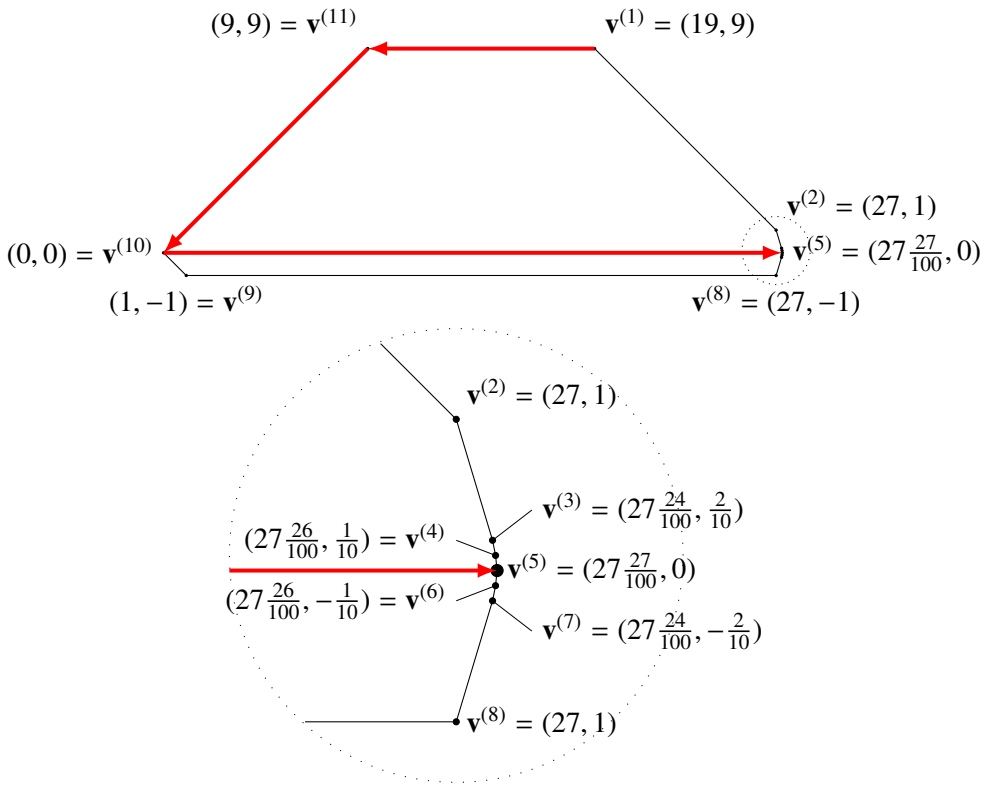

Figure 2.9: A polytope with a feasible maximal backwards circuit walk of length three.

Every other feasible maximal circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(5)}$ has length at least four. To see this, we illustrate all possible combination of first (dashed) and second feasible maximal circuit steps in Figure 2.10. From none of these second step points we can reach $\mathbf{v}^{(5)}$ in only one additional step, except from the point $\mathbf{v}^{(10)}$ in the top left picture. But this yields the backwards circuit walk

depicted in Figure 2.9. Observe that all second steps that end in the upper edge $(\mathbf{v}^{(11)}, \mathbf{v}^{(1)})$ have coordinates $(x, 9)^\intercal$ for an *integral* x, and thus we cannot reach $\mathbf{v}^{(5)} = (27\frac{27}{100}, 0)^\intercal$ by applying the circuit $(-1, 1)^\intercal$ at these points. The final sketch is a zoomed-in view on the bottom right picture. It illustrates all possible second steps after applying $(-1, 1)^\intercal$ at $\mathbf{v}^{(1)}$ for a convenient verification of the fact that we cannot get to $\mathbf{v}^{(5)}$ with only one more step.

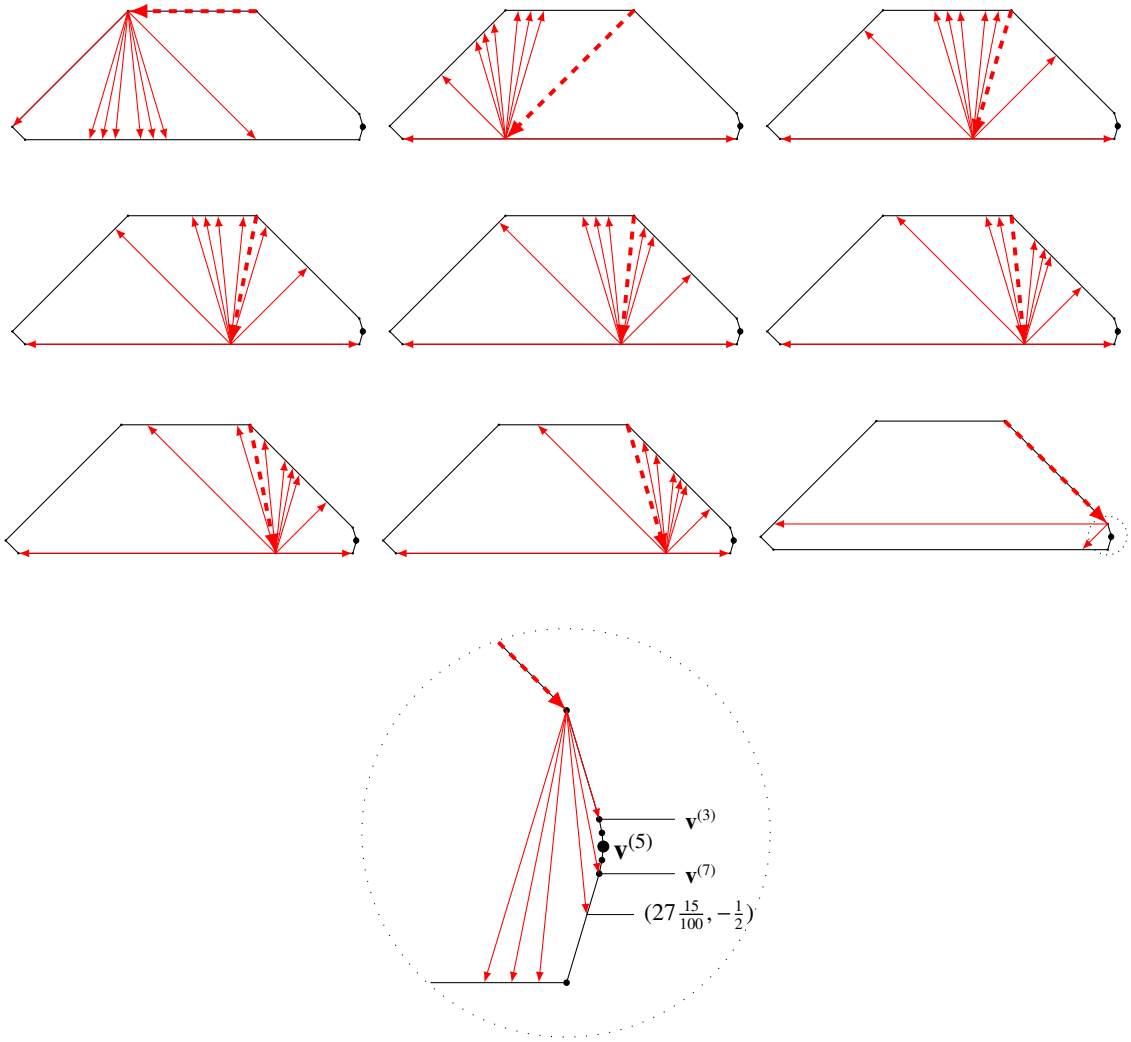

Figure 2.10: Possible combinations of first and second feasible maximal circuit steps from $\mathbf{v}^{(1)}$.

□

Likewise, one may obtain a shorter circuit walk by allowing oneself to use a circuit twice, and thus $\mathcal{CD}_{fm} < \mathcal{CD}_{fmr}$.

Lemma 2.8 ($\mathcal{CD}_{fmr} > \mathcal{CD}_{fm}$)**.** *For $n = 2$, there is a polytope with a pair of vertices for which every optimal feasible maximal circuit walk is repetitive. Hence the distances $\mathcal{CD}_{fmr}$ and $\mathcal{CD}_{fm}$ differ in this case.*

Proof. We consider the polytope on nine vertices depicted in Figure 2.11. Note that there are two edges e_0 and e_7 with direction $(1,0)^\intercal$, an edge e_1 with direction $(1,-1)^\intercal$, an edge e_6 with direction $(1,1)^\intercal$, and the edge e_8 with direction $(0,1)^\intercal$. Further, in the right part there are four steeper edges: e_2 with direction $(1,-4)^\intercal$, e_3 with direction $(1,-5)^\intercal$, e_4 with direction $(1,5)^\intercal$, e_5 with direction $(1,4)^\intercal$.

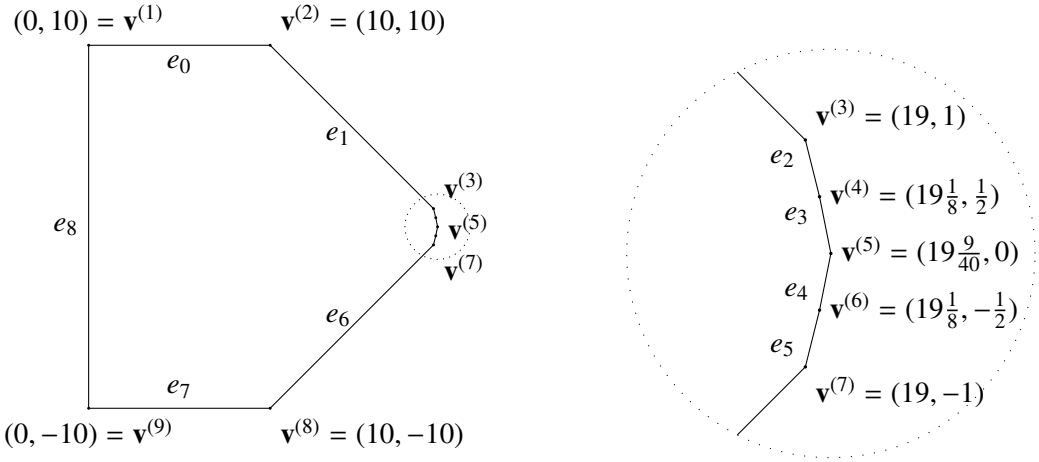

Figure 2.11: The polytope for the proof of Lemma 2.8.

There is a feasible maximal circuit walk of length three from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(5)}$ that is repetitive.

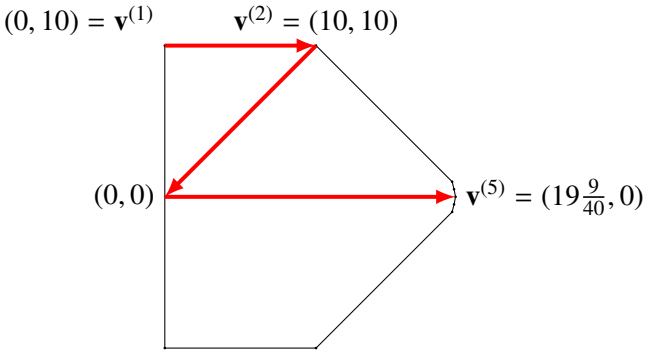

Figure 2.12: A feasible maximal repetitive circuit walk of length three.

Every other circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(5)}$ has length at least four. Therefore, we illustrate all possible combination of first (dashed) and second steps in Figure 2.13. From none of these second step points we can reach $\mathbf{v}^{(5)}$ in only one additional step, except from the point $(0,0)$ in the first picture. But this is the repetitive circuit walk from Figure 2.12. For those points for which it might not be immediately obvious that we cannot reach $\mathbf{v}^{(5)} = (19\frac{9}{40}, 0)$ along any edge direction in only one more step we added the coordinates for a convenient verification.

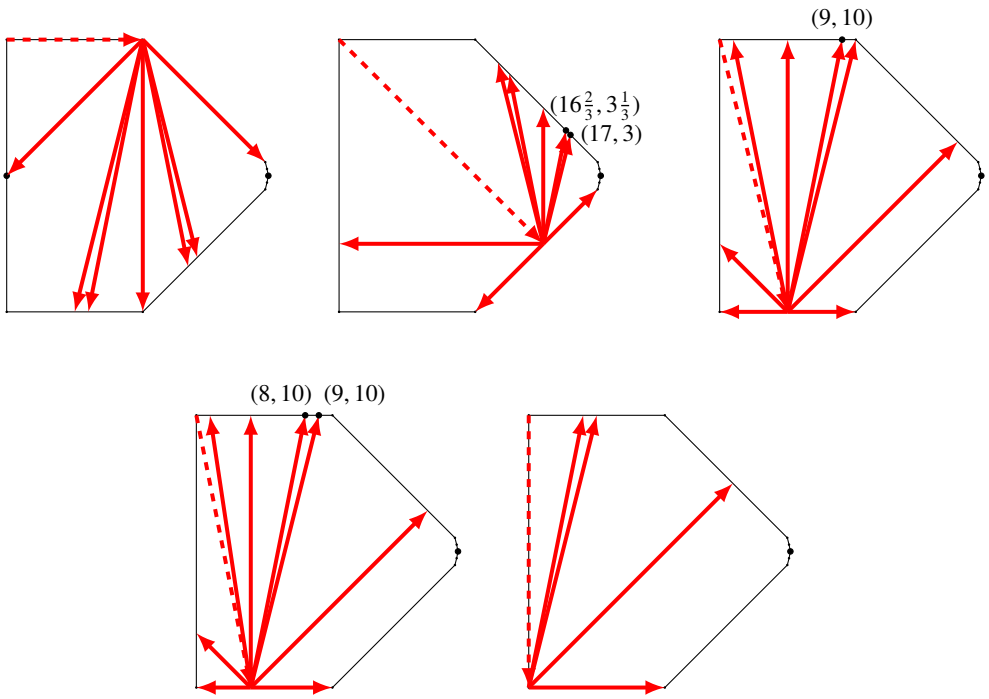

Figure 2.13: Possible combinations of first and second feasible maximal circuit steps from $\mathbf{v}^{(1)}$.

□

We also can show that the first and the second layer are related not only by the obvious weak inequalities, but by $\mathcal{CD}_{efm(b)(r)} > \mathcal{CD}_{fm(b)(r)}$.

Lemma 2.9 ($\mathcal{CD}_{efm(b)(r)} > \mathcal{CD}_{fm(b)(r)}$)**.** *For $n = 2$, there is a polytope with a pair of vertices for which every optimal feasible maximal circuit walk is not an edge walk, and there is such a walk that is non-repetitive and non-backwards. Hence the distances $\mathcal{CD}_{efm}$ and $\mathcal{CD}_{fm}$, the distances $\mathcal{CD}_{efmb}$ and $\mathcal{CD}_{fmb}$, and the distances $\mathcal{CD}_{efmr}$ and $\mathcal{CD}_{fmr}$ differ in this case.*

Proof. In the polytope below, an optimal edge walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ along the edges has length three, while there is a feasible maximal non-repetitive non-backwards circuit walk of length two.

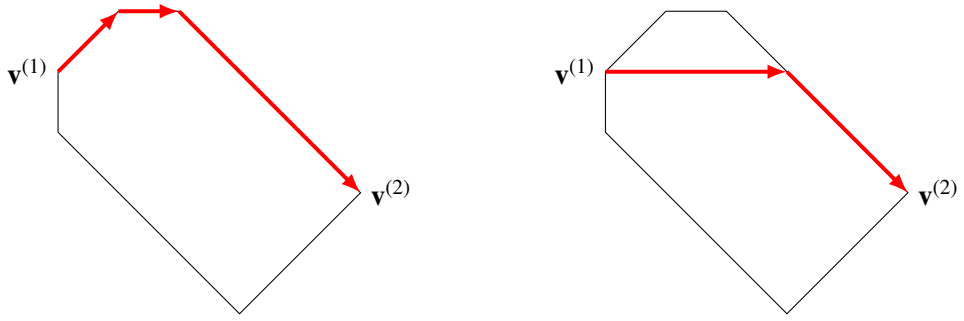

Figure 2.14: An optimal edge walk and an optimal feasible maximal walk.

□

In the third and lower layers of the hierarchy we drop the maximality condition. This may again reduce the distance of vertices.

Lemma 2.10 ($\mathcal{CD}_{fm(b)(r)} > \mathcal{CD}_{f(b)(r)}$)**.** *For $n = 2$, there is a polytope with a pair of vertices for which every optimal feasible circuit walk is not maximal, and there is such a walk that is non-repetitive and non-backwards. Hence the distances $\mathcal{CD}_{fm}$ and $\mathcal{CD}_{f}$, the distances $\mathcal{CD}_{fmb}$ and $\mathcal{CD}_{fb}$, and the distances $\mathcal{CD}_{fmr}$ and $\mathcal{CD}_{fr}$ differ in this case.*

Proof. In the polytope below, an optimal feasible maximal walk from $\mathbf{v}^{(2)}$ to $\mathbf{v}^{(1)}$ has length at least three: No matter which circuit direction we apply at $\mathbf{v}^{(2)}$ with maximum length, we cannot get to $\mathbf{v}^{(1)}$ in just one additional step.

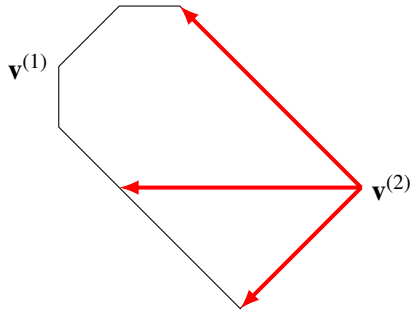

Figure 2.15: Possible first feasible maximal circuit steps at $\mathbf{v}^{(2)}$.

On the other hand, there is a feasible non-repetitive non-backwards circuit walk of length two.

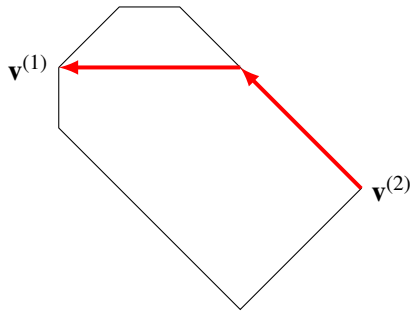

Figure 2.16: A feasible circuit walk of length two.

□

Further, we can prove that requiring a non-repetitive walk may increase the distance of a feasible walk, i.e. $\mathcal{CD}_{fr} > \mathcal{CD}_{f}$ and $\mathcal{CD}_{fbr} > \mathcal{CD}_{fb}$. Thereto, we need a polytope in dimension at least four as the following lemma tells us. In Lemma 2.12 we then show that such a polytope indeed exists.

Lemma 2.11. *For $n \leq 3$, every optimal feasible circuit walk is non-repetitive. Hence the distances $\mathcal{CD}_{f}$ and $\mathcal{CD}_{fr}$ coincide in this case.*

Proof. Repetitive circuit walks obviously must have length at least three. If every optimal circuit walk between two vertices is repetitive, then any feasible non-repetitive circuit walk must have length at least four. But in dimension ≤ 3 there always is a sign-compatible and thus feasible non-repetitive circuit walk of length at most three (for this, see also Lemmas 2.14 and 2.18). □

Lemma 2.12 ($\mathcal{CD}_{fr(b)} > \mathcal{CD}_{f(b)}$)**.** *For $n = 4$, there is a polytope with a pair of vertices for which every optimal feasible circuit walk is repetitive (but non-backwards). Hence the distances $\mathcal{CD}_f$ and $\mathcal{CD}_{fr}$ and the distances $\mathcal{CD}_{fb}$ and $\mathcal{CD}_{frb}$ differ in this case.*

Proof. Let the polytope

$$P = \left\{ \mathbf{x} \in \mathbb{R}^4 : \ \mathbf{l} \leq B\mathbf{x} \leq \mathbf{u} \right\}$$

be defined by

$$B = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}, \quad \mathbf{l} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ -\infty \\ -\infty \\ -\infty \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} 3/2 \\ 1 \\ 1 \\ 1 \\ 2 \\ 2 \\ 2 \end{pmatrix}.$$

The rows of the matrix B define the directions of 11 hyperplanes bounding the polytope.

The *vertices* are the intersections of four of these hyperplanes, in case this intersection is a single point that is contained in P. A simple computation shows that we have 23 vertices $\left(\{0,1\}^4 \backslash \{(1,0,0,0)^\intercal\}\right) \cup \left(\left\{\frac{3}{2}\right\} \times \left\{0, \frac{1}{2}\right\}^3\right)$. In particular $\mathbf{v}^{(1)} := (0,0,0,0)^\intercal$ and $\mathbf{v}^{(2)} := (1,1,1,1)^\intercal$ are vertices of P.

The *circuits* are the potential edge directions for varying $\mathbf{l}$ and $\mathbf{u}$, that is, they are given by the intersection of three hyperplanes, in case this intersection is 1-dimensional. Again it is not hard to compute that these directions are $\pm\mathbf{e}^i$ (where $\mathbf{e}^i$ is the i'th unit vector) and $\pm\left(\{1\} \times \{0,-1\}^3\right)$ and hence constitute the set of circuits.

Claim: Every feasible circuit walk from $\mathbf{v}^{(1)} = (0,0,0,0)^\intercal$ to $\mathbf{v}^{(2)} = (1,1,1,1)^\intercal$ of length at most three is repetitive.

Proof of claim: We investigate how we can go to $\mathbf{v}^{(2)} = (1,1,1,1)^\intercal$ in at most three feasible circuit steps. Observe that at any point we cannot apply circuits that violate the lower bounds of 0 or the upper bounds of 1, respectively $\frac{3}{2}$. Hence as a first feasible circuit step, we can only apply $\mathbf{e}^1$ or without loss of generality $\mathbf{e}^2$.

Applying $\mathbf{e}^2$ yields a point $(0, x_2, 0, 0)^\intercal$ with $x_2 \leq 1$. In the second step we can either apply without loss of generality $\mathbf{e}^3$, giving $(0, x_2, x_3, 0)^\intercal$ with $x_2, x_3 \leq 1$, or we apply $\mathbf{e}^1$ or $(1,-1,0,0)^\intercal$ giving $(x_1, x_2, 0, 0)^\intercal$ with $x_1 \leq \frac{3}{2}$, $x_2 \leq 1$. From neither of these points we can go to $\mathbf{v}^{(2)} = (1,1,1,1)^\intercal$ with one more circuit step: We cannot increase the first and the last component at the same time, nor increase the last two components by one simultaneously without decreasing the first component to, respectively under $\frac{1}{2}$.

Applying $\mathbf{e}^1$ as a first step yields a point $(x_1, 0, 0, 0)^\intercal$ with $x_1 \leq \frac{3}{2}$. In the next step we can either increase only one component, without loss of generality only the second one (by applying $\mathbf{e}^2$ or $(-1, 1, 0, 0)^\intercal$), giving $(x_1, x_2, 0, 0)^\intercal$ with $x_1 \leq \frac{3}{2},\ x_2 \leq 1$, but as before we cannot reach $\mathbf{v}^{(2)}$ in one more circuit step. Otherwise in the second step we increase at least two components (by applying $(-1, 1, 1, 1)^\intercal$ or without loss of generality $(-1, 1, 1, 0)^\intercal$), giving $(x_1, x_2, x_3, x_4)^\intercal$ with $x_1 + x_2 = x_1 + x_3 = x_1 + x_4 \leq \frac{3}{2},\ x_2, x_3, x_4 \leq 1$ (respectively, $(x_1, x_2, x_3, 0)^\intercal$ with $x_1 + x_2 = x_1 + x_3 \leq \frac{3}{2},\ x_2, x_3 \leq 1$). In particular we know that $x_2 = x_3 < 1$ or $x_1 < 1$. Hence to reach to $\mathbf{v}^{(2)} = (1, 1, 1, 1)^\intercal$ in one more step, we have to increase the second and third component simultaneously to 1 (which decreases the first component to $\leq \frac{1}{2}$), or we have to increase the first one without decreasing any other component (that is, we apply $\mathbf{e}^1$ *again*). This proves our claim.

On the other hand, applying the circuits $\mathbf{e}^1$, $(-1, 1, 1, 1)^\intercal$ and $\mathbf{e}^1$ with step length one each is indeed a feasible non-backwards circuit walk of length three from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$. □

In contrast, we only know $\mathcal{CD}_{fb} \geq \mathcal{CD}_f$ and $\mathcal{CD}_{fbr} \geq \mathcal{CD}_{fr}$. In other words, we do not know whether there is a polyhedron with a pair of vertices for wich every optimal feasible edge walk is backwards. These are the only weak inequalities in the hierarchy and we conjecture that these are strict as well. In Lemma 2.13 we explain why a polytope proving this conjecture has to be of dimension five or higher.

Lemma 2.13. *For $n \leq 4$, every optimal feasible circuit walk is non-backwards. Hence the distances $\mathcal{CD}_f$ and $\mathcal{CD}_{fb}$ coincide in this case.*

Proof. We first show that if an optimal feasible circuit walk is backwards then it has length at least four. Clearly it has length at least three. Assume there is a polytope with vertices $\mathbf{v}^{(1)}$ and $\mathbf{v}^{(2)}$ that are connected by a feasible circuit walk of length three that is backwards:

$$\mathbf{v}^{(1)} = \mathbf{y}^{(0)}, \quad \mathbf{y}^{(1)} = \mathbf{y}^{(0)} + \alpha_1 \mathbf{g}^1, \quad \mathbf{y}^{(2)} = \mathbf{y}^{(1)} + \alpha_2 \mathbf{g}^2, \quad \mathbf{y}^{(3)} = \mathbf{y}^{(2)} + \alpha_3(-\mathbf{g}^1) = \mathbf{v}^{(2)} .$$

But then there is a feasible circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ of length two, which is

$$\mathbf{v}^{(1)} = \mathbf{y}^{(0)}, \quad \bar{\mathbf{y}}^{(1)} = \mathbf{y}^{(0)} + (\alpha_1 - \alpha_3)\, \mathbf{g}^1, \quad \bar{\mathbf{y}}^{(2)} = \bar{\mathbf{y}}^{(1)} + \alpha_2 \mathbf{g}^2 \qquad \text{if } \alpha_1 \geq \alpha_3 ,$$

respectively,

$$\mathbf{v}^{(1)} = \mathbf{y}^{(0)}, \quad \bar{\mathbf{y}}^{(1)} = \mathbf{y}^{(0)} + \alpha_2 \mathbf{g}^2, \quad \bar{\mathbf{y}}^{(2)} = \bar{\mathbf{y}}^{(1)} + (\alpha_3 - \alpha_1)\, (-\mathbf{g}^1) \qquad \text{if } \alpha_1 < \alpha_3 .$$

Clearly these circuit walks satisfy $\bar{\mathbf{y}}^{(2)} = \mathbf{v}^{(2)}$ and are indeed feasible by convexity of the polytope. Therefore, a feasible backwards circuit walk of length three cannot be optimal.

Now in case every optimal circuit walk between two vertices is backwards, any feasible non-backwards circuit walk must have length at least five. But in dimension $n \leq 4$ there always is a sign-compatible and thus feasible non-backwards circuit walk of length at most four (see Lemmas 2.14 and 2.18). □

We conclude the feasible circuit walks with the sign-compatible ones, i.e. $\mathcal{CD}_{fs}$. Unlike the many combinations where the $\geq$-relation is clear from imposing additional or less constraints, it is not obvious for $\mathcal{CD}_{fs} \geq \mathcal{CD}_{fbr}$.

Lemma 2.14. *Any optimal sign-compatible circuit walk is feasible, non-backwards and non-repetitive.*

Proof. It follows immediately from the definition of sign-compatible walks that these are feasible. In fact, the steps of a sign-compatible circuit walk can be applied in arbitrary order, yielding feasible sign-compatible walks again. Hence, by reordering we can assume that all steps that use a circuit $\pm\mathbf{g}^i$ are applied consecutively. But then these multiple steps could be combined into a single circuit step and this yields a shorter circuit walk. □

For proving the strict inequality, again it is enough to state a polyhedron for which $\mathcal{CD}_{fs} > \mathcal{CD}_{fbr}$ holds for two of its vertices.

Lemma 2.15 ($\mathcal{CD}_{fs} > \mathcal{CD}_{fbr}$)**.** *For $n = 3$, there is a polytope with a pair of vertices for which every optimal feasible circuit walk is not sign-compatible, and there is such a walk that is non-repetitive and non-backwards. Hence the distance $\mathcal{CD}_{fs}$ differs from $\mathcal{CD}_{fbr}$, $\mathcal{CD}_{fr}$, $\mathcal{CD}_{fb}$, and $\mathcal{CD}_{f}$.*

Proof. Consider the polytope

$$P = \left\{ \mathbf{x} \in \mathbb{R}^3 : \mathbf{l} \leq B\mathbf{x} \leq \mathbf{u} \right\}$$

defined by

$$B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}, \quad \mathbf{l} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ -\infty \\ -\infty \end{pmatrix}, \quad \mathbf{u} = \begin{pmatrix} \infty \\ 1 \\ 1 \\ 2 \\ 2 \end{pmatrix}.$$

All possible edge directions $\mathbf{g}$ of P are given by

$$\pm\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \pm\begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \pm\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ -1 \\ -1 \end{pmatrix},$$

and the corresponding vectors $B\mathbf{g}$ are

$$\pm\begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \\ 1 \end{pmatrix}, \pm\begin{pmatrix} 0 \\ 1 \\ 0 \\ 1 \\ 0 \end{pmatrix}, \pm\begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ -1 \\ 0 \\ 0 \\ 1 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ 0 \\ -1 \\ 1 \\ 0 \end{pmatrix}, \pm\begin{pmatrix} 1 \\ -1 \\ -1 \\ 0 \\ 0 \end{pmatrix}.$$

We want to perform circuit walks from $\mathbf{v}^{(1)} = (0,0,0)^\intercal$ to $\mathbf{v}^{(2)} = (1,1,1)^\intercal$. We have $B\left(\mathbf{v}^{(2)} - \mathbf{v}^{(1)}\right) = (1,1,1,2,2)^\intercal$. Hence only the unit vectors $\mathbf{e}^1$, $\mathbf{e}^2$ and $\mathbf{e}^3$ can be applied in sign-compatible walks. Thus an optimal feasible sign-compatible walk from $\mathbf{v}^{(1)} = (0,0,0)^\intercal$ to $\mathbf{v}^{(2)} = (1,1,1)^\intercal$ has length at least three, as clearly we cannot connect these vertices using only two of the unit vectors.

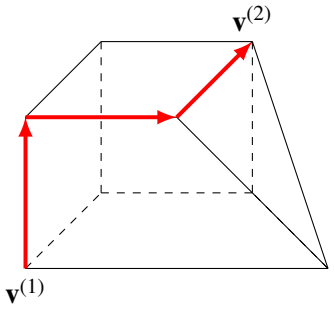

Figure 2.17: A feasible sign-compatible circuit walk of length three.

On the other hand, there is a feasible non-repetitive non-backwards circuit walk of length two that is not sign-compatible.

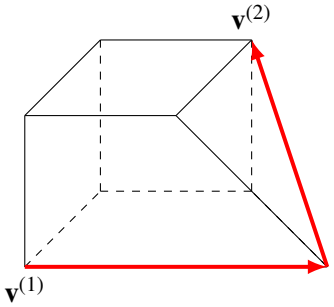

Figure 2.18: A feasible circuit walk of length two that is not sign-compatible.

□

In the final part of the hierarchy, shown in the lowest horizontal layer of the table, we do not even require feasibility. Here we only have to consider $\mathcal{CD}$: We already saw that $\mathcal{CD} = \mathcal{CD}_b = \mathcal{CD}_r$, and all other properties imply feasibility. We finally demonstrate that going infeasible indeed can be of advantage.

Lemma 2.16 ($\mathcal{CD}_f > \mathcal{CD}$)**.** *For $n = 3$, there is a polytope with a pair of vertices for which no optimal soft circuit walk is feasible, and there is such an (infeasible) optimal walk that is sign-compatible. Hence the distance $\mathcal{CD}$ differs from the distances $\mathcal{CD}_{fs} = \mathcal{CD}_s$, $\mathcal{CD}_{fbr}$, $\mathcal{CD}_{fb}$, $\mathcal{CD}_{fr}$, and $\mathcal{CD}_f$.*

Proof. The polytope below is obtained from a cube by cutting off six of its vertices using three pairs of hyperplanes, and keeping an opposite pair $\mathbf{v}^{(1)}$, $\mathbf{v}^{(2)}$ of vertices as depicted. Assume

the center of gravity of the cube is $\mathbf{0}$. Then the normals of these hyperplanes are equal to the coordinates of the vertices cut off. The 'depth' of the cuts is arbitrarily small.

Any feasible circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ has length at least three: To see this we illustrate the directions of all possible first steps at $\mathbf{v}^{(1)}$ (red) and all possible last steps to $\mathbf{v}^{(2)}$ (green) of a feasible circuit walk. Note that these steps are not necessarily maximal.

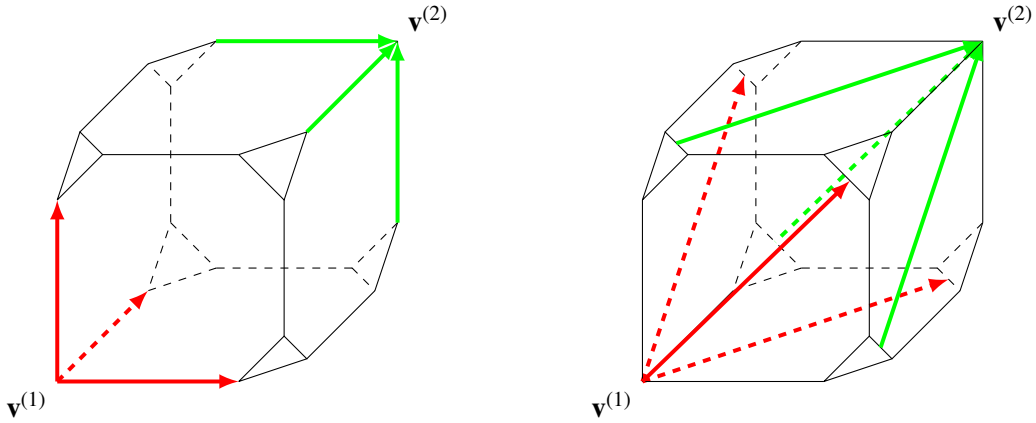

Figure 2.19: Possible first and last steps of a circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$.

Clearly there is no point that can be reached in a single step from $\mathbf{v}^{(1)}$ and from which one can reach $\mathbf{v}^{(2)}$ in a single step (at the same time). Hence any feasible circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ has length at least three.

On the other hand, there is a circuit walk of length two from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ when allowing to go infeasible:

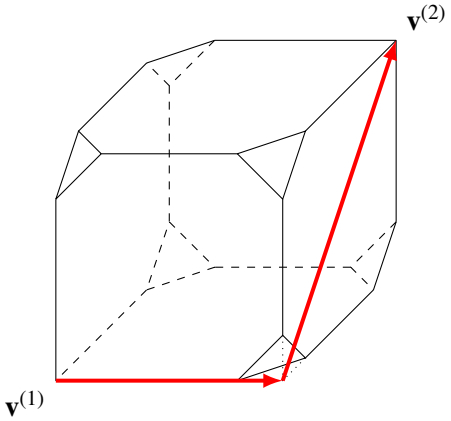

Figure 2.20: A circuit walk of length two that is not feasible.

□

2.3 Upper bounds on circuit diameters

In the previous section we studied how the several classes of circuit diameters bound each other. We now ask for general upper bounds on the respective diameters. Our main result is the following.

Theorem 2.17. *Let $P = \{ \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \}$ with $A \in \mathbb{Z}^{m_A \times n}$, $B \in \mathbb{Z}^{m_B \times n}$ be a polyhedron in $\mathbb{R}^n$. Then the circuit diameters CD_f, CD_{fb}, CD_{fr}, CD_{fbr}, CD_{fs}, and CD are bounded above by*

$$\min \{ \operatorname{rank}(A) - n + m_B\ ,\ n - \operatorname{rank}(A) \}\ .$$

The theorem follows immediately from the next lemma and transitivity of the relations in the hierarchy Figure 2.5.

Lemma 2.18. *Let $P = \{ \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d} \}$ with $A \in \mathbb{Z}^{m_A \times n}$, $B \in \mathbb{Z}^{m_B \times n}$ be a polyhedron in $\mathbb{R}^n$ and let $\mathbf{v}^{(1)}, \mathbf{v}^{(2)}$ be two of its vertices. Then there is a feasible sign-compatible circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$ of length at most* $\min \{ \operatorname{rank}(A) - n + m_B\ ,\ n - \operatorname{rank}(A) \}$.

Proof. It suffices to consider the case $\operatorname{rank}\binom{A}{B} = \operatorname{rank}(A) + \operatorname{rank}(B)$. Otherwise the representation of P has redundant rows in the matrix B and the bound derived below may only become lower.

Let $\sim \in \{=, \leq, \geq\}^{m_B}$ such that its i-th component $\sim_i$ is defined as

$$\sim_i = \begin{cases} \leq & \text{if } \left(B(\mathbf{v}^{(2)} - \mathbf{v}^{(1)}\right)_i < 0 \\ = & \text{if } \left(B(\mathbf{v}^{(2)} - \mathbf{v}^{(1)}\right)_i = 0 \\ \geq & \text{if } \left(B(\mathbf{v}^{(2)} - \mathbf{v}^{(1)}\right)_i > 0 \end{cases}.$$

Then $\mathbf{v}^{(2)} - \mathbf{v}^{(1)} \in \{\mathbf{x} \in \mathbb{R}^n : B\mathbf{x} \sim 0\} =: C_\sim$. This is a polyhedral rational cone in which all elements are pairwise sign-compatible. Observe that $B\mathbf{v}^{(i)} \leq \mathbf{d}$ and at least $n - \operatorname{rank}(A)$ linear independent inequalities of this kind are tight. Hence $(B(\mathbf{v}^{(2)} - \mathbf{v}^{(1)}))_i = 0$ for at least $2(n - \operatorname{rank}(A)) - m_B$ (linearly independent) inequalities if $m_B \leq 2(n - \operatorname{rank}(A))$ (and possibly for none if $m_B \geq 2(n - \operatorname{rank}(A))$). Hence $C_\sim$ has dimension at most $n - (2(n - \operatorname{rank}(A)) - m_B) = 2\operatorname{rank}(A) - n + m_B$ if $m_B \leq 2(n - \operatorname{rank}(A))$ (and dimension at most n if $m_B \geq 2(n - \operatorname{rank}(A))$).

Let $H_\sim$ be the unique minimal generating set of $C_\sim \cap \ker(A)$ over $\mathbb{R}$, where the components of each vector in $H_\sim$ are scaled to integers with greatest common divisor one. Then all elements in $H_\sim$ are circuits. Note that $\dim(C_\sim \cap \ker(A)) = \dim(C_\sim) - \operatorname{rank}(A)$ due to $\operatorname{rank}\binom{A}{B} = \operatorname{rank}(A) + \operatorname{rank}(B)$ and hence $\dim(C_\sim \cap \ker(A)) \leq \operatorname{rank}(A) - n + m_B$ if $m_B \leq 2(n - \operatorname{rank}(A))$ (and $\dim(C_\sim \cap \ker(A)) \leq n - \operatorname{rank}(A)$ if $m_B \geq 2(n - \operatorname{rank}(A))$). By Caratheodory's Theorem, $\mathbf{v}^{(2)} - \mathbf{v}^{(1)} \in C_\sim \cap \ker(A)$ can thus be written as a combination of at most $\operatorname{rank}(A) - n + m_B$ (respectively, $n - \operatorname{rank}(A)$) of the generators contained in $H_\sim$. □

For two common representations of polyhedra using only the matrix A or B, respectively, we can refine the diameter bounds as follows.

Corollary 2.19. *Let $P = \{\, \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b},\, \mathbf{x} \geq \mathbf{0} \,\}$ be a polyhedron in $\mathbb{R}^n$. Then the circuit diameters CD_f, CD_{fb}, CD_{fr}, CD_{fbr}, CD_{fs}, and CD are bounded above by*

$$\min \{\, \operatorname{rank}(A)\,,\ n - \operatorname{rank}(A) \,\}\,.$$

Proof. Note that $B = -I_n$ and thus $m_B = n$. Then Theorem 2.17 immediately implies the claim. □

Corollary 2.20. *Let $P = \{\, \mathbf{x} \in \mathbb{R}^n : B\mathbf{x} \leq \mathbf{d} \,\}$ be a polyhedron in $\mathbb{R}^n$, $B \in \mathbb{Z}^{m_B \times n}$. Then the circuit diameters CD_f, CD_{fb}, CD_{fr}, CD_{fbr}, CD_{fs}, and CD are bounded above by*

$$\min \{\, m_B - n\,,\ n \,\}\,.$$

Proof. Note that there is no matrix A in the description of the polyhedron, so the claim follows from Theorem 2.17 using $\operatorname{rank}(A) = 0$. □

As has been pointed out several times, we are particularly interested in the bound posed by the Hirsch conjecture. For circuit diameters defined on circuit walks using arbitrary length steps we get the following result.

Corollary 2.21. *For any d-dimensional polyhedron with f facets, the circuit diameters CD_f, CD_{fb}, CD_{fr}, CD_{fbr}, CD_{fs}, and CD are bounded above by $f - d$.*

Proof. Observe that $f \leq m_B$ and $d = \dim(P) = n - \operatorname{rank}(A)$. Thus the circuit diameters are bounded above by $m_B - (n - \operatorname{rank}(A)) \leq f - d$ by Theorem 2.17. □

Recall that this bound does not hold for the combinatorial diameter CD_{efm} [27] and its validity is open for the circuit diameter CD_{fm}. Moreover, for both diameters the existence of upper bounds that are polynomial in f and d is unresolved.

2.4 The circuit distance hierarchy in dimension two

We conclude this chapter on the circuit hierarchy with a discussion of the different notions of circuit distances in dimension $n = 2$, where we use the notation CD^2 etc. Here the hierarchy collapses into fewer distinct categories, see Figure 2.21. We prove this in the following Theorem 2.22, together with all possible values of the respective circuit distances.

After this, we will investigate that the combinatorial diameter and the circuit diameter CD^2_{fm} can differ significantly in Lemma 2.24.

Theorem 2.22. *For $n = 2$ the circuit hierarchy collapses as depicted in Figure 2.21.*

Additionally, for a polygon on k vertices we have that

$$\begin{aligned} \mathcal{CD}^2_{efmb} &\in \{1, \dots, k-3\} \ \ (k \geq 5) \\ \mathcal{CD}^2_{efm} = \mathcal{CD}^2_{efmr} &\in \left\{1, \dots, \left\lfloor \frac{k}{2} \right\rfloor\right\} \\ \mathcal{CD}^2_{fm(b)(r)} &\in \left\{1, \dots, \left\lfloor \frac{k}{2} \right\rfloor\right\} \\ \mathcal{CD}^2_{f} = \mathcal{CD}^2_{fr} = \mathcal{CD}^2_{fb} = \mathcal{CD}^2_{fbr} = \mathcal{CD}^2_{fs} = \mathcal{CD}^2 &\in \{1, 2\} \end{aligned}$$

and there are polygons with pairs of vertices that attain the maximal distances in the ranges claimed above.

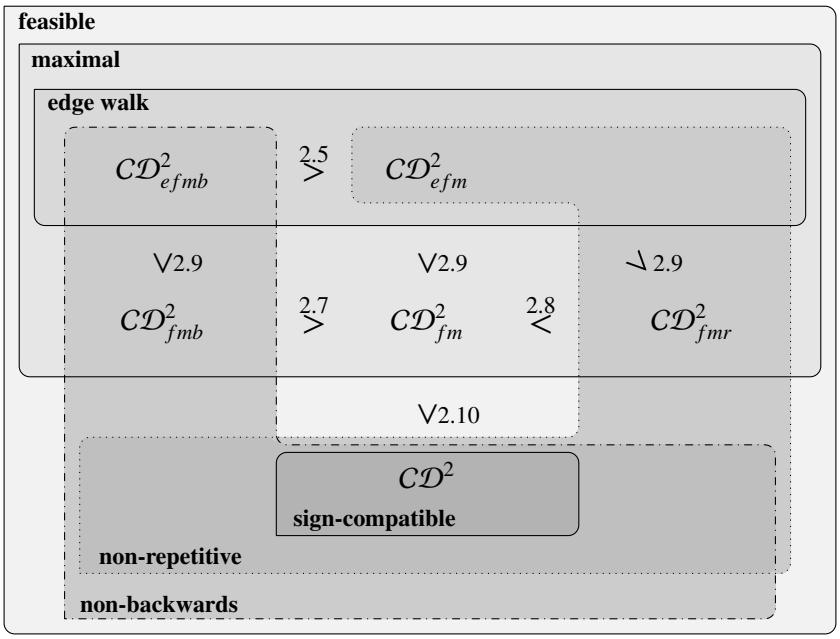

Figure 2.21: The hierarchy of circuit distances in dimension two.

Proof. Note that Lemma 2.5, Lemma 2.7, Lemma 2.8, Lemma 2.9 and Lemma 2.10 prove distinctness of the respective circuit distances in dimension two. Hence the strict inequalities claimed in the circuit hierarchy in Figure 2.21 hold (again the numbers near the inequality symbols refer to these lemmas). For

$$\mathcal{CD}^2_{f} = \mathcal{CD}^2_{fr} = \mathcal{CD}^2_{fb} = \mathcal{CD}^2_{fbr} = \mathcal{CD}^2_{fs} = \mathcal{CD}^2 \in \{1, 2\}$$

it is enough to recall Lemma 2.18: In dimension two any two vertices are connected by a feasible sign-compatible circuit walk of length at most two. Further, if there is a circuit step that directly connects two vertices, this circuit walk of length one trivially satisfies any property.

In dimension two there are no repetitive edge walks and hence $\mathcal{CD}^2_{efm} = \mathcal{CD}^2_{efmr}$. The statement on the possible values of circuit distance is obvious.

For a polygon on k vertices with a pair of vertices with $CD^2_{efmb} = k - 3$ is readily derived from the one given in Figure 2.6 in Lemma 2.5 by putting $k-4$ vertices 'to the left' of $\mathbf{v}^{(1)}$ and $\mathbf{v}^{(2)}$. The upper bound of $k - 3$ follows from the fact that an optimal circuit walk that is backwards uses at least three steps.

We clearly have $CD^2_{fm(b)(r)} \leq \left\lfloor \frac{k}{2} \right\rfloor$. In Lemma 2.23 we show that this value indeed can be attained. □

To complete the proof of Theorem 2.22 we still have to show that there are polygons with vertices that have feasible maximal circuit distance $\left\lfloor \frac{k}{2} \right\rfloor$. For the sake of a clean presentation, we provide the proof for k even. It can readily be extended to odd k by adding another vertex.

Lemma 2.23. *Let k be even. Then there is a polygon on k vertices with circuit diameter CD^2_{fm} equal to $\frac{k}{2}$.*

Proof. Let k even be given. We construct a polygon on k vertices $\mathbf{v}^{(0)}, \ldots, \mathbf{v}^{(k-1)}$, with edges $e_i = \left(\mathbf{v}^{(i)}, \mathbf{v}^{(i+1)}\right)$ for $i = 0, \ldots, k-1$ (where $\mathbf{v}^{(k)} := \mathbf{v}^{(0)}$) such that $CD^2_{fm}\left(\mathbf{v}^{(0)}, \mathbf{v}^{(\frac{k}{2})}\right) = \frac{k}{2}$. In particular, the corresponding optimal maximal circuit walk will be an edge walk.

Note that for $n = 2$ the actual edges of the polygon are in fact all potential edge directions. Thus, for a simpler wording in the following, we will talk about 'walking along edges' or 'in direction of an edge'.

We now come to the construction of the polygon P. First of all we fix the edge directions and hence the set of circuits associated with P. To this end we choose $\frac{k}{2}$ slopes $0 > s_0 > s_1 > s_2 > \ldots > s_{\frac{k}{2}-1}$ arbitrarily. In the upcoming construction we assign edge e_0 slope $-s_0$; edge e_{k-1} slope s_0, and for $i = 1, \ldots, \frac{k}{2} - 1$ we assign e_i slope s_i and e_{k-1-i} slope $-s_i$. This will produce a polygon of shape as depicted in Figure 2.22.

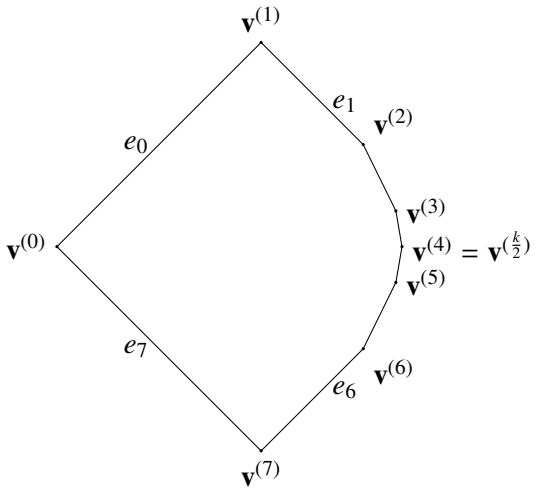

Figure 2.22: Sketch of the polygon for $k = 8$.

Observe that the slopes of the edges from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(\frac{k}{2})}$ iteratively become steeper, just as the slopes of the edges from $\mathbf{v}^{(k-1)}$ to $\mathbf{v}^{(\frac{k}{2})}$. Further the polygon will be symmetric with respect to the

first coordinate axis (which we call x_1-axis from now on).

It remains to arrange the vertices. We do this iteratively, fixing a pair of vertices 'opposite' $\mathbf{v}^{(i)}, \mathbf{v}^{(k-i)}$ in each step such that in P the following property **(*)** is satisfied:

(*) Every maximal feasible circuit walk starting at $\mathbf{v}^{(0)}$ of length at most $\frac{k}{2}$ that contains a point $\mathbf{v}'$ with larger x_1-coordinate than $\mathbf{v}^{(i)}$ or $\mathbf{v}^{(k-i)}$ (or equivalently, that hits an edge e_j with $i \leq j \leq k-1-j$) must contain $\mathbf{v}^{(i)}$ or $\mathbf{v}^{(k-i)}$.

This will immediately imply that the circuit distance $\mathcal{CD}^2_{fm}$ from $\mathbf{v}^{(0)}$ to $\mathbf{v}^{(\frac{k}{2})}$ is $\frac{k}{2}$: Every circuit walk of length at most $\frac{k}{2}$ from $\mathbf{v}^{(0)}$ to $\mathbf{v}^{(\frac{k}{2})}$ does reach a $\mathbf{v}'$ with larger x_1-coordinate than every $\mathbf{v}^{(i)}$ for all $i = 1, \ldots, \frac{k}{2} - 1$. (We will informally call this 'going beyond $\mathbf{v}^{(i)}$'.) Hence by **(*)**, any such circuit walk must contain a vertex from each of these $\frac{k}{2} - 1$ pairs of vertices. This takes at least $\frac{k}{2} - 1$ steps, and it takes one additional step to reach the target vertex $\mathbf{v}^{(\frac{k}{2})}$.

Construction of initial vertices: Fix $\mathbf{v}^{(0)} = (0, 0)$. Let edges e_0, respectively e_{k-1}, start at $\mathbf{v}^{(0)}$ and end in a point $\mathbf{v}^{(1)}$ on e_0 and a point $\mathbf{v}^{(k-1)}$ on e_{k-1} such that $\mathbf{v}^{(1)}$ and $\mathbf{v}^{(k-1)}$ have identical x_1-coordinates.

(*) holds for the pair $\mathbf{v}^{(1)}, \mathbf{v}^{(k-1)}$: At $\mathbf{v}^{(0)}$ we can only apply circuit steps with directions e_0 or e_{k-1} (any other direction is too steep). As we apply maximal steps, the second point of any circuit walk is either $\mathbf{v}^{(1)}$ or $\mathbf{v}^{(k-1)}$.

Construction of a new pair of vertices: Let the vertices $\mathbf{v}^{(0)}, \mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(i)}, \mathbf{v}^{(k-i)}, \ldots \mathbf{v}^{(k-1)}$, $i < \frac{k}{2} - 1$, be constructed and satisfy **(*)**. We now construct the vertices $\mathbf{v}^{(i+1)}$, $\mathbf{v}^{(k-i-1)}$ together with the incident edges e_i, e_{k-1-i}.

1. Let edges with directions e_i (respectively e_{k-1-i}) start at $\mathbf{v}^{(i)}$ (respectively $\mathbf{v}^{(k-i)}$). Let $\mathbf{w}^{(i)}$ be their intersection (which has x_2-coordinate 0). This defines a polygon P_i.

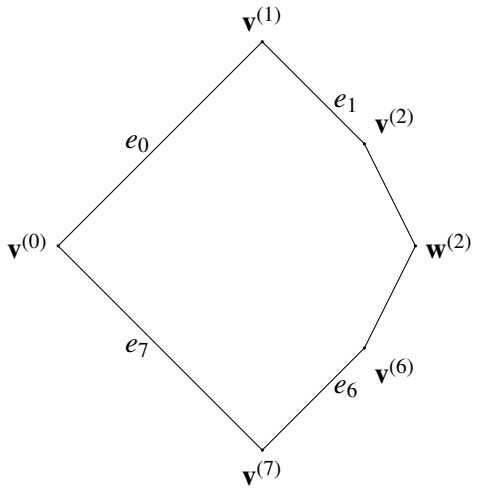

Figure 2.23: The polygon P_2 for $k = 8$.

2. In P_i consider all feasible maximal circuit walks of length at most $\frac{k}{2}$ that begin at $\mathbf{v}^{(0)}$ and do not walk along the (actual) edge e_i (respectively e_{k-1-i}) to the vertex $\mathbf{w}^{(i)}$. Then none of these walks contains $\mathbf{w}^{(i)}$: A step that hits $\mathbf{w}^{(i)}$ is not allowed to go along the edges we just inserted by definition and we cannot reach $\mathbf{w}^{(i)}$ from e_0 (respectively e_{k-1}) in one circuit step by construction. Hence it must start at an edge e_j with without loss of generality $0 < j < i$. But then there would have been a feasible maximal circuit walk of length at most $\frac{k}{2}$ in P_j that goes beyond $\mathbf{v}^{(j)}$ which contradicts the definition of $\mathbf{v}^{(j)}$ in P_j.

3. Among all points contained in all of these circuit walks, there are points that have a largest x_1-value. These points lie on the edges e_i (respectively e_{k-1-i}) by construction. We now set $\mathbf{v}^{(i+1)}$ to be such a point on e_i (respectively $\mathbf{v}^{(k-1-i)}$ on e_{k-1-i}). This yields a pair of vertices $\mathbf{v}^{(i+1)}, \mathbf{v}^{(k-1-i)}$ of identical x_1-value and with $\mathbf{v}^{(i+1)}, \mathbf{v}^{(k-1-i)} \neq \mathbf{w}^{(i)}$ (with the same arguments as before).

We have to show that $\mathbf{v}^{(i+1)}$ and $\mathbf{v}^{(k-i-1)}$ satisfy **(*)** in P. Therefore, consider a maximal feasible circuit walk in P starting at $\mathbf{v}^{(0)}$ and of length at most $\frac{k}{2}$ that goes beyond $\mathbf{v}^{(i+1)}$. This walk in P translates to a walk in P_i and clearly these walks in P and P_i coincide until they go beyond $\mathbf{v}^{(i+1)}$ (in both P and P_i) by applying some circuit $\mathbf{g}^j$ at some point $\mathbf{y}^{(j)}$ in the respective circuit walks. Let $\mathbf{y}^{(j+1)}$ be the subsequent point in the circuit walk in P, respectively $\bar{\mathbf{y}}^{(j+1)}$ in P_i. In particular these $\mathbf{y}^{(j+1)}$ and $\bar{\mathbf{y}}^{(j+1)}$ have a larger x_1-value than $\mathbf{v}^{(i+1)}$. By construction of $\mathbf{v}^{(i+1)}$ we can only go beyond $\mathbf{v}^{(i+1)}$ in at most $\frac{k}{2}$ circuit steps in P_i when going along the (actual) edge, without loss of generality along e_i. Hence without loss of generality $\mathbf{g}^j$ is the edge direction e_i and $\mathbf{y}^{(j)} \in e_i$. Thus we have $\mathbf{y}^{(j)} = \mathbf{v}^{(i+1)}$ as we apply maximal steps, in particular the vertex $\mathbf{v}^{(i+1)}$ is contained in the circuit walk in P.

Construction of the final vertex: Set $\mathbf{v}^{(\frac{k}{2})} := \mathbf{w}^{(i)}$ for $i = \frac{k}{2} - 1$. This concludes the construction of a polygon P with property **(*)**. □

Recall that Theorem 2.22 tells us all possible values of the respective circuit diameters. These results immediately imply that there are polyhedra that have constant circuit distances $\mathcal{CD}^2_{(f(b)(r)(s))}$, while the other notions grow linear.

We finally want to investigate, how the combinatorial diameter and the circuit diameter $\mathcal{CD}^2_{fm}$ can differ. The following lemma tells us that this gap can be significant: The former diameter can grow linear while the latter remains constant.

Lemma 2.24. *Let P be a regular polygon on k vertices. Then for the combinatorial diameter we have*

$$\mathcal{CD}^2_{efm(b)(r)} = \begin{cases} \frac{k-1}{2} & \text{if } k \text{ odd} \\ \frac{k}{2} & \text{if } k \text{ even} \end{cases},$$

while the circuit diameter is given by

$$\mathcal{CD}^2_{fm(b)(r)} = \begin{cases} 1 & \text{if } k \text{ odd} \\ 2 & \text{if } k \text{ even} \end{cases}.$$

Proof. For $\mathcal{CD}^2_{efm(b)(r)}$ the claim is obvious.

To determine the circuit distances $\mathcal{CD}_{fm(b)(r)}$, let $\mathbf{v}^{(1)}, \ldots, \mathbf{v}^{(k)}$ be the vertices of the polygon and $(\mathbf{v}^{(i)}, \mathbf{v}^{(i+1)})$ its edges (where $\mathbf{v}^{(k+1)} := \mathbf{v}^{(1)}$). Let k be odd. It suffices to show that from $\mathbf{v}^{(1)}$ we can reach any other vertex in just a single circuit step. For this, it is enough to see $\mathbf{v}^{(2i)} = \mathbf{v}^{(1)} + \alpha \cdot (\mathbf{v}^{(i+1)} - \mathbf{v}^{(i)})$ for some α and $\mathbf{v}^{(2i+1)} = \mathbf{v}^{(1)} + \alpha' \cdot (\mathbf{v}^{(\frac{k+1}{2}+i)} - \mathbf{v}^{(\frac{k+1}{2}+i+1)})$ for some α'. In other words, the directions $\mathbf{v}^{(2i)} - \mathbf{v}^{(1)}$ and $\mathbf{v}^{(i+1)} - \mathbf{v}^{(i)}$, respectively the directions $\mathbf{v}^{(2i+1)} - \mathbf{v}^{(1)}$ and $\mathbf{v}^{(\frac{k+1}{2}+i)} - \mathbf{v}^{(\frac{k+1}{2}+i+1)}$, are parallel.

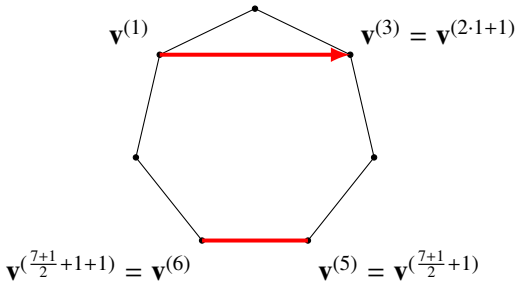

Figure 2.24: A circuit step from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(3)}$ and the edge of corresponding direction.

Now let k be even. First observe that there are always two collinear edges and hence not all pairs of vertices can be connected by a single circuit step. Hence the diameter is at least two. As before, we have $\mathbf{v}^{(2i)} = \mathbf{v}^{(1)} + \alpha \cdot (\mathbf{v}^{(i+1)} - \mathbf{v}^{(i)})$ for some α. In case we want to walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2i+1)}$ we first go to $\mathbf{v}^{(2i)}$ and then along edge $(\mathbf{v}^{(2i)}, \mathbf{v}^{(2i+1)})$, as depicted in the walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(5)}$ in Figure 2.4. Hence the regular k-polygon for k even has diameter two with respect to $\mathcal{CD}^2_{fm(b)(r)}$.

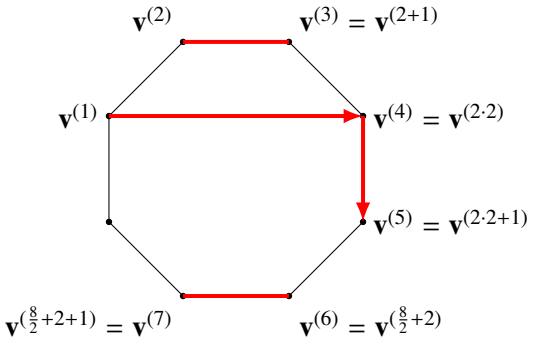

Figure 2.25: Optimal circuit walks from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(4)}$ and $\mathbf{v}^{(5)}$ and edges of corresponding direction.

□

Chapter 3

Circuit diameter bounds for transportation polytopes

Classical or two-way transportation problems (also called *Hitchcock problems* [19]) model the minimum-cost transportation of a good from M suppliers to N demand locations, where each of these $M + N$ points sends or receives, respectively, a specified amount of the product. They are among the most fundamental problems in mathematical programming, operations research and statistics, see for example [12, 13, 22, 29] and references therein.

As the standard transportation problem requires the optimization of a linear objective function, the combinatorial diameter of transportation polytopes is of particular interest. For $M\times N$ transportation polytopes, the Hirsch conjecture essentially claims an upper bound of $M + N - 1$. It holds for $M = 2$ [12] and for a special class of 0, 1-transportation polytopes, the so-called partition polytopes, which even satisfy a much lower bound [2]. However, it is still open whether the Hirsch conjecture is true for general $M\times N$ transportation polytopes. For $M \geq 3$ the best upper bounds are linear: A bound of $8(M + N - 2)$ is presented in [7] and improved to $4(M + N - 1)$ with a sketch of the proof in [12].

In the following we prove lower and upper bounds for the four main categories of circuit distances introduced in the previous chapter. Our lower bound result is common to all notions of circuit diameters.

Theorem 3.1 (Lower bounds for transportation polytopes)**.** *For all M, N, there is an $M\times N$ transportation polytope for which any circuit diameter is at least the minimum of $M + N - 1$ and $(M - 1)(N - 1)$.*

Note that we have $(M - 1)(N - 1) \leq M + N - 1$ for all $M \leq N$ except $M = 2$ or $M = N = 3$. In particular, this lower bound tells us that $M\times N$ transportation polytopes with $N > M \geq 3$ are 'at least *Hirsch-sharp*', that is, there are polytopes whose combinatorial diameter attains the bound of $M + N - 1$ claimed by the Hirsch conjecture.

On the contrary, proving the upper bound of $M + N - 1$ for all circuit diameters equally is much more involved. We tackle this problem by investigating the combinatorial diameter and $C\mathcal{D}_{fm}$ for small instances of transportation problems ($M = 2, 3$), while in the general $M\times N$ case we look at the less restrictive notions of circuit diameters for which proving upper bounds turns out to be far

more practicable. In fact, we get validity of the bound induced by the Hirsch conjecture, and we already know that this bound is tight.

Theorem 3.2 (Upper bounds for transportation polytopes)**.** *For $M\times N$ transportation polytopes, the weak and the soft circuit diameter $C\mathcal{D}_f$ and $C\mathcal{D}$ are bounded above by $M + N - 1$.*

We then restrict ourselves to $2\times N$ and $3\times N$ transportation polytopes. For the combinatorial diameter (we use $C\mathcal{D}_e$ in this chapter) we obtain the following results.

Theorem 3.3 (Hirsch conjecture for $M = 2, 3$)**.**

(i) The monotone Hirsch conjecture holds for $2\times N$ transportation polytopes. Further, the (monotone) combinatorial diameter $C\mathcal{D}_e$ is bounded above by $M + N - 2 = N$.

(ii) The Hirsch conjecture holds for $3\times N$ transportation polytopes. In particular, the combinatorial diameter $C\mathcal{D}_e$ is bounded above by $M + N - 1 = N + 2$.

Both bounds are tight. This follows from Theorem 3.1 for $M = 3$, and we present a $2\times N$ transportation polytopes with combinatorial diameter N in Example 3.9. Note that the upper bound of N also tells us that $2\times N$ transportation polytopes are not *Hirsch sharp*, as the Hirsch bound of $N + 1$ cannot be attained. Therefore, Theorem 3.1 (i) improves the bound implied by the validity of the Hirsch conjecture for $M = 2$ [12].

For the circuit diameter $C\mathcal{D}_{fm}$ we get similar bounds.

Theorem 3.4 (Circuit diameter bounds for $M = 2, 3$)**.**

(i) For $2\times N$ transportation polytopes, the circuit diameter $C\mathcal{D}_{fm}$ is bounded above by $N - 1$.

(ii) For $3\times N$ transportation polytopes, the circuit diameter $C\mathcal{D}_{fm}$ is bounded above by $N + 2$.

By Theorem 3.1 these bounds are tight as well. Note that for $2\times N$ transportation polytopes our general upper bound on $C\mathcal{D}_{fm}$ is better than that for combinatorial diameter, while the bounds for the $3\times N$ case coincide. Actually, Theorem 3.4 (ii) was already implied by Theorem 3.3. However, we consider the $C\mathcal{D}_{fm}$ case explicitly as we present a more specific proof here. It provides us with a stronger statement about the actual circuit distances and this approach might be useful for tackling $M > 3$ as well.

The chapter is joint work with Steffen Borgwardt, Jesús A. De Loera and Jake Miller [4]. It is structured as follows: In Section 3.1 we recall the necessary background on transportation polytopes and present our notation and tools for the discussion. In particular, we explain what the respective circuit walks look like for transportation polytopes. We then turn to the upper bounds on circuit diameters. We begin with the proof of Theorem 3.2 in Section 3.2, followed by proving Theorem 3.3 in Section 3.3 and Theorem 3.4 in Section 3.4. This chapter ends with the investigation of the lower bounds in Section 3.5, where we show that Theorem 3.1 holds.

3.1 Preliminaries

An $M\times N$ transportation problem has M supply points and N demand points to be met. Each supply point holds a quantity $u_i > 0$ and each demand point needs a quantity $v_j > 0$. Let $y_{ij} \geq 0$ denote the flow from supply point i to demand point j. Then the set of *feasible (flow) assignments* $\mathbf{y} \in \mathbb{R}^{M\times N}$ can be described as

$$\begin{aligned} \sum_{j=1}^{N} y_{ij} &= u_i \qquad i = 1, ..., M, \\ \sum_{i=1}^{M} y_{ij} &= v_j \qquad j = 1, ..., N, \\ y_{ij} &\geq 0 \qquad i = 1, ..., M,\ j = 1, ..., N. \end{aligned}$$

The set of solutions to these constraints constitutes a *transportation polytope*. Here the vectors $\mathbf{u}$ and $\mathbf{v}$ are called the *marginals* or *margins* for the transportation polytope.

When discussing an $M\times N$ transportation problem, it is common practice to think of the supply and demand points as nodes in the complete bipartite graph $K_{M,N}$. We denote the nodes corresponding to the supply points $\{\mathfrak{s}_1, \ldots, \mathfrak{s}_M\}$ and the nodes corresponding to the demand points $\{\mathfrak{d}_1, \ldots, \mathfrak{d}_N\}$. For every feasible solution $\mathbf{y} \in \mathbb{R}^{M\times N}$ we define the *support graph* $B(\mathbf{y})$ as the subgraph of $K_{M,N}$ that contain precisely the edges $\left\{\{\mathfrak{s}_i, \mathfrak{d}_j\} \ : \ y_{ij} > 0, i \in \{1, \ldots, M\}, j \in \{1, \ldots, N\}\right\}$ of non-zero flow. We use this representation throughout the chapter to visualize our methods.

With the term *assignment* we either refer to the vector $\mathbf{y} \in \mathbb{R}^{M\times N}$ itself or just to the edge set of its support graph $B(\mathbf{y})$, depending on the context. Note that $B(\mathbf{y})$ is directly derived from $\mathbf{y}$. We typically denote assignments by the capital letters O (for 'original'), C (for 'current'), and F (for 'final'). If we want to refer to the actual flow assignment as a vector explicitly, we write for example $\mathbf{y}^O = (y^O_{11}, \ldots, y^O_{MN})$. We use $|O|$ to refer to the number of edges of the support graph.

Observe that for general transportation polytopes the support graph is not necessarily connected, but this is the case for non-degenerate transportation polytopes: An $M\times N$ transportation polytope is *non-degenerate* if every vertex has exactly $M + N - 1$ non-negative entries. This is the case if and only if there are no non-empty proper subsets $I \subsetneq \{1, \ldots, M\}$ and $J \subsetneq \{1, \ldots, M\}$ such that $\sum_{i\in I} u_i = \sum_{j\in J} v_j$, see [29]. Note that for each degenerate $M\times N$ transportation polytope there is a non-degenerate $M\times N$ transportation polytope of the same or larger combinatorial diameter $C\mathcal{D}_e$ [29]. Therefore, it suffices to consider non-degenerate transportation polytopes to prove upper bounds on $C\mathcal{D}_e$. We exploit this in Section 3.3.

In contrast, we cannot assume non-degeneracy when exhibiting other notions of circuit distance as it is not clear whether for every degenerate $M\times N$ transportation polytope there is a perturbed non-degenerate $M\times N$ transportation polytope bounding the respective circuit diameters of the original one from above.

When studying circuit distances, the *vertices* of the polytope are of special interest. For transportation polytopes they can be characterized in terms of the support graphs: A feasible point $\mathbf{y}$ is a vertex if and only if its support graph contains no cycles, that is, $B(\mathbf{y})$ is a spanning forest.

Observe that a vertex $\mathbf{y}$ is uniquely determined by (the edge set of) its support graph $B(\mathbf{y})$ and the vertices $\mathbf{y}$ of non-degenerate transportation polytopes are given by spanning trees (see for example [22]). For a vertex O (or $\mathbf{y}^O$) we distinguish two kinds of demand nodes:

1. *Leaf demands* are those demand points which are leaves in $B(\mathbf{y}^O)$. When we say *leaf edges* we refer to edges incident to leaf demands. (This differs from the standard notion of leaf edges used in graph theory!)

2. *Mixed demands* are those demand points which have degree at least two in $B(\mathbf{y}^O)$. We denote the set of mixed demands by D_m^O and the *mixed edges* are given by $E_m^O := \left\{\{\mathfrak{s}_i, \mathfrak{d}_j\} : \mathfrak{d}_j \in D_m^O\right\}$.

Note that for a vertex O of a non-degenerate $2\times N$ transportation polytope we always have $|D_m^O| = 1$ and $|E_m^O| = 2$, in the non-degenerate $3\times N$ case either $|D_m^O| = 1$ and $|E_m^O| = 3$ or $|D_m^O| = 2$ and $|E_m^O| = 4$. The three configurations are illustrated in Figure 3.1 (mixed edges are bold). Here the sets D_m^O of mixed demands are $\{\mathfrak{d}_4\}$, $\{\mathfrak{d}_4\}$ and $\{\mathfrak{d}_4, \mathfrak{d}_6\}$, respectively.

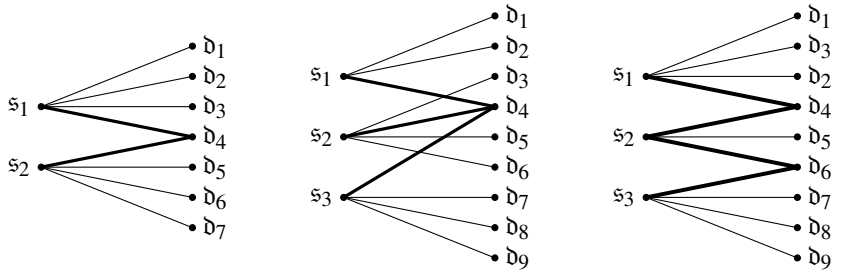

Figure 3.1: Vertices of non-degenerate $2\times N$ transportation polytopes always have exactly one mixed demand point (left), while vertices of non-degenerate $3\times N$ transportation polytopes can have either one or two (middle, right).

For every supply point $\mathfrak{s}_i$ we denote the vertices adjacent to $\mathfrak{s}_i$ in an assignment O (its neighborhood) by

$$N_i^O := \left\{\mathfrak{d}_j : y_{ij}^O > 0\right\} = \left\{\mathfrak{d}_j : \{\mathfrak{s}_i, \mathfrak{d}_j\} \in O\right\} .$$

We continue with characterizing the *actual* edges of the polytope in terms of the support graphs.

Proposition 3.5 (see [22])**.** *Let O and C be two assignments of an $M\times N$ transportation polytope. Then the corresponding vertices are connected by an edge if and only if $O \cup C$ contains a unique cycle.*

This unique cycle describes an edge direction of the transportation polytope. It is easy to see that every cycle of $K_{M,N}$ can appear as an edge of some $M\times N$ transportation polytope if we choose suitable margins. Thus, the set of circuits of an $M\times N$ transportation polytope consists of all even simple cycles of the form $(\mathfrak{s}_{i_1}, \mathfrak{d}_{j_1}, \mathfrak{s}_{i_2}, \mathfrak{d}_{j_2}, \ldots, \mathfrak{s}_{i_k}, \mathfrak{d}_{j_k})$. Applying such a circuit at a (feasible) point $\mathbf{y}$ corresponds to changing the flow on the edges of $K_{M,N}$: We increase flow on all edges $\{\mathfrak{s}_{i_l}, \mathfrak{d}_{j_l}\}$ and decrease flow on all edges $\{\mathfrak{d}_{i_l}, \mathfrak{s}_{j_{l+1}}\}$ by the same arbitrary amount, the *s*tep length. (For a shorter wording, we will often say that we *increase* or *decrease* edges.)

This in particular ensures that the 'margin equations' defining the transportation polytope remain satisfied. However, we do not necessarily remain feasible as applying a circuit possibly decreases the flow on an edge $\{\mathfrak{d}_{i_l}, \mathfrak{s}_{j_{l+1}}\}$ below its lower bound of zero. So if we want to remain feasible, the step length at $\mathbf{y}$ can be at most the minimum over all $y_{j_l i_{l+1}}$. This implies that the circuit steps with respect to the four different concepts are as follows:

- $\mathcal{CD}$: We can apply any circuit with any step length.
- $\mathcal{CD}_f$: We can apply any circuit $(\mathfrak{s}_{i_1}, \mathfrak{d}_{j_1}, \mathfrak{s}_{i_2}, \mathfrak{d}_{j_2}, \ldots, \mathfrak{s}_{i_k}, \mathfrak{d}_{j_k})$ for which $y_{j_l i_{l+1}} > 0$ for all l. The step length is at most the minimum over all $y_{j_l i_{l+1}}$.
- $\mathcal{CD}_{fm}$: We can apply any circuit $(\mathfrak{s}_{i_1}, \mathfrak{d}_{j_1}, \mathfrak{s}_{i_2}, \mathfrak{d}_{j_2}, \ldots, \mathfrak{s}_{i_k}, \mathfrak{d}_{j_k})$ for which $y_{j_l i_{l+1}} > 0$ for all l. The step length equals the minimum over all $y_{j_l i_{l+1}}$.
- $\mathcal{CD}_e$: By Proposition 3.5, an edge step (*pivot*) goes from an assignment O to an assignment C that differs from O in exactly one edge. Being at a vertex O, such a step can be constructed by inserting an arbitrary edge $\{\mathfrak{s}_i, \mathfrak{d}_j\} \notin O$ into O. This closes an even cycle, which describes the circuit (edge direction) we apply. Again we alternately increase and decrease along this cycle by the minimum over all $y_{j_l, i_{l+1}}$, where we increase on the edge we inserted. Due to non-degeneracy (which we can assume here) this deletes exactly one edge and hence leads to an assignment C that is a neighboring vertex.

Observe that for $\mathcal{CD}_e$ every circuit step inserts one edge and deletes one edge and hence the corresponding support graphs always remain cycle free. In contrast to this, $\mathcal{CD}_{fm}$ can insert multiple edges while deleting at least one edge, so that there can be cycles. In weak circuit walks $\mathcal{CD}_f$ we can insert multiple edges and we do not have to delete an edge at all. Finally, infeasible points can appear in soft circuit walks $\mathcal{CD}$. That is why we do not consider a support graph in this case. Also recall that we can assume connectivity of the support graphs due to non-degeneracy only for the edge walks.

For sake of notation, we distinguish two types of distances from an assignment O to a fixed assignment F. We will use $\mathcal{CD}^O$, $\mathcal{CD}_e^O$, etc. to denote the respective circuit distances from O and F, while the *edge distance* $|O\backslash F|$ is simply the number of edges that are in O, but not in F. Note that $O\backslash F$ consists of those edges that have *to be deleted* (or are *to delete*) when walking from O to F.

Clearly we have $\mathcal{CD}_e^O \geq |O\backslash F|$: By applying a single pivot at an assignment O, one obtains an assignment C which has at most one additional edge (the new, inserted one) in common with F. In contrast, we can have $\mathcal{CD}_{fm}^O < |O\backslash F|$ as we will see in Example 3.6. This example was first mentioned in [7]. It illustrates a situation that is crucial for proving the Hirsch conjecture for $M = 2$ and for understanding the combinatorial diameter of transportation polytopes: At some point in our edge walk from the initial to the final assignment we have to delete an edge that actually exists in the final assignment!

Example 3.6. *Consider the edge walk from an assignment O to an assignment F in Figure 3.2. The nodes are labeled with the margin values, the edges with the current flow; the bold edges highlight the circuit we apply and the dashed edges are those we insert*

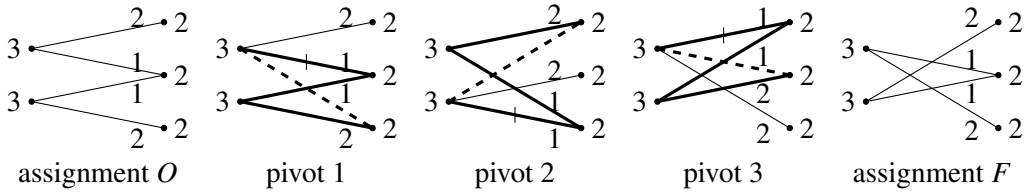

Figure 3.2: An edge walk from vertex O to vertex F of length three.

No matter which edge we insert in the first step, we have to delete an edge that is contained in F. Thus we need at least two more steps as we still have to insert two edges from F. Therefore, the edge walk above is a walk of minimum length and we have that $\mathcal{CD}_e$ is strictly larger than $|O \backslash F|$.

In contrast, we can go from O to F in only one circuit step of type $\mathcal{CD}_{fm}$ (and thus also $\mathcal{CD}_f$ or $\mathcal{CD}$): As we allow to go through the interior of the polytope, we can insert and delete two edges in just one step.

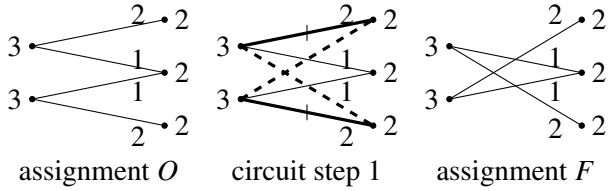

Figure 3.3: A feasible maximal circuit walk from vertex O to vertex F of length one.

We conclude this introductory section with a remark on the bound posed by the Hirsch conjecture. It is closely related to the diameter bound of $M + N - 1$ in Theorem 3.1 and 3.2: We know that the dimension of an $M \times N$ transportation polytope equals $(M - 1)(N - 1)$ [22], and for non-degenerate transportation polytopes the number of facets is equal to $M \cdot N - k$, where k is the number of *critical edges*. This follows immediately from Theorem 2 in [22]. We call an edge $\{\mathfrak{s}_i, \mathfrak{d}_j\}$ *critical* for a transportation polytope if the edge $\{\mathfrak{s}_i, \mathfrak{d}_j\}$ exists in every support graph, that is, $y_{ij} > 0$ for all solutions $\mathbf{y}$. With this we can state the Hirsch conjecture as

Conjecture 1 (Hirsch conjecture for transportation polytopes)**.** *The combinatorial diameter of an $M \times N$ transportation polytope is at most $M + N - 1 - k$, where k is the number of critical edges of the transportation polytope.*

Thus, the Hirsch conjecture would imply the upper bound of $M+N-1$. Observe that $M+N-1-k$ is precisely the number of edges in which two assignments of the corresponding transportation polytope can differ. In particular, for proving the Hirsch conjecture for $M \times N$ transportation polytopes it is enough to show that there is a sequence of pivots that inserts the edges in F one after another and no inserted edge is deleted again.

3.2 Upper bounds on circuit diameters of $M\times N$ transportation polytopes

One of the first upper bounds on the combinatorial diameter of $M\times N$ transportation polytopes was shown in [29], a quadratic upper bound of $M{\cdot}N$. Meanwhile, this was improved to linear upper bounds, the best one currently known is $8(M+N-2)$ [7]. By transitivity this translates to an upper bound on $C\mathcal{D}_{fm}$, $C\mathcal{D}_f$ and $C\mathcal{D}$.

However, for $C\mathcal{D}_f$ and $C\mathcal{D}$ we can show the much stronger bound of $M + N - 1$. This is an immediate consequence of Theorem 2.17 from the previous chapter.

Lemma 3.7. *For an $M\times N$ transportation polytope, the weak circuit diameter $C\mathcal{D}_f$ is bounded above by $M + N - 1$.*

Proof. A transportation polytope can be written as $P = \{\mathbf{y} \in \mathbb{R}^{M\times N} : A\mathbf{y} = \binom{\mathbf{u}}{\mathbf{v}}, \mathbf{y} \geq \mathbf{0}\}$, where the matrix $A \in \mathbb{Z}^{(M+N)\times(M\cdot N)}$ is the node-edge incidence matrix of $K_{M,N}$. Observe that A has row rank $M+N-1$ since of the $M+N$ margin equalities one is redundant and can be derived from the others. We further have $m_B = M{\cdot}N$ inequalities $y_{ij} \geq 0$. Thus, $C\mathcal{D}_f \leq C\mathcal{D}_{fs} \leq M + N - 1 - M{\cdot}N + M{\cdot}N$ by Theorem 2.17. □

For the circuit diameter $C\mathcal{D}_{fm}$ and the combinatorial diameter $C\mathcal{D}_e$ of transportation polytopes this bound seems too ambitious, as already indicated by the long history of attempts in proving the Hirsch conjecture. Instead we focus on investigating small instances, which might shed some light on this question. In the following two sections we have a look at $2\times N$ and $3\times N$ transportation polytopes. In these cases we are indeed able to prove the Hirsch bound!

However, the Hirsch conjecture remains open for $M\times N$ transportation polytopes for $M \geq 4$. We want to point out that some of the concepts we apply in the following can easily be adapted to the $M\times N$ case, whereas several arguments are specific for $M = 2, 3$. In particular, the lengthy case distinctions are not likely to be a reasonable approach for larger M.

3.3 Upper bounds on the combinatorial diameter

In this section we cover Theorem 3.3. Before turning to the actual proofs, we introduce of a marking system that is at the core of our approach. We use it for showing that in the $2\times N$ case we get an upper bound of N on the monotone diameter (Theorem 3.3 (i)) and for proving validity of the Hirsch conjecture for $3\times N$ transportation polytopes (Theorem 3.3 (ii)).

3.3.1 Basic concepts for the proofs

For proving upper bounds on the combinatorial diameter we will construct a walk from an initial assignment O to an assignment F by iteratively inserting edges from F into the current assignment, just as we did in Figure 3.2. In this process we now distinguish marked and unmarked edges in the support graph of the current assignment. In the beginning every edge is unmarked. Any pivot may

only insert edges from F and we mark every edge we insert. We are also allowed to mark certain edges (from F) that already exist in the current assignment without applying a pivot. Therein, an edge $\{\mathfrak{s}_i, \mathfrak{d}_j\}$ can be marked only if

1. $\mathfrak{d}_j$ is a leaf demand in F, or
2. $\mathfrak{d}_j$ is a mixed demand in F and all leaf edges (in F) incident to $\mathfrak{s}_i$ already exist in O and are marked.

We will consider the possible configurations of marked and unmarked edges in the mixed part of O that can appear throughout the process. For each such configurations we specify an edge from F to mark (after possibly inserting it). In this we never delete a marked edge. Following this approach we will need at most $|F|$ steps to get to the final assignment.

Like this we prove our upper bounds of $N + 1$ and $N + 2$, respectively, on the combinatorial diameter of $2\times N$ and $3\times N$ transportation polytopes. By refining our arguments we then even show the slightly stronger Hirsch conjecture for $M = 2, 3$ and in the case of $2\times N$ transportation polytopes we improve the diameter bound by one to N in general.

Before starting with the proofs, let us outline some general conventions, situations and arguments that frequently appear in our analysis.

In the sketches throughout this section, marked edges are drawn in bold while unmarked edges are illustrated using thin lines. Edges that are possibly marked are depicted as a plain line with a dashed bold line over it.

When talking about 'mixed' or 'leaf' edges without referring to a specific assignment, we always mean that the edges are mixed or leaves in the final assignment F. For example, '$\mathfrak{s}_i$ has all its leaf edges' means that all edges that are leaf edges in F and incident to $\mathfrak{s}_i$ in F also exist in O.

Recall that $\{\mathfrak{s}_i, \mathfrak{d}_j\}$ is an 'edge to increase' if $y^O_{ij} < y^F_{ij}$ and an 'edge to decrease' if $y^O_{ij} > y^F_{ij}$. Clearly, if there is an edge to increase incident to a node $\mathfrak{s}_i$ or $\mathfrak{d}_j$ in O, there also must be an edge to decrease incident to this node and vice versa. Further observe that edges to increase can only be edges to insert or edges that are mixed in O, while edges to decrease are edges we have to delete or edges that exist in O and are mixed in F. These principles are frequently used in our proofs.

Another important observation is the following: If there exist marked edges $\{\mathfrak{s}_{i_1}, \mathfrak{d}_j\}, \{\mathfrak{s}_{i_2}, \mathfrak{d}_j\} \in O$, then $\mathfrak{s}_{i_1}$ and $\mathfrak{s}_{i_2}$ already have all their leaf demands: $\{\mathfrak{s}_{i_1}, \mathfrak{d}_j\}$ and $\{\mathfrak{s}_{i_2}, \mathfrak{d}_j\}$ are marked and thus in F. In particular these edges are mixed in F. But these mixed edges can be marked only if all leaf edges incident to $\mathfrak{s}_{i_1}$ and $\mathfrak{s}_{i_2}$, respectively, already exist in O. So in particular we have that $O = F$ if E^O_m (mixed edges in O) is completely marked.

We finally present a lemma that will be useful for finding an edge to (possibly insert and) mark in many configurations that appear in our proofs. In this lemma we assume that the 'partially marked' assignment was obtained by the rules described above. 'An edge $e \in E^O_m$ is an even number of edges away from $\mathfrak{s}_i$ in E^O_m' means that if we consider the path with edges in E^O_m, starting at node $\mathfrak{s}_i$ and ending with the edge e, this path has an even number of edges.

Lemma 3.8. *Let $O \neq F$ be two assignments in a non-degenerate $2\times N$ or $3\times N$ transportation polytope, with O partially marked. Assume there is some $\mathfrak{s}_i$, such that all marked edges in E_m^O are an even number of edges away from $\mathfrak{s}_i$ in E_m^O. Then after (at most) one pivot, we may mark some edge in O.*

Proof. First note that the condition that all marked edges in the mixed part E_m^O are an even number of edges away from $\mathfrak{s}_i$ in E_m^O implies that a pivot that inserts an edge incident to $\mathfrak{s}_i$ cannot delete a marked edge in E_m^O as they are all increased. Thus, when inserting some edge $\{\mathfrak{s}_i, \mathfrak{d}_j\}$ incident to $\mathfrak{s}_i$, we only have to worry about deleting marked edges that are not mixed in O, that is, about the unique edge incident to $\mathfrak{d}_j$ in O. We then proceed according to the following rules, applied in that precise order:

1. If there is an unmarked leaf edge $\{\mathfrak{s}_i, \mathfrak{d}_j\}$, we insert it (if necessary) and mark it. Since this is a leaf edge in F, the decreasing edge incident to $\mathfrak{d}_j$ is not in F and thus not marked in O.

2. Else $\mathfrak{s}_i$ has all its leaf edges. If there is an edge $\{\mathfrak{s}_i, \mathfrak{d}_j\} \in E_m^F \cap O$, we mark it.

3. Else all edges incident to $\mathfrak{s}_i$ in E_m^O are unmarked (they are an odd number of edges away from $\mathfrak{s}_i$) and we have to delete these edges (otherwise we would have marked one of them in 1. or 2.). Since there is at least one such edge, there also is an edge incident to $\mathfrak{s}_i$ to increase. This must be an edge we still have to insert and it is mixed in F (otherwise we would have inserted it in 1.).

   If there is only one mixed edge $\{\mathfrak{s}_i, \mathfrak{d}_j\}$ to insert, this is the only mixed edge incident to $\mathfrak{s}_i$ in F (a second mixed edge would already exist in O. But then either we would have marked it in 2., or it would have already been marked and the lemma could not have been applied). We insert and mark this mixed edge $\{\mathfrak{s}_i, \mathfrak{d}_j\}$. Let $\{\mathfrak{s}_k, \mathfrak{d}_j\}$ (for some $i \neq k$) be the edge incident to $\mathfrak{d}_j$ in O. If this pivot deleted $\{\mathfrak{s}_k, \mathfrak{d}_j\}$, then $\mathfrak{s}_i$ would be incident to an edge to decrease ($\{\mathfrak{s}_i, \mathfrak{d}_j\}$ is mixed in F, but a leaf after this pivot), but $\mathfrak{s}_i$ would not be incident to an edge to increase, a contradiction. Hence no marked edge will be deleted.

4. Else there are two mixed edges incident to $\mathfrak{s}_i$ to insert (let these edges be $\{\mathfrak{s}_i, \mathfrak{d}_1\}$ and $\{\mathfrak{s}_i, \mathfrak{d}_2\}$; note this does not happen in the $2\times N$ case). If $\mathfrak{d}_1$ or $\mathfrak{d}_2$ is not incident to a marked edge in O, we insert the respective edge $\{\mathfrak{s}_i, \mathfrak{d}_j\}$ and mark it.

   Otherwise, both $\mathfrak{d}_1$ and $\mathfrak{d}_2$ are incident to marked edges, which necessarily are mixed in F. Without loss of generality let $\mathfrak{s}_i = \mathfrak{s}_3$ and let $\{\mathfrak{s}_1, \mathfrak{d}_1\}$, $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ be the marked mixed edges. Then we know that $\mathfrak{s}_1$ and $\mathfrak{s}_2$ (and also $\mathfrak{s}_3$) already have all their leaf edges and these are marked. We claim that in O we have without loss of generality one of the following configurations:

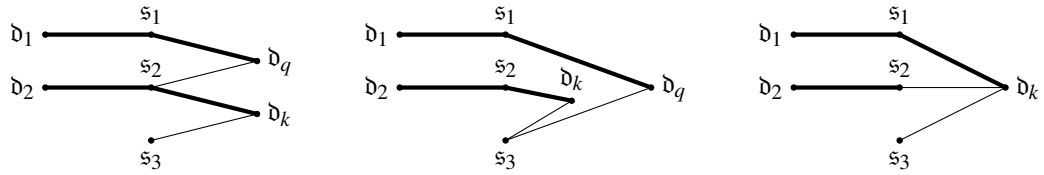

$\mathfrak{d}_1$ and $\mathfrak{d}_2$ are leaf demands in O

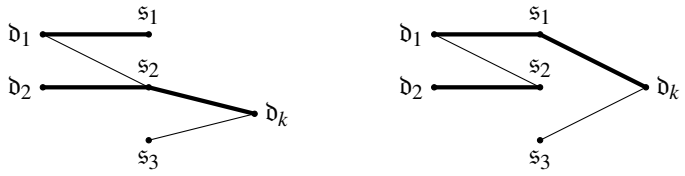

Without loss of generality $\mathfrak{d}_1$ is mixed demand in O

To see that these are the only possible configurations, first observe that not both $\mathfrak{d}_1$ and $\mathfrak{d}_2$ can be mixed in O. Clearly the mixed demands $\mathfrak{d}_q$ and $\mathfrak{d}_k$ in O must be leaves in F. Thus, the leaf edges incident to $\mathfrak{d}_q$ and $\mathfrak{d}_k$ already exist in O and must be marked. As all marked edges in E_m^O must be an even number of edges away from $\mathfrak{s}_3$, the marked edges incident to $\mathfrak{d}_q$ and $\mathfrak{d}_k$ can only be the ones depicted above. Note that the last configuration cannot even occur since the marked edge $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ would be an odd number of edges away from $\mathfrak{s}_3$.

In all other cases we can insert $\{\mathfrak{s}_3, \mathfrak{d}_1\}$: This is clear for the fourth configuration. For the configurations in the first row we have to show that $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ cannot be deleted. But if it was deleted, there would be an edge incident to $\mathfrak{s}_1$ to insert, but no edge to decrease.

This proves the claim. □

The proof of Lemma 3.8 essentially describes an algorithm to decide which pivot to use or simply which edge to mark. We refer to using this algorithm as *applying Lemma 3.8*.

3.3.2 $2\times N$ transportation polytopes

Validity of the Hirsch conjecture for a $2\times N$ transportation polytope, which was already proven in [12], implies a general upper bound of $N + 1$ on the combinatorial diameter. We now refine this bound by one to N, which is actually tight. Then we show that the monotone Hirsch conjecture holds with an upper bound of N as well, as for any linear function we can find a corresponding edge walk that is non-decreasing.

First recall that it suffices to consider non-degenerate polytopes for proving upper bounds on the combinatorial diameter. Let $O \neq F$ be two assignments of a non-degenerate $2\times N$ transportation polytope. These are spanning trees of $K_{2,N}$, consisting of $N + 1$ edges. It is not hard to see that every two assignments have at least two edges in common and thus O and F differ by at most $N - 1$ edges (see also Lemma 3.17). The following Example 3.9 shows that there are assignments

that indeed satisfy $|O \backslash F| = N - 1$. Even more, this example provides us with a family of $2\times N$ transportation polytopes with combinatorial diameter N. Thus the upper bound we prove in the following is tight in the sense that for all $N \geq 3$ there is a $2\times N$ transportation polytope with combinatorial diameter equal to N.

Example 3.9. *Consider an instance of a transportation problem with $u_1 = u_2 = 2N - 3$, $v_1 = 2N - 4$, and $v_j = 2$ for all $j = 2, \ldots, N$. This defines a non-degenerate transportation polytope. Now take the following two assignments O and F.*

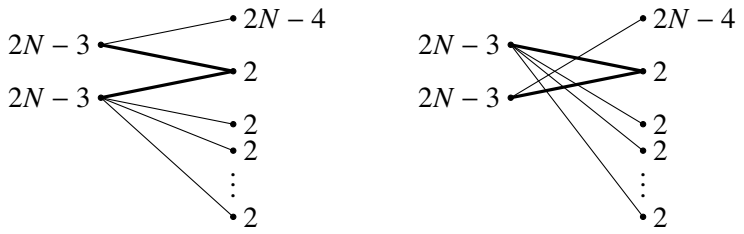

Figure 3.4: Two assignments with combinatorial distance N.

In the first step, no matter which edge from $F \backslash O$ we insert, we delete an edge that is contained in F. More precisely, the edge we delete is mixed in F and in O (the mixed edges have minimum flow of one among all edges in O). Thus, inserting any edge into O creates a new assignment that still differs from F in $N - 1$ edges. Hence we need at least $N - 1$ more pivots to reach the final assignment, resulting in a total of at least N steps.

Note that in the above example we had $D_m^O = D_m^F$ and the configuration is similar to the 2×3 transportation polytope in Example 3.6 from Section 3.1. This is the situation we have to take special care of when proving the Hirsch conjecture in the $2\times N$ case.

To show the upper bound of N, we use the following lemma about our marking system in the $2\times N$ case.

Lemma 3.10. *Let $O \neq F$ be two assignments in a non-degenerate $2\times N$ transportation polytope and let O be partially marked. Suppose there is some edge $e \in O \cap E_m^F$ such that*

1. *e is a leaf edge in O, or*
2. *$E_m^O = E_m^F$ and the other mixed edge in O is marked.*

Then applying Lemma 3.8 cannot delete e.

Proof. Let $e = \{\mathfrak{s}_2, \mathfrak{d}_1\}$. If e is a leaf edge in O and Lemma 3.8 inserts an edge that is a leaf in F, then e is not involved in the pivot and thus will not be deleted. Otherwise we either have that $\{\mathfrak{s}_2, \mathfrak{d}_1\}$ is a leaf edge in O and we insert the other mixed edge $\{\mathfrak{s}_1, \mathfrak{d}_1\}$, or we have $E_m^O = E_m^F$ and $\{\mathfrak{s}_1, \mathfrak{d}_1\} \in E_m^F$ is already marked. In the latter case we must be applying Lemma 3.8 to $\mathfrak{s}_2$ since $\mathfrak{s}_1$ is incident to a marked edge in the mixed part of O.

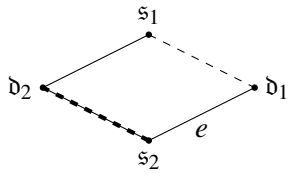

Inserting the mixed edge $\{\mathfrak{s}_1, \mathfrak{d}_1\}$

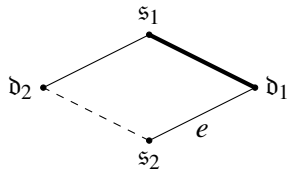

Inserting an edge when $E_m^O = E_m^F$

These are all cases that can appear when applying Lemma 3.8. In both configurations we know that $\mathfrak{s}_1$ must have all its leaf edges (since we either insert a mixed edge incident to $\mathfrak{s}_1$ or this edge is already marked). It remains to show that $\{\mathfrak{s}_2, \mathfrak{d}_1\}$ will not be deleted. But if it was deleted, $\mathfrak{s}_1$ would have an edge to decrease ($\{\mathfrak{s}_1, \mathfrak{d}_2\} \in O \backslash F$) but none to increase. □

We now can prove

Lemma 3.11. *The combinatorial diameter of a* $2\times N$ *transportation polytope is at most* N.

Proof. Starting with a partially marked assignment O, we show that after (at most) one pivot we may obtain a new assignment (which we will also call O) with one more edge from the final assignment F marked. The marked mixed edges in O look without loss of generality as one of the following cases:

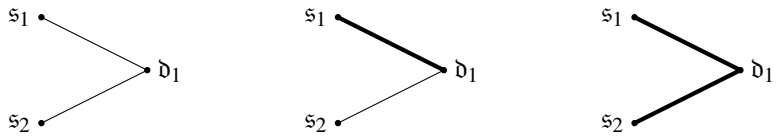

In the first two cases we may apply Lemma 3.8 to $\mathfrak{s}_2$ to obtain a marking after (at most) one pivot. In the third case, $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_1\}$ are both marked in O and thus mixed edges in F. In particular, we have $E_m^O = E_m^F$. Further we know that $\mathfrak{s}_1$ and $\mathfrak{s}_2$ both have all their leaf edges in F since they are incident to mixed edges that are marked. Hence we have $O = F$.

This immediately implies that the diameter is at most $N + 1$, since there are $|F| = N + 1$ edges to mark in F and each marking takes at most one pivot. To prove the diameter bound of N, we only have to show there always is some edge we can mark without applying a pivot. If there is a leaf edge $e \in F$ such that $e \in O$ when we start our marking process, we can mark e immediately and our proof is complete. Otherwise, no such edge exists. But then we have $E_m^O = E_m^F$: We know that $|O \cap F| \geq 2$ ($|O \cap F| = |O| + |F| - |O \cup F| \geq 2(N + 1) - 2N = 2$ since $|O \cup F| \leq 2N$) and there are no shared edges that are leaves in F. Hence the two shared edges are mixed in F and thus also mixed in O.

So let $E_m^O = E_m^F = \{e_1, e_2\}$. If neither e_1 nor e_2 is ever deleted when we perform our pivots, then we will be able to mark them both without performing a pivot and the diameter of the associated polytope will be at most $N - 1$. If one of them is deleted at some point during our process, without loss of generality say e_1, then e_2 becomes a leaf edge. Then e_2 will never be deleted: Whenever we apply Lemma 3.8, e_2 either remains a leaf edge or it becomes a mixed edge again by inserting e_1. But then e_1 will be marked. Hence e_2 will not be deleted by Lemma 3.10. Therefore, we will

eventually be able to mark e_2 without a pivot and the diameter of the associated polytope is at most N. This completes the proof. □

Observe that every critical edge can be marked without applying a pivot, as such edges exist in every assignment. Thus the proof of Lemma 3.11 actually tells us that the diameter is bounded above by $\min\{N, N+1-k\}$, where k is the number of critical edges. Therefore, $2\times N$ transportation polytopes satisfies the Hirsch conjecture with an upper bound of N as stated in Theorem 3.3 (i).

The monotone Hirsch conjecture

We now turn to a stronger version of the Hirsch conjecture. For an $M\times N$ transportation polytope P with k critical edges, we can state it as follows.

Conjecture 2 (The monotone Hirsch conjecture)**.** *Given any cost vector $\mathbf{s} = (s_{11}, \ldots, s_{MN})^\intercal \in \mathbb{R}^{M\times N}$ and any vertex $\mathbf{y}^O$ of P, there always is an edge walk of length at most $M+N-k$ from $\mathbf{y}^O$ to a vertex $\mathbf{y}^F = \arg\max_{\mathbf{y}\in P} \mathbf{s}^\intercal\mathbf{y}$ that visits vertices in a sequence of non-decreasing objective function values.*

We refer to the diameter of a polytope with respect to such a non-decreasing sequence of objective function values as the *monotone diameter*. To prove our upper bounds on this diameter, we only have to refine the previous proofs to make sure we construct a monotone edge walk.

But first note that for $2\times N$ transportation polytopes the vector $\mathbf{s} = (s_{11}, \ldots, s_{2N})^\intercal \in \mathbb{R}^{2\times N}$ already tells us what the maximizing vertex $\mathbf{y}^F$ looks like. Note that the assumption $s_{1i} - s_{2i} \geq s_{1(i+1)} - s_{2(i+1)}$ for all $i \leq N-1$ is no restriction, as it can be achieved by simply re-indexing the demands $\mathfrak{d}_j$.

Lemma 3.12. *Let $\mathbf{s} \in \mathbb{R}^{2\times n}$ satisfy $s_{1i} - s_{2i} \geq s_{1(i+1)} - s_{2(i+1)}$ for all $i \leq N-1$ and let j be a maximal index such that $\sum_{i=1}^{j-1} v_i < u_1$. Then the assignment $\mathbf{y}^F$ defined by*

- $y^F_{1i} = v_i$ *for* $i < j$, $y^F_{1j} = u_1 - \sum_{i=1}^{j-1} v_i$, *and* $y^F_{1i} = 0$ *for* $i > j$,
- $y^F_{2i} = 0$ *for* $i < j$, $y^F_{2j} = v_j - y^F_{1j}$, *and* $y^F_{2i} = v_i$ *for* $i > j$,

is an optimizer for $\max \mathbf{s}^\intercal\mathbf{y}$.

Proof. The index j and the values y^F_{1j}, y^F_{2j} are well-defined since there is no index j' with $\sum_{i=1}^{j'} v_i = u_1$ due to non-degeneracy of the transportation polytope. $\mathfrak{d}_j$ is the unique mixed demand in the assignment.

It suffices to prove that the reduced costs of all pivots possible at $\mathbf{y}^F$ are non-positive. Such a pivot corresponds to adding an edge incident to a demand $\mathfrak{d}_i$ for $i \neq j$. Since all $\mathfrak{d}_i$ with $i < j$ are incident to $\mathfrak{s}_1$ and all $\mathfrak{d}_i$ with $i > j$ are incident to $\mathfrak{s}_2$, the reduced costs satisfy

$$(s_{2i} - s_{1i}) + (s_{1j} - s_{2j}) \leq 0 \quad \text{due to } s_{1j} - s_{2j} \leq s_{1i} - s_{2i} \text{ for } i < j$$

and

$$(s_{1i} - s_{2i}) + (s_{2j} - s_{1j}) \leq 0 \quad \text{due to } s_{1j} - s_{2j} \geq s_{1i} - s_{2i} \text{ for } i > j.$$

□

We now show that the edge walk to the final vertex can be chosen as a sequence of non-decreasing objective function values. We split this proof into two parts, one for $D_m^O = D_m^F$ (which implies $E_m^O = E_m^F$ in the $2\times N$ case) and one for $D_m^O \neq D_m^F$. We begin with $D_m^O = D_m^F$.

Lemma 3.13. *Let* $\mathbf{s} \in \mathbb{R}^{2\times N}$, $\mathbf{y}^F = \arg\max_{\mathbf{y}\in P} \mathbf{s}^\intercal \mathbf{y}$. *Suppose we have* $O \neq F$ *but* $D_m^O = D_m^F$. *Then any pivot inserting an edge from* $F\backslash O$ *is nondecreasing.*

Proof. Let $D_m^O = D_m^F = \{\mathfrak{d}_j\}$ and let without loss of generality $\mathbf{s}$ be as in Lemma 3.12, i.e. $\mathbf{s} \in \mathbb{R}^{2\times N}$ satisfies $s_{1i} - s_{2i} \geq s_{1(i+1)} - s_{2(i+1)}$ for all $i \leq N-1$. Then inserting an edge from $F\backslash O$ means that we insert an edge $\{\mathfrak{s}_1, \mathfrak{d}_i\}$ with $i < j$ or an edge $\{\mathfrak{s}_2, \mathfrak{d}_i\}$ with $i > j$.

In the first case we have $s_{1i} - s_{2i} \geq s_{1j} - s_{2j}$ and thus obtain reduced costs of

$$(s_{1i} - s_{2i}) + (s_{2j} - s_{1j}) \geq 0\,.$$

In the second case we have $s_{1i} - s_{2i} \leq s_{1j} - s_{2j}$ and thus reduced costs

$$(s_{2i} - s_{1i}) + (s_{1j} - s_{2j}) \geq 0\,.$$

□

Lemma 3.14. *Let* $\mathbf{s} \in \mathbb{R}^{2\times N}$, $\mathbf{y}^F = \arg\max_{\mathbf{y}\in P} \mathbf{s}^\intercal \mathbf{y}$. *Suppose we have* $O \neq F$ *and* $D_m^O \neq D_m^F$. *Then there is some* $\mathfrak{s}_i$ *for which we may apply Lemma 3.8 such that the corresponding pivot is non-decreasing.*

Proof. Let again, without loss of generality, $\mathbf{s} \in \mathbb{R}^{2\times N}$ satisfy $s_{1i} - s_{2i} \geq s_{1(i+1)} - s_{2(i+1)}$ for all $i \leq N-1$. We have $D_m^F = \{\mathfrak{d}_j\}$, while $D_m^O = \{\mathfrak{d}_q\}$ for some $q \neq j$.

First consider the case $q < j$. Note that $q < j$ implies that $\{\mathfrak{s}_2, \mathfrak{d}_q\} \notin F$ and thus the mixed incident to $\mathfrak{s}_2$ in O is unmarked. Hence we may apply Lemma 3.8 to $\mathfrak{s}_2$. Thus, when we apply a pivot, we insert an edge $\{\mathfrak{s}_2, \mathfrak{d}_p\} \in F\backslash O$ for which $p \geq j \geq q$. Such a pivot has reduced costs

$$(s_{2p} - s_{1p}) + (s_{1q} - s_{2q}) \geq 0\,,$$

and thus a corresponding pivot obtained by applying Lemma 3.8 will be non-decreasing. The case $q > j$ follows analogously with the roles of $\mathfrak{s}_1$ and $\mathfrak{s}_2$ switched. □

Finally, we obtain the desired statement.

Lemma 3.15. *A* $2\times N$ *transportation polytope has monotone diameter at most* N.

Proof. For a given $\mathbf{s} \in \mathbb{R}^{2\times N}$, let $\mathbf{y}^F = \arg\max_{\mathbf{y}\in P} \mathbf{s}^\intercal \mathbf{y}$. Let F be the corresponding maximal assignment and O be the original assignment. By Lemma 3.11, it is possible to arrive at F after at most

N pivot steps – using the approach outlined in its proof. It suffices to see that one always can find a non-decreasing pivot.

If we have $D_m^O = D_m^F$ this follows from Lemma 3.13, as then all pivots inserting edges in $F \setminus O$ are non-decreasing. Otherwise there is a non-decreasing next pivot that adheres to the process in the proof of Lemma 3.11 by Lemma 3.14. This proves the claim. □

3.3.3 $3\times N$ transportation polytopes

Next we show that the Hirsch conjecture is true for the $3\times N$ transportation polytopes. It claims a bound of $N + 2 - k$ on the combinatorial diameter, where k is the number of critical edges. But before turning to the details of the proof, let us give a top-level view.

Proof of Theorem 3.3 (ii). Let P be a non-degenerate $3\times N$–transportation polytope with k critical edges. Let O be some initial vertex assignment and F be a final vertex assignment. We will present an algorithm that constructs an edge walk from O to F by choosing an edge from F to insert (if needed) and mark at each step of the edge walk. We start with all edges unmarked. Throughout the whole process, the conditions of our marking system (see Section 3.3.1) will always be satisfied, in particular no marked edge is ever deleted. Thus we need $|F|$ markings to reach our final assignment. This requires at most $|F| - k = 3 + N - k - 1 = N + 2 - k$ insertions (pivots), since the k critical edges exist in every assignment an thus can be marked at some point without applying a pivot. Hence the Hirsch conjecture holds for $3\times N$ transportation polytopes. □

The cusp of this argument is the lengthy algorithm for choosing which edge to mark (after possibly inserting it) presented below. Like in the proof of Lemma 3.11 we consider the possible configurations of marked and unmarked mixed edges in O in a case-by-case analysis. Therein we apply Lemma 3.8 whenever possible to easily prove the existence of a valid marking. However, there are several cases where it cannot be used, and this complicates the algorithm considerably. Further we must be wary of the following two configurations of marked mixed edges.

$D_m^O = 2$, Case 2c $\qquad$ $D_m^O = 2$, Case 3a

Figure 3.5: Two difficult configurations.

We avoid the first situation ($D_m^O = 2$, Case 2c) altogether, and we allow to enter the second of case ($D_m^O = 2$, Case 3a) only if $\mathfrak{d}_2$ satisfies some special requirements.

Algorithm for the proof of Theorem 3.3 (ii)

For the sake of a simple wording, we always refer to the current assignment as O. We start out with a copy of our initial assignment with no edges marked and iteratively (insert and) mark edges as explained in the following. We structure our investigation of the possible configurations of the mixed edges in our current assignment by the number of mixed demands $|D_m^O|$, which is equal to one or two, and the number of marked edges in E_m^O. Note that our labeling of the supply and demand nodes is no restriction, as it always can be achieved by a simple renaming.

Case $|D_m^O| = 1$

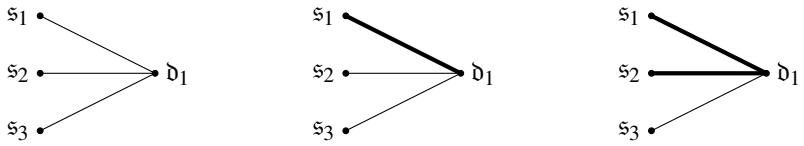

0 marked edges | 1 marked edge | 2 marked edges

If our mixed edges look like any of these cases, there is some $\mathfrak{s}_i$ such that we can apply Lemma 3.8 and mark an edge after (at most) one pivot. Otherwise all mixed edges are marked:

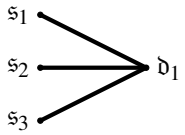

Then, by assumption, all $\mathfrak{s}_i$ have all their leaf edges and $E_m^O = E_m^F$. Hence $O = F$.

Case $|D_m^O| = 2$

0 or 1 marked edges

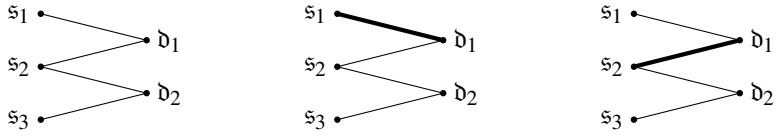

Case 0 | Case 1a | Case 1b

We may apply Lemma 3.8 to $\mathfrak{s}_3$ in any of the three configurations above, and mark an edge after (at most) one pivot.

2 marked edges

Case 2a.

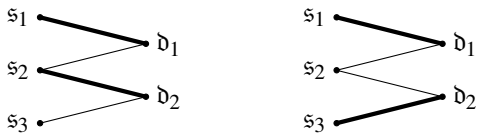

Case 2a(i) Case 2a(ii)

We can apply Lemma 3.8 to $\mathfrak{s}_3$ in case 2a(i) and to $\mathfrak{s}_2$ in case 2b(ii). Thus we mark one edge after (at most) one pivot.

Case 2b.

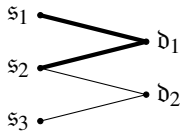

We do the following.

1. If $\mathfrak{d}_2$ is a leaf demand in F, we know that $\{\mathfrak{s}_3, \mathfrak{d}_2\} \in F$ ($\mathfrak{s}_1$ and $\mathfrak{s}_2$ already have all their leaf edges since they are incident to marked mixed edges), we mark it.

2. Else $\mathfrak{d}_2$ is a mixed demand in F. If $\{\mathfrak{s}_2, \mathfrak{d}_2\} \in F$, we mark it.

   Otherwise $\{\mathfrak{s}_1, \mathfrak{d}_2\}, \{\mathfrak{s}_3, \mathfrak{d}_2\} \in F$. We insert $\{\mathfrak{s}_1, \mathfrak{d}_2\}$:

   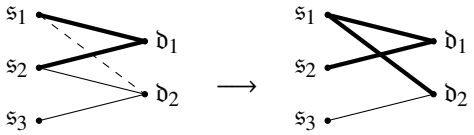

   This pivot deletes $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ (otherwise there is an edge incident to $\mathfrak{s}_2$ to delete, but no edge to increase).

Case 2c.

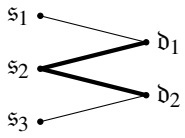

This case will never occur when we stick to the algorithm given here: The edge we mark always either is the one inserted or it already exists. Thus, to enter case 2c, one of edges $\{\mathfrak{s}_2, \mathfrak{d}_1\}$, $\{\mathfrak{s}_2, \mathfrak{d}_2\}$

would have already been marked in the previous assignment and it would have been in the mixed part of this previous assignment.

Applying Lemma 3.8 always marks an edge incident to some $\mathfrak{s}_i$ which is not currently not incident to a marked edge in the mixed part of the assignment. Thus we do not enter this case from case $|D_m^O| = 1$ or from $|D_M^O| = 2$ for 0 marked edges, 1 marked edge, or case 2a, where we apply Lemma 3.8.

Further, note that once there are two marked edges incident to the same $\mathfrak{d}_j$, they will forever be marked and in the mixed part of the assignment. Hence we never end up in case 2c from case 2b or any of the 3 marked edge cases.

3 marked edges

Case 3a.

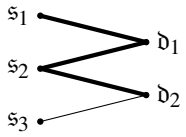

First note the following: When we enter this configuration, $\mathfrak{d}_2$ must be a mixed demand in F. Now observe that at least one of the edges $\{\mathfrak{s}_2, \mathfrak{d}_1\}$ or $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ must have been a marked mixed edge in the previous assignment (cf. case 2c). Hence we never enter case 3a when applying Lemma 3.8, as Lemma 3.8 always marks an edge incident to some $\mathfrak{s}_i$ with no edges marked in the mixed part of the assignment. Therefore we will never end up in this case from $|D_m^O| = 1$ or $|D_M^O| = 2$ for 0 marked edges, 1 marked edge, or case 2a.

Next, from case 2b,1. we enter case 3b. From case 2b,2. we in fact enter case 3a, but by assumption of case 2b,2 $\mathfrak{d}_2$ is a mixed edge as claimed. From case 3b we either enter the case with 4 marked edges, or we will enter case 3a (from case 3b,2. if $\{\mathfrak{s}_3, \mathfrak{d}_3\}$ is unmarked). But then we have that the claimed demand is a mixed demand in F.

On top of this we will show at the end of the discussion of case 3a that we either remain in this case with the same mixed edges and thus $\mathfrak{d}_2$ still is a mixed demand, or we enter another case.

Using this knowledge, we do the following.

1. If there is an unmarked leaf edge $\{\mathfrak{s}_3, \mathfrak{d}_3\}$, we insert $\{\mathfrak{s}_3, \mathfrak{d}_3\}$ (if necessary) and mark it. If $\{\mathfrak{s}_2, \mathfrak{d}_3\} \in O$, clearly no marked edges are deleted. Else we have $\{\mathfrak{s}_1, \mathfrak{d}_3\} \in O$.

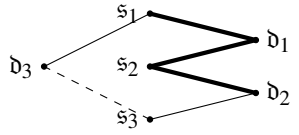

   But then our pivot cannot delete $\{\mathfrak{s}_2, \mathfrak{d}_1\}$. Suppose it does, and let C be the assignment obtained by deleting $\{\mathfrak{s}_2, \mathfrak{d}_1\}$. In C we would need to decrease $\{\mathfrak{s}_1, \mathfrak{d}_1\}$, and thus there is some

edge $\{\mathfrak{s}_1, \mathfrak{d}_j\}$ incident to $\mathfrak{s}_1$ to increase. However since $\mathfrak{s}_1$ already has all its leaf edges and they are marked and leaves in O, this would imply that $\{\mathfrak{s}_1, \mathfrak{d}_j\}$ is mixed in F. But $\mathfrak{d}_1$ and $\mathfrak{d}_2$ are the only mixed demands in F and clearly $\mathfrak{d}_2 \neq \mathfrak{d}_j$, a contradiction.

2. Else $\mathfrak{s}_3$ has all its leaf edges and these are marked. Since $\mathfrak{d}_2$ is a mixed demand in F, we have that $\{\mathfrak{s}_3, \mathfrak{d}_2\} \in F$ so we mark it.

Finally, as announced at the beginning of the case, we want to show the case we can enter by these steps. If step 2. is applied we enter the case with 4 marked edges. If step 1. is applied we arrive at one of the follow configurations:

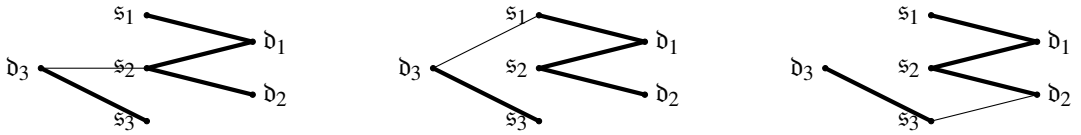

In the first two figures we have a configuration of type case 3b. In the last figure we have a configuration that is of type case 3a, but $\mathfrak{d}_2$ still is a mixed demand as claimed.

Case 3b.

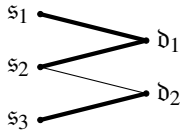

We do the following.

1. If $\mathfrak{d}_2$ is a mixed demand in F, we mark $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ or insert and mark $\{\mathfrak{s}_1, \mathfrak{d}_2\}$. The latter pivot deletes $\{\mathfrak{s}_2, \mathfrak{d}_2\}$, as otherwise there would be an edge incident to $\mathfrak{s}_2$ to delete, but no edge to increase. Note we can do these markings since $\mathfrak{s}_1$ and $\mathfrak{s}_2$ necessarily have all their leaf edges.

2. Else $\mathfrak{d}_2$ is a leaf demand in F. Note there are no edges incident to $\mathfrak{s}_3$ to delete, as for such an edge $\{\mathfrak{s}_3, \mathfrak{d}_3\}$, $\mathfrak{d}_3$ must be a leaf demand in F, but $\mathfrak{s}_1$ and $\mathfrak{s}_2$ already have all their leaf edges. As there is an edge incident to $\mathfrak{s}_3$ to increase (the leaf edge $\{\mathfrak{s}_3, \mathfrak{d}_2\}$), but no more edges to delete, there must be some leaf edge $\{\mathfrak{s}_3, \mathfrak{d}_3\} \in O$ to decrease which is mixed in F.

   We insert the other mixed edge $\{\mathfrak{s}_1, \mathfrak{d}_3\}$ or $\{\mathfrak{s}_2, \mathfrak{d}_3\}$ and mark it.

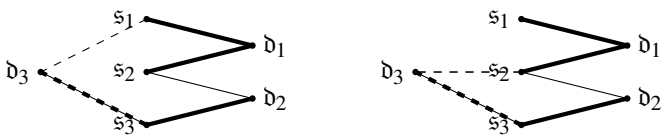

$\{\mathfrak{s}_1, \mathfrak{d}_3\} \in F$ $\{\mathfrak{s}_3, \mathfrak{d}_2\} \in F$

If $\{\mathfrak{s}_3, \mathfrak{d}_3\}$ is marked in O, the assignments O and F differed in exactly one edge and we just inserted the last edge from F. Thus we arrive at F after this pivot by Proposition 3.5. Otherwise $\{\mathfrak{s}_3, \mathfrak{d}_3\}$ is unmarked. Then if we insert $\{\mathfrak{s}_2, \mathfrak{d}_3\}$, clearly no marked edges are deleted. If instead we insert $\{\mathfrak{s}_1, \mathfrak{d}_3\}$, we know that $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ is not deleted as otherwise $\mathfrak{s}_2$ would have two edges to decrease but none to increase.

4 marked edges

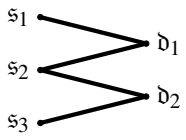

All $\mathfrak{s}_i$ are incident to marked mixed edges and thus have all their leaf edges. Further we must have $E_m^O = E_m^F$ and hence $O = F$.

This concludes the proof of Theorem 3.3 (ii).

3.4 Upper bounds on the circuit diameter

This section is dedicated to proving Theorem 3.4 on the circuit diameter $C\mathcal{D}_{fm}$. In fact, we prove the even stronger statement that the circuit distance from an assignment O to an assignment F is at most $|O \backslash F|$. In terms of the support graphs this means that we can delete at least one edge from $O \backslash F$ in every maximal circuit step.

Recall that we cannot exclude the degenerate case as it is not clear whether for every degenerate transportation polytope there is a perturbed non-degenerate transportation polytope bounding the circuit diameter of the original one above. For transportation polytopes this means that the support graphs are not necessarily connected, and thus the vertices are described by spanning forests instead of spanning trees.

3.4.1 $2\times N$ transportation polytopes

We begin with the $2\times N$ part of Theorem 3.4. First observe that we can describe a circuit step by two disjoint edges from the current assignment on which we want to *decrease* flow, one edge $\{\mathfrak{s}_1, \mathfrak{d}_j\}$ incident to $\mathfrak{s}_1$ and one edge $\{\mathfrak{s}_2, \mathfrak{d}_l\}$ incident to $\mathfrak{s}_2$ since this implies that we increase on $\{\mathfrak{s}_1, \mathfrak{d}_l\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_j\}$ (dashed). In particular, the latter edges are inserted if they do not exist in the current assignment.

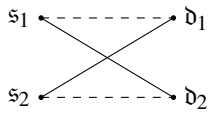

Figure 3.6: A circuit induced by the edges to decrease (solid).

Lemma 3.16. *Let O and F be two vertices of a $2\times N$ transportation polytope. Then the circuit distance $\mathcal{CD}_{fm}$ from O to F is at most $|O\backslash F|$. Further, if $|F| = N$ then the circuit distance $\mathcal{CD}_{fm}$ from O to F is at most $|O\backslash F| - 1$.*

Proof. If $O \neq F$, there must be an edge in $O\backslash F$ that we have to delete. We show that there is a circuit step that deletes such an edge in $O\backslash F$ and does not insert any edge not contained in F.

Case 1: There are edges incident to both $\mathfrak{s}_1$ and $\mathfrak{s}_2$ to delete.
Let $\{\mathfrak{s}_1, \mathfrak{d}_j\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_l\}$ be those edges. We apply the pivot that reduces flow on these particular edges. This deletes at least one of these edges and increases $\{\mathfrak{s}_1, \mathfrak{d}_l\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_j\}$, both of which must be in F (as both $\mathfrak{d}_l$ and $\mathfrak{d}_j$ must be connected to some supply point in F).

Case 2: One supply point still has an edge to delete while the other does not.
Without loss of generality, assume there is an edge $\{\mathfrak{s}_1, \mathfrak{d}_l\}$ incident to $\mathfrak{s}_2$ to delete, but there is no edge incident to $\mathfrak{s}_1$ to delete. Then we have to increase and hence there is an edge $\{\mathfrak{s}_1, \mathfrak{d}_j\}$ incident to $\mathfrak{s}_1$ to decrease but no edge to delete. Thus $\{\mathfrak{s}_1, \mathfrak{d}_j\}$ is a mixed edge in F. Since $\mathfrak{s}_1$ is incident to at most one mixed edge in F, this is the only edge incident to $\mathfrak{s}_1$ to decrease. We apply the pivot that decreases $\{\mathfrak{s}_1, \mathfrak{s}_j\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_l\}$. Assume it would delete $\{\mathfrak{s}_1, \mathfrak{d}_j\}$. Then there would be no more edges incident to $\mathfrak{s}_1$ to decrease, but an edge to increase ($\{\mathfrak{s}_1, \mathfrak{d}_j\}$ would have to be inserted again). Thus $\{\mathfrak{s}_2, \mathfrak{d}_l\}$ is the only edge that is deleted.

Note in both cases, we delete at least one edge from $O\backslash F$, and never delete an edge from F. Also, we only insert edges from F. Hence we will have $|O\backslash F| = 0$ after at most $|O\backslash F|$ circuit steps, at which point $O = F$.

Finally, we look at the case $|F| = N$. First note that $|F| = N$ implies that the final assignment is not connected and hence there is no mixed demand in F. In order to show that the circuit distance is at most $|O\backslash F| - 1$, it is enough to find a circuit step that deletes two edges from $O\backslash F$ at once. Therefore, consider the last step of a circuit walk constructed as described above. Recall that we already showed that we never delete an edge that is contained in F.

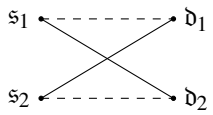

Then last circuit step; dashed edges are increasing, solid are decreasing

Suppose that one of the edges being decreased was not deleted, without loss of generality say $\{\mathfrak{s}_1, \mathfrak{d}_2\}$. Then since this is the last circuit step we have $\{\mathfrak{s}_1, \mathfrak{d}_2\}, \{\mathfrak{s}_2, \mathfrak{d}_2\} \in F$. This implies that F has a mixed demand $\mathfrak{d}_2$, but this is a contradiction as there are no mixed demands in F. Hence this

last circuit step indeed deletes two edges from $O\backslash F$. Combining this with the above arguments we get a circuit distance of $|O\backslash F| - 1$ from O to F if $|F| = N$. □

Now $C\mathcal{D}_{fm} \leq N - 1$ is a consequence of the following simple observation.

Lemma 3.17. *Let O and F be two vertices of a (possibly degenerate) $2\times N$ transportation polytope P. Then either $|O\backslash F| \leq N - 1$ or $|O\backslash F| = N$ and $|F| = N$*

Proof. We have $|O\backslash F| + |F| = |O \cup F| \leq 2N$ (this is the number of edges of $K_{2,N}$) and $|F| \geq N$ (a spanning forest in $K_{2,N}$ must have at least N edges). Thus we get $|O\backslash F| \leq 2N - |F| \leq 2N - N = N$. Further, if $|O\backslash F| = N$ then we have $|F| \leq 2N - |O\backslash F| = 2N - N = N$. As always $|F| \geq N$, this implies that we have $|F| = N$ if $|O\backslash F| = N$. □

Theorem 3.1 already tells us that this bound is tight for all N. However, in the $2\times N$ case one can easily state an explicit example of a transportation polytope with $C\mathcal{D}_{fm} = N - 1$ for arbitrary N.

Example 3.18. *Consider the (non-degenerate) transportation polytope given by margins $u_1 = u_2 = 2N - 1$, $v_1 = 2N$, $v_j = 2$ for $j = 2, \ldots, N$. The two assignments below have circuit distance $C\mathcal{D}_{fm}$ equal to $N - 1$: Every circuit step can insert at most one edge incident to $\mathfrak{s}_1$ and we have to add $N - 1$ such edges.*

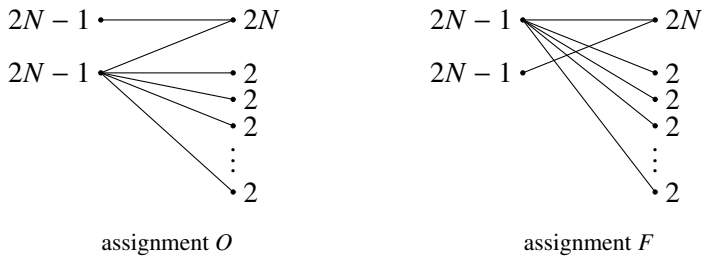

Figure 3.7: Two assignments with circuit distance $N - 1$.

This gives as an example of transportation polytopes with circuit diameter $C\mathcal{D}_{fm} = N - 1$. Note that we also have combinatorial diameter $N - 1$ since $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_1\}$ are critical edges.

3.4.2 $3\times N$ transportation polytopes

We now turn to the $3\times N$ part of Theorem 3.4. The discussion for $M = 3$ will be more involved than the $2\times N$ case (note that the circuits cannot be not characterized that easily anymore). However, our general approach can still be applied: We will show that there always is a maximal circuit step that deletes an edge from $O\backslash F$ while only inserting edges from F. Like this we prove an upper bound of $|O\backslash F| \leq M + N - 1 = M + 2$ on the circuit distance from an assignment O to an assignment F. Note that this is a slightly stronger statement that the one for the combinatorial distances as there are assignments with $C\mathcal{D}_{fm} \leq |O\backslash F| < C\mathcal{D}_e$. However, by Theorem 3.2 the best

possible general upper bound on the circuit diameter $C\mathcal{D}_{fm}$ is $M + 2$, which coincides with the general upper bound on the combinatorial diameter of $3\times N$ transportation polytopes.

Theorem 3.4 (ii) follows immediately from the following Lemma.

Lemma 3.19. *Let O and F be two vertices of a $3\times N$ transportation polytope. Then the circuit distance $C\mathcal{D}_{fm}$ from O to F is at most $|O\backslash F|$.*

Proof. It is enough to prove that there is a circuit step that deletes an edge from $O\backslash F$, while increasing (and hence inserting) only edges that exist in F. Applying such steps consecutively leads to a circuit walk of length at most $|O\backslash F|$.

With the following case-by-case discussion, we present a procedure for finding such a circuit step.

Case 1: There are edges incident to $\mathfrak{s}_1$, $\mathfrak{s}_2$ and $\mathfrak{s}_3$ to delete.
Let without loss of generality $\{\mathfrak{s}_i, \mathfrak{d}_i\} \in O\backslash F$ for $i = 1, 2, 3$. Then there must be edge $\{\mathfrak{s}_{k_i}, \mathfrak{d}_i\} \in F$ to increase for all i. The union of these six edges ($\{\mathfrak{s}_i, \mathfrak{d}_i\} \in O\backslash F$ and $\{\mathfrak{s}_{k_i}, \mathfrak{d}_i\} \in F$) on (at most) six vertices $\mathfrak{s}_1, \mathfrak{s}_2, \mathfrak{s}_3, \mathfrak{d}_1, \mathfrak{d}_2, \mathfrak{d}_3$ must contain a cycle that describes a circuit that only decreases edges in $O\backslash F$ and only increases edges in F. Thus we found a desired circuit step.

Case 2: Two supply nodes have edges to delete while the third has none.
Without loss of generality let the two nodes with edges to delete be $\mathfrak{s}_1$ and $\mathfrak{s}_3$ and let the third node be $\mathfrak{s}_2$. Then every edge incident to $\mathfrak{s}_2$ to delete must be a mixed edge in F and thus there are *at most* two such edges. The circuit step we apply must not delete these edges!

1. If there are no edges incident to $\mathfrak{s}_2$ to decrease, we can find $\{\mathfrak{s}_1, \mathfrak{d}_1\}, \{\mathfrak{s}_3, \mathfrak{d}_3\} \in O\backslash F$ to delete such that there exist edges $\{\mathfrak{s}_3, \mathfrak{d}_1\}, \{\mathfrak{s}_1, \mathfrak{d}_3\} \in F$ to increase. We can apply the circuit induced by these four edges; it does not contain $\mathfrak{s}_2$.

2. Assume there is exactly one edge incident to $\mathfrak{s}_2$ to decrease (we call this edge $\{\mathfrak{s}_2, \mathfrak{d}_2\}$). Let without loss of generality $\{\mathfrak{s}_1, \mathfrak{d}_1\}, \{\mathfrak{s}_3, \mathfrak{d}_3\} \in O\backslash F$; in particular these edges have to be decreased. For each $i = 1, 2, 3$ choose an edge $\{\mathfrak{s}_{k_i}, \mathfrak{d}_i\}$ to be increased, but select $\mathfrak{s}_{k_i} = \mathfrak{s}_2$ only if there are no other options. These six edges on at most six vertices induce a circuit that increases only edges to increase and decreases edge that have to be decreased.

   If this circuit does not contain $\mathfrak{s}_2$, we can apply it and are done. If this circuit contains $\mathfrak{s}_2$, we need to show that applying it does not delete $\{\mathfrak{s}_2, \mathfrak{d}_2\}$. Suppose it does and let C be the new assignment. Then $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ needs to be increased (inserted) again in C. Hence (in C) $\mathfrak{s}_2$ must be incident to some edge to decrease. Since $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ was the only such edge in O, the decreasing edge in C can only be the edge we increased in the previous circuit step, which is $\{\mathfrak{s}_2, \mathfrak{d}_i\}$ for some $i = 1, 3$. But in C there can be no edge to increase incident to $\mathfrak{d}_i$: In O, $\{\mathfrak{s}_2, \mathfrak{d}_i\}$ was the only such edge (otherwise we would not have chosen $\{\mathfrak{s}_{k_i}, \mathfrak{d}_i\} = \{\mathfrak{s}_2, \mathfrak{d}_i\}$), but now it must be decreased. Thus the only edge incident to $\mathfrak{d}_i$ that possibly has to be increased C is the one we decreased in the circuit step from O to C. But this is not an edge to increase as $\{\mathfrak{s}_i, \mathfrak{d}_i\} \notin F$. Therefore, $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ cannot be deleted when applying this circuit.

3. Otherwise there are two edges incident to $\mathfrak{s}_2$ to decrease (let these edges be $\{\mathfrak{s}_2, \mathfrak{d}_{21}\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_{23}\}$). If there is a circuit that does not involve $\mathfrak{s}_2$, that is, there are $\{\mathfrak{s}_1, \mathfrak{d}_1\}, \{\mathfrak{s}_3, \mathfrak{d}_3\} \in O \backslash F$ such that $\{\mathfrak{s}_3, \mathfrak{d}_1\}, \{\mathfrak{s}_1, \mathfrak{d}_3\} \in F$, we apply this circuit.

   Else we may assume without loss of generality that for all edges $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ to delete, we have $\{\mathfrak{s}_3, \mathfrak{d}_1\} \notin F$. Recall that $\{\mathfrak{s}_2, \mathfrak{d}_{21}\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_{23}\}$ are mixed edges in F. Thus we have two more mixed edges in F, without loss of generality $\{\mathfrak{s}_1, \mathfrak{d}_{21}\}, \{\mathfrak{s}_3, \mathfrak{d}_{23}\} \in F$. Now, having to decrease $\{\mathfrak{s}_2, \mathfrak{d}_{21}\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_{23}\}$ in O implies that we have to increase $\{\mathfrak{s}_1, \mathfrak{d}_{21}\}$ and $\{\mathfrak{s}_3, \mathfrak{d}_{23}\}$ (as $\{\mathfrak{s}_1, \mathfrak{d}_{23}\}, \{\mathfrak{s}_3, \mathfrak{d}_{21}\} \notin F$). Further, as we have to delete $\{\mathfrak{s}_3, \mathfrak{d}_3\}$, there must be $\{\mathfrak{s}_i, \mathfrak{d}_3\} \in F$ to increase for $i = 1$ or $i = 2$. Thus, in O we are in one of the situations depicted below, where dashed edges have to be increased, while solid edges have to be decreased; bold edges are in F, while non-bold edges are in $O \backslash F$. Recall that for all edges $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ to delete we must have $\{\mathfrak{s}_2, \mathfrak{d}_1\} \in F$ by assumption.

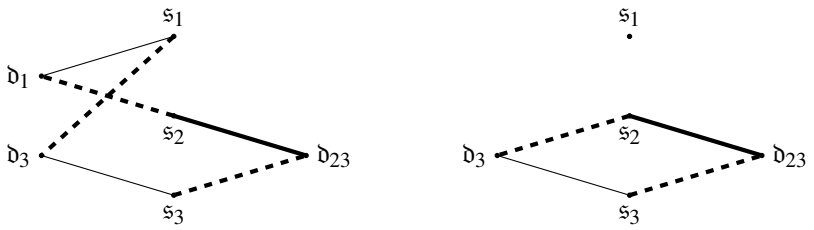

Two configurations of edges in O

   Note that each of these configurations induces a circuit. We claim that we can apply these circuits without deleting an edge from F. So assume that applying one of these circuits leads to an assignment C with $\{\mathfrak{s}_2, \mathfrak{d}_{23}\} \notin C$. Then $\{\mathfrak{s}_2, \mathfrak{d}_{23}\}$ has to be increased again in C and thus $\{\mathfrak{s}_3, \mathfrak{d}_{23}\}$ must be decreased (observe that $\{\mathfrak{s}_1, \mathfrak{d}_{23}\} \notin O$ as in this case we could have applied a circuit not involving $\mathfrak{s}_2$). But then there is no edge incident to $\mathfrak{s}_3$ to increase in C: Such an edge was to be increased in O as well, or it was decreased by the previous circuit step. The latter was the case for $\{\mathfrak{s}_3, \mathfrak{d}_3\} \notin F$ which must not be increased at any time. We claim that the edge $\{\mathfrak{s}_3, \mathfrak{d}_{23}\}$ (the edge that now has to be decreased) was the *only* edge incident to $\mathfrak{s}_3$ to increase in O. Clearly, there can be no other mixed edge (in F) incident to $\mathfrak{s}_3$ to increase (in O). So assume that $\mathfrak{s}_3$ has a leaf edge $\{\mathfrak{s}_3, \mathfrak{d}_j\}$ to increase. Then for $i = 1$ or $i = 2$ we have that $\{\mathfrak{s}_i, \mathfrak{d}_j\} \in O$ must be decreased, in particular deleted. But there is no edge incident to $\mathfrak{s}_2$ to delete, and we cannot have $i = 1$ as we assumed that for all $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ to delete, we have $\{\mathfrak{s}_3, \mathfrak{d}_1\} \notin F$. Thus, there is no edge incident to $\mathfrak{s}_3$ to increase in C and therefore $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ was not deleted. Hence we can apply one of the circuits depicted above.

Case 3: There is only one supply node with an edge to delete.

Without loss of generality there are edges incident to $\mathfrak{s}_1$ to delete, but not to $\mathfrak{s}_2$ and $\mathfrak{s}_3$. Recall that $\mathfrak{s}_2$ and $\mathfrak{s}_3$ appear in our circuits only if they have edges to be decreased at all. If existing, such

edges must not be deleted and hence must be mixed edges in F. Note that there has to be such an edge for at least one of $\mathfrak{s}_2$, $\mathfrak{s}_3$ since $\mathfrak{s}_1$ has an edge to delete.

1. There is at most one edge to decrease incident to either of $\mathfrak{s}_2$ and $\mathfrak{s}_3$ (let $\{\mathfrak{s}_2, \mathfrak{d}_2\}$, $\{\mathfrak{s}_3, \mathfrak{d}_3\}$ be those edges, if existing). Let $\{\mathfrak{s}_1, \mathfrak{d}_1\} \in O \backslash F$ be an edge to delete. First observe that there is at most one edge incident to $\mathfrak{d}_1$ to increase (Otherwise we would have $\{\mathfrak{s}_2, \mathfrak{d}_1\}, \{\mathfrak{s}_3, \mathfrak{d}_1\} \in F$ are mixed edge to increase in O. Thus the both mixed edges to decrease, $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ and $\{\mathfrak{s}_3, \mathfrak{d}_3\}$, must exist in O. But then these four mixed edges would form a cycle in F.). Let without loss of generality $\{\mathfrak{s}_2, \mathfrak{d}_1\}$ be this one edge to increase. Then the mixed edge $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ to decrease must exist and hence $\{\mathfrak{s}_1, \mathfrak{d}_2\}$ or $\{\mathfrak{s}_3, \mathfrak{d}_2\}$ has to be increased. Thus, we have one of the following two situations in O (edges illustrated as before).

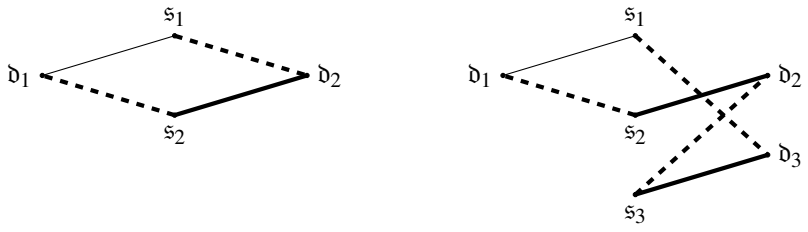

Two configurations of edges in O

   For the second configuration, observe that $\{\mathfrak{s}_2, \mathfrak{d}_3\}$ cannot be edge incident to $\mathfrak{d}_3$ to increase as this would create a cycle in F. We claim that applying the circuit induced by the respective constellation does not delete edges that are in F. Therefore, let C be the new assignment.

   If $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ was deleted, it has to be increased in C. The edge incident to $\mathfrak{s}_2$ to decrease could only be the previously increased edge $\{\mathfrak{s}_2, \mathfrak{d}_1\}$. But in C there can be no edge incident to $\mathfrak{d}_1$ to increase: $\{\mathfrak{s}_1, \mathfrak{d}_1\} \notin F$, $\{\mathfrak{s}_2, \mathfrak{d}_1\}$ has to be decreased, and $\{\mathfrak{s}_3, \mathfrak{d}_1\}$ was not an edge to increase in O and thus is no such edge in C (it was not affected by the pivot).

   Similarly, if $\{\mathfrak{s}_3, \mathfrak{d}_3\} \notin C$ (which can only happen if we apply the latter circuit), it has to be increased in C, and the edge incident to $\mathfrak{s}_3$ to decrease could only be the previously increased edge $\{\mathfrak{s}_3, \mathfrak{d}_2\}$. Thus, $\{\mathfrak{s}_2, \mathfrak{d}_2\}$ has to be increase (observe that $\{\mathfrak{s}_1, \mathfrak{d}_2\} \notin F$). But as before, there can be no edge incident to $\mathfrak{s}_2$ to decrease.

2. There are two edges incident to either $\mathfrak{s}_2$ or $\mathfrak{s}_3$ to decrease (without loss of generality let $\{\mathfrak{s}_2, \mathfrak{d}_1\}$ and $\{\mathfrak{s}_2, \mathfrak{d}_3\}$ be those edges). Since these are mixed edges, we have without loss of generality that $\{\mathfrak{s}_1, \mathfrak{d}_1\}$ and $\{\mathfrak{s}_3, \mathfrak{d}_3\}$ are the remaining mixed edges in F. As $\{\mathfrak{s}_2, \mathfrak{d}_3\}$ has to be decreased, $\{\mathfrak{s}_3, \mathfrak{d}_3\}$ has to be increased. Thus there must be an edge incident to $\mathfrak{s}_3$ to decrease and this edge must be mixed in F by assumption of case 3. But this implies that F has at least five mixed edges, a contradiction.

Case 4: There are no edges to delete.
In particular $O \backslash F = \emptyset$ and thus $O = F$.

This proves the claim. □

3.5 Lower bounds on circuit diameters of $M\times N$ transportation polytopes

We now turn to lower bounds on the circuit diameters. We will construct instances of transportation polytopes with assignments that differ by $\min\{(M-1)(N-1), M+N-1\}$ edges and show that by a perturbation of the margins we can ensure that every circuit walk from one assignment to the other one takes at least $|O\backslash F|$ steps. Note that this is always true for the combinatorial walks, as we delete exactly one edge in every step.

The following Lemma immediately implies Theorem 3.1. In the proof we outline a general principle of constructing a set of margins with the property mentioned above. Recall that we have $M+N-1 \leq (M-1)(N-1)$ for all $M \leq N$ but $M=2$ or $M=N=3$.

Lemma 3.20. *For all M, N, there is an $M\times N$ transportation polytope with soft circuit diameter CD at least* $\min\{\,(M-1)(N-1),\ M+N-1\,\}$.

Proof. We begin by constructing an $M\times N$ transportation polytope P with margins $\mathbf{u}, \mathbf{v}$ and two vertices O, F such that there is a sign-compatible circuit walk from O to F that uses exactly $k = \min\{(M-1)(N-1), M+N-1\}$ linearly independent circuits. Recall that $\dim(P) = (M-1)(N-1)$ and that the matrix defining an $M\times N$ transportation polytope has row rank $M+N-1$.

Consider the set of circuits depicted in Figure 3.8, where the dashed edges are the ones being increased. They are sign-compatible and linearly independent, as all of them use at least one edge no other circuit uses. It is not difficult to check that for all $M \leq N$, there exist at least k such circuits and that neither the edges to decrease form a cycle nor the edges to increase do.

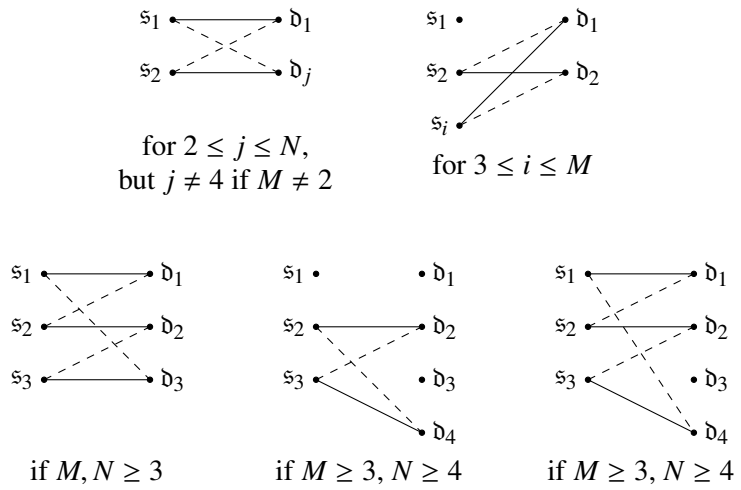

Figure 3.8: A set of sign-compatible, linearly independent circuits.

Given a set of k such circuits $\mathbf{g}^i$, define the margins $\mathbf{u}, \mathbf{v}$ of P to component-wisely be the number of circuits in this set the respective point lies on. Then the set of edges to decrease induce a vertex O as they do not form a cycle, and the set of edges to increase induce a vertex F. Note that we can

write $\mathbf{y}^F - \mathbf{y}^O = \sum_{i=1}^{k} \alpha_i \mathbf{g}^i$ for appropriate $\alpha_i > 0$.

Let us now introduce the perturbation of the margins without loss of generality 'along $\mathbf{g}^1$'. This defines a new polytope P' as follows: We derive new margins $\mathbf{u}', \mathbf{v}'$ from $\mathbf{u}, \mathbf{v}$ by choosing an ϵ and then setting $u'_j = u_j + \epsilon$ and $v'_j = v_j + \epsilon$ for all supply and demand points incident with the circuit $\mathbf{g}^1$. For ϵ sufficiently small, the same support graphs $B(\mathbf{y}^O)$, $B(\mathbf{y}^F)$ still induce vertices $\mathbf{y}'^O$, $\mathbf{y}'^F$ of P' and we have $\mathbf{y}'^F - \mathbf{y}'^O = (\alpha_1 + \epsilon)\mathbf{g}^1 + \sum_{i=2}^{k} \alpha_i \mathbf{g}^i$. Observe that such a perturbation can be done with respect to any subset of the $\mathbf{g}^i$; here each $\mathbf{g}^i$ gets its own small ϵ_i. We actually can do this by a sequence of perturbations along one circuit each time, as they are sign-compatible.

We now use such perturbations to get a transportation polytope P' with vertices $\mathbf{y}'^O$ and $\mathbf{y}'^F$ that have soft circuit distance at least k. As the $\mathbf{g}^i$ are linearly independent, no strict subset suffices to be able to walk from y^O to y^F in P. Now assume that there is another set of (up to) $k-1$ circuits $\mathbf{g}'^1, \dots, \mathbf{g}'^{k-1}$ such that $\mathbf{y}^O - \mathbf{y}^F$ is in the linear subspace spanned by $\mathbf{g}'^1, \dots, \mathbf{g}'^{k-1}$. There is a $\mathbf{g}^i$ which is linearly independent from $\mathbf{g}'^1, \dots, \mathbf{g}'^{k-1}$. Perturbing P along such a $\mathbf{g}^i$ yields an infinite set of polytopes P' for which $\mathbf{y}'^O - \mathbf{y}'^F$ is not in the span of $\mathbf{g}'^1, \dots, \mathbf{g}'^{k-1}$. As there is only a finite number of sets of up to $k-1$ circuits (recall these only depend on M and N, not the margins), there is a perturbation (or rather a sequence of perturbations) along the $\mathbf{g}^i$ such that we obtain a polytope P' for which $\mathbf{y}'^F - \mathbf{y}'^O$ is not in the linear subspace spanned by any set of (up to) $k-1$ circuits. This proves the claim. □

Let us demonstrate what a perturbation as described in Lemma 3.20 looks like.

Example 3.21. *Recall Example 3.6 and the circuit walk taking just a single circuit step. We call this circuit* $\mathbf{g}$.

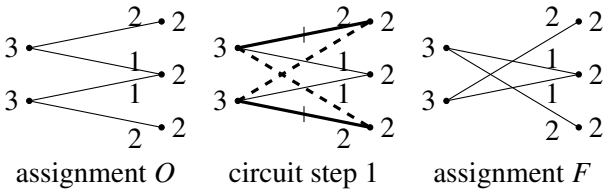

Figure 3.9: A circuit walk from vertex O to vertex F using only one circuit step $\mathbf{g}$.

The 2×3 *transportation polytope in the example has dimension two. Our goal is to come up with a perturbation of the margins such that any (not necessarily feasible) circuit walk from O to F has to use two steps. In particular we want to rule out going from O to F by only applying* $\mathbf{g}$*!*

To do so, consider the two circuits $\mathbf{g}^1$, $\mathbf{g}^2$ *depicted in Figure 3.10. The dashed edges are the ones that are increased (Note that they are not sign-compatible, but this not relevant for our simple example, as we only want to demonstrate a perturbation along a single circuit.). They are linearly independent, as they both share only the edges* $\{\mathfrak{s}_1, \mathfrak{d}_2\}$ *and* $\{\mathfrak{s}_2, \mathfrak{d}_2\}$*. Applying both of them with step length two transfers O to F via an infeasible circuit walk of length two.*

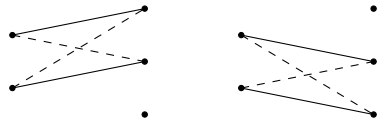

Figure 3.10: Two circuits $\mathbf{g}^1, \mathbf{g}^2$ in a 2×3 transportation polytope.

We now apply the construction described in the second part of the proof of Lemma 3.20 to make it impossible to go from O to F using only $\mathbf{g}$. Note that $\mathbf{g}$ is linearly independent both from $\mathbf{g}^1$ and from $\mathbf{g}^2$, so it does not matter which of the two circuits we pick for the construction. We use $\mathbf{g}^1$ and choose a sufficiently small $\epsilon > 0$. We then add ϵ to all nodes incident to $\mathbf{g}^1$. This is depicted in Figure 3.11.

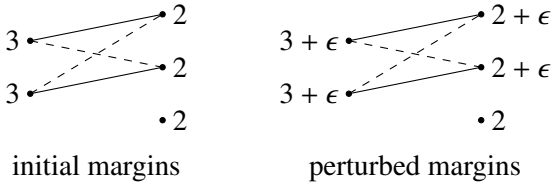

Figure 3.11: A perturbation of margins along circuit g^1.

Then the original support graphs $B(\mathbf{y}^O)$ and $B(\mathbf{y}^F)$ again induce vertices O' and F' in the new polytope. We show the corresponding $\mathbf{y}^{O'}$ and $\mathbf{y}^{F'}$ in Figure 3.12. Clearly, now an application of $\mathbf{g}$ cannot transfer O' to F', as $y^{O'}_{11} \neq y^{O'}_{23}$.

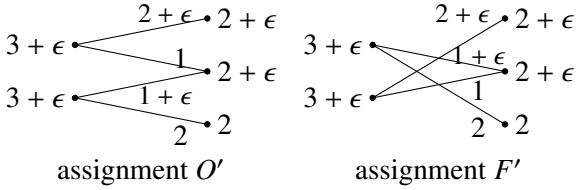

Figure 3.12: The assignments corresponding to $B(\mathbf{y}^O)$, $B(\mathbf{y}^F)$ in the perturbed polytope.

.

By this perturbation, we ruled out the possibility of having a 'shortcut' via circuit $\mathbf{g}$. In a larger example, one would now continue with a sequence of perturbations of smaller and smaller perturbation values to rule out any further shorter circuit walks . Choosing each ϵ sufficiently small, one does not reintroduce any of the shorter circuit walk at a later point.

This lower bound result concludes our discussion of the circuit diameters of transportation polytopes.

Chapter 4

Circuit diameter bounds for dual network flow polyhedra

Dual transportation polyhedra, associated with the transportation problems we investigated in the previous chapter, are one of the few classes of polyhedra for which the Hirsch conjecture is known to be true. This was proven by Balinski in [1]. He even showed that these polyhedra are Hirsch-sharp. In other words, there are instances of dual transportation polyhedra whose combinatorial diameter (we use $\mathcal{CD}_e$ in this chapter) attains the Hirsch bound, which is quadratic in the number of nodes of the associated bipartite graph. In contrast, we will prove linear upper bounds for all less restrictive notions of circuit diameters in Theorem 4.1. Like this we obtain a family of polyhedra for which circuits walks through the interior of the polyhedron are indeed much shorter than edge walks.

Theorem 4.1 (Upper bounds for dual transportation polyhedra)**.** *For a dual transportation polyhedron defined on a bipartite graph on $|V|$ nodes,*

(i) the circuit diameter $\mathcal{CD}_{fm}$ is bounded above by $|V| - 2$ and

(ii) the circuit diameters $\mathcal{CD}_{fs}$, $\mathcal{CD}_f$ and $\mathcal{CD}$ are bounded above by $|V| - 3$.

One natural generalization of transportation problems are min-cost flow problems (also transshipment problems), which allow for transportation between any pair of nodes, not only from supply to demand points. They are defined on arbitrary directed graphs $G = (V, E)$ instead of undirected bipartite ones. For the polyhedra associated with the dual problems we prove the following upper and lower bounds for the main categories of circuit distances.

Theorem 4.2 (Upper bounds for dual network flow polyhedra)**.** *For a dual network flow polyhedron $P_{G,\mathbf{c}}$*

(i) the combinatorial diameter $\mathcal{CD}_e$ is bounded above by both, $|E| \cdot (|V| - 1)$ and $\frac{1}{3}|V|^3$,

(ii) the circuit diameter $\mathcal{CD}_{fm}$ is bounded above by $\frac{1}{2}|V| \cdot (|V| - 1)$, and

(iii) the circuit diameters $\mathcal{CD}_{fs}$, $\mathcal{CD}_{f(b)(r)}$ and $\mathcal{CD}$ are bounded above by both, $|V| - 1$ and $|E| - |V| + 1$.

The quadratic bound (ii) strengthens a cubic one implied by a result in [10]. Note that, roughly speaking, we add a factor of $|V|$ on the bounds for $\mathcal{CD}_e$ and $\mathcal{CD}_{fm}$ from the special case of dual transportation polyhedra. When complementing our discussion and asking for lower bounds, we will see that this generalized setup is indeed much more involved and thus we cannot expect same bounds as in the bipartite case. Most notably, the upper bounds from Theorem 4.1 can be violated for all notions of circuit diameters.

Theorem 4.3 (Lower bounds for dual network flow polyhedra)**.** *For a dual network flow polyhedron $P_{G,\mathbf{c}}$*

(i) the combinatorial diameter $\mathcal{CD}_e$ and circuit diameter $\mathcal{CD}_{fm}$ are bounded below by $\frac{4}{3}|V|-4$, and

(ii) the circuit diameters $\mathcal{CD}_{fs}$, $\mathcal{CD}_{f(b)(r)}$ and $\mathcal{CD}$ are bounded below by both, $|V|-1$ and $|E|-|V|+1$

in the sense that for all values of $|V|$ and $|E|$ there are dual network flow polyhedra that attain the respective bounds.

Note that (ii) tells us that the upper bounds from Theorem 4.2 (iii) are tight. For the first statement (i) we will exhibit a family of dual network flow polyhedra that violate this upper bound on $\mathcal{CD}_{fs}$ and thus also the upper bound for the bipartite case by an arbitrary additive constant. This indicates why we get weaker upper bounds for the combinatorial diameter and for $\mathcal{CD}_{fm}$ for the more general polyhedra in Theorem 4.2.

This chapter is structured as follows: In Section 4.1 we formally introduce dual network flow polyhedra and characterize their vertices, edges, and circuits, which reveal a lot of combinatorial structure. We further prepare some tools for the proofs of our diameter bounds. Next we consider the special case of dual transportation polyhedra in Section 4.2, where we prove Theorem 4.1. We then finally turn to the general case of dual transportation polyhedra. In Section 4.3 we prove the upper bounds stated in Theorem 4.2 before investigating the lower bounds from Theorem 4.3 in Section 4.4. The results on the combinatorial diameter and the circuit diameter $\mathcal{CD}_{fm}$ were published in [5] and [6], together with Steffen Borgwardt and Raymond Hemmecke.

4.1 Preliminaries

We now introduce the dual network flow polyhedra associated with uncapacitated min-cost $\mathbf{b}$-flow problems as a generalization of transportation problems. Let $G = (V, E)$ be a directed connected graph on node set $V = \{0, \ldots, |V|-1\}$. A directed edge $e \in E$ with node i as its tail and node j as its head is denoted $e = (i, j)$, or simply $e = ij$.

Let $A \in \{-1, 0, 1\}^{|V|\times|E|}$ be the node-arc incidence matrix of G, where $a_{ie} = -1$ and $a_{je} = 1$ if the arc $e = ij$ is contained in E. Let $\mathbf{b} \in \mathbb{R}^{|V|}$. An uncapacitated $\mathbf{b}$-flow on G is given by any solution $\mathbf{x} \in \mathbb{R}^{|E|}$ to $A\mathbf{x} = \mathbf{b},\ \mathbf{x} \geq \mathbf{0}$, that is, b_i is the total (incoming minus outgoing) flow through a node

$i \in V$. If $b_i > 0$ the node i is called a supply, while nodes with $b_i < 0$ are called demands. Note that we allow for several supply and demand nodes and that supply nodes can also have incoming flow and there can be outgoing flow at the demand nodes.

For a cost function $\mathbf{c}\colon E \to \mathbb{R}_+$ indexed by the edges, the *uncapacitated min-cost* $\mathbf{b}$*-flow problem* on G (also *transshipment problem*) and its dual are given by

$$\min\left\{\mathbf{c}^\intercal\mathbf{x}:\ A\mathbf{x}=\mathbf{b},\ \mathbf{x}\geq\mathbf{0},\ \mathbf{x}\in\mathbb{R}^{|E|}\right\}$$

and

$$\max\left\{\mathbf{u}^\intercal\mathbf{b}:\ A^\intercal\mathbf{u}\leq\mathbf{c},\ \mathbf{u}\in\mathbb{R}^{|V|}\right\}.$$

The *dual network flow polyhedron* associated to some graph G and vector $\mathbf{c}\in\mathbb{R}^{|E|}$ is the feasible region of the associated dual linear program. We write it as

$$P_{G,\mathbf{c}}=\left\{\mathbf{u}\in\mathbb{R}^{|V|}:-u_a+u_b\leq c_{ab}\ \forall\ ab\in E,\ u_0=0\right\},$$

where we put $u_0 = 0$ as is standard to make $P_{G,\mathbf{c}}$ pointed.

Throughout this chapter we will exploit the special structure of dual network flow polyhedra $P_{G,\mathbf{c}}$ by describing their vertices, edges and circuits in terms of subgraphs of the defining graph G. This allows us to visualize our concepts when proving bounds on the respective circuit diameters. In all figures, the labels inside a node represent the respective values of u_i, while labels next to the nodes refer to the node's identifier i.

For $\mathbf{u}\in P_{G,\mathbf{c}}$ we denote by $G(\mathbf{u})$ the graph with nodes V and with edges $ab\in E$ for which $-u_a+u_b\leq c_{ab}$ is tight. Recall that for proving upper bounds on the combinatorial diameter $C\mathcal{D}_e$ it is enough to consider non-degenerate polyhedra, while for the other categories of circuit diameters we have to cover the degenerate case as well. Thus, in the former case we can exploit the following observation: If a dual network flow polyhedron is non-degenerate, the graphs $G(\mathbf{u})$ do not contain cycles. Therefore, note that a cycle corresponds to a set of linearly dependent inequalities that are all tight at the same time, but in non-degenerate polyhedra there are no such over-determined points.

The *vertices* of the polyhedron play a fundamental role in our diameter studies. A vertex of $P_{G,\mathbf{c}}$ is determined by a set of $|V|-1$ linearly independent inequalities $-u_a+u_b\leq c_{ab}$ that are tight (recall that we already set $u_0=0$). The edges associated with these $|V|-1$ inequalities describe a *spanning tree* of G. Thus, $\mathbf{u}\in P_{G,\mathbf{c}}$ is a vertex if and only if $G(\mathbf{u})$ is a spanning subgraph of G and *any* spanning tree T with $E(T)\subseteq E(G(\mathbf{u}))$ uniquely determines the vertex.

Next we want to describe the *edge directions (circuits)* of the polyhedron. An edge e of the polyhedron $P_{G,\mathbf{c}}$ is characterized by $|V|-2$ linearly independent inequalities that are tight. These tight inequalities correspond to edges in the graphs $G(\mathbf{u})$ of elements $\mathbf{u}\in e$. $G(\mathbf{u})$ thus consists of two connected components if $\mathbf{u}$ is not a vertex of $P_{G,\mathbf{c}}$. Let R and S be the respective node sets and assume without loss of generality that $0\in R$. When going along the edge e of the polyhedron,

we keep all edges within R and S, respectively, as the corresponding inequalities remain tight. But this means that we must change all components u_i with $i \in S$ by the same amount ϵ, while all components u_i with $i \in R$ remain unchanged (as we have $u_0 = 0$ and $0 \in R$). Thus, the edge direction (or circuit) is describe by a partition $V = R \dot\cup S$ for which the respective node sets R and S are non-empty and connected in the underlying undirected graph. We apply the circuit $R \dot\cup S$, $0 \in R$, by increasing or decreasing, respectively, all elements in S by the same value. If at some point another inequality $-u_a + u_b \leq c_{ab}$ becomes tight, we applied a maximal circuit step and a new edge ab was inserted in the corresponding graph. Observe that $a \in R$, $b \in S$ if we increased S and $b \in R$, $a \in S$ if we decreased S. This also implies that we cannot increase S if there is an edge from R to S in $G(\mathbf{u})$, and we cannot decrease S if there is an edge from S to R (at least if we want to stay feasible).

The arguments above further tell us that two vertices $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ are connected by an edge of the polyhedron if and only if the subgraph of G with edge set $E\left(G(\mathbf{u}^{(1)})\right) \cap E\left(G(\mathbf{u}^{(2)})\right)$ consists of exactly two connected components. These components describe the common edge.

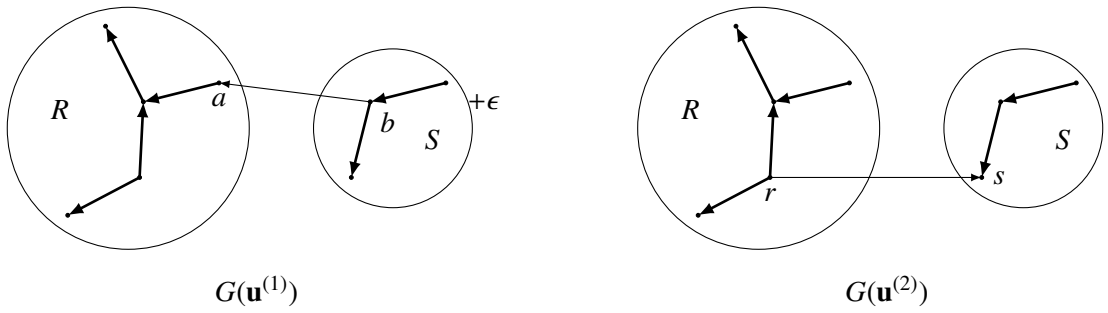

Figure 4.1: Two neighboring vertices of a dual network flow polyhedron.

We saw that every edge is characterized by a partition $V = R \dot\cup S$ with R and S connected. Conversely, it is not hard to see that every such partition describes an edge of $P_{G,\mathbf{c}}$ for a suitable choice of $\mathbf{c}$: Let E_R and E_S edge sets of spanning trees of the respective connected components and set $c_{ab} = 0$ for all $ab \in E_R \cup E_S$ and $c_{ab} = 1$ otherwise. This proves the following

Lemma 4.4 (Circuits of $P_{G,\mathbf{c}}$)**.** *The set of circuits C_G, associated to the matrix defining the polyhedron $P_{G,\mathbf{c}}$, consists of the vectors*

$$g_i = \begin{cases} 0, & \text{if } i \in R, \\ 1, & \text{if } i \in S, \end{cases} \tag{4.1}$$

for node sets $R, S \subseteq V$ such that the underlying undirected subgraphs of the respective node sets are connected and satisfy $R \cup S = V$, $R \cap S = \emptyset$, and $0 \in R$.

Recall that for proving upper bounds on the circuit diameter of a specific polyhedron $P_{G,\mathbf{c}}$ we may assume that none of the inequalities $-u_a + u_b \leq c_{ab}$ is redundant (otherwise we could remove such an edge ab from G, leaving the polyhedron the same but making the set of circuits smaller and thus the circuit diameter potentially bigger).

We continue with a tool we use for the proofs in Sections 4.2.2 and 4.3. The idea of *contracting edges* simplifies the construction of the circuit walks that prove our upper bounds on the combinatorial distance $C\mathcal{D}_e$ and on the circuit distance $C\mathcal{D}_{fm}$ (see Theorem 4.2).

Assume that we are at a feasible point $\mathbf{y}$ of a polyhedron $P_{G,\mathbf{c}}$ from which we want to construct a circuit walk to some vertex $\mathbf{u}$. Suppose that $E(G(\mathbf{y}))$ and $E(G(\mathbf{u}))$ already have an edge ab in common. Then we would like to keep this edge on the remaining circuit walk. Therefore, the difference between u_a and u_b has to remain constant, which means that in every circuit step given by $V = R \dot\cup S$, a and b are assigned either both to R or both to S. To simplify this proceeding, we will interpret a and b as one node in the following way: We *contract* the edge ab and continue our circuit walk in a smaller polyhedron defined on a graph with one node less and adjusted edge set.

Geometrically this corresponds to intersecting the dual network flow polyhedron with the hyperplane $\left\{\mathbf{u} \in \mathbb{R}^{|V|} \colon\ -u_a + u_b = c_{ab}\right\}$. This defines a face of the polyhedron, which is a dual network flow polyhedron in its own right. We then continue the circuit walk in this face. More formally, let ab be the common edge in $G = (V, E)$. The new polyhedron $P_{G',c'}$ is defined by a new graph $G' = (V', E')$ and a new vector c' as follows (for a simple notation we use $c_{ij} = \infty$ if $ij \notin E$):

$$
\begin{aligned}
V' =& V \backslash \{b\} \\
E' =& \{ij : ij \in E \text{ and } i, j \neq a, b\} \\
&\cup \{aj : aj \in E \text{ or } bj \in E\} \cup \{ia : ia \in E \text{ or } ib \in E\} \\
c'_{ij} =& \begin{cases} c_{ij} & \text{for } \ i, j \neq a,\ ij \in E' \\ \min\left\{c_{aj}, c_{bj} + c_{ab}\right\} & \text{for } \ i = a,\ aj \in E' \\ \min\{c_{ia} + c_{ab}, c_{ib}\} & \text{for } \ j = a,\ ia \in E' \end{cases}
\end{aligned}
$$

For the definition of $\mathbf{c}'$, observe the following: Assume that ab exists in $G(\mathbf{u})$ (i.e. $-u_a + u_b = c_{ab}$) and there are edges $aj, bj \in E$ for some j. Suppose we decrease both u_a and u_b. Then $-u_a + u_j \leq c_{aj}$ will become tight before $-u_b + u_j \leq c_{bj}$ if and only if $c_{aj} \leq c_{bj} + c_{ab}$. Hence, when keeping ab, the latter case will never occur and thus only the first inequality is relevant. On the other hand, $c_{aj} > c_{bj} + c_{ab}$ implies that only $-u_b + u_j \leq c_{bj}$ can become tight, and thus we only need to consider this inequality, but we have to adjust the value for $\mathbf{c}$ (observe $u_b = u_a + c_{ab}$). The other case is analogous. So we know that every circuit walk in $P_{G',c'}$ admits a circuit walk in $P_{G,c}$ that keeps the edge ab such that we can continue the walk in the smaller polyhedron.

We close this introductory sections with a fundamental observation on the graphs of the elements of a dual network flow polyhedron: We show that the existence of a feasible point whose graph contains a certain edge ab implies the non-existence of feasible points whose graphs contain another directed path from a to b.

Lemma 4.5. *Let $P_{G,\mathbf{c}}$ be a dual network flow polyhedron. Let $v_0v_k \in E$ such that in G there is another directed path $\mathcal{P}$ from v_0 to v_k, i.e. there are nodes $v_0, v_1, \ldots, v_k \in V$, $k \geq 2$, such that $v_iv_{i+1} \in E$ for all $i = 0, \ldots, k-1$.*

Assume there is a feasible point $\mathbf{w} \in P_{G,c}$ *with* $v_0v_k \in E(G(\mathbf{w}))$ *and let* $\mathbf{u} \in P_{G,c}$ *with* $\mathcal{P} \subseteq G(\mathbf{u})$. *Then also* $v_0v_k \in E(G(\mathbf{u}))$. *Thus, there can be no such* $\mathbf{u}$ *if* $P_{G,\mathbf{c}}$ *is non-degenerate.*

Proof. The feasible point $\mathbf{w} \in P_{G,c}$ satisfies

$$c_{v_0v_k} = -w_{v_0} + w_{v_k} = \sum_{i=0}^{k-1} \left(-w_{v_i} + w_{v_{i+1}}\right) \leq \sum_{i=0}^{k-1} c_{v_iv_{i+1}} \ .$$

$\mathbf{u} \in P_{G,c}$ satisfies $-u_{v_i} + u_{v_{i+1}} = c_{v_iv_{i+1}}$ for $i = 0, \ldots, k-1$ and $-u_{v_0} + u_{v_k} \leq c_{v_0v_k}$. We get

$$\sum_{i=0}^{k-1} c_{v_iv_{i+1}} = \sum_{i=0}^{k-1} \left(-u_{v_i} + u_{v_{i+1}}\right) = -u_{v_0} + u_{v_k} \leq c_{v_0v_k} \leq \sum_{i=0}^{k-1} c_{v_iv_{i+1}}.$$

Hence, all inequalities must be satisfied with equality and we get $-u_{v_0} + u_{v_k} = c_{v_0v_k}$, that is, $v_0v_k \in E(G(\mathbf{u}))$. □

Observe that Lemma 4.5 can easily be generalized to a stronger statement: Assume there is a feasible point whose graph contains a directed path from some node v_0 to some node v_k. Then every point of the dual network flow polyhedron whose graph contains another directed v_0, v_k -path must contain the first path as well. This can only happen in the degenerate case.

4.2 The bipartite case: Circuit diameter bounds for dual transportation polyhedra

Before investigating the circuit diameters of general dual network flow polyhedra, we deal with an interesting special case in this section.

An $M\times N$ transportation problem (see Chapter 3) is defined on an undirected bipartite graph. It can be seen as a $\mathbf{b}$-flow problem defined on a directed graph with all edges pointing from the set of supply nodes $V_1 = \{0, \ldots, M-1\}$ to the set of demand nodes $V_2 = \{M, \ldots, M+N-1\}$. The nodes $i \in V_1$ have positive total flow $b_i > 0$ (supply), while for the nodes $i \in V_2$ we have $b_i < 0$ (demand). Then, similar to the definition of dual network flow polyhedra, a *dual transportation polyhedron* associated to a bipartite graph G on $M\times N$ nodes is given by some vector $\mathbf{c} \in \mathbb{R}^{|E|}$ via

$$P_{G,\mathbf{c}} = \left\{ \mathbf{u} \in \mathbb{R}^{M+N} : -u_a + u_b \leq c_{ab} \ \forall\, ab \in E \text{ with } a \in V_1, b \in V_2,\ u_0 = 0 \right\} .$$

Note that all results presented in Section 4.1 translate to this setting.

For dual $M\times N$ transportation polyhedra, Balinski proved an upper bound of $(M-1)(N-1)$ on the combinatorial diameter $C\mathcal{D}_e$ in [1]. This bound is tight in the sense that for all M, N there are dual transportation polyhedra with combinatorial diameter $(M-1)(N-1)$ and thus $C\mathcal{D}_e$ can be linear in the number of edges and quadratic in the number of nodes of the underlying graph. On the contrary, we will prove that all less restrictive circuit distances are linear in the number of vertices, as stated in Theorem 4.1. To this end, we present a structural results on the vertices of

dual transportation polyhedra that is specific to the bipartite case in Section 4.2.1. In Section 4.2.2 we then prove our upper bounds, part (i) and part (ii) of Theorem 4.1.

4.2.1 Preliminary results

As in the general case, each vertex of a dual transportation polyhedron $P_{G,\mathbf{c}}$ is described by a spanning tree of G with edges corresponding to the inequalities $-u_a + u_b \leq c_{ab}$ that are tight at the vertex. We show that for every pair of vertices of a dual transportation polyhedron there are such spanning trees that have at least two edges in common. Using this result we can refine our proofs in Section 4.2.2. This will improve our upper bounds on the circuit distances.

Lemma 4.6. *Let $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ be two vertices of $P_{G,\mathbf{c}}$. Then there are spanning trees T_1 and T_2 of G with $E(T_i) \subseteq E(G(\mathbf{u}^{(i)}))$ such that $|E(T_1) \cap E(T_2)| \geq 2$.*

Proof. Let $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ be distinct vertices. As we can translate $P_{G,\mathbf{c}}$, we may assume without loss of generality that $\mathbf{u}^{(1)} = \mathbf{0}$. Clearly, this mere shift does not change any structure, in particular it does not change $G(\mathbf{u}^{(1)})$ and $G(\mathbf{u}^{(2)})$. For better readability, we split the vectors $\mathbf{u}^{(i)}$, $i = 1, 2$, into two vectors $\mathbf{v}^{(i)}$ and $\mathbf{w}^{(i)}$ for the components belonging to V_1 and V_2, respectively. Let $G_i := G(\mathbf{u}^{(i)})$. Now assume that $E(G_1) \cap E(G_2) = \emptyset$.

As G_1 is connected, there must be an edge $(v_0, w_0) \in G_1$. As $(v_0, w_0) \notin G_2$, we have $-v_0^{(2)} + w_0^{(2)} < -v_0^{(1)} + w_0^{(1)} = 0$ and hence $w_0^{(2)} < v_0^{(2)}$.

As G_2 is connected, there must be an edge $(v_1, w_0) \in G_2$. As $(v_1, w_0) \notin G_1$, we must have $0 = -v_1^{(1)} + w_0^{(1)} < -v_1^{(2)} + w_0^{(2)}$ and hence $v_1^{(2)} < w_0^{(2)}$.

Again, as G_1 is connected, there must be an edge $(v_1, w_1) \in G_1$. As $(v_1, w_1) \notin G_2$, we have $-v_1^{(2)} + w_1^{(2)} < -v_1^{(1)} + w_1^{(1)} = 0$ and hence $w_1^{(2)} < v_1^{(2)}$.

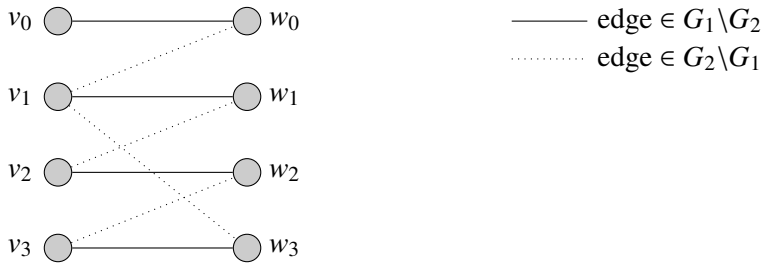

Figure 4.2: An alternating path of edges that eventually close an even cycle.

Continuing like this we create a path with edges alternately from $G_1 \setminus G_2$ and $G_2 \setminus G_1$. As there are only finitely many nodes, eventually some v_i (or w_j) is selected a second time and we close a cycle. But then we have

$$v_i^{(2)} > w_i^{(2)} > v_{i+1}^{(2)} > \ldots > v_k^{(2)} = v_i^{(2)}$$

(or $w_j^{(2)} > \ldots > w_j^{(2)}$), a contradiction. Hence we must have $E(G_1) \cap E(G_2) \neq \emptyset$.

Now assume that $E(G_1)$ and $E(G_2)$ have only one edge in common. Let (v_0, w_0) be this edge. Then $v_0^{(2)} = w_0^{(2)} = 0$. As G_2 is connected, there must be an edge $(v_0, w_i) \in G_2$ or an edge

$(v_i, w_0) \in G_2$. Without loss of generality, assume we have $(v_1, w_0) \in G_2$. As $(v_1, w_0) \notin G_1$, we have $0 = -v_1^{(1)} + w_0^{(1)} < -v_1^{(2)} + w_0^{(2)}$ and hence $v_1^{(2)} < w_0^{(2)}$.

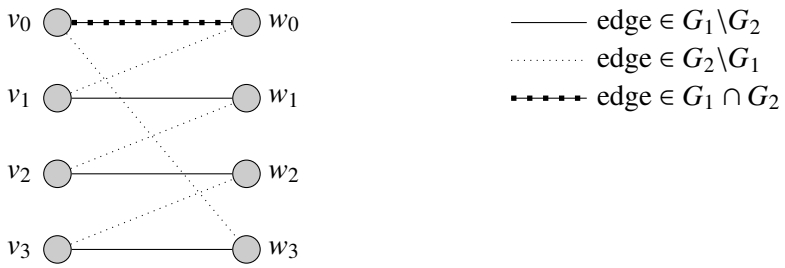

Figure 4.3: An alternating path of edges with one common edge that eventually closes an even cycle.

Continuing like this we again create a path. The first edge is in $G_1 \cap G_2$, while from the second edge on the edges are alternately from $G_2 \setminus G_1$ and $G_1 \setminus G_2$. Again, as there are only finitely many nodes eventually some v_i (or w_j) is selected a second time and we close a cycle. If this node is not v_0, we get the same contradiction as before. Otherwise we have

$$0 = v_1^{(2)} = w_1^{(2)} > v_2^{(2)} > w_2^{(2)} > \ldots > v_1^{(2)} ,$$

again a contradiction. Hence we must have $|E(G_1) \cap E(G_2)| > 1$.

We now obtain the spanning trees T_1 and T_2 with $|E(T_1) \cap E(T_2)| \geq 2$ by taking two edges from $E(G_1) \cap E(G_2)$ and extending this set to the edge sets of a spanning trees T_i of G by adding edges from the respective G_i. □

4.2.2 Linear upper bounds on circuit distances

We first present and prove the core part of Theorem 4.1 (i).

Lemma 4.7. *For a dual transportation polyhedron defined on a bipartite graph on $|V| = M + N$ nodes, the circuit diameter $\mathcal{CD}_{fm}$ is bounded from above by $|V| - 1$.*

Proof. Let $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ be two vertices of $P_{G,\mathbf{c}}$, given by spanning trees T_1 and T_2 of G with $E(T_1) \subseteq E(G(\mathbf{u}^{(1)}))$ and $E(T_2) \subseteq E(G(\mathbf{u}^{(2)}))$. We will show how to construct a feasible maximal circuit walk $\mathbf{u}^{(1)} = \mathbf{y}^{(0)}, \ldots, \mathbf{y}^{(k)} = \mathbf{u}^{(2)}$, such that $G(\mathbf{y}^{(i)})$ has at least i edges in common with T_2. This immediately implies $k \leq |V| - 1$ which proves the claim. It should be noted that the subgraphs $G(\mathbf{y}^{(i)})$ may not be connected, since our circuit walk possibly goes through the interior of $P_{G,\mathbf{c}}$ along (potential) edge directions.

Given a point $\mathbf{y}^{(i)} \neq \mathbf{u}^{(2)}$ of the circuit walk, let $C = G(V(C), E(C))$ be the connected component of $(V, E(G(\mathbf{y}^{(i)})) \cap E(T_2))$ containing the node 0. Possibly, C consists only of the node 0. As $\mathbf{y}^{(i)} \neq \mathbf{u}^{(2)}$, we must have $C \neq T_2$ and thus there is some node $s \in V$ which is not in C, but which is connected to C via some edge rs in T_2. We now construct an edge direction $\mathbf{g}$ from C_G such

that $\mathbf{y}^{(i+1)} := \mathbf{y}^{(i)} + \alpha\mathbf{g}$ arises from a maximal step along $\mathbf{g}$ and such that $(V, E(G(\mathbf{y}^{(i+1)})) \cap E(T_2))$ contains C and the edge rs from T_2. Starting from $\mathbf{y}^{(0)}$ and repeating this process iteratively, we see that $G(\mathbf{y}^{(i)})$ has at least i edges in common with T_2, which implies the result.

To construct $\mathbf{g}$, we need to define $R, S \subseteq V$ that describe the edge direction from C_G. Without loss of generality we will assume that $s \in V_2$. The case $s \in V_1$ works analogously by merely switching the roles of V_1 and V_2 and hence by switching the roles of $\epsilon\mathbf{g}$ and $-\epsilon\mathbf{g}$ below.

(a) All nodes from C are assigned to R.

(b) All nodes from $V_2 \setminus \{ s \}$ which are connected to C by an edge in E, are assigned to R.

(c) All nodes $t \in V \setminus R$ that are connected to s by a path in G that does not contain a node in R, are assigned to S.

(d) All remaining nodes are assigned to R.

As G is connected, this construction leads to sets R and S that are nonempty, satisfy $R \dot\cup S = V$, and define connected components of G that are connected by the edge $rs \in E$. Hence, R and S define an element $\mathbf{g} \in C_G$ via Equation (4.1). Observe that $s \in V_2$ as $s \notin V(C)$ and $0 \in R$. We wish to include the edge rs into our graph, that is, we wish to make the inequality $-u_r + u_s \leq c_{rs}$ tight at $\mathbf{y}^{(i+1)}$. Thus we have to increase increasing the component $y_s^{(i)}$. Therefore, we *add* $\epsilon\mathbf{g}$ to $\mathbf{y}^{(i)}$. (If $s \in V_1$, we *subtract* $\epsilon\mathbf{g}$ from $\mathbf{y}^{(i)}$.) We choose as ϵ the smallest non-negative number such that an inequality $-u_a + u_b \leq c_{ab}$ with $a \in R$ and $b \in S$ becomes tight. Note that $\epsilon = 0$ is not excluded, but we show that this will never happen. In fact, we show that the edge ab (on which $-u_a + u_b \leq c_{ab}$ becomes tight) is exactly the edge rs that we wish to include.

Assume now on the contrary that $ab \neq rs$. Note that by construction at steps (b) and (c) we must have $b = s$, as all edges from R to $S \cap V_2$ have s as their end point and these are exactly the edges on which an inequality may become tight when walking along direction $\mathbf{g} \in C_G$. Hence we must have $a \neq r$. Observe that $\mathbf{y}^{(i+1)} := \mathbf{y}^{(i)} + \epsilon\mathbf{g}$ and $\mathbf{u}^{(2)}$ agree in their components in $V(C)$, that is, $u_c^{(2)} = y_c^{(i+1)}$ for all $c \in V(C)$, because $G(\mathbf{y}^{(i)} + \epsilon\mathbf{g})$ and T_2 coincide on the edges in C and $0 \in V(C)$. Since $\mathbf{y}^{(i+1)} \in P_{G,\mathbf{c}}$ and since $as \in E(G(\mathbf{y}^{(i+1)}))$ and $rs \notin E(G(\mathbf{y}^{(i+1)}))$, we have

$$-u_a^{(2)} + y_s^{(i+1)} = c_{as} \text{ but } -u_r^{(2)} + y_s^{(i+1)} < c_{rs}.$$

On the other hand, since $\mathbf{u}^{(2)} \in P_{G,\mathbf{c}}$ and since $as \notin E(T_2)$ and $rs \in E(T_2)$, we have

$$-u_a^{(2)} + u_s^{(2)} < c_{as} \text{ but } -u_r^{(2)} + u_s^{(2)} = c_{rs}.$$

From $-u_a^{(2)} + y_s^{(i+1)} = c_{as}$ and $-u_a^{(2)} + u_s^{(2)} < c_{as}$ we conclude $y_s^{(i+1)} > u_s^{(2)}$, whereas $-u_r^{(2)} + y_s^{(i+1)} < c_{rs}$ and $-u_r^{(2)} + u_s^{(2)} = c_{rs}$ imply $y_s^{(i+1)} < u_s^{(2)}$. This contradiction shows $a = r$ and the claim is proved. □

Lemma 4.6 now implies the following strengthening of Lemma 4.7 and thus proves Theorem 4.1 (i).

Lemma 4.8. *For a dual transportation polyhedron defined on a bipartite graph on $|V| = M + N$ nodes, the circuit diameter $\mathcal{CD}_{fm}$ is bounded from above by $|V| - 2$.*

Proof. The proof is analogous to the proof of Lemma 4.7. We merely have to observe that by Lemma 4.6 we can choose spanning trees T_1 and T_2 that have one edge in common and that we may assume without loss of generality that this edge has 0 as one of its endpoints. Thus we only have to add at most $|V| - 2$ edges in at most $|V| - 2$ steps to reach $\mathbf{u}^{(2)}$. □

The following example illustrates the circuit steps constructed in the above proof of Lemma 4.7.

Example 4.9. *We consider the bipartite graph on node sets $V_1 = \{0, 1, 2, 3\}$ and $V_2 = \{4, 5, 6, 7\}$ with edge set a depicted below.*

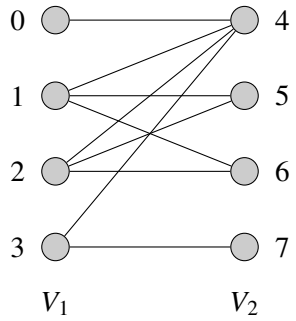

Figure 4.4: The graph that defines $P_{G,\mathbf{c}}$.

Let $\mathbf{c} \in \mathbb{R}^{|E|}$ be defined as $c_e = 1$ if $e = (2, 4)$ or $e = (1, 6)$ and set $c_e = 0$ otherwise. Then $\mathbf{y}^{(i)}$ and $\mathbf{u}^{(2)}$ illustrated in Figure 4.5 are elements of the dual transportation polyhedron $P_{G,\mathbf{c}}$, and $\mathbf{u}^{(2)}$ is a vertex as its graph is a spanning tree. Recall that the labels inside the vertices of the graphs are the values of the corresponding components of the vectors $\mathbf{y}^{(i)}, \mathbf{u}^{(2)} \in \mathbb{R}^{|V|}$.

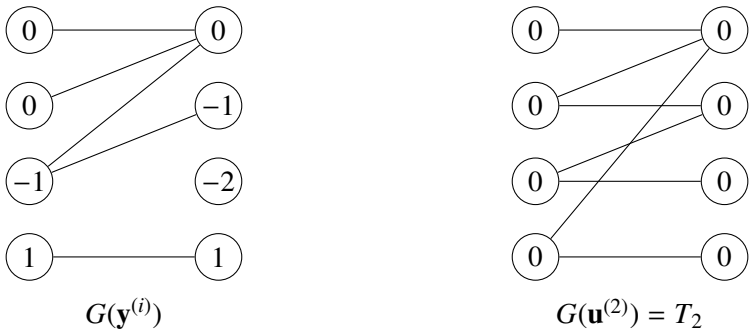

Figure 4.5: Two elements of a dual transportation polyhedron.

We now describe a circuit step at $\mathbf{y}^{(i)}$ on a circuit walk to the vertex $\mathbf{u}^{(2)}$ as applied in the proof of Lemma 4.7. The connected component C of $(V, E(G(\mathbf{y}^{(i)})) \cap E(T_2))$ containing the node 0 consists of the nodes $0, 1, 4$ and edges $(0, 4), (1, 4)$. There is a node $5 \notin V(C)$ that is directly connected to C via the edge $(1, 5)$ from T_2. We want to insert this edge, so we have $r = 1$, $s = 5$. We construct a

corresponding circuit given by $R \dot\cup S = V$ by assigning each node to one of the sets, following the rules described in the proof of Lemma 4.7 (see also Figure 4.6):

(a) $0, 1, 4$ *are assigned to R (as nodes in C).*

(b) 6 *is assigned to R (as a node in $V_2 \setminus \{ s \}$ which is connected to C by an edge in E).*

(c) $2, 5$ *are assigned to S (as nodes in $V \setminus R$ that are connected to s by a path in G not containing nodes in R).*

(d) $3, 7$ *are assigned to R (as the remaining nodes).*

Hence we get $R = \{0, 1, 3, 4, 6, 7\}$ and $S = \{2, 5\}$. To insert the edge $(1, 5)$, we increase all components in S by one. Like this we obtain $\mathbf{y}^{(i+1)}$ as the successive point in the circuit walk. Note that this step deletes all edges from S to R.

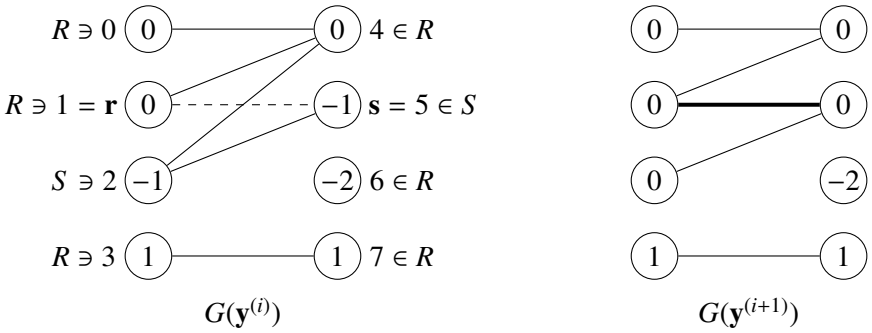

Figure 4.6: A maximal feasible circuit step that inserts $(1, 5)$.

In Lemma 4.8 we proved an upper bound of $|V| - 2$ on the circuit diameter $C\mathcal{D}_{fm}$. In fact, we also proved this bound for the more restrictive circuit diameter $C\mathcal{D}_{fmbr}$. To see this, note that the sequence of circuit steps applied in the proof of Lemma 4.7 is non-repetitive and non-backwards: Every circuit step induced by $R \dot\cup S$ inserts an edge $e = ab$, that is, the inequality $-u_a + u_b \leq c_{ab}$ becomes tight and it remains tight for the remaining circuit walk. Thus, we neither apply the same direction again (like this we would violate the inequality), nor do we apply the circuit in reverse direction (this would delete the edge ab again).

In contrast, our circuit walk is not necessarily sign-compatible. Therefore, observe that being sign-compatible would mean that the term $-u_a + u_b$ is either only increase or only decreased during the entire circuit walk. In particular, if there is some edge e that is inserted and removed (or the other way around), then this circuit walk is not sign-compatible. But the algorithm in the proof of Lemma 4.7 can yield a circuit walk that deletes and then reinserts an edge:

Example 4.10 (A non-sign-compatible circuit step)**.** *Consider the bipartite graph on node sets $V_1 = \{0, 1, 2\}$ and $V_2 = \{3, 4, 5\}$ with edge set E as depicted in Figure 4.7. The edges are labeled with the values c_{ab} of the vector $\mathbf{c} \in \mathbb{R}^{|E|}$ defining the dual transportation polyhedron $P_{G,\mathbf{c}}$.*

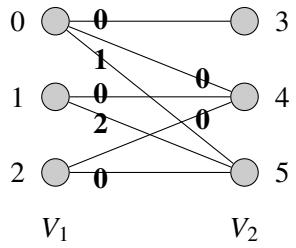

Figure 4.7: The graph and the values of $\mathbf{c} \in \mathbb{R}^{|E|}$ defining $P_{G,\mathbf{c}}$.

Obviously, the following two are vertices of $P_{G,\mathbf{c}}$.

Figure 4.8: Vertices of a dual transportation polyhedron.

At $\mathbf{y}^{(0)} = \mathbf{u}^{(1)}$ *we have* $V(C) = \{0, 3\}$ *for the connected component* C *of* $(V, E(G(\mathbf{y}^{(i)})) \cap E(T_2))$ *containing the node* 0*. According to the algorithm in the proof of Lemma 4.7 we can only choose* $e = (0, 4)$ *to insert. Thus we have* $R = \{0, 3, 5\}$ *(by rules (a) and (b)) and* $S = \{1, 2, 4\}$ *(by rule (c)). We now increase all components in* S *until the first inequality becomes tight, which must be the one corresponding to* $(0, 4)$ *as shown in the proof of Lemma 4.7. But this deletes* $(2, 5)$*, an edge that is contained in the graph of the final vertex. We have to reinsert* $(2, 5)$*, and thus the circuit walk cannot be sign-compatible.*

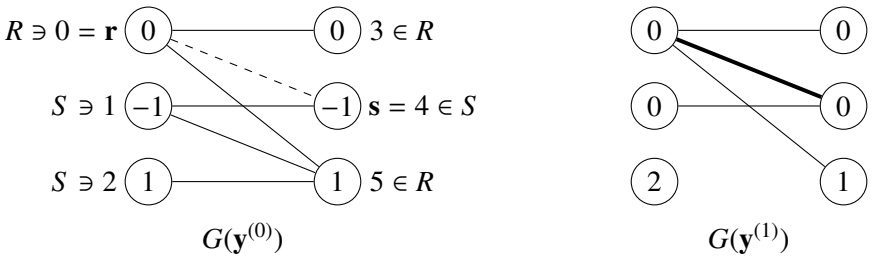

Figure 4.9: A circuit step that deletes an edge that has to be reinserted.

However, we can easily prove an even better upper bound of $|V| - 3$ on $C\mathcal{D}_{fs}$ as a corollary of Theorem 2.17. Note that by transitivity this result readily translates to all less restrictive notions of circuit distances, in particular to $C\mathcal{D}_f$ and $C\mathcal{D}$.

Lemma 4.11. *For a dual transportation polyhedron defined on a bipartite graph on* $|V| = M + N$ *nodes, the circuit diameter with respect to* $C\mathcal{D}_{fs}$ *is bounded from above by* $|V| - 3$.

Proof. Let $\mathbf{u}^{(1)}, \mathbf{u}^{(2)}$ be two vertices of a dual transportation polyhedron $P_{G,\mathbf{c}}$ defined on $G = (V, E)$. By Lemma 4.6, there are $ab, a'b' \in E(G(\mathbf{u}^{(1)})) \cap E(G(\mathbf{u}^{(2)}))$ with $ab \neq a'b'$. Now consider

$$P' := \{ \mathbf{u} \in \mathbb{R}^{|V|} : u_0 = 0,\ -u_a + u_b = c_{ab},\ -u_{a'} + u_{b'} = c_{a'b'},\ -u_r + u_s \leq c_{rs}\ \forall rs \in E \backslash \{ab, a'b'\} \}\,.$$

Then $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ are vertices of P'. Observe that $P' \subseteq \mathbb{R}^{|V|}$ is defined by three independent equations and $|E| - 2$ inequalities. Thus, by Theorem 2.17 every two vertices are connected by a sign-compatible circuit walk of length at most $\min\{3 - |V| + |E| - 2, |V| - 3\} \leq |V| - 3$. The claim follows by observing that the set of circuits associated with P' is contained in the set of circuits associated with $P_{G,\mathbf{c}}$ and thus the circuit walk in P' is a circuit walk in $P_{G,\mathbf{c}}$ as well. □

This completes the proof of Theorem 4.1 (ii).

4.3 Upper bounds on circuit diameters

We now turn to the more general framework of dual network flow polyhedra. As already outlined in the begining of this chapter we cannot expect similar upper bounds as in the bipartite case, as for every category of circuit diameter there are examples that violate these upper bounds. Especially for the more restrictive notions it is unlikely to find such bounds, as we can exceed the former bound by an arbitrary constant, see Lemma 4.18 in section 4.4. Roughly speaking, when turning to general graphs we have to add a factor of $|V|$ on the previous bounds for the combinatorial diameter and for $C\mathcal{D}_{fm}$, yielding quadratic bounds. In contrast, the upper bounds on $C\mathcal{D}_{fs}$ and all less restrictive diameters remain linear in the number of vertices. All these bounds are stated in Theorem 4.2 and proved in this section.

We begin with the strongest category, the combinatorial diameter. Recall that for proving upper bounds on the combinatorial diameter of polyhedra it is enough to consider non-degenerate polyhedra and thus we can assume that for every vertex $\mathbf{u}$, $G(\mathbf{u}) = T$ is a spanning tree.

Lemma 4.12. *For a dual network flow polyhedron $P_{G,\mathbf{c}}$, the combinatorial diameter $C\mathcal{D}_e$ is bounded from above by* $\min\{\ (|V| - 1) \cdot |E|\ ,\ \frac{|V|^3}{6}\ \}$.

Proof. Let $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ be two vertices of a non-degenerate dual network flow polyhedron $P_{G,\mathbf{c}}$, given by spanning trees $T_1 = G(\mathbf{u}^{(1)})$ and $T_2 = G(\mathbf{u}^{(2)})$. We construct an edge walk from $\mathbf{u}^{(1)}$ to $\mathbf{u}^{(2)}$ as follows: Being at a vertex $\mathbf{y}$ of $P_{G,\mathbf{c}}$ with spanning tree $T = G(\mathbf{y})$, we choose an edge $rs \in T_2 \backslash T$ we wish to insert. We show how to construct an edge walk of length at most $|E|$ that leads to a vertex $\bar{\mathbf{y}}$ for which $rs \in E(G(\bar{\mathbf{y}}))$, that is, our specified edge is added to the corresponding spanning tree. Then we contract this edge (see Section 4.1) to ensure that we do not delete it again. Starting with $\mathbf{y} = u^{(1)}$ and repeating this for all $|V| - 1$ edges in T_2 proves the claimed bound of $(|V| - 1) \cdot |E|$.

Now, let $\mathbf{y}$ be the current vertex in our edge walk and let $T = G(\mathbf{y})$ be the corresponding spanning tree. We choose an arbitrary edge $rs \in T_2$ we wish to insert. Given a spanning tree T and the node s we distinguish forward and backward edges in $E(T)$: We see s as the root of the tree T. Then

every edge in $E(T)$ lies on a unique path starting at s (this path is independent of the directions of the edges). We call the edges pointing away from s *backward edges*, the edges pointing towards s *forward edges*.

In T there is a unique path (undirected) from r to s. Let e be the last backward edge on this path. Note that by Lemma 4.5 such an edge must exist. Let R and S be the node sets of the connected components of $T - e$ such that $r \in R$ and $s \in S$. Observe that in particular all nodes from which we can reach s on a directed path in the spanning tree T are assigned to S (and these nodes form an arborescence of forward edges with root s).

We wish to include the edge rs in our graph, that is, we wish to make the inequality $-u_r+u_s \leq c_{rs}$ tight. Without loss of generality we assume $0 \in R$, therefore we add an ϵ to all components y_i of $\mathbf{y}$ with $i \in S$. (If $0 \notin R$, we would subtract ϵ from all components y_i with $i \in R$.) We choose as ϵ the smallest non-negative number such that any inequality $-u_a + u_b \leq c_{ab}$ with $a \in R$ and $b \in S$ becomes tight. Due to non-degeneracy there is only one such inequality. This creates a new feasible point $\mathbf{y}'$, which is indeed a neighboring vertex of $\mathbf{y}$ by construction.

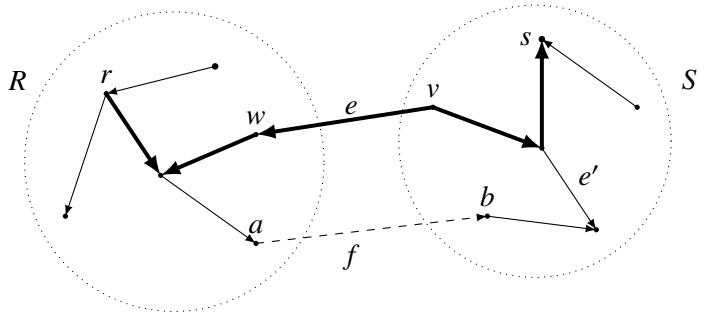

Figure 4.10: A circuit step that inserts ab and deletes vw when aiming at inserting rs.

So, in this edge step $e = vw$ is deleted and $f = ab$ is inserted. If we inserted $f = rs$, we contract this edge and start over again, aiming to insert another edge $r's'$ from $E(T_2)$. Otherwise we consider the path connecting r and s in the new spanning tree T'. As before the last backward edge e' defines sets R' and S' and we repeat the same procedure until eventually rs is inserted. It remains to prove that this indeed happens after at most $|E|$ steps. It is enough to show that the deleted edge $e = vw$ is not inserted again: As there is a directed path from v to s in $G(\mathbf{y})$, v and all nodes on this path will always be assigned to S (in particular, no edge on this path is deleted). As only edges from R to S are inserted, $e = vw$ with $v \in S$ cannot be reinserted. This proves the claimed upper bound $(|V| - 1) \cdot E$.

To see the upper bound $\frac{|V|^3}{6}$, we only have to change the way we count the number of steps that we need to insert the edge rs in a current underlying graph on i nodes (the number of nodes decreases with every contraction): Note that the current graph has at most $i \cdot (i - 1)$ edges, and in particular at most $\binom{i}{2}$ edges $e = vw$ with $v \in S$ and $w \in R$. As we only insert edges from R to S, this tells us an upper bound of $\binom{i}{2}$ steps until rs inserted. After contracting this edge, we start this process again on a graph with $i - 1$ nodes. Hence we obtain an edge walk of length at most

$$\sum_{i=2}^{|V|}\binom{i}{2} = \sum_{i=2}^{|V|}\left(\frac{1}{2}i\cdot(i-1)\right) = \frac{1}{2}\left(\sum_{i=2}^{|V|}i^2 - \sum_{i=2}^{|V|}i\right)$$
$$= \frac{1}{2}\left(\frac{|V|(|V|+1)(2|V|+1)}{6} - \frac{|V|(|V|+1)}{2}\right) = \frac{|V|^3-|V|}{6} \leq \frac{|V|^3}{6}.$$

□

We continue with the circuit distance $\mathcal{CD}_{fm}$ for which we prove a quadratic bound as well, but this time it is quadratic in the number of nodes. This strengthens the cubic bound implied by Corollary 5 in [10]. Recall that we cannot simply assume a non-degenerate polyhedron here.

Lemma 4.13. *For a dual network flow polyhedron $P_{G,\mathbf{c}}$, the circuit diameter $\mathcal{CD}_{fm}$ is bounded from above by $\frac{|V|\cdot(|V|-1)}{2}$.*

Proof. Let $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ be two vertices of the polyhedron $P_{G,\mathbf{c}}$. Let T_2 be a spanning tree with $E(T_2) \subseteq E(G(\mathbf{u}^{(2)}))$. We construct a feasible maximal circuit walk from $\mathbf{u}^{(1)}$ to $\mathbf{u}^{(2)}$ as follows:

Being at a point $\mathbf{y} \in P_{G,\mathbf{c}}$ of our circuit walk, we choose an edge $rs \in T_2 \backslash E(G(\mathbf{y}))$ we wish to insert. We construct a feasible maximal circuit walk to a point $\bar{\mathbf{y}} \in P_{G,\mathbf{c}}$ with $rs \in E(G(\bar{\mathbf{y}}))$. This walk has length at most $i-1$, where i is the number of nodes in the current underlying graph. As in the proof of Lemma 4.12, we then contract it to make sure that we do not delete it when continuing our circuit walk. We start with $\mathbf{y} = \mathbf{u}^{(1)}$ and repeat this procedure for all $|V|-1$ edges in $E(T_2)$. As the number of nodes decreases after every contraction this yields our quadratic upper bound of $\sum_{i=1}^{|V|-1} i = \frac{1}{2}(|V|\cdot(|V|-1))$.

Now, let $\mathbf{y}$ be a feasible point in the circuit walk. Let $rs \in E(T_2)\backslash E(G(\mathbf{y}))$ be an arbitrary edge we wish to insert, that is, we have to make $-u_r + u_s \leq c_{rs}$ tight. To this end, we construct a circuit direction that increases the component y_s. This circuit is given by $R \dot\cup S = V$ for the node sets R and S constructed by the following sequence of rules:

1. r is assigned to R.
2. s is assigned to S.
3. All nodes from $V\backslash\{r\}$ from which s can be reached on a directed path using edges in $E(G(\mathbf{y}))$ are assigned to S. (These edges form an arborescence with root s.)
4. All nodes $t \in V\backslash S$ that are connected to r in the underlying undirected graph are assigned to R.
5. All remaining nodes are assigned to S.

Observe that from s we cannot reach r on a directed path in $E(G(\mathbf{y}))$ by Lemma 4.5, and thus the sets R and S are well-defined and these sets clearly satisfy all the conditions to define a circuit. Let $\mathbf{g}$ be the corresponding circuit direction defined via Equation (4.1). Without loss of generality we assume that $0 \in R$. The case $0 \in S$ works analogously by merely switching the roles of R and S and subtracting $\epsilon\mathbf{g}$ to decrease y_r.

We now apply the maximal circuit step given by $\mathbf{g}$, that is, the next point in our circuit walk is $\mathbf{y}' := \mathbf{y} + \epsilon\mathbf{g}$, where ϵ is the smallest non-negative number such that an inequality $-u_a + u_b \leq c_{ab}$ with $a \in R$ and $b \in S$ becomes tight (observe that there could be multiple such inequality, as we do not assume non-degeneracy of the polyhedron $P_{G,\mathbf{c}}$). In any case the gap in between $-u_r + u_s$ and its upper bound c_{rs} becomes smaller. If rs was indeed inserted, we contract the edge and continue in a smaller polyhedron.

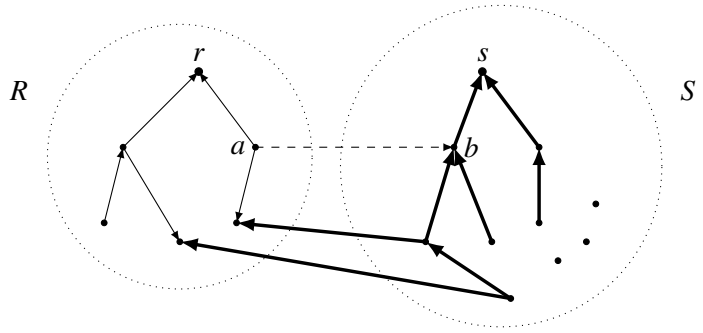

Figure 4.11: A circuit step that inserts ab and deletes two edges from S to R when aiming at inserting rs.

Otherwise, the inserted edge extends the arborescence by at least the node a. We again apply a circuit step by constructing sets R' and S' for $\mathbf{y}'$ as before, which inserts rs or extends the arborescence further. Continuing like this after at most $i - 2$ steps all nodes but r are contained in the arborescence (if rs was not already inserted). Then the next step must add rs by Lemma 4.5. Thus, inserting rs takes at most $i - 1$ maximal circuit steps. □

Note that the circuit walks in the proof of Lemma 4.13 are non-backwards and non-repetitive, that is, we do not choose any partition $R \dot\cup S$ twice: Within the process of inserting some rs, all $R \dot\cup S$ are distinct since the arborescence is extended in every step and thus each set S contains at least one node (the node a) no previous S contained. To see that in the overall circuit walk no $R \dot\cup S$ is chosen twice, observe that when inserting rs we always have $r \in R$ and $s \in S$, but when aiming at inserting the next edge $r's'$, r and s will always be assigned to the same set.

Finally, we turn to the weak and soft circuit diameters. We get a linear upper bound on $C\mathcal{D}_{fs}$ and thus on all less restrictive circuit distances.

Lemma 4.14. *For a dual network flow polyhedron $P_{G,\mathbf{c}}$, the circuit diameter $C\mathcal{D}_{fs}$ is bounded from above by* $\min\{\, |V| - 1\,,\ |E| - |V| + 1 \,\}$.

Proof. We have $P_{G,\mathbf{c}} = \left\{ \mathbf{u} \in \mathbb{R}^{|V|} :\ u_0 = 0,\ A^\intercal \mathbf{u} \leq \mathbf{c} \right\}$, where $A^\intercal \in \mathbb{R}^{|E| \times |V|}$ is the transpose of the node-edge incidence matrix of G. Observe that there is only one equation defining the polyhedron and thus the corresponding matrix has rank one. Hence by Theorem 2.17 we get $C\mathcal{D}_{fs} \leq \min\{\, 1 - |V| + |E|\,,\ |V| - 1 \,\}$. □

This concludes the proof of Theorem 4.2.

We finally want to point out that these diameter bounds also hold for dual network flow polyhedra defined on directed graphs that are not connected: To make the polyhedron pointed, we set for each connected component the value of one variable to zero (just as we fixed $u_0 = 0$ for graphs with just one connected component). Then the algorithmic approaches described in the proofs of Lemma 4.12 and Lemma 4.13 and the arguments in the proof of Lemma 4.14 can be applied to each connected component individually, yielding even better bounds on the combinatorial diameter and the circuit diameter $C\mathcal{D}_{fm}$.

4.4 Lower bounds on circuit diameters

We now complement our discussion and turn to lower bounds on the circuit diameters by constructing dual network flow polyhedra that have at least a certain diameter. We first prove that all circuit distances are bounded below by $|V| - 1$. Then we consider $C\mathcal{D}_{fm}$ more thoroughly and show that it can in fact exceed $|V| - 1$ by an arbitrary constant. To this end we will introduce a special glueing construction for graphs that allows us to simply add up the circuit distances of the associated dual network flow polyhedra.

But we begin with Theorem 4.3 (ii), that is, we show the lower bound of $|V|-1$ on the soft circuit distance $C\mathcal{D}$. In fact, this result tells us that the bipartite upper bound of $|V| - 2$ does not hold for any of the circuit distance categories in the more general setting of arbitrary graphs. Further, this shows that our upper bound on the weak and soft circuit distances, Theorem 4.2 (iii), is tight.

Lemma 4.15. *For any n, there are dual network flow polyhedra defined on graphs $G = (V, E)$ on $|V| = n$ nodes that have (soft) circuit diameter at least $|V| - 1 = |E| - |V| + 1$.*

Proof. We provide two examples of dual network flow polyhedra that satisfy the claim.

1) Paths. Let n be given. Consider the polyhedron $P_{G^p,\mathbf{c}^p}$ defined on the following graph G^p, edges labeled with the respective values of $\mathbf{c}^p$.

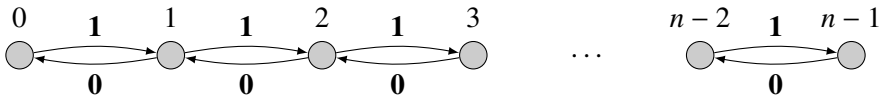

Figure 4.12: The graph G^p and the corresponding values of $\mathbf{c}^p$.

The following two spanning trees correspond to vertices of $P_{G^p,\mathbf{c}^p}$.

Figure 4.13: Two vertices of $P_{G^p,\mathbf{c}^p}$.

Clearly, the circuits as connected subsets of nodes are of the form $\{0, \dots, i\} \dot\cup \{i + 1, \dots, n - 1\}$ for some $i = 0, \dots n - 2$. Any such element affects only the edges $(i, i + 1)$ and $(i + 1, i)$ (or,

equivalently, the respective inequalities $-u_i + u_{i+1} \leq c_{i,i+1}$ and $-u_{i+1} + u_i \leq c_{i+1,i}$). Thus, we need at least $n-1$ circuit steps to reverse the $n-1$ edges in which the vertices depicted in Figure 4.13 differ.

2) Stars. Let n be given. Consider the polyhedron $P_{G^s,\mathbf{c}^s}$ defined on the following graph G^s, edges labeled with the respective values of $\mathbf{c}^s$.

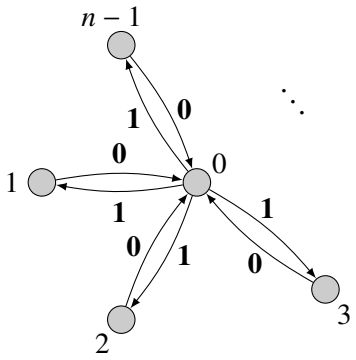

Figure 4.14: The graph G^s and the corresponding values of $\mathbf{c}^s$.

The following two spanning trees correspond to vertices of $P_{G^s,\mathbf{c}^s}$.

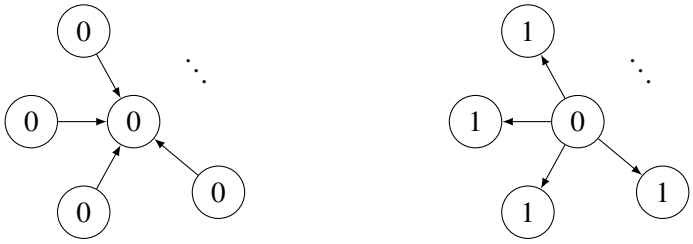

Figure 4.15: Two vertices of $P_{G^s,\mathbf{c}^s}$.

Clearly, the circuits as connected subsets of nodes are of the form $\{i\} \dot\cup (V \setminus \{i\})$ for some $i = 1, \ldots n-1$. Any such element affects only the edges $(0, i)$ and $(i, 0)$ (or, equivalently, the respective inequalities $-u_0 + u_i \leq c_{0,i}$ and $-u_i + u_0 \leq c_{i,0}$). Thus, we need at least $n-1$ circuit steps to reverse the $n-1$ edges in which the vertices depicted in Figure 4.15 differ. □

We now turn to the circuit distance $C\mathcal{D}_{fm}$ again. We begin with a small example for which $C\mathcal{D}_{fm}$ is even larger than $|V| - 1$, the lower bound we just proved in Lemma 4.15. More precisely, we get a circuit distance of $|V| = 4$. What seems just like a minor refinement in fact includes a tremendous observation: For general dual transportation polyhedra, there may not be a maximal circuit step that inserts an edge from the spanning tree representing the target vertex. Note that in the undirected bipartite case we are always able to apply such a step and the proof of Lemma 4.7 relied on this idea.

Example 4.16 ($C\mathcal{D}_{fm} = |V|$)**.** *Consider the dual network flow polyhedron $P_{G,\mathbf{c}}$ associated with the graph on four nodes depicted in Figure 4.16; the edges are labeled with the respective values of* $\mathbf{c}$*.*

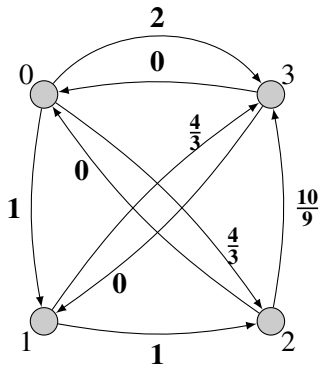

Figure 4.16: A graph G on four nodes and the values of $\mathbf{c} \in \mathbb{R}^{|E|}$.

Observe that the polyhedron $P_{G,\mathbf{c}}$ is non-degenerate (there can be no cycle of tight inequalities). The following two spanning trees correspond to vertices $\mathbf{u}^{(1)}$ and $\mathbf{u}^{(2)}$ of $P_{G,\mathbf{c}}$.

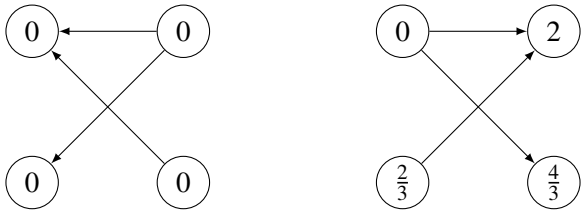

$T_1 = G(\mathbf{u}^{(1)})$ $\qquad$ $T_2 = G(\mathbf{u}^{(2)})$

Figure 4.17: Two vertices of $P_{G,\mathbf{c}}$.

These two vertices are connected via the following edge walk of length four. Hence their circuit distance $C\mathcal{D}_{fm}$ is at most four.

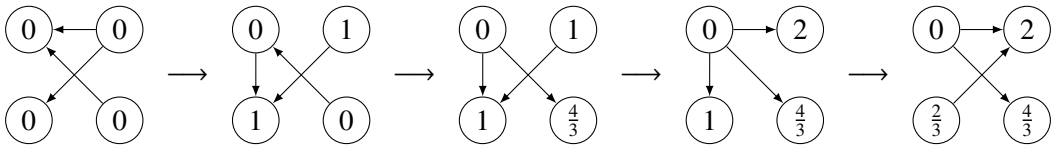

Figure 4.18: A maximal feasible circuit walk from $\mathbf{u}^{(1)}$ to $\mathbf{u}^{(2)}$.

In Figure 4.16 we now illustrate all maximal feasible circuit steps we could apply at $\mathbf{u}^{(1)}$, leading to points $\mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(6)}$. The respective circuits are stated below the graphs as subsets $S \subseteq V$ we increase or decrease, respectively (thus we have $0 \notin S$ and $0 \in R = V \backslash S$). Note that $S = \{1, 2\}$ is not applicable. Observe that in each case the inserted (bold) edge is not contained in $E(T_2)$.

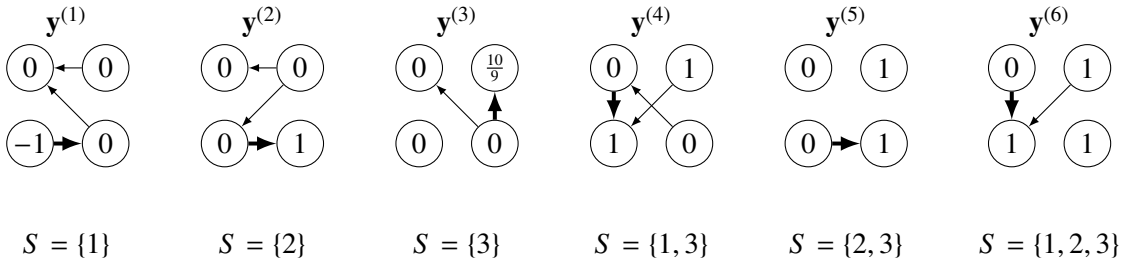

Figure 4.19: All feasible maximal circuit steps at $\mathbf{u}^{(1)}$.

This already shows that for arbitrary graphs we cannot insert an edge from the target tree in every step. However, we still have to verify that the circuit distance from $\mathbf{u}^{(1)}$ *to* $\mathbf{u}^{(2)}$ *is indeed* $|V| = 4$. *For this concrete example it can be done in a lengthy consideration of all options to apply two maximal feasible circuit steps.*

However, we can also use a far more elementary observation on circuit walks: It would be sufficient to have that in the remaining circuit walk we never insert two edges (from $E(T_2)$*) in one single step. Now, even if the latter property would not hold for a given* $\mathbf{c}$*, we could always achieve this with a slight perturbation:*

A finite number of linear conditions on the right-hand sides $\mathbf{c}$ *can guarantee that a certain maximal circuit step inserts at most (and thus exactly) one edge. Therefore, finitely many such conditions can guarantee that in all circuit walks of a fixed finite length k no maximal circuit step inserts two edges at the same time. These finitely many conditions only exclude right-hand sides* $\mathbf{c}$ *that lie in the union of a finite number of hyperplanes. In particular, we can get a* $\mathbf{c}$ *with the desired property by a perturbation.*

We now use this graph satisfying '$\mathcal{CD}_{fm} = |V|$' to obtain family of graphs with associated polyhedra for which $\mathcal{CD}_{fm}$ exceeds the number of nodes $|V|$ by an arbitrary constant. This reinforces the observation we cannot expect similar bounds for $\mathcal{CD}_{fm}$ (in terms of $|V|$) as in the bipartite case.

To this end, we now come to the *glueing construction* for graphs: If we glue k graphs together at a single, arbitrary node, we obtain a larger graph and the circuit diameter of the polyhedra associated with the component graphs just sum up to the circuit diameter of the polyhedron associated with the larger graph obtained by glueing.

More formally, let $G_i = (V_i, E_i)$, $i = 1, \dots, k$ be k connected directed graphs. For every graph choose an arbitrary node $v_0^i \in V_i$. By glueing the graphs together at the v_0^i, joining them to one node v_0, we construct a new graph $G = (V, E)$, with node and edge set given by

$$V := \{v_0\} \cup \bigcup_{i=1}^{k} \left(V_i \setminus \left\{ v_0^i \right\} \right) ,$$

$$E := \bigcup_{i=1}^{k} \left(\left\{ ab : \ ab \in E_i, a, b \neq v_0^i \right\} \cup \left\{ v_0 b : v_0^i b \in E_i \right\} \cup \left\{ a v_0 : a v_0^i \in E_i \right\} \right).$$

We depict the graphs G_i by highlighting the nodes v_0^i, while all remaining nodes and edges are

represented by a loop:

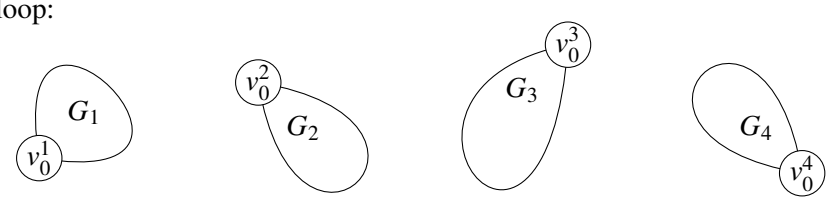

Figure 4.20: Four graphs in simplified description.

Glueing these 4 graphs together yields a graph G that can be illustrated as follows:

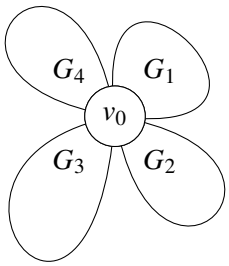

Figure 4.21: A graph obtained by glueing four graphs.

Now the following lemma tells us that the circuit diameter of the dual network flow polyhedra associated with the respective graphs are directly related.

Lemma 4.17. *Let $P_{G_i,\mathbf{c}^i}$, $i = 1, \ldots, k$ be arbitrary dual network flow polyhedra with circuit diameter equal to (at least) d_i, with respect to any circuit distance $\mathcal{CD}_\sim$. Let G be the graph obtained by glueing these k graphs together, and define $c \in \mathbb{R}^{|E|}$ by $c_{lj} = c^i_{lj}, lj \in E_i$, where we use $v^i_0 = v_0$.*

Then $P_{G,\mathbf{c}}$ has circuit diameter $\mathcal{CD}_\sim$ (at least) $\sum_{i=1}^k d_i$.

Proof. Let a circuit direction of $P_{G,\mathbf{c}}$ be given by a partition $V = R \dot\cup S$. Assume without loss of generality $v_0 \in R$. Then $S \subseteq V_i \backslash \left\{v^i_0\right\}$ for some $i \in \{1, \ldots, k\}$, as the node set S must be connected in the underlying graph and $v_0 \notin S$.

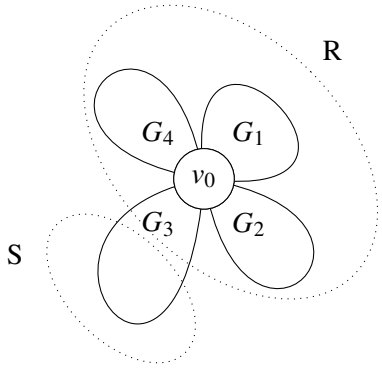

Figure 4.22: A circuit direction on a graph obtained by glueing.

Therefore, every step of a circuit walk modifies only variables of one single component G_i,

such that every edge walk circuit walk in $P_{G,\mathbf{c}}$ of length d' directly translates into k circuit walks in $P_{G_1,\mathbf{c}^1}, \ldots, P_{G_k,\mathbf{c}^k}$ of length $d'_1, \ldots, d'_k$ with $\sum_{i=1}^{k} d'_i = d'$ and vice versa. □

We now apply the glueing constructions to the graph on four nodes we introduced in Example 4.16. Denote by G^k the graph obtained by glueing k copies of that graph. Then G^k is a graph on $3k+1$ nodes and by Lemma 4.17 it has diameter $C\mathcal{D}_{fm}$ at least $4k$. This provides us with a family of graphs G^k with associated dual network flow polyhedra that admit a circuit diameter of at least $4k = |V| + k - 1$. Hence we exceed the number of nodes by $k-1$, which can be chosen arbitrarily big. In particular we arbitrarily violate the diameter bound for bipartite graphs.

This further yields a family of polyhedra whose circuit diameter approaches $\frac{4}{3}|V|$ and we get the following lower bound statement for the circuit diameter $C\mathcal{D}_{fm}$ of dual network flow polyhedra.

Lemma 4.18. *For any $n \geq 4$, there is a dual network flow polyhedron defined on a graph $G = (V, E)$ on $|V| = n$ nodes that has circuit diameter $C\mathcal{D}_{fm}$ at least $\frac{4}{3}|V| - 4$.*

Proof. For $n = 3k+1$ with $k \in \mathbb{Z}$ the claim follows by choosing $G = G^k$, as $k = \frac{|V|-1}{3}$ and we have circuit diameter $C\mathcal{D}_{fm}$ at least $4k$. If $n = 3k+2$ ($n = 3k+3$) we simply add one leaf (two leaves) to G^k. Then $k = \frac{|V|-2}{3}$ ($k = \frac{|V|-3}{3}$) and $C\mathcal{D}_{fm}$ is still at least $4k$. □

This concludes the proof of Theorem 4.3 and our investigation of the circuit diameters of dual network flow polyhedra.

Outlook

In this thesis, we were concerned with the concept of circuit diameters. We presented several fundamental results, most notably our hierarchy of circuit diameters, and we demonstrated the usefulness of our ideas by investigating two special classes of polyhedra.

However, as we first introduced this whole topic just recently in [3] and [5], there are still many open questions and several possibilities for developing the ideas further and in new directions. In this final chapter we highlight some of them.

First of all, the results presented in this thesis immediately give rise to the following questions:

- **?** Are the weak circuit distance categories $C\mathcal{D}_{fb(r)}$ and $C\mathcal{D}_{f(r)}$ distinct in general?
  (Recall that these are the only weak inequalities in the hierarchy and an example that proves strict inequality has to be in dimension five or higher.)

- **?** For general $M\times N$ *transportation polytopes*, can we improve the linear upper bound of $8(M + N - 2)$ for the combinatorial diameter or for the circuit diameter, respectively? Does the Hirsch conjecture (Hirsch bound) hold?

- **?** For dual *network flow polyhedra*, is there a linear upper bound on the circuit diameter, or even on the combinatorial diameter? Does the Hirsch bound of $|E| - |V| + 1$ hold?
  On the other hand, can we find stronger lower bounds?

Another straightforward continuation of this work is the investigation of the hierarchy for further classes of polyhedra. Are there more classes for which the hierarchy collapses? Are there polyhedra for which the respective circuit diameters differ widely?

Aside from that, there arise several interesting questions concerning the 'original' circuit diameter $C\mathcal{D}_{fm}$. First of all, it is still open whether there is a general polynomial upper bound. In particular we would like to know whether the Hirsch bound is always satisfied, as already conjectured in [5]:

Conjecture 3. *For any d-dimensional polyhedron with f facets the circuit diameter $C\mathcal{D}_{fm}$ is bounded above by $f - d$.*

Recall that the Hirsch conjecture does not hold for the combinatorial diameter [21, 27]. However, it is open whether the counterexamples from [21] and [27] give rise to a counterexample for Conjecture 3 as well. The following question may give some indication.

? Given a polyhedron, can it always be turned into a polyhedron of same combinatorial structure for which the circuit diameter equals its combinatorial diameter?
In other words, is there a perturbation such that every shortest feasible maximal circuit walk actually goes along the edges of the polyhedron?

In Lemma 2.23 in Section 2.4 we saw that this is true in dimension two: For any k there is a polygon on k vertices for which the circuit diameter and the combinatorial diameter coincide.

However, even the effects of perturbing only the right-hand sides are hard to predict as demonstrated in Example 2.1.3 in Section 2.1.3. But understanding these effects could be extremely useful: A positive answer to the following question would simplify proving upper bounds on the circuit diameter since we could assume non-degeneracy just as we did for the combinatorial diameter.

? When perturbing the right-hand sides, is the maximal circuit diameter attained by a non-degenerate polyhedron?

But there are also much more far-reaching questions. Up to now, we were only concerned with the circuit distance itself but we did not consider algorithmic aspect such as the (efficient) computation of optimal circuit walks, neither did we study how to apply the different concepts of circuit walks in augmentation algorithms. In particular, there is still little understanding in how to exploit the idea of going infeasible.

We conclude the thesis with a short initial discussion of a natural extension of the circuit diameter to an integral analog – the Graver diameter.

The integral circuit diameter

In his seminal paper [14], Graver did not only treat the continious case, but also provided a universal integral test set, by now also known at the *Graver basis*.

Definition (See Def. 4.3 in [14])**.** *The* Graver basis $\mathcal{G}(A,B)$ *associated with matrices A and B consists of those non-zero vectors* $\mathbf{g} \in \ker_{\mathbb{Z}}(A) = \{\mathbf{z} \in \mathbb{Z}^n : A\mathbf{z} = \mathbf{0}\}$*, for which $B\mathbf{g}$ is $\sqsubseteq$-minimal in the set* $\{ B\mathbf{y} : \mathbf{y} \in \ker_{\mathbb{Z}}(A) \backslash \{\mathbf{0}\} \}$.

For $\mathbf{v}, \mathbf{w} \in \mathbb{R}^d$ *we have* $\mathbf{w} \sqsubseteq \mathbf{v}$ *(the partial ordering on $\mathbb{R}^d$), if $v_i w_i \geq 0$ and $|w_i| \leq |v_i|$ for all $i = 1, \ldots, d$.*

It is not hard to see that $C(A,B) \subseteq \mathcal{G}(A,B)$. Note that we can have $C(A,B) \neq \mathcal{G}(A,B)$. Graver bases admit a representation property for the integral elements in the kernel of A similar to the representation property for circuits (see Theorem 4.5 in [14]): Every element $\mathbf{v} \in \ker_{\mathbb{Z}}(A)$ can be written as a positive linear $\sqsubseteq$-compatible sum of elements in $\mathcal{G}(A,B)$, i.e. $\mathbf{v} = \sum_{i=1}^{k} \alpha_i \mathbf{g}^i$ and we have $\mathbf{g}^i \in \mathcal{G}(A,B)$, $B\mathbf{g}^i \sqsubseteq B\mathbf{v}$, and $\alpha_i \in \mathbb{Z}_+$ for all $i = 1, \ldots, k$. It follows immediately that the Graver basis is an optimality certificate for linear objectives (see Corollary 4.6 in [14]). Further, is is not hard to see that every edge direction of the integer hull $P_I = \operatorname{conv}(P \cap \mathbb{Z}^n) = \operatorname{conv}(\{\mathbf{x} \in \mathbb{Z}^n : A\mathbf{x} = \mathbf{b},\ B\mathbf{x} \leq \mathbf{d}\})$ is contained in the Graver basis [14].

The well-studied common setup for Graver bases is $\mathcal{G}(A) \coloneqq \mathcal{G}(A, I_n)$: We know that $\mathcal{G}(A)$ also provides optimality certificates for the minimization of separable convex objective functions over the lattice points of a polyhedron [23], that at most polynomially many (in the binary encoding length of the input data) Graver-best augmentation steps are needed in order to reach an optimal solution [18], and that N-fold separable-convex integer linear programs can be solved in polynomial time [11, 16, 17]. For a more thorough introduction to the theory of Graver bases and for more references on this topic, see for example [9, 25].

It is not hard to see that the definitions of circuit walks, circuit distances and circuit diameters from Section 2.1 readily translate to the integral case. Note that an *integral circuit walk* or *Graver walk* connects two vertices $\mathbf{v}^{(1)}$, $\mathbf{v}^{(2)}$ of the integer hull P_I of the polyhedron and that all points of a circuit walk $\mathbf{v}^{(1)} = \mathbf{y}^{(0)}, \dots, \mathbf{y}^{(k)} = \mathbf{v}^{(2)}$ must satisfy $\mathbf{y}^{(i)} \in \mathbb{Z}^n$. Then the Graver diameter is the maximum length of a shortest Graver walk between any two vertices of P_I, denoted $\mathcal{GD}$ etc.

We close this brief outlook with a demonstration of the challenges that come with integrality. Therefore we compare $\mathcal{CD}_{fm}$ and $\mathcal{GD}_{fm}$. The following fundamental examples show that they cannot be related via '$\leq$' or '$\geq$' in general!

Observation. *We can have $\mathcal{GD}_{fm} < \mathcal{CD}_{fm}$ as the Graver basis might contain additional elements.*

We consider a polygon on four vertices. Clearly we have $\mathcal{CD}_{fm} = 2$ as the circuits coincide with the directions of the edges $\binom{1}{0}$ and $\binom{1}{-2}$. In contrast, the Graver basis contains two more vectors $\binom{0}{1}$ and $\binom{1}{-1}$. Therefore, opposite vertices are connected by a Graver walk of length one.

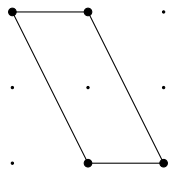

Figure 23: A polytope of Graver diameter one.

Observation. *We can have $\mathcal{GD}_{fm} > \mathcal{CD}_{fm}$ as the combinatorial structure of a polytope might differ from the combinatorial structure of its integer hull.*

Let us have a look at the polygon on three vertices in Figure 24, whose integer hull actually has four vertices. Then we have $\mathcal{CD}_{fm} = 1$ as any two vertices of the polygon are connected by an edge. However, opposite vertices of the integer hull cannot be connected by a Graver walk of length one as neither $\binom{2}{1}$ nor $\binom{2}{-1}$ belong to the Graver basis.

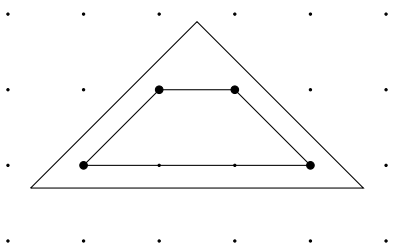

Figure 24: A polytope of Graver diameter two.

Observation. *We can have $\mathcal{GD}_{fm} > \mathcal{CD}_{fm}$, even if $P = P_I$, as the optimal circuit walk is not necessarily integral, i.e. not a Graver walk.*

We consider the two-dimensional polytope on six vertices illustrated in Figure 25. It is not hard to check that any two vertices are connected by a circuit walk of length at most two, thus $\mathcal{CD}_{fm} = 2$. We now have a closer look at a circuit walk from $\mathbf{v}^{(1)}$ to $\mathbf{v}^{(2)}$.

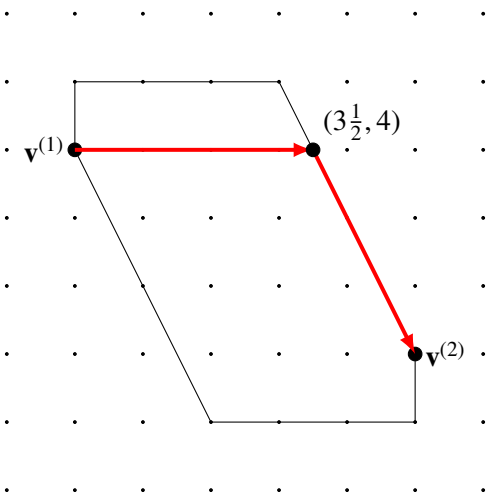

Figure 25: A circuit walk of length two.

Note that this circuit walk is not integral! The Graver basis associated with this polyhedron consists of the edge directions and one additional element, $\binom{1}{-1}$. Thus, there are four possible first Graver steps at $\mathbf{v}^{(1)}$, but none of these first steps allows to reach $\mathbf{v}^{(2)}$ with only one more Graver step. Hence we have Graver diameter three.

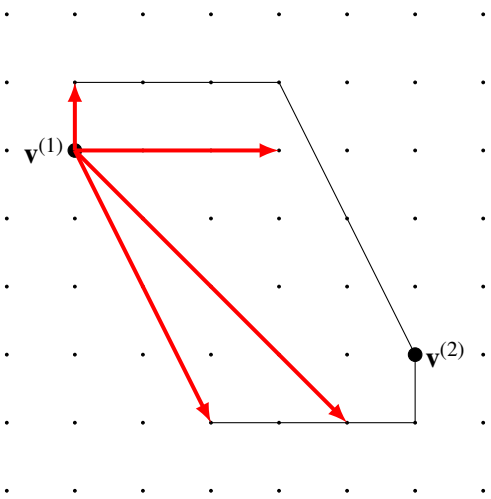

Figure 26: Maximal Graver steps at $\mathbf{v}^{(1)}$.

Bibliography

[1] M. L. Balinski. The Hirsch conjecture for dual transportation polyhedra. *Mathematics of Operations Research*, 9(4):629–633, 1984.

[2] S. Borgwardt. On the diameter of partition polytopes and vertex-disjoint cycle cover. *Mathematical Programming, Series A*, 141:1–20, 2013.

[3] S. Borgwardt, J. A. De Loera, and E. Finhold. Edges vs Circuits: a Hierarchy of Diameters in Polyhedra. eprint arXiv:1409.7638, 2014.

[4] S. Borgwardt, J. A. De Loera, E. Finhold, and J. Miller. The Hierarchy of Circuit Diameters and Transportation Polytopes. eprint arXiv:1411.1701, 2014.

[5] S. Borgwardt, E. Finhold, and R. Hemmecke. On the circuit diameter of dual transportation polyhedra. eprint arXiv:1405.3184, to appear in *SIAM Journal on Discrete Mathematics*.

[6] S. Borgwardt, E. Finhold, and R. Hemmecke. Quadratic diameter bounds for dual network flow polyhedra. eprint arXiv:1408.4184, 2014.

[7] G. Brightwell, J. Heuvel, and L. Stougie. A Linear Bound on the Diameter of the Transportation Polytope. *Combinatorica*, 26:133–139, 2006.

[8] G. Dantzig. *Linear Programming and Extensions*. Princeton Univ. Press, 1963.

[9] J. A. De Loera, R. Hemmecke, and M. Köppe. *Algebraic and Geometric Ideas in the Theory of Discrete Optimization*, volume 14 of *MOS-SIAM Series on Optimization*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2013.

[10] J. A. De Loera, R. Hemmecke, and J. Lee. Augmentation Algorithms for Linear and Integer Linear Programming. eprint arXiv:1408.3518, 2014.

[11] J. A. De Loera, R. Hemmecke, S. Onn, and R. Weismantel. *n*-fold integer programming. *Discrete Optimization*, 5:231–241, 2008. In Memory of George B. Dantzig.

[12] J. A. De Loera and E. D. Kim. Combinatorics and Geometry of Transportation Polytopes: An Update. eprint arXiv:1307.0124, 2013.

[13] J. A. De Loera, E. D. Kim, S. Onn, and F. Santos. Graphs of transportation polytopes. *Journal of Combinatorial Theory - Series A*, 116:1306–1325, 2009.

[14] J. E. Graver. On the foundations of linear and integer linear programming I. *Mathematical Programming*, 9:207–226, 1975.

[15] R. Hemmecke. On the positive sum property and the computation of graver test sets. *Mathematical Programming*, 96:247–269, 2003.

[16] R. Hemmecke, M. Köppe, and R. Weismantel. Graver basis and proximity techniques for block-structured separable convex integer minimization problems. eprint arXiv:1207.1149, to appear in *Mathematical Programming*.

[17] R. Hemmecke, S. Onn, and L. Romanchuk. N-Fold integer programming in cubic time. *Mathematical Programming*, 137:325–341, 2013.

[18] R. Hemmecke, S. Onn, and R. Weismantel. A polynomial oracle-time algorithm for convex integer minimization. *Mathematical Programming*, 126:97–117, 2011.

[19] F. L. Hitchcock. The distribution of a product from several sources to numerous localities. *Journal of Mathematics and Physics*, 20:224–230, 1941.

[20] E. D. Kim and F. Santos. An update on the Hirsch conjecture. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 112(2):73–98, 2010.

[21] V. Klee and D. W. Walkup. The d-step conjecture for polyhedra of dimension $d < 6$. *Acta Mathematica*, 117:53–78, 1967.

[22] V. Klee and C. Witzgall. Facets and vertices of transportation polyhedra. *Mathematics of the Decision Science, Part 1. Lectures in Applied Mathematics*, 11:257–282, 1968.

[23] K. Murota, H. Saito, and R. Weismantel. Optimality criterion for a class of nonlinear integer programs. *Operations Research Letters*, 32:468–472, 2004.

[24] D. Naddef. The Hirsch conjecture is true for $(0, 1)$-polytopes. *Mathematical Programming*, 45:109–110, 1989.

[25] S. Onn. *Convex Discrete Optimization*. Zurich Lectures in Advanced Mathematics. European Mathematical Society, 2010.

[26] R. T. Rockafellar. The Elementary Vectors of a Subspace of R^n. *Combinatorial Mathematics and its Applications*, pages 104–127, 1969, 2011.

[27] F. Santos. A counterexample to the Hirsch conjecture. *Annals of Mathematics*, 176:383–412, 2012.

[28] B. Sturmfels and R. R. Thomas. Variation of cost functions in integer programming. *Mathematical Programming*, 77:357–387, 1997.

[29] V. A. Yemelichev, M. M. Kovalëv, and M. K. Kravtsov. *Polytopes, graphs and optimisation*. Cambridge University Press, Cambridge, 1984.