

Fiber graphs

Dissertation

zur Erlangung des akademischen Grades

doctor rerum naturalium
(Dr. rer. nat.)

von Tobias Windisch, M.Sc. (hons)
geb. am 12.10.1987 in Bamberg

genehmigt durch die Fakultät für Mathematik
der Otto-von-Guericke-Universität Magdeburg

Gutachter: Jun.-Prof. Dr. Thomas Kahle
Otto-von-Guericke-Universität Magdeburg

Prof. Seth Sullivant, PhD
North Carolina State University

eingereicht am: 22.07.2016

Verteidigung am: 09.12.2016

Abstract

A fiber graph is a graph on the integer points of a polytope whose edges come from a set of allowed moves. Fiber graphs are given implicitly which makes them a useful tool in many applications of statistics and discrete optimization whenever an exploration of vast discrete structures is needed. The first part of this thesis discusses the graph-theoretic structure of fiber graphs with a particular focus on their diameter and edge-expansion. We define the fiber dimension of a simple graph as the smallest dimension where it can be represented as a fiber graph and prove an upper bound on the fiber dimension that only depends on the chromatic number of the graph. In the second part, random walks on fiber graphs are studied and it is shown that, when a fixed set of moves is used, rapid mixing is impossible. In order to improve mixing rates for fiber walks in fixed dimension, we evaluate possible adaptations of the set of moves, one that adds a growing number of linear combinations of moves to the set of allowed moves and one that allows arbitrary lengths of single moves. We show that both methods lead to spectral expanders in fixed dimension. Finally, the parity binomial edge ideal of a graph is introduced. Unlike the binomial edge ideal, it does not have a square-free Gröbner bases and is radical if only if the graph is bipartite or the characteristic of the ground field is not two. We compute the universal Gröbner basis and the minimal primes and show that both encode combinatorics of even and odd walks.

Zusammenfassung

Ein Fasergraph ist ein Graph auf den ganzzahligen Vektoren eines Polytops, dessen Kanten aus einer Menge zugelassener Richtungsvektoren entstehen. Fasergraphen sind implizit gegeben und deshalb ein wichtiges Werkzeug in vielen Anwendungen der Statistik und Optimierung, wann immer riesige diskrete Strukturen untersucht werden. Der erste Teil dieser Arbeit beschäftigt sich mit der graphen-theoretischen Struktur von Fasergraphen mit besonderem Augenmerk auf deren Durchmesser und Kanten-Expansion. Wir definieren die Faserdimension eines einfachen Graphen als die kleinste Dimension, in der er als Fasergraph dargestellt werden kann, und beweisen eine obere Schranke der Faserdimension, die nur von der chromatischen Zahl des Graphen abhängt. Im zweiten Teil werden Zufallsbewegungen auf Fasergraphen untersucht und es wird gezeigt, dass mit fixierten Richtungsvektoren eine schnelle Konvergenz nicht möglich ist. Um die Konvergenzrate zu verbessern, untersuchen wir mögliche Anpassungen der zugelassenen Richtungsvektoren: eine, die eine wachsende Zahl an Linearkombinationen der Richtungen hinzufügt und eine zweite, die fixierte Richtungen beliebiger Länge zulässt. Wir zeigen, dass beide Methoden spektrale Expander in fixierter Dimension liefern. Zum Schluss wird das paritätäre binomische Kantenideal eines Graphen vorgestellt. Dieses hat, anders als das binomische Kantenideal, keine quadratfreie Gröbnerbasis und ist radikal genau dann, wenn der Graph bipartit oder die Charakteristik des Grundkörpers ungleich zwei ist. Wir berechnen die universelle Gröbnerbasis sowie die minimalen Primideale und zeigen, dass beide eine Kombinatorik von geraden und ungeraden Pfaden kodieren.

Acknowledgements

First and foremost, I am indebted to my advisor, Thomas Kahle, for his excellent guidance and enduring support. During all these years, his door was open for me at any time and he never failed to be reachable. He taught me so extremely many things about mathematics and beyond, not alone because he was always prepared to give an unprepared blackboard talk about almost any topic I came up with. All my knowledge about binomial ideals, Markov bases, 3D-printing, handball, and crypto currencies is due to him.

Research is impossible without money, and I would like to thank all the institutions that were willing to support me with theirs: the *German National Academic Foundation* (Studienstiftung des deutschen Volkes), the *TopMath* program at TU Munich, and the *German Academic Exchange Service* (Deutscher Akademischer Austauschdienst). I also would like to thank all the current and former members of the *Institute of Algebra and Geometry* at OvGU Magdeburg and the research group on *Applied Geometry and Discrete Mathematics* at TU Munich for the very inspiring working atmosphere and the many great talks, discussions, and coffee breaks.

I thank my co-authors Raymond Hemmecke, Thomas Kahle, Camilo Sarmiento, and Caprice Stanley for their work and effort, and for the great time I had while working on our projects.

My interests in algebraic statistics and fiber graphs in particular was triggered by Raymond Hemmecke, and I thank him for all his advice and support during my time at TU Munich. I am indebted to Seth Sullivant for pointing me to the hit-and-run algorithm and for initiating the joint work with Caprice Stanley on compressed fiber graphs. I owe much to Alexander Engström who had always a friendly ear for my questions and who encouraged me to pursue them further. Mathematical thanks go also to Gennadiy Averkov, Johannes Hofscheier, and Benjamin Nill for helping me with my various questions on polytopes and to the innumerable many people I talked to on all the conferences I visited over the last years. I also want to express my gratitude to Sonja Petrović, Seth Sullivant, and Ruriko Yoshida for their great hospitality during my US visit.

This thesis could not have been created without the many magnificent open-source software packages that helped me finding my theorems. Thus, I am grateful to the developers of `4ti2` [1], `LatTE` [9], `Macaulay2` [58], `Normaliz` [20], `polymake` [57], `R` [98], `sage` [34], and many more for their effort to make mathematics available to computers.

I thank all my friends, particularly Linus and Markus, for their tireless persistence to bring me back to life whenever I was buried in mathematics. Finally, I thank Katharina for her unconditional love and her support. In my life as a PhD student, geographically teared between Magdeburg, Immenstadt, Bamberg, and Munich, mentally spread among algebra, combinatorics, and statistics, in planes and trains, up hill and down dale, she was always my universal constant and the source of my motivation.

Contents

1	Introduction	1
1.1	Random walks	2
1.2	Fiber graphs	6
1.3	Running examples	9
1.4	Fiber walks in statistics	10
2	Graph properties of fiber graphs	15
2.1	Connectedness and connectivity	15
2.2	Bounds on the diameter	18
2.3	Graph degrees	22
2.4	Edge-expansions	24
3	The fiber dimension of a graph	30
3.1	Embeddings	30
3.2	Fiber dimension one	33
3.3	Complete graphs	35
3.4	Distinct pair-sum polytopes	36
3.5	Computational aspects	38
4	Symmetric fiber walks	40
4.1	Fixed Markov bases	40
4.2	Adapted Markov bases	44
4.3	Varying constraint matrices	47
5	Heat-bath walks	49
5.1	Spectral analysis	50
5.2	Augmenting Markov bases	56
5.3	Best move selection	61
6	Parity binomial edge ideals	63
6.1	Markov bases	64
6.2	Universal Gröbner basis	66
6.3	Minimal primes	72
6.4	Radicality	76
	Bibliography	79
	Index	87

1 Introduction

The assessment of statistical models, based on expectation and experience, tries to reduce the complexity of our ambient world and makes questions about nature amenable to algorithms and mathematics. The prevalent working scheme in inferential statistics is to draw conclusions based on finitely many independent observations about the whole set of interests. Frequently, observations are represented as multi-way *contingency tables* and the class of *log-linear* statistical models describes how their attributes relate among each other [28]. Given an observed contingency table on the one and a statistical model on the other hand, an intruding question is how well the model explains the observed data, that is determining its *goodness-of-fit*. The running engine that many exact goodness-of-fit tests for log-linear models have under their hood is a random walk on a *fiber graph*. Essentially, a fiber graph is a graph on the lattice points of a polytope where two nodes are adjacent if their difference lies in a set of allowed moves. Being graphs on lattice points, the combinatorial outreach of fiber graphs goes far beyond statistics as they appear naturally in discrete optimization [31] and commutative algebra, particularly in the context of toric ideals [109] and matrix gradings [79]. Regardless of the application, an interesting situation is when fiber graphs are connected, in which case the set of moves is called a *Markov basis*. The seminal paper of Diaconis and Sturmfels [39] triggered a lot of research on fiber graphs that was mostly dedicated to the determination of Markov bases for a variety of statistical models [72, 112, 110, 111] and the design of algorithms for their computation [42, 61, 41]. Since Markov bases can be determined with tools from commutative algebra, research on fiber graphs is a topic of *algebraic statistics*, a field that studies statistical models with algebraic methods.

Due to the implicit structure of fiber graphs, Markov bases give rise to implementable and irreducible random walks that can approximate any probability distribution on the underlying set of lattice points. The number of steps needed to approximate a given distribution sufficiently is the *mixing time* of the random walk. As a general impediment of any Markov chain Monte Carlo approach, there is a priori no information on the mixing time of a random walk available and its exact determination remains computationally unfeasible. To put hands on mixing times, a combinatorial understanding of fiber graphs is necessary and it is the goal of this thesis to take a first step towards their graph-theoretic understanding.

This thesis summarizes and unifies the author's work on questions related to fiber graphs [119, 120, 108] and parity binomial edge ideals [73]. Being at the non-empty intersection of commutative algebra, statistics, graph theory, and discrete geometry, a lot of different concepts, notations, and terminologies are required in this thesis. Chapter 1 is devoted to briefly set up the very basic definitions that are used throughout and it provides pointers to the literature for more details on the particular topic. More specific concepts are introduced in the respective chapters locally. Typically, the node set of a fiber graph is a *fiber* $\mathcal{F}_{A,b} := \{u \in \mathbb{N}^d : Au = b\}$ of an integer matrix $A \in \mathbb{Z}^{m \times d}$ and a right-hand side $b \in \mathbb{Z}^m$, and our leading principle is to study fiber graphs that originate from the same integer matrix, but for varying right-hand sides. In the first part of this thesis, we study how much a priori information about the graph structure of fiber graph can be read off the input only. Our results concern the asymptotic character of

diameter and edge-expansion of fiber graphs for sequences of right-hand sides $(b_i)_{i \in \mathbb{N}}$ in $\mathbb{N}A$. We state properties on $(b_i)_{i \in \mathbb{N}}$ and the moves that imply a decline of the edge-expansion as $\frac{1}{i}$ and a growth of the diameter linearly in i . Every simple graph can be written as a fiber graph, but not all structures are possible in any dimension. This motivates our investigation of the *fiber dimension* of a graph in Chapter 3, that is the smallest dimension in which a graph can be represented as fiber graph. The second part of this thesis deals with random walks on fiber graphs. First, we use our results from Chapter 2 to prove that symmetric fiber walks cannot mix rapidly in fixed dimension. A canonical adaption of Markov bases is stated and shown to be faster than fiber walks with conventional Markov bases asymptotically for large right-hand sides. A different adaption is shown in Chapter 5, where *heat-bath* walks on *compressed* fiber graphs are examined. We prove that the diameter of compressed fiber graphs is bounded from above by a constant when the right-hand side varies (Theorem 2.2.17) and that under additional assumptions on the Markov basis, the heat-bath walk mixes rapidly on them (Theorem 5.2.9). Finally, the *parity binomial edge ideal* \mathcal{I}_G of a graph $G = (V, E)$ is studied in Chapter 6, which encodes the adjacency relations into a binomial ideal in $\mathbb{k}[x_v, y_v : v \in V]$. We determine a prime decomposition and prove radicality for parity binomial edge ideals when $\text{char}(\mathbb{k}) \neq 2$ in terms of combinatorial invariants of the graph G . Moreover, the universal Gröbner basis of \mathcal{I}_G is stated in Section 6.2 and shown to support one part of a recent conjecture from [8].

Basic notations

The natural numbers are $\mathbb{N} := \{0, 1, 2, \dots\}$. For any $n \in \mathbb{N}$, we set $[n] := \{m \in \mathbb{N} : 1 \leq m \leq n\}$ and we use $\mathbb{N}_{>n}$ and $\mathbb{N}_{\geq n}$ to denote the subsets of \mathbb{N} whose elements are strictly greater and greater than n respectively. For any $u \in \mathbb{Z}^d$, the symbols $u^+, u^- \in \mathbb{N}^d$ denote the unique support disjoint vectors that fulfill $u = u^+ - u^-$. The $d \times d$ identity matrix is denoted by I_d and the unit vectors of \mathbb{Q}^d are e_1, \dots, e_d . Let $\mathcal{S} \subseteq \mathbb{Q}^d$ and $i \in \mathbb{Q}$, then the i -th *dilation* of \mathcal{S} is the set $i \cdot \mathcal{S} := \{i \cdot s : s \in \mathcal{S}\}$. For another subset $\mathcal{T} \subseteq \mathbb{Q}$, we let $\mathcal{T} \cdot \mathcal{S} = \{t \cdot s : t \in \mathcal{T}, s \in \mathcal{S}\}$.

Let $(a_i)_{i \in \mathbb{N}}$ and $(b_i)_{i \in \mathbb{N}}$ be two sequences in \mathbb{Q} , then $(a_i)_{i \in \mathbb{N}} \in \mathcal{O}(b_i)_{i \in \mathbb{N}}$ if there exist $i_0 \in \mathbb{N}$ and $C \in \mathbb{Q}_{>0}$ such that $|a_i| \leq C \cdot |b_i|$ for all $i \geq i_0$. Similarly, $(a_i)_{i \in \mathbb{N}} \in \Omega(b_i)_{i \in \mathbb{N}}$ if there exist $i_0 \in \mathbb{N}$ and $C \in \mathbb{Q}_{>0}$ such that $|a_i| \geq C \cdot |b_i|$ for all $i \geq i_0$. The sequence $(a_i)_{i \in \mathbb{N}}$ is a subsequence of $(b_i)_{i \in \mathbb{N}}$ if there is a strongly increasing sequence $(i_k)_{k \in \mathbb{N}}$ in \mathbb{N} such that $a_{i_k} = b_k$ for all $k \in \mathbb{N}$.

A *graph* is always undirected and can have multiple loops. Occasionally, we point to the node set of a graph $G = (V, E)$ with $V(G) := V$ and to its edge set with $E(G) := E$. The *adjacency matrix* of G is denoted by $A^G \in \mathbb{N}^{|V| \times |V|}$ and for a node $v \in V$, $\text{deg}_G(v)$ denotes the number of edges in E incident to v . A (u, v) -*walk* in G of *length* r is a sequence $(w_1, \dots, w_{r+1}) \in V^{r+1}$ with $w_1 = u$ and $w_{r+1} = v$ such that $\{w_k, w_{k+1}\} \in E$ for all $k \in [r]$. A *path* is a walk where the intermediate nodes w_2, \dots, w_r are distinct and different from u and v . A *cycle* (*circuit*) is an (u, v) -walk (path) with $u = v$. For a set of nodes $S \subseteq V$, $G[S]$ denotes the *induced subgraph* on S . The *complete graph* and the *circuit graph* on n nodes are K_n and C_n respectively. The *complete r -partite graph* on node classes of size $n_1, \dots, n_r \in \mathbb{N}$ is denoted K_{n_1, \dots, n_r} .

1.1 Random walks

In this and the remaining sections, the notation is mainly borrowed from the excellent textbooks [40] on graph theory and [81] on Markov chains. When graphs meet probability theory, then walks become random walks. Different than their name let one suggest, random walks on

graphs are very deterministic and tangible in mathematics:

Definition 1.1.1. Let $G = (V, E)$ be a graph. A map $\mathcal{W} : V \times V \rightarrow [0, 1]$ is a *random walk* on G if for all distinct $s, t \in V$ with $\{s, t\} \notin E$, $\mathcal{W}(s, t) = 0$ and if for all $v \in V$, $\sum_{u \in V} \mathcal{W}(v, u) = 1$.

Let $G = (V, E)$ be a graph. When there is no ambiguity, a random walk is represented as a $|V| \times |V|$ -matrix, for instance when it is clear how the elements of V are ordered. Let $\mathcal{W} : V \times V \rightarrow [0, 1]$ and $\mathcal{W}' : V \times V \rightarrow [0, 1]$ be maps, then their *product* $\mathcal{W} \circ \mathcal{W}' : V \times V \rightarrow [0, 1]$ is

$$(\mathcal{W} \circ \mathcal{W}')(u, v) = \sum_{w \in V} \mathcal{W}(u, w) \cdot \mathcal{W}'(w, v).$$

Representing \mathcal{W} and \mathcal{W}' as matrices, then $\mathcal{W} \circ \mathcal{W}' = \mathcal{W} \cdot \mathcal{W}'$ is precisely the product of matrices. Let $\pi : V \rightarrow \mathbb{R}$ be a map, then we similarly define $(\mathcal{W} \circ \pi)(u) = \sum_{v \in V} \mathcal{W}(u, v) \cdot \pi(v)$ and $(\pi \circ \mathcal{W})(u) = \sum_{v \in V} \pi(v) \cdot \mathcal{W}(v, u)$ to be the multiplications of a matrix with a vector. We let \mathcal{W}^0 be the map that sends $(u, v) \in V \times V$ to 1 if $u = v$ and to 0 otherwise, and define $\mathcal{W}^t := \mathcal{W} \circ \mathcal{W}^{t-1}$ for $t \in \mathbb{N}_{\geq 1}$ recursively. With a random walk \mathcal{W} , the node set of its underlying graph can be explored randomly by selecting for any starting node $v \in V$ a node w randomly from V according to the distribution $\mathcal{W}(v, \cdot)$ and iterating the procedure at the new node. For any $t \in \mathbb{N}$, the quantity $\mathcal{W}^t(v, u)$ is then the probability that the random walk that starts at v is at u after t steps. The assumptions in Definition 1.1.1 guarantee that two distinct nodes that are subsequently visited are adjacent in G and hence this method produces walks randomly.

Remark 1.1.2. To perform a random walk on a graph $G = (V, E)$, we do not need an explicit description, or list, of its nodes V and edges E . It suffices to have a local understanding of the graph, for instance an algorithm that computes for any $v \in V$ its neighborhood in G . In computer science, graphs with that property are called *implicit graphs* [66, Definition 2.3].

We now define the *simple walk* on a graph which is, as the name suggests, a very simple random walk, mainly because it selects at every step uniformly from the respective neighborhoods:

Definition 1.1.3. The *simple walk* on $G = (V, E)$ is the map \mathcal{S}_G on $V \times V$ defined by

$$\mathcal{S}_G(u, v) = \begin{cases} \frac{A^G(u, v)}{\deg_G(u)}, & \text{if } \{u, v\} \in E \\ 0, & \text{otherwise} \end{cases}.$$

In general, every random walk $\mathcal{W} : V \times V \rightarrow [0, 1]$ comes along with a *discrete-time Markov chain* whose state space is the node set of its underlying graph [18, Section 1]: For any initial distribution $\pi_0 : V \rightarrow [0, 1]$ and any $t \in \mathbb{N}$, a probability mass function π_t on V is given by $\pi_t(v) := \sum_{u \in V} \pi_0(u) \cdot \mathcal{W}^t(v, u)$ and it is not hard to show that $(\pi_t)_{t \in \mathbb{N}}$ is a Markov chain. The following properties on the random walk ensure convergence – in a sense that is made precise later – of its attached Markov chain $(\pi_t)_{t \in \mathbb{N}}$ (see also Theorem 1.1.5):

Definition 1.1.4. Let $G = (V, E)$ be a graph and \mathcal{W} a random walk on G . The random walk \mathcal{W} is *symmetric* if \mathcal{W} is a symmetric map and *aperiodic* if for all $v \in V$, $\gcd\{t \in \mathbb{N}_{>0} : \mathcal{W}^t(v, v) > 0\} = 1$. A random walk \mathcal{W} is *irreducible* if for all $v, u \in V$, there exists $t \in \mathbb{N}$ such that $\mathcal{W}^t(v, u) > 0$ and *reducible* otherwise. A random walk \mathcal{W} is *reversible* if there exists a probability mass function $\mu : V \rightarrow [0, 1]$ such that $\mu(u) \cdot \mathcal{W}(u, v) = \mu(v) \cdot \mathcal{W}(v, u)$ for all $u, v \in V$. A probability mass function $\pi : V \rightarrow [0, 1]$ is a *stationary distribution* of \mathcal{W} if $\pi \circ \mathcal{W} = \pi$.

Irreducibility of random walks is an important and desired property since irreducible random walks have a unique stationary distribution [81, Corollary 1.17]. If the random walk is symmetric additionally, then this distribution is the uniform distribution on the node set of the underlying graph. It is immediate from Definition 1.1.4 that connectedness of the underlying graph is a necessary condition to construct irreducible random walks. The next theorem shows that an irreducible and aperiodic random walk *converges* to its stationary distribution. To specify what the convergence of a random walk is, let us set up a distance measure on the set of mass functions that is suitable for statistical purposes: The *total variation distance* of two mass functions π and π' on a finite set V is defined by $\|\pi - \pi'\|_{\text{TV}} := \frac{1}{2} \sum_{v \in V} |\pi(v) - \pi'(v)|$. An equivalent, and more statistically motivated, definition of the total variation distance is in [81, Proposition 4.2].

Theorem 1.1.5. *Let $G = (V, E)$ be a graph and $\mathcal{W} : V \times V \rightarrow [0, 1]$ be an irreducible and aperiodic random walk with stationary distribution π . Then there exists $C \in \mathbb{Q}$ and $\alpha \in (0, 1)$ such that $\max_{v \in V} \|\mathcal{W}^t(v, \cdot) - \pi\|_{\text{TV}} \leq C \cdot \alpha^t$ for all $t \in \mathbb{N}$.*

Proof. This is [81, Theorem 4.9]. □

Putting Theorem 1.1.5 in other words: Random walks can be used to draw samples from the distribution they converge to. To make this approximation applicable in practice, the number of steps that are needed to be sufficiently close to the stationary distribution needs to be known:

Definition 1.1.6. Let $G = (V, E)$ be a graph and \mathcal{W} an irreducible and aperiodic random walk with positive stationary distribution π . The *mixing time* of \mathcal{W} is the map $\mathcal{T}_{\mathcal{W}} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{N}$,

$$\mathcal{T}_{\mathcal{W}}(\epsilon) := \min\{t \in \mathbb{N} : \max_{v \in V} \|\mathcal{W}^s(v, \cdot) - \pi\|_{\text{TV}} \leq \epsilon \text{ for all } s \geq t\}$$

In statistics, it is common practice to denote the mixing time of a random walk \mathcal{W} by $\mathcal{T}_{\mathcal{W}}(0.25)$, which typically suffices to approximate the stationary distribution of \mathcal{W} sufficiently well in applications (see also [81, Section 4.5]). The general definition of the mixing time as stated in Definition 1.1.6 is cumbersome and intricate to work with in practice. To derive an equivalent convergence measurement that additionally takes the combinatorial structure of the random walk into account, we need a few definitions: A scalar $\lambda \in \mathbb{R}$ is an *eigenvalue* of a random walk \mathcal{W} if there exists a map $\pi : V \rightarrow \mathbb{R}$ not identically zero such that $\mathcal{W} \circ \pi = \lambda \cdot \pi$. In this case, π is an *eigenfunction* of \mathcal{W} . The absolute value of all eigenvalues of a random walk is smaller than one [81, Lemma 12.1] and hence the following definition is well:

Definition 1.1.7. Let G be a graph on n nodes and \mathcal{W} a random walk on G . The eigenvalues of \mathcal{W} are denoted by $\lambda_1(\mathcal{W}), \dots, \lambda_n(\mathcal{W})$ so that $1 = \lambda_1(\mathcal{W}) \geq \lambda_2(\mathcal{W}) \geq \dots \geq \lambda_n(\mathcal{W}) \geq -1$ is fulfilled. The *second largest eigenvalue modulus* of \mathcal{W} is $\lambda(\mathcal{W}) := \max\{\lambda_2(\mathcal{W}), -\lambda_n(\mathcal{W})\}$.

Remark 1.1.8. Let G be a graph and \mathcal{W} be a random walk on G that is not irreducible, then the eigenvalue $\lambda_1(\mathcal{W}) = 1$ has multiplicity greater than 2 and hence $\lambda(\mathcal{W}) = 1$. If \mathcal{W} is irreducible, then the Perron-Frobenius theorem [95, 55] implies that $\lambda_1(\mathcal{W})$ is a simple eigenvalue of \mathcal{W} .

Theorem 1.1.9. *Let $G = (V, E)$ be a graph and \mathcal{W} be a reversible and irreducible random walk on G with stationary distribution π , then for all $\epsilon > 0$,*

$$\log\left(\frac{1}{2\epsilon}\right) \cdot \left(\frac{1}{1 - \lambda(\mathcal{W})} - 1\right) \leq \mathcal{T}_{\mathcal{W}}(\epsilon) \leq \log\left(\frac{1}{\epsilon \cdot \min_{v \in V} \pi(v)}\right) \cdot \frac{1}{1 - \lambda(\mathcal{W})}.$$

Proof. The first inequality is [81, Remark 13.7] and the second is [81, Theorem 12.3]. Although our definition of mixing time $\mathcal{T}_{\mathcal{W}}(\epsilon)$ (Definition 1.1.6) is an upper bound to the mixing time $t_{\text{mix}}(\epsilon)$ as defined in [81, Section 4.5], it is immediate from their proof of [81, Theorem 12.3] that the upper bound is valid for both definitions. \square

The second largest eigenvalue modulus is a measurement of the convergence rate: Theorem 1.1.9 says that the closer $\lambda(\mathcal{W})$ is to 1, the slower is the convergence to the stationary distribution. However, it is precarious to assign the adjectives *fast* and *slow* to the mixing behaviour of a single random walk. Instead, the mixing time has to be in relation to the size of its state space asymptotically. It is common to express rapid mixing of random walks in terms of the following spectral property [3, 69, 25, 66, 115]:

Definition 1.1.10. For any $i \in \mathbb{N}$, let $G_i = (V_i, E_i)$ be a graph and let \mathcal{W}_i be a random walk on G_i . The sequence $(\mathcal{W}_i)_{i \in \mathbb{N}}$ is *rapidly mixing* if there is a polynomial $p \in \mathbb{Q}_{\geq 0}[t]$ such that

$$\lambda(\mathcal{W}_i) \leq 1 - \frac{1}{p(\log |V_i|)}$$

for all $i \in \mathbb{N}$. The sequence $(\mathcal{W}_i)_{i \in \mathbb{N}}$ is an *expander* if there exists $\delta > 0$ such that for all $i \in \mathbb{N}$,

$$\lambda(\mathcal{W}_i) \leq 1 - \delta.$$

Due to Theorem 1.1.9, being rapidly mixing or an expander is equivalently expressed in terms of the mixing time (Remark 1.1.12). The name *expander* relates to the fact that their *edge-expansion* (Definition 2.4.1) can strictly be bounded away from zero (Proposition 4.1.9).

Example 1.1.11. The simple walk on the complete graph K_n has eigenvalues $\{1, -\frac{1}{n-1}\}$ and hence $(\mathcal{S}_{K_n})_{n \in \mathbb{N}}$ is an expander. It is not hard to see that the spectrum of the simple walk on the circuit graph C_n is $\{\cos(\frac{2\pi k}{n}) : 0 \leq k \leq n-1\}$ and thus $\lambda(\mathcal{S}_{C_n}) \geq \cos(\frac{2\pi}{n})$. For $n \rightarrow \infty$, $\cos(\frac{2\pi}{n})$ tends faster to 1 than $1 - \frac{1}{p(\log n)}$ for all $p \in \mathbb{Q}[t]$ and hence $(\mathcal{S}_{C_n})_{n \in \mathbb{N}}$ is not rapidly mixing.

Remark 1.1.12. More generally, a sequence $(\mathcal{W}_i)_{i \in \mathbb{N}}$ of random walks on graphs $G_i = (V_i, E_i)$ with stationary distribution $\pi_i : V_i \rightarrow [0, 1]$ is *rapidly mixing* when there is a polynomial p such that $\mathcal{T}_{\mathcal{W}_i}(\epsilon) \leq p(\log(\epsilon^{-1}), \log(\min_{v \in V_i} \pi_i(v)^{-1}))$ (see for instance [3, Section 5]). Since we focus in this thesis exclusively on symmetric random walks, that is π_i is the uniform distribution, the spectral property of Definition 1.1.10 is a pragmatic, but equivalent, reformulation. Roughly speaking, a symmetric random walk mixes rapidly if only a logarithmic part of the graph nodes has to be traversed. Observe that in this framework, the number of computations needed for a single transition are not taken into account at all. Thus, when we assume that all graphs G_i are given implicit (Remark 1.1.2) and that \mathcal{W}_i needs at most $q(\epsilon, i)$ many computations to sample locally from its induced distributions on the graph neighborhoods, then \mathcal{W}_i generates elements from V_i uniformly with at most $q(\epsilon, i) \cdot p(\log \epsilon^{-1}, \log |V_i|)$ many computations.

Remark 1.1.13. Let $G = (V, E)$ be a graph, then the stationary distribution of \mathcal{S}_G is the map $u \mapsto \deg_G(u) \cdot (2|E|)^{-1}$. If G is *d-regular*, that is if all its nodes are incident to d edges, then $\mathcal{S}_G = \frac{1}{d}A^G$ is symmetric and hence its stationary distribution is the uniform distribution on V .

With Theorem 1.1.9, analyzing the mixing time of a random walk boils down to a linear algebra problem. However, the dimensions of the matrix we want to compute the eigenvalues from is not given a priori and the number of its entries grows quadratically in the number of

nodes of the graph. A powerful and popular tool to bound the mixing time of a random walk and that respects the underlying graph structure is the *conductance* of the random walk which has first been used in a statistical context in [106].

Definition 1.1.14. Let $G = (V, E)$ be a graph and $\mathcal{W} : V \times V \rightarrow [0, 1]$ be a random walk on G with stationary distribution $\pi : V \rightarrow [0, 1]$. The *conductance* of \mathcal{W} is

$$\Phi(\mathcal{W}) := \min \left\{ \frac{\sum_{i \in S} \sum_{j \in V \setminus S} \pi(i) \mathcal{W}(i, j)}{\pi(S)} : S \subseteq V, 0 < \pi(S) \leq \frac{1}{2} \right\}.$$

In [69], it was shown that $1 - 2\Phi(\mathcal{W}) \leq \lambda_2(\mathcal{W}) \leq 1 - \frac{1}{2}\Phi(\mathcal{W})^2$. In particular, if $\lambda_2(\mathcal{W}) = \lambda(\mathcal{W})$, then the conductance is another equivalent measurement of the convergence rate. However, when $\lambda(\mathcal{W}) > \lambda_2(\mathcal{W})$, then these bounds are not valid for $\lambda(\mathcal{W})$ and no statement can be made. A way to circumvent this issue is to manually increase the rejection probability of the walk:

Definition 1.1.15. Let $G = (V, E)$ be a graph and \mathcal{W} a random walk on G . The *lazy version* of \mathcal{W} is the random walk $\mathcal{L}(\mathcal{W}) = \frac{1}{2}(\mathcal{W} + I_n)$.

Remark 1.1.16. Since for any graph $G = (V, E)$, the eigenvalues of symmetric random walks $\mathcal{W} : V \times V \rightarrow [0, 1]$ on G are within $[-1, 1]$, all eigenvalues of their lazy versions are in $[0, 1]$.

1.2 Fiber graphs

A *polytope* is a set of the form $\text{conv}_{\mathbb{Q}}(\mathcal{F}) \subset \mathbb{Q}^d$ for a finite set $\mathcal{F} \subset \mathbb{Q}^d$. If $\mathcal{F} \subset \mathbb{Z}^d$, then \mathcal{P} is a *lattice polytope*. The key player of this thesis are graphs on *saturated* sets, i.e. sets which are the lattice points of a polytope. In particular, a finite set $\mathcal{F} \subset \mathbb{Z}^d$ is saturated if and only if $\text{conv}_{\mathbb{Q}}(\mathcal{F}) \cap \mathbb{Z}^d = \mathcal{F}$. To get started, let us recall how we construct graphs on lattice points [39]:

Definition 1.2.1. Let $\mathcal{F} \subseteq \mathbb{Z}^d$ and $\mathcal{M} \subseteq \mathbb{Z}^d$ be sets. Then $\mathcal{F}(\mathcal{M})$ is the graph on \mathcal{F} where two nodes $u, v \in \mathcal{F}$ are adjacent if $u - v \in \pm\mathcal{M}$. If \mathcal{F} is saturated, then $\mathcal{F}(\mathcal{M})$ is called *fiber graph*.

We have stated Definition 1.2.1 in its full generality, but we restrict here to the case where the involved sets \mathcal{F} and \mathcal{M} are finite, and leave the infinite case for further investigations. The choice of the name in Definition 1.2.1 needs some clarification. It is hard to track when and where the notion *fiber graph* has entered algebraic statistic. In the last two decades, it has become a collective term that stands for graphs on sets $(v + \mathcal{L}) \cap \mathbb{N}^d$, where $v \in \mathbb{N}^d$ and $\mathcal{L} \subseteq \mathbb{Z}^d$ is a *lattice*, that is a set of the form $\{\sum_{i=1}^r \lambda_i w_i : \lambda_1, \dots, \lambda_r \in \mathbb{Z}\}$ for fixed $w_1, \dots, w_r \in \mathbb{Z}^d$ (see for instance [86]). When $\mathcal{L} \cap \mathbb{N}^d = \{0\}$, then $(v + \mathcal{L}) \cap \mathbb{N}^d$ is a finite set for all $v \in \mathbb{Z}^d$. If \mathcal{L} is *saturated*, i.e. if there is an integer matrix $A \in \mathbb{Z}^{m \times d}$ for some $m \in \mathbb{N}$ such that $\mathcal{L} = \ker_{\mathbb{Z}}(A)$, then $(v + \mathcal{L}) \cap \mathbb{N}^d$ is a saturated set. In this case, the set $(v + \mathcal{L}) \cap \mathbb{N}^d$ is the *fiber* $\mathcal{F}_{A,b} := \{u \in \mathbb{N}^d : Au = b\}$ of A and the *right-hand side* $b := Av \in \mathbb{Z}^m$. To work with fibers, we postulate the following throughout:

Convention. Matrices $A \in \mathbb{Z}^{m \times d}$ that define fibers are assumed to have a non-trivial kernel and to satisfy $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$.

Clearly, only for right-hand sides in the *affine semigroup* of A , denoted by $\mathbb{N}A := \{Au : u \in \mathbb{N}^d\}$, the fiber is non-empty. Although the node sets of the graphs in Definition 1.2.1 do not need to be fibers, the next lemma justifies why the name is appropriate in spite of that.

Lemma 1.2.2. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be a saturated set and $\mathcal{M} \subset \mathbb{Z}^d$ be a finite set. There exists a matrix $A \in \mathbb{Z}^{k \times m}$, $b \in \mathbb{N}A$, and $\mathcal{M}' \subset \ker_{\mathbb{Z}}(A)$ such that $\mathcal{F}(\mathcal{M}) \cong \mathcal{F}_{A,b}(\mathcal{M}')$.*

Proof. Translation of \mathcal{F} does not change the graph structure of $\mathcal{F}(\mathcal{M})$ and thus we can assume that $\mathcal{F} \subset \mathbb{N}^d$. Let $\mathcal{P} \subset \mathbb{Q}_{\geq 0}^d$ be a polytope with $\mathcal{F} = \mathcal{P} \cap \mathbb{N}^d$. Since \mathcal{P} is a rational polytope, there exists $B \in \mathbb{Z}^{n \times d}$ with $\dim \ker_{\mathbb{Z}}(B) = 0$, $n \geq d$, and $b \in \mathbb{Z}^n$ such that $\mathcal{P} = \{x \in \mathbb{Q}_{\geq 0}^d : Bx \leq b\}$. Consider the injective and affine map

$$\phi : \mathbb{Q}^d \rightarrow \mathbb{Q}^{d+n}, x \mapsto \begin{pmatrix} x \\ b - Bx \end{pmatrix}$$

and let $\mathcal{P}' := \phi(\mathcal{P}) = \{(x, y)^T \in \mathbb{Q}_{\geq 0}^{d+n} : Bx + y = b\}$. Since the kernel of B is trivial, ϕ induces a bijection from $\mathcal{F} = \mathcal{P} \cap \mathbb{N}^d$ to $\mathcal{F}' := \mathcal{P}' \cap \mathbb{N}^{d+n}$. For $\mathcal{M}' := \{(m, -Bm)^T : m \in \mathcal{M}\}$, we have $\mathcal{F}'(\mathcal{M}') \cong \mathcal{F}(\mathcal{M})$. With $A = (B, I_n) \in \mathbb{Z}^{n \times (d+n)}$, we get $\mathcal{F}_{A,b} = \mathcal{P}'$ and $\mathcal{M}' \subset \ker_{\mathbb{Z}}(A)$. \square




Figure 1.1: A fiber graph in \mathbb{Q}^2 .

When working with fiber graphs, the overall goal is to make them connected, for instance, in order to run irreducible random walks on them. The authors of [39] coined the following concept:

Definition 1.2.3. Let $\mathcal{M} \subset \mathbb{Z}^d$ be a finite set. Then \mathcal{M} is a *Markov basis* for a finite set $\mathcal{F} \subset \mathbb{Z}^d$ if $\mathcal{F}(\mathcal{M})$ is connected and \mathcal{M} is a *Markov basis* for a collection \mathfrak{F} of finite subsets of \mathbb{Z}^d if \mathcal{M} is a Markov basis for all $\mathcal{F} \in \mathfrak{F}$. For a norm $\|\cdot\|$ on \mathbb{R}^d , $\|\mathcal{M}\| := \max_{m \in \mathcal{M}} \|m\|$.

This definition of a Markov basis is a slight extension of that given in [43, Definition 1.1.12]. Their definition can easily be recovered by plugging in the collection of fibers for an integer matrix $A \in \mathbb{Z}^{m \times d}$, that is the set $\mathfrak{F}_A := \{\mathcal{F}_{A,b} : b \in \mathbb{N}A\}$. For simplicity, we call a Markov basis for \mathfrak{F}_A just a Markov basis for A . In the literature, a Markov basis for $\mathcal{F}_{A,b}$ is often called a *Markov subbasis* [26, 99]. We discuss in Section 2.1 how Markov bases for this type of collections are computed. Observe that, in general, finite Markov bases do not have to exist:

Example 1.2.4. Let $\mathcal{F}_i := \{0, i\} \subset \mathbb{Z}$ for $i \in \mathbb{N}$, then clearly the collection $(\mathcal{F}_i)_{i \in \mathbb{N}}$ cannot have a finite Markov basis. But trivially, all collections of saturated sets in \mathbb{Z} have a finite Markov basis, namely $\{1\}$. This fails to be true in \mathbb{Z}^2 , where $\mathcal{F}_i = \{(0, 0), (1, i)\} \subset \mathbb{Z}^2$ is saturated for every $i \in \mathbb{N}$ and every Markov basis of $(\mathcal{F}_i)_{i \in \mathbb{N}}$ is a superset of $\{(1, i) : i \in \mathbb{N}\}$.

Convention. Let $(w_1, \dots, w_r) \in \mathcal{F}^r$ be a path in $\mathcal{F}(\mathcal{M})$ from $u = w_1$ to $v = w_r$, then the difference of subsequent elements $m_i := w_{i+1} - w_i$ is a move from $\pm\mathcal{M}$. For brevity, we frequently just call the sum $v = u + \sum_{i=1}^r m_i$ a path from u to v , where we implicitly assume that the partial sums satisfy $u + \sum_{i=1}^l m_i \in \mathcal{F}$ for all $l \in [r]$.

In the joint work [108] with Caprice Stanley, the *compression* of a graph from Definition 1.2.1 was introduced, which allows edges of arbitrarily length in the directions of the moves used:

Definition 1.2.5. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets. The *compression* of the graph $\mathcal{F}(\mathcal{M})$ is the graph $\mathcal{F}^c(\mathcal{M}) := \mathcal{F}(\mathbb{Z} \cdot \mathcal{M})$.

Clearly, the compressed version of a fiber graph is connected if and only if the fiber graph itself is connected. In Chapter 5, random walks on this type of graphs are studied and it is shown that they converge rapidly in fixed dimension. The diameter of compressed fiber graphs is discussed in Section 2.2, where the following set of moves becomes important:

Definition 1.2.6. Two vectors $u, v \in \mathbb{Z}^d$ are *sign-compatible* if $u_i \cdot v_i \geq 0$ for all $i \in [d]$. We write $u \sqsubseteq v$ if u and v are sign-compatible and if $|u_i| \leq |v_i|$ for all $i \in [d]$. The *Graver basis* $\mathcal{G}_{\mathcal{L}}$ of a lattice $\mathcal{L} \subseteq \mathbb{Z}^d$ is the set of all \sqsubseteq -minimal elements in \mathcal{L} . The *Graver basis* of a matrix $A \in \mathbb{Z}^{m \times d}$ is the Graver basis of the saturated lattice $\ker_{\mathbb{Z}}(A)$.

The Graver basis is always finite and there are many good reasons why they are an especially nice set of moves, some of these reasons are shown in Proposition 2.1.3, Lemma 2.2.14, and Proposition 2.1.8. We refer to [31, Chapter 2] for more nice and not so nice Graver facts.




Figure 1.2: Compressing graphs.

We now merge fiber graphs with random walks, shortly *fiber walks*. Given a finite set $\mathcal{F} \subset \mathbb{Z}^d$ together with a Markov basis $\mathcal{M} \subset \mathbb{Z}^d$, one simply can use the simple walk on $\mathcal{F}(\mathcal{M})$ to explore \mathcal{F} . However, the simple walk may not be aperiodic, for instance when the underlying graph is bipartite. In this case, the simple walk commutes back and forth between its color classes and does not converge to its stationary distribution. This also happens for fiber graphs:

Proposition 1.2.7. Let $\mathcal{F} \subset \mathbb{Z}^d$ be a saturated set and let $\mathcal{M} \subset \mathbb{Z}^d$ be a Markov basis for \mathcal{F} with $\dim(\mathcal{F}) = |\mathcal{M}|$. Then $\mathcal{F}(\mathcal{M})$ is bipartite.

Proof. Let $k := \dim(\mathcal{F}) \leq d$. Since \mathcal{M} is a Markov basis of \mathcal{F} , $\dim(\mathcal{F}) = \dim(\text{span}_{\mathbb{Q}}\{\mathcal{M}\})$ and thus we can write $\mathcal{M} = \{m_1, \dots, m_k\}$. The assumption on the dimension says that \mathcal{M} is linearly independent. Let $v \in \mathcal{F}$ and let $v + \sum_{i=1}^k \lambda_i m_i + \sum_{i=1}^k -\mu_i m_i = v$ be a circuit in $\mathcal{F}(\mathcal{M})$ of length $r = \sum_{i=1}^k (\lambda_i + \mu_i)$. Linear independence gives $\lambda_i = \mu_i$ for all $i \in [k]$ and thus r is even. \square

Bipartiteness of fiber graphs cannot always read off as easy as in Proposition 1.2.7 and can be quite hidden subtly. A feasible way to construct a fiber walk that converges without any more assumptions on \mathcal{F} and \mathcal{M} – beside connectedness of $\mathcal{F}(\mathcal{M})$ – is to manually add rejections:

Definition 1.2.8. Let \mathcal{F} and $\mathcal{M} \subset \mathbb{Z}^d$ be finite set. The *simple fiber walk* is the simple walk on the graph obtained from $\mathcal{F}(\mathcal{M})$ by adding $|\{m \in \pm\mathcal{M} : v + m \notin \mathcal{F}\}|$ many loops to all $v \in \mathcal{F}$.

We do not require the set \mathcal{F} in Definition 1.2.8 to be saturated, although we restrict to this case most of the time. With the additional halting states, we obtain:

Proposition 1.2.9. Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite and non-empty set and $\mathcal{M} \subset \mathbb{Z}^d$ a Markov basis for \mathcal{F} . The simple fiber walk on $\mathcal{F}(\mathcal{M})$ is irreducible, aperiodic, symmetric, reversible, and its stationary distribution is the uniform distribution on \mathcal{F} .

Proof. The random walk is irreducible and symmetric since $\mathcal{F}(\mathcal{M})$ is connected and $|\pm\mathcal{M}|$ -regular (Remark 1.1.13). Thus, it suffices to show that $\mathcal{F}(\mathcal{M})$ has one aperiodic state to show that all states are aperiodic due to [81, Lemma 1.6]. Choose $v \in \mathcal{F}$ and $m \in \mathcal{M}$ arbitrarily. Since \mathcal{F} is finite, let $\lambda \in \mathbb{N}$ be the largest natural number such that $v + \lambda m \in \mathcal{F}$. Then m cannot be applied on $v + \lambda m$ and thus there is a positive probability that the simple fiber walk stays at $v + \lambda m$. Symmetry of the simple fiber walk implies immediately its reversibility and that the uniform distribution is its unique stationary distribution. \square

For a finite set of moves \mathcal{M} , the neighborhood of a given node $v \in \mathcal{F}$ in $\mathcal{F}(\mathcal{M})$ can be enumerated by going through all moves $m \in \pm\mathcal{M}$ and by checking whether $v + m \in \mathcal{F}$ holds. We emphasize that, by definition, Markov bases are finite sets, which is important from an algorithmic perspective. The efficiency of the decision on membership $v + m \in \mathcal{F}$ depends on how the set \mathcal{F} is given. Typically, \mathcal{F} is given implicitly in \mathcal{H} -representation, as in the case $\mathcal{F} = \mathcal{F}_{A,b}$, and here we can decide over membership efficiently. As we have seen in Remark 1.1.2, the capability to enumerate the neighborhood of any node suffices to perform random walks on this graph.

Remark 1.2.10. For finite sets $\mathcal{F}, \mathcal{M} \subset \mathbb{Z}^d$, the simple fiber walk on $\mathcal{F}(\mathcal{M})$ is implemented as follows: At a given node $v \in \mathcal{F}$, select uniformly an element $m \in \pm\mathcal{M}$ and stay at v if $v + m \notin \mathcal{F}$, or walk to $v + m \in \mathcal{F}$ otherwise (which may also be v if $m = 0$).




Figure 1.3: The simple fiber walk with unit vectors in a convex polygon after 2500 steps.

Remark 1.2.11. The *Metropolis-Hastings*-methodology (see Definition 1.4.2) allows to modify the simple fiber walk so that it converges to any given probability distribution on \mathcal{F} .

Remark 1.2.12. A fiber walk that is not restricted to some finite set in \mathbb{Z}^d , but uses a finite set of moves, is often called a *lattice walk*. The asymptotic counting of the number of lattice walks in \mathbb{Z}^d , or \mathbb{N}^d , under some fixed parameters such as length, start node, or end node is a topic of combinatorics [78]. The correspondence between *Catalan strings* and *Dyck paths* is just one of the many paradigms to describe combinatorial structures as lattice walks. Recently, the proof of *Gessel's walk conjecture* in [75] received a lot of attention. It states that there are $\mathcal{O}(16^n)$ *Gessel walks* of length $2n$, that are lattice walks in \mathbb{N}^2 whose start and end nodes are the origin and that use the moves $\{\pm(1,1), \pm(1,0)\}$. Another aspect is the sampling of lattice walks themselves. For instance, the uniform sampling of certain 2-dimensional lattice walks is discussed in [85].

1.3 Running examples

This section introduces matrices that appear frequently in this thesis.

Example 1.3.1. For $k \in \mathbb{N}$, let $\mathbf{1}_k$ be the k -dimensional vector with all entries equal to 1 and let

$$H_k := \begin{bmatrix} I_k & I_k & \mathbf{0} & \mathbf{0} & -\mathbf{1}_k & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & I_k & I_k & \mathbf{0} & -\mathbf{1}_k \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & 1 & 1 \end{bmatrix} \in \mathbb{Z}^{(2k+1) \times (4k+2)}. \quad (1.3.1)$$

The linear independent set of moves $\mathcal{R}_k \subset \mathbb{Z}^{4k+2}$ that consists of

$$(0, \dots, 0, 1, \dots, 1, 0, \dots, 0, -1, \dots, -1, 1, -1)^T$$

and $e_i - e_{k+i}$ for $i \in \{1, \dots, k, 2k+1, \dots, 3k\}$ is a Markov basis for H_k [62, Theorem 3]. An explicit description of the Graver basis of H_k is in [62, Theorem 2].

Example 1.3.2. The node-edge incidence matrix of the complete bipartite graph $K_{n,m}$ is denoted $A_{n,m} \in \{0, 1\}^{(m+n) \times m \cdot n}$. For instance,

$$A_{2,3} = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}.$$

In a statistical context, $A_{n,m}$ is the constraint matrix of the $n \times m$ *independence model*, which is briefly discussed in Example 1.4.7. We refer to [43, Chapter 1] and [39] for more information about the statistics behind. Elements in the kernel of $A_{n,m}$ are represented as $n \times m$ contingency tables whose row and column sums are zero. Among these, the easiest ones are the *basic moves* $\mathcal{M}_{n,m}$ which are all elements in the orbit of

$$\begin{bmatrix} 1 & -1 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & & \\ \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & & \cdots & 0 \end{bmatrix} \in \mathbb{Z}^{n \times m}$$

under the group action of $S_n \times S_m$ on the rows and columns. In particular, $|\mathcal{M}_{n,m}| = 2 \cdot \binom{n}{2} \cdot \binom{m}{2}$. It is well-known that the basic moves form a Markov basis for $A_{n,m}$ that is minimal in the sense that a removal of any element takes away the Markov basis property [43, Proposition 1.2.2].

Example 1.3.3. For any $d \in \mathbb{N}$, let $A_d = (1, \dots, 1) \in \mathbb{Z}^{1 \times d}$. For any $b \in \mathbb{N}$, the elements of $\mathcal{F}_{A_d, b}$ corresponds to the monomials in $\mathbb{k}[x_1, \dots, x_d]$ of degree b and hence $|\mathcal{F}_{A_d, b}| = \binom{b+d-1}{d-1}$. It is easy to see that the set $\mathcal{M}_d := \{e_1 - e_k : 2 \leq k \leq d\}$ is a Markov basis for A_d .

1.4 Fiber walks in statistics

Let Ω be a finite set and $\pi : \Omega \rightarrow [0, 1]$ be a probability mass function. Many problems in statistics are stated equivalently as the problem of estimating the expected value of a function

$f : \Omega \rightarrow \mathbb{R}^m$ under π , i.e. $\mathbb{E}_\pi(f) = \sum_{\omega \in \Omega} f(\omega)\pi(\omega)$. Given a sequence $(\omega_i)_{i \in \mathbb{N}}$ of elements in Ω that are drawn independently from π , the law of large numbers ensures that almost surely,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \left(\sum_{i=0}^{t-1} f(\omega_i) \right) = \mathbb{E}_\pi(f). \quad (1.4.1)$$

Typically, it is computationally expensive to sample from π directly. The *Markov chain Monte Carlo* approach works around this issue by constructing a connected graph $G = (\Omega, E)$ and an irreducible and aperiodic random walk $\mathcal{W} : \Omega \times \Omega \rightarrow [0, 1]$ on G that has stationary distribution π . The Ergodic theorem [81, Theorem 4.16] ensures that for any starting vector $\omega_0 \in \Omega$ and any sequence $(\omega_i)_{i \in \mathbb{N}}$ that is obtained by an execution of the random walk \mathcal{W} starting at ω_0 , equation (1.4.1) holds almost surely. Again, the second largest eigenvalue modulus of \mathcal{W} is an indication of how fast the sum converges to $\mathbb{E}_\pi(f)$. Recall that the *variance* of the function f is

$$\text{Var}_\pi(f) := \frac{1}{2} \sum_{u \in \Omega} \sum_{v \in \Omega} (f(u) - f(v))^2 \pi(u)\pi(v).$$

Theorem 1.4.1. *Let $\mathcal{W} : \Omega \times \Omega \rightarrow [0, 1]$ be a reversible random walk with stationary distribution $\pi : \Omega \rightarrow [0, 1]$ and let $\omega_0 \in \Omega$. Suppose that $(\omega_i)_{i \in \mathbb{N}}$ is obtained by an execution of \mathcal{W} starting at ω_0 and let $\epsilon > 0$ and $\delta > 0$. If $t_0 \geq \mathcal{T}_\mathcal{W}(\frac{\epsilon}{2})$ and $t \geq \frac{4 \cdot \text{Var}_\pi(f)}{\delta^2 \cdot \epsilon} \cdot \frac{1}{1 - \lambda(\mathcal{W})}$, then the probability that*

$$\left| \frac{1}{t} \left(\sum_{i=0}^{t-1} f(\omega_{t_0+i}) \right) - \mathbb{E}_\pi(f) \right| \geq \delta$$

is at most ϵ .

Proof. This is [81, Theorem 12.19]. □

Theorem 1.4.1 suggests to omit the first t_0 samples obtained by the random walk \mathcal{W} . This is often called a *burn-in* in the literature. Although the set Ω is finite, an enumeration of its elements is frequently unfeasible in practice. Instead, we typically have to deal with an implicit description of its elements, as $\mathcal{F}_{A,b}$ for given A and b (see Example 1.4.6). Even if we are able to construct a graph and a random walk on Ω , the respective stationary distribution is certainly not π . The *Metropolis-Hastings* methodology helps to modify a given random walk so that it converges to any positive mass function on the node set Ω :

Definition 1.4.2. Let $G = (\Omega, E)$ be a graph, \mathcal{W} a random walk on G , and $\pi : \Omega \rightarrow [0, 1]$ a positive mass function. The *Metropolis-Hastings walk* $\mathcal{M}_{\mathcal{W},\pi}$ is the random walk on G defined by

$$\mathcal{M}_{\mathcal{W},\pi}(u, v) := \begin{cases} \mathcal{W}(u, v) \cdot \min \left\{ 1, \frac{\pi(v)\mathcal{W}(v,u)}{\pi(u)\mathcal{W}(u,v)} \right\}, & \text{if } u \neq v \\ 1 - \sum_{w \in \Omega \setminus \{u\}} \mathcal{W}(u, w) \cdot \min \left\{ 1, \frac{\pi(w)\mathcal{W}(w,u)}{\pi(u)\mathcal{W}(u,w)} \right\}, & \text{if } u = v \end{cases}$$

The Metropolis-Hastings walk was introduced in [60], which in turn arises as a generalization of the results in the classic paper [89]. Many properties of the original random walk are invariant under this deformation and an execution of the Metropolis-Hastings walk is possible despite the fact that π may be expensive to evaluate (Remark 1.4.4).

Proposition 1.4.3. *Let $G = (\Omega, E)$ be a graph, let \mathcal{W} be an irreducible and aperiodic random walk on G , and let $\pi : \Omega \rightarrow [0, 1]$ be a positive probability mass function, then $\mathcal{M}_{\mathcal{W},\pi}$ is irreducible, aperiodic, reversible with respect to π , and has stationary distribution π .*

Proof. This is [37, Lemma 1.1]. □

Remark 1.4.4. Given a random walk \mathcal{W} , the Metropolis-Hastings walk $\mathcal{M}_{\mathcal{W},\pi}$ is performed as follows: Suppose the random walk is at node $u \in \Omega$, then we sample $v \in \Omega$ according to $\mathcal{W}(u, \cdot)$ and compute $p := \min \left\{ 1, \frac{\pi(v)\mathcal{W}(v,u)}{\pi(u)\mathcal{W}(u,v)} \right\}$. In a second step, we walk to v with probability p and we stay at u with probability $1 - p$. Basically, an instance of the Metropolis-Hastings walk is just an instance of \mathcal{W} that is enriched with an additional rejection probability which depends on the current state and the proposal drawn by \mathcal{W} . In many cases, where a direct evaluation of π is not possible, π is known up to a constant factor, that is

$$\pi(u) = \frac{g(u)}{\sum_{\omega \in \Omega} g(\omega)}$$

for some function $g : \Omega \rightarrow \mathbb{R}_{\geq 0}$ that is easy to evaluate (as in Example 1.4.5 and Example 1.4.6). In this case, the rejection step in a Metropolis-Hastings walk needs only to evaluate the ration $\frac{\pi(v)}{\pi(u)} = \frac{g(v)}{g(u)}$, where the normalizing constant cancels.

Algorithm 1 puts all pieces together and shows the common workflow of a *Markov chain Monte Carlo* approach to approximate $\mathbb{E}_{\pi}(f)$. We finish this section with explicit examples showing the importance of fiber walks in statistics.

Algorithm 1 Markov chain Monte Carlo

Input: Ω (a finite set), $\pi : \Omega \rightarrow [0, 1]$ (mass function), $f : \Omega \rightarrow \mathbb{R}^m$, $t \in \mathbb{N}$

- 1: **procedure** MCMC:
 - 2: Construct a connected (implicit) graph $G = (\Omega, E)$
 - 3: Construct an irreducible and aperiodic random walk \mathcal{W} on G
 - 4: Perform $\mathcal{M}_{\mathcal{W},\pi}$ to obtain samples $\omega_1, \dots, \omega_t \in \Omega$
 - 5: **RETURN** $\frac{1}{t} \sum_{i=1}^t f(\omega_i)$
-

Example 1.4.5 (Ratio counting). Let $\mathcal{F} \subset \mathbb{Z}^d$ be a saturated set, $\mathcal{F}' \subseteq \mathcal{F}$ be a subset, and $\mathcal{M} \subset \mathbb{Z}^d$ be a Markov basis for \mathcal{F} . With Algorithm 1, we can approximate the ratio $\frac{|\mathcal{F}'|}{|\mathcal{F}|}$. Let $\pi : \mathcal{F} \rightarrow [0, 1]$ be the uniform distribution and $f := \mathbf{1}_{\mathcal{F}'} : \mathbb{Q}^d \rightarrow \{0, 1\}$ be the indicator function of \mathcal{F}' , then $|\mathcal{F}| \cdot \mathbb{E}_{\pi}(f) = |\mathcal{F}'|$. The simple fiber walk on $\mathcal{F}(\mathcal{M})$ is symmetric and has π as stationary distribution, so in this case, we even do not have to use the modified Metropolis-Hastings walk.

Example 1.4.6 (Goodness-of-fit). Let $\Omega = [d]$ be a finite set and $\pi \in [0, 1]^d$ a probability mass function on Ω . In many practical problems, the distribution π is expensive to evaluate or unknown, but sampling from π is easy. Given n independent samples $x_1, \dots, x_n \in \Omega$ from π , a typical question in statistical inference is whether the true distribution π belongs to a given statistical model. For instance, the *log-linear model* on Ω defined by a matrix $A \in \mathbb{Q}^{m \times d}$ with $(1, \dots, 1) \in \text{rowspan}(A)$ is the set

$$\mathcal{P}_A := \left\{ \theta \in (0, 1)^d : \theta_1 + \dots + \theta_d = 1 \wedge (\log \theta_1, \dots, \log \theta_d) \in \text{rowspan}(A) \right\}.$$

A broadly accepted way to evaluate the *goodness-of-fit* of log-linear models is by doing an exact conditional test (see also [39, 43] and references therein), which we explain now. For $i \in [d]$, let $u_i^{\text{obs}} := |\{j \in [n] : x_j = i\}|$ be the frequency count of i within the samples. The frequency counts

of n independent samples from Ω is multinomial distributed over $\{v \in \mathbb{N}^d : \|v\|_1 = n\}$ and when $\pi \in \mathcal{P}_A$, then the probability to observe $v \in \mathbb{N}^d$ with $\|v\|_1 = n$ is

$$\frac{n!}{v_1! \cdots v_d!} \pi_1^{v_1} \cdots \pi_d^{v_d} = \frac{n!}{v_1! \cdots v_d!} \exp(\alpha^T Av)$$

with $(\log \pi_1, \dots, \log \pi_d) = \alpha^T A$. Assume that the statistical model \mathcal{P}_A comes with a *test statistics*, that is a map $X : \mathbb{N}^d \rightarrow \mathbb{R}$ which measures the extremeness of an observed frequency count within the model \mathcal{P}_A . Then, the elements of $\mathcal{X} := \{v \in \mathbb{N}^d : \|v\|_1 = n\} \cap \{v \in \mathbb{N}^d : X(v) \geq X(u^{\text{obs}})\}$ are the frequencies that are more extreme than u^{obs} . Now, the probability of observing a frequency count under π that is more extreme than u^{obs} is

$$\sum_{v \in \mathcal{X}} \frac{n!}{v_1! \cdots v_d!} \exp(\alpha^T Av). \quad (1.4.2)$$

Thus, if this value is small, it is very unlikely to observe u^{obs} and the conclusion we draw is that the hypothesis $\pi \in \mathcal{P}_A$ is false. However, the value in (1.4.2) depends on α , which in turn relies on the unknown distribution π and hence this quantity cannot be computed. Instead, we exploit the fact that A is a *sufficient statistics* for the multinomial distribution on $\{v \in \mathbb{N}^d : \|v\|_1 = n\}$ induced by \mathcal{P}_A . That is, when conditioning on the subset $\mathcal{F}_{A,b}$ where $b := Au^{\text{obs}}$, then

$$\frac{\sum_{v \in \mathcal{X} \cap \mathcal{F}_{A,b}} \frac{n!}{v_1! \cdots v_d!} \exp(\alpha^T Av)}{\sum_{w \in \mathcal{F}_{A,b}} \frac{n!}{w_1! \cdots w_d!} \exp(\alpha^T Aw)} = \frac{\sum_{v \in \mathcal{X} \cap \mathcal{F}_{A,b}} \frac{1}{v_1! \cdots v_d!}}{\sum_{w \in \mathcal{F}_{A,b}} \frac{1}{w_1! \cdots w_d!}} \quad (1.4.3)$$

and the unknown parameter α cancels. The quantity in (1.4.3) is the *conditional p-value* of the test and in practice, the hypothesis $\pi \in \mathcal{P}_A$ is rejected when the conditional p -value is below a threshold of 0.05. The sum in equation (1.4.3) runs over $\mathcal{F}_{A,b}$ and is thus impossible to evaluate in practice. However, an approximation of the p -value with Algorithm 1 is applicable. First, define a probability mass function $\tilde{\pi} : \mathcal{F}_{A,b} \rightarrow [0, 1]$ by

$$\tilde{\pi}(v) = \frac{\frac{1}{v_1! \cdots v_d!}}{\sum_{w \in \mathcal{F}_{A,b}} \frac{1}{w_1! \cdots w_d!}}$$

and let $f : \mathcal{F}_{A,b} \rightarrow \{0, 1\}$ be the indicator function of $\mathcal{X} \cap \mathcal{F}_{A,b}$, then the conditional p -value is $\mathbb{E}_{\tilde{\pi}}(f)$. After computing a Markov basis $\mathcal{M} \subset \mathbb{Z}^d$ for $\mathcal{F}_{A,b}$ (for example with Proposition 2.1.1), a connected graph on $\mathcal{F}_{A,b}$ is obtained and we can use the Metropolis-Hastings walk as a modification of the simple fiber walk to sample from $\mathcal{F}_{A,b}$ according to $\tilde{\pi}$. In practice, X is often

$$X(v) := \sum_{i=1}^d \frac{(v_i - n \cdot \tilde{\theta}_i)^2}{n \cdot \tilde{\theta}_i}$$

where $\tilde{\theta} \in \mathcal{P}_A$ is a *maximum likelihood estimator* for the samples $x_1, \dots, x_n \in \Omega$, that is an element of (the possibly empty set) $\arg \max_{\theta \in \mathcal{P}_A} \sum_{i=1}^n \log \theta_{x_i}$. Since the function $\theta \mapsto \sum_{i=1}^n \log \theta_{x_i}$ is convex on the open set \mathcal{P}_A , the maximum likelihood estimator $\tilde{\theta}$ is unique when it exists.

Example 1.4.7. Let $\Omega = [n] \times [m]$ and let $\pi = (\pi_{1,1}, \dots, \pi_{1,m}, \pi_{2,1}, \dots, \pi_{n,m})$ be a probability mass function on Ω . Let $\pi_i^{(1)} := \sum_{j=1}^m \pi_{i,j}$ and $\pi_j^{(2)} := \sum_{i=1}^n \pi_{i,j}$ be the marginal probabilities. We want to test whether $\pi_{i,j} = \pi_i^{(1)} \cdot \pi_j^{(2)}$ holds for all $(i, j) \in [n] \times [m]$, that is

whether $\pi^{(1)}$ is stochastically independent of $\pi^{(2)}$. It is not hard to show that this is true if and only if π is an element of the log-linear model defined by $A_{n,m}$ from Example 1.3.2. As a special log-linear model, testing on stochastic independence can be done as in Example 1.4.6. The frequency count of n observations $(v_1, w_1), \dots, (v_n, w_n) \in \Omega$ is typically represented as an $n \times m$ contingency table $u^{\text{obs}} \in \mathbb{N}^{n \times m}$. The entries of the image of u^{obs} under $A_{n,m}$ are the rows sums r_1, \dots, r_n and the column sums c_1, \dots, c_m of the table u . A maximum likelihood estimator is then $\tilde{\theta}_{i,j} := \frac{r_i \cdot c_j}{n \cdot n}$, provided that $r_i > 0$ and $c_j > 0$ [43, Example 2.1.2].

Remark 1.4.8. If $\Omega = [n_1] \times \dots \times [n_m]$, then more complex relations among the m features can be modeled with *hierarchical models* [43, Chapter 1.2] which also belong to the class of log-linear models. Here, the relations among the items are represented by a simplicial complex Γ on $[m]$. The independence model is a special case, namely for the simplicial complex $\Gamma = \{\emptyset, \{1\}, \dots, \{m\}\}$.

Remark 1.4.9. *Discrete reaction networks* are used to model the dynamic behaviour of chemical reactions [94, 19, 4]. The discrete nature of atoms and atomic reactions makes the methods of discrete mathematics applicable. Typically, there are d chemical species and with the number of atoms of species i equal to u_i , every state of the network is represented as an element $u \in \mathbb{N}^d$. The dynamics of the atom numbers proceeds only along elementary reactions $m \in \mathbb{Z}^d$ which is applied to a state u only if $u + m \in \mathbb{N}^d$. In many situations, there are only finitely many elementary reactions $\mathcal{M} \subset \mathbb{Z}^d$ available to the system and the probability for every reaction to be applied depends on the entries of u only. There may also be *irreversible* reactions, that are elements $m \in \mathcal{M}$ such that $-m \notin \mathcal{M}$. In particular, reachability in these networks is not symmetric and hence is more subtle than it may first seem. Additionally, $\mathcal{M} \cap \mathbb{N}^d \setminus \{0\}$ is not required to be empty and the lattice spanned by \mathcal{M} is not saturated a priori. In the language of Definition 1.2.1, this setup deals with directed fiber graphs on infinite node sets. Typical questions have a purely combinatorial character, as the reachability of a state starting at a given state, or are analytic, as whether the random walk is positive recurrent or transient.

2 Graph properties of fiber graphs

In this section, the toolbox for our further investigation is developed. After surveying known results on the connectedness and connectivity of fiber graphs (Section 2.1), we turn our attention to their diameter (Section 2.2). For a given collection \mathfrak{F} , we detect properties on the set of moves \mathcal{M} that allow to bound the diameter of $\mathcal{F}(\mathcal{M})$ linearly in $\max\{\|u - v\| : u, v \in \mathcal{F}\}$ from below and above for all $\mathcal{F} \in \mathfrak{F}$. As a consequence, the diameter of $\mathcal{F}_{A,i,b}(\mathcal{M})$ for fixed $A \in \mathbb{Z}^{m \times d}$, $\mathcal{M} \subset \mathbb{Z}^d$, and non-trivial $b \in \mathbb{N}A$ grows at least and at most linearly as i varies. On the other hand, the diameter of compressed fiber graphs cannot be arbitrarily large. We prove that for any matrix $A \in \mathbb{Z}^{m \times d}$ and any Markov basis $\mathcal{M} \subset \mathbb{Z}^d$, the diameter of the compressed fiber graphs on $\mathcal{F}_{A,b}$ is bounded from above by a constant as b varies (Theorem 2.2.17). Finally, we show that under certain assumptions on a sequence $(b_i)_{i \in \mathbb{N}}$ in $\mathbb{N}A$, the edge-expansion of the corresponding fiber graphs declines essentially as $\mathcal{O}(\frac{1}{i})_{i \in \mathbb{N}}$.

2.1 Connectedness and connectivity

Diaconis and Sturmfels have shown in [39] that the combinatorics necessary to decide the connectedness of fiber graphs is encoded in a binomial ideal attached to them. To set up their machinery, let $\mathbb{k}[x_1, \dots, x_d]$ be the polynomial ring over a field \mathbb{k} and $u \in \mathbb{Z}^d$ be an integer vector. The unique decomposition $u = u^+ - u^-$ allows to attach the binomial

$$x_1^{u_1^+} \cdots x_d^{u_d^+} - x_1^{u_1^-} \cdots x_d^{u_d^-} \in \mathbb{k}[x_1, \dots, x_d]$$

to u . Similarly, we attach to any set $\mathcal{L} \subset \mathbb{Z}^d$ the ideal $\mathcal{I}_{\mathcal{L}} \subset \mathbb{k}[x_1, \dots, x_d]$ generated by all binomials coming from vectors in \mathcal{L} . If \mathcal{L} is a lattice, then $\mathcal{I}_{\mathcal{L}}$ is the *lattice ideal* of \mathcal{L} (see [31, Chapter 11]). The following is due to Diaconis and Sturmfels:

Proposition 2.1.1. *Let $\mathcal{L} \subset \mathbb{Z}^d$ be a lattice with $\mathcal{L} \cap \mathbb{N}^d = \{0\}$ and $\mathcal{M} \subset \mathcal{L}$ a finite set. Then \mathcal{M} is a Markov basis for the collection $\{(\mathcal{L} + u) \cap \mathbb{N}^d : u \in \mathbb{N}^d\}$ if and only if $\mathcal{I}_{\mathcal{M}} = \mathcal{I}_{\mathcal{L}}$.*

Proof. This is a straightforward extension of [39, Theorem 3.1] to non-saturated lattices (see also [31, Lemma 11.3.3]). \square

A special lattice ideal is the *toric ideal* of a matrix $A \in \mathbb{Z}^{m \times d}$ which is the ideal $\mathcal{I}_A := \mathcal{I}_{\ker_{\mathbb{Z}}(A)}$. Proposition 2.1.1 says that a set $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ is a Markov basis for A if and only if $\mathcal{I}_{\mathcal{M}} = \mathcal{I}_A$. Hilbert's basis theorem [76, Corollary 2.13] on the other hand says that every ideal in $\mathbb{k}[x_1, \dots, x_d]$ is finitely generated and hence there always exists a finite Markov basis for A . The “only if” direction of Proposition 2.1.1 enables access for a practical computation of Markov bases by computing a finite binomial generating set of \mathcal{I}_A . In commutative algebra, a desirable generating set for a given ideal is a *Gröbner basis* [109, Chapter 1]. There exists a huge literature on Gröbner bases, including many excellent textbooks focusing on the fascinating theory around [76, 31, 109]. Many algorithms that compute a Markov basis have one or more Gröbner basis computations under the hood, like the *saturation algorithm* [65], the *elimination algorithm for toric ideals* [16],

or the *Project-and-Lift algorithm* [61, Section 3]. The latter algorithm is implemented in the software `4ti2` [1] and tends to be the fastest method in practice.

Proposition 2.1.1 is often called the *fundamental theorem of algebraic statistics* and along with it comes the fundamental *problem* of algebraic statistics: The only known way to compute a Markov basis for $\mathcal{F}_{A,b}$ is to compute a Markov basis \mathcal{M} for the whole collection \mathfrak{F}_A . As a consequence, \mathcal{M} contains a lot of moves that are redundant for an irreducible random walk on the fiber $\mathcal{F}_{A,b}$ of interest, but whose pure presence increases the rejection rate of any fiber walk (see also Proposition 4.3.2). To get around that problem, a truncated version of Buchberger’s algorithm for toric ideals was introduced in [113] that computes for given $b \in \mathbb{N}A$ a Gröbner basis for $\mathcal{F}_{A,b'}$ with $b - b' \in \mathbb{N}A$. This was further relaxed in [86]. Another approach is to write a matrix A as a *toric fiber product* of “easier” matrices B_1 and B_2 and to lift and glue Markov bases of B_1 and B_2 to a Markov basis of A afterwards [111]. This technique turns out to be highly beneficial for many statistical models, especially hierarchical ones [50, 101].

Remark 2.1.2. In [101], the notion of an *inequality Markov basis* was introduced, which is a Markov basis for sets of the form $\{u \in v + \mathcal{L} : Bu \leq b\}$ for a fixed lattice $\mathcal{L} \subset \mathbb{Z}^d$, a fixed integer matrix $B \in \mathbb{Z}^{m \times d}$, and varying $v \in \mathbb{Z}^d$ and $b \in \mathbb{Z}^m$. The computation of an inequality Markov basis for \mathcal{L} and B can be reduced to a Markov basis computation in the sense of Proposition 2.1.1.

The connectedness of fiber graphs is subtle and highly sensitive to small changes of the saturated set or the set of moves. For instance, when \mathcal{M} is a Markov basis for \mathcal{F} , then connectedness does not carry over to subsets $\mathcal{F}' \subset \mathcal{F}$ in general, not even saturated ones. Similarly, adding a single row to a constraint matrix $A \in \mathbb{Z}^{m \times d}$ can have huge effects on the size of a minimal Markov basis. The Graver basis, on the other hand, is a very robust set of moves in the following sense:

Proposition 2.1.3. *Let $A \in \mathbb{Z}^{m \times d}$, then \mathcal{G}_A is Markov basis for $\{v \in \mathbb{Z}^d : Av = b, l \leq v \leq u\}$ for all $b \in \mathbb{N}A$ and $l, u \in \mathbb{Z}^d$.*

Proof. This is [31, Lemma 3.2.4]. □

Remark 2.1.4. Let $B \in \mathbb{Z}^{m \times d}$ such that $\mathcal{F}_b := \{u \in \mathbb{Z}^d : Bu \leq b\}$ is bounded for all $b \in \mathbb{Z}^m$. There are many ways to compute a Markov basis for the collection $\mathfrak{F} := \{\mathcal{F}_b : b \in \mathbb{Z}^m\}$. One way is by Proposition 2.1.3: The projection onto the first d coordinates of the Graver basis of the matrix $(B, I_m) \in \mathbb{Z}^{m \times (d+m)}$ is a Markov basis for \mathfrak{F} . Another way is to use [101, Lemma 6], which says that when \mathcal{M}' is a Markov basis for the lattice $\mathbb{Z} \cdot B$ generated by the columns of B , then $\{m \in \mathbb{Z}^d : Bm \in \mathcal{M}'\}$ is an inequality Markov basis for \mathbb{Z}^d and B (Remark 2.1.2) and hence a Markov basis for \mathfrak{F} . In particular, given a polytope $\mathcal{P} \subset \mathbb{Q}^d$ in \mathcal{H} -representation, a Markov basis for the collection $\{(i \cdot \mathcal{P}) \cap \mathbb{Z}^d : i \in \mathbb{N}\}$ can be computed.

The theory around Markov bases hosts many open questions, and they typically range from the computation of Markov bases for certain statistical models to the investigation of bounds on the degrees of Markov binomials. The following question arises while working on this thesis and aims in a new direction: Does there exist a universal constant on the number of moves needed to connect any given saturated set in fixed dimension? Expressed with quantifiers:

Question 2.1.5. *Is there for all $d \in \mathbb{N}$ a constant C_d such that any saturated set $\mathcal{F} \subset \mathbb{Z}^d$ has a Markov basis $\mathcal{M} \subset \mathbb{Z}^d$ with $|\mathcal{M}| \leq C_d$?*

A trivial lower bound is $C_d \geq d$, since any Markov basis of a saturated sets has at least $\dim(\mathcal{F})$ many moves. For $d = 1$, the unifying set $\{1\}$ gives $C_1 := 1$. For $d = 2$, however, $C_2 > 2$ must

be true: Figure 2.1 shows a saturated set $\mathcal{F} \subset \mathbb{Z}^2$ where $|\mathcal{M}| \geq 3$ for all Markov bases $\mathcal{M} \subset \mathbb{Z}^2$ for \mathcal{F} . In particular, there does not always exist a Markov basis with linearly independent moves. Currently, we do not know whether minimal Markov bases in \mathbb{Z}^2 can be arbitrarily large.




Figure 2.1: A saturated set where every Markov basis has at least 3 moves.

For the remainder of this section, we survey results on the connectivity of fiber graphs, mainly based on the joint work [62] with Raymond Hemmecke. We first need a few definitions. Let $G = (V, E)$ be a graph, then G is k -node-connected if $|V| > k$ and if for all $X \subseteq V$ with $|X| < k$, the induced graph on $V \setminus X$ is connected. Similarly, the graph G is k -edge-connected if $|E| > k$ and if for all $X \subseteq E$ with $|X| < k$, the graph $(V, E \setminus X)$ is connected. The *edge-connectivity* (*node-connectivity*) of G is the largest natural number $k \in \mathbb{N}$ such that G is k -edge-connected (k -node-connected). Clearly, the node-connectivity of any graph is bounded from above by its edge-connectivity, which in turn is bounded from above by the minimal degree of the graph. We refer to [40] for more details on the connectivity of a graph. Let $A \in \mathbb{Z}^{m \times d}$, then a set $\mathcal{R} \subset \ker_{\mathbb{Z}}(A)$ is a (*reduced*) *Gröbner basis* of A if the set of corresponding binomials in $\mathbb{k}[x_1, \dots, x_d]$ is a (*reduced*) Gröbner basis of \mathcal{I}_A . For instance, the set of basic moves $\mathcal{M}_{n,m}$ is a Gröbner basis of $A_{n,m}$. It was shown in [97] that for $r \in \mathbb{N}_{>2}$ and $b_r := (r, \dots, r)^T \in \mathbb{N}^{2n}$, the node-connectivity and the minimal degree of $\mathcal{F}_{A_{n,n}, b_r}(\mathcal{M}_{n,n})$ are both $\binom{n}{2}$. There, it also was conjectured that this extends to the general case:

Conjecture 2.1.6. *Let $A \in \mathbb{Z}^{m \times d}$ and let $\mathcal{R} \subset \ker_{\mathbb{Z}}(A)$ be a reduced Gröbner basis of A . There exists $N \in \mathbb{N}^m$ so that the node-connectivity of $\mathcal{F}_{A,b}(\mathcal{R})$ equals its minimal degree for all $b \geq N$.*

Let us briefly recall the construction of the counter-example to Conjecture 2.1.6 from [62]. Let $A_1 \in \mathbb{Z}^{m \times d_1}$ and $A_2 \in \mathbb{Z}^{m \times d_2}$ and define

$$A_1 \times A_2 := \begin{bmatrix} A_1 & A_2 \\ 0 & A_2 \end{bmatrix}.$$

For any $b_i \in \mathbb{N}A_i$, we have the cartesian decomposition $\mathcal{F}_{A_1 \times A_2, (b_1+b_2) \times b_2} = \mathcal{F}_{A_1, b_1} \times \mathcal{F}_{A_2, b_2}$. In fact, given additionally two sets of moves $\mathcal{M}_i \subset \ker_{\mathbb{Z}}(A_i)$, we show in Section 3.1 that the fiber graph on $\mathcal{F}_{A_1 \times A_2, (b_1+b_2) \times b_2}$ with the set of moves $\mathcal{M}_1 \times \{0\} \cup \{0\} \times \mathcal{M}_2$ is the *cartesian product* of the graphs $\mathcal{F}_{A_1, b_1}(\mathcal{M}_1)$ and $\mathcal{F}_{A_2, b_2}(\mathcal{M}_2)$ (Proposition 3.1.9). Since $\ker_{\mathbb{Z}}(A_1 \times A_2) = \ker_{\mathbb{Z}}(A_1) \times \ker_{\mathbb{Z}}(A_2)$, linear algebra does not see a difference between $A_1 \times A_2$ and the diagonal matrix that has A_1 and A_2 as blocks. In particular, the Graver basis satisfies $\mathcal{G}_{A_1 \times A_2} = \mathcal{G}_{A_1} \times \mathcal{G}_{A_2}$ and Gröbner and Markov bases of $A_1 \times A_2$ decompose in the very same way. When $A_2 = I_m$ is the identity matrix, then all its fibers are single points and the fiber graph on $\mathcal{F}_{A_1 \times A_2, (b_1+b_2) \times b_2} = \mathcal{F}_{A_1, b_1} \times \{b_2\}$ is isomorphic to $\mathcal{F}_{A_1, b_1}(\mathcal{M}_1)$, for arbitrary $b_2 \in \mathbb{N}^m$. This allows to 'push' the entries of the right-hand side of $A_1 \times A_2$ beyond any given bound, without changing the graph structure:

Proposition 2.1.7. *Let $A \in \mathbb{Z}^{m \times d}$, $b \in \mathbb{N}A$, $\mathcal{M} \subset \mathbb{Z}^d$, and $c = (1, \dots, 1)^T \in \mathbb{Z}^m$. There exists $C \in \mathbb{N}$ such that for all $i \geq C$, $\mathcal{F}_{A,b}(\mathcal{M}) \cong \mathcal{F}_{A \times I_m, (b+i \cdot c) \times i \cdot c}(\mathcal{M} \times \{0\})$.*

Proof. This is [62, Theorem 1]. □

Rephrasing Proposition 2.1.7: Any given fiber graph is isomorphic to a fiber graph with arbitrarily large right-hand side entries. This provides a new view on the behaviour of graph-theoretic properties of fiber graphs and the way we should make conjectures about them. For instance, to disprove Conjecture 2.1.6, it suffices to find just a single right-hand side where the node-connectivity of the corresponding fiber graph differs from its minimal degree. It turns out, a particular right-hand side for the matrix from Example 1.3.1 settles this case:

Proposition 2.1.8. *The node-connectivity of $\mathcal{F}_{H_k, e_{2k+1}}(\mathcal{R}_k)$ is 1 and its minimal degree is k .*

Proof. For any graph, the node-connectivity is bounded from above by its edge-connectivity and bounded from below by 1 for connected graphs. Thus, this is precisely [62, Corollary 5.1]. \square

The set of moves \mathcal{R}_k is not only a Markov basis of H_k , but also a reduced Gröbner basis [62, Theorem 3]. Together with Proposition 2.1.7, $H_k \times I_{2k+1}$ is a counter-example to Conjecture 2.1.6. An appealing question is when the connectivity of fiber graphs is best-possible, that is, equal to its minimal-degree. This may be achieved by either adding or removing moves from the set of allowed moves, since both operations affect the connectivity and the minimal degree at the same time. For H_k , adding more structural moves makes the edge-connectivity best-possible.

Proposition 2.1.9. *The edge-connectivity of all fibers of H_k is best-possible when using the Graver basis as moves.*

Proof. This is [62, Theorem 4]. \square

Exploring the connectivity of fiber graphs is an important question in further understanding their structure. We are convinced that the connectivity of Graver fiber graphs is best-possible, i.e. that the answer to the following question is 'no':

Question 2.1.10. *Is there a matrix $A \in \mathbb{Z}^{m \times d}$ and $b \in \mathbb{N}^A$ such that the minimal degree of the fiber graph $\mathcal{F}_{A,b}(\mathcal{G}_A)$ is strictly larger than its edge-connectivity?*

2.2 Bounds on the diameter

Let G be a graph, then the *distance* $\text{dist}_G(u, v)$ between two distinct nodes u and v which are contained in the same connected component of G is the number of edges in a shortest (u, v) -path. We set $\text{dist}_G(u, v) := \infty$ if u and v are disconnected. The *diameter* of G , denoted $\text{diam}(G)$, is the maximal distance that appears between any pair of its nodes. In this section, we determine lower and upper bounds on the diameter of fiber graphs and their compressed counterparts. The results of this sections come, with small modifications, from [108, Section 3].

Lemma 2.2.1. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite and non-empty sets. Then for any norm $\|\cdot\|$,*

$$\text{diam}(\mathcal{F}(\mathcal{M})) \geq \frac{1}{\|\mathcal{M}\|} \cdot \max\{\|u - v\| : u, v \in \mathcal{F}\}.$$

Proof. If $\mathcal{F}(\mathcal{M})$ is not connected, then the statement holds trivially, so assume that \mathcal{M} is a Markov basis for \mathcal{F} . Let $u', v' \in \mathcal{F}$ such that $\|u' - v'\| = \max\{\|u - v\| : u, v \in \mathcal{F}\}$ and let $m_1, \dots, m_r \in \mathcal{M}$ so that $u' = v' + \sum_{i=1}^r m_i$ is a path of minimal length, then $\|u' - v'\| \leq r \cdot \|\mathcal{M}\|$ and the claim follows from $\text{diam}(\mathcal{F}(\mathcal{M})) \geq \text{dist}_{\mathcal{F}(\mathcal{M})}(u', v') = r$. \square

Remark 2.2.2. Let $\mathcal{P} \subset \mathbb{Q}^d$ be a lattice polytope and $\mathcal{F} := \mathcal{P} \cap \mathbb{Z}^d$ its attached saturated set. For any non-zero $l \in \mathbb{Z}^d$, the l -width of \mathcal{P} is $\text{width}_l(\mathcal{P}) := \max\{(u-v)^T l : u, v \in \mathcal{P}\}$. Since $\text{width}_l(\mathcal{P}) = \max\{u^T l : u \in \mathcal{P}\} - \min\{u^T l : u \in \mathcal{P}\}$ and since the maximum and minimum is attained at a vertex of \mathcal{P} , we have $\text{width}_l(\mathcal{P}) = \max\{(u-v)^T l : u, v \in \mathcal{F}\}$. For all $l \in \{-1, 0, 1\}^d$ and $u, v \in \mathcal{F}$, we have $(u-v)^T l \leq \|u-v\|_1$ and thus $\text{width}_l(\mathcal{P}) \leq \max\{\|u-v\|_1 : u, v \in \mathcal{F}\}$. Let $u', v' \in \mathcal{F}$ such that $\|u' - v'\|_1 = \max\{\|u-v\|_1 : u, v \in \mathcal{F}\}$ and let $l'_i := \text{sign}(u'_i - v'_i)$ for $i \in [d]$, then $\|u' - v'\|_1 = (u' - v')^T \cdot l' \leq \text{width}_{l'}(\mathcal{P}) \leq \max\{\|u-v\|_1 : u, v \in \mathcal{F}\} = \|u' - v'\|_1$. The *lattice width* of \mathcal{P} is $\text{width}(\mathcal{P}) := \min\{\text{width}_l(\mathcal{P}) : l \in \mathbb{Z}^d \setminus \{0\}\}$ and thus Lemma 2.2.1 gives

$$\|\mathcal{M}\|_1 \cdot \text{diam}(\mathcal{F}(\mathcal{M})) \geq \text{width}(\text{conv}_{\mathbb{Q}}(\mathcal{F})).$$

Given a collection \mathfrak{F} of saturated sets in \mathbb{Z}^d , Lemma 2.2.1 says that there exists a constant $C \in \mathbb{Q}_{>0}$ such that $\text{diam}(\mathcal{F}(\mathcal{M})) \geq C \cdot \max\{\|u-v\| : u, v \in \mathcal{F}\}$ holds for all $\mathcal{F} \in \mathfrak{F}$. We now investigate when a similar upper bound on $\text{diam}(\mathcal{F}(\mathcal{M}))$ holds, which is not always the case:

Example 2.2.3. Let $i \in \mathbb{N}$ and let $\mathcal{F}_i \subset \mathbb{Z}^3$ be the 'pyramid' from Figure 2.2, where i denotes the number of integers on its largest slice. Since the description of \mathcal{F}_i in symbols is quite vacuous and technical, we omit it here. Clearly, \mathcal{F}_i is a saturated set and $\mathcal{M} = \{(1, 0, 0), (0, 0, 1), (0, 1, 1)\}$ is a Markov basis for \mathcal{F}_i for all $i \in \mathbb{N}$. It is straightforward to check that $\mathcal{F}_i(\mathcal{M})$ is a path of length $|\mathcal{F}_i| - 1 = \frac{(i+1) \cdot i}{2} + i - 2$ and that $\max\{\|u-v\| : u, v \in \mathcal{F}_i\} = \frac{i-3}{2} \leq C \cdot i$. Since $\text{diam}(\mathcal{F}_i(\mathcal{M})) \geq C' \cdot i^2$, the diameter grows quadratically in the 1-norm distance.




Figure 2.2: The pyramid for $i = 7$.

The following concept is well-established (see for instance [6, Chapter 6]) in the theory of Markov bases and helps us to derive an upper bound on the diameter of fiber graphs:

Definition 2.2.4. Let $\mathcal{M} \subset \mathbb{Z}^d$ be finite. Then \mathcal{M} is *norm-reducing* for a finite set $\mathcal{F} \subset \mathbb{Z}^d$ if for all $u, v \in \mathcal{F}$, there is $m \in \mathcal{M}$ such that $u + m \in \mathcal{F}$ and $\|u + m - v\|_1 < \|u - v\|_1$. The set \mathcal{M} is *norm-reducing* for a collection \mathfrak{F} of finite sets of \mathbb{Z}^d if \mathcal{M} is norm-reducing for all $\mathcal{F} \in \mathfrak{F}$.

Any norm-reducing set of moves is a Markov basis [6, Proposition 6.1] and it is well-known that the Graver basis of any matrix $A \in \mathbb{Z}^{m \times d}$ is norm-reducing for \mathfrak{F}_A [6, Proposition 6.4]. An example of a Markov basis that is norm-reducing and strictly contained in the Graver basis is $\mathcal{M}_{n,m}$ from Example 1.3.2. Between being norm-reducing and being a Markov basis, there is much space left for the following property:

Definition 2.2.5. Let \mathfrak{F} be a collection of finite subsets of \mathbb{Z}^d and $\mathcal{M} \subset \mathbb{Z}^d$ a finite set. Then $\mathcal{M} \subset \mathbb{Z}^d$ is *norm-like* for \mathfrak{F} if there exists a constant $C \in \mathbb{N}_{>0}$ such that for all $\mathcal{F} \in \mathfrak{F}$ and all $u, v \in \mathcal{F}$, $\text{dist}_{\mathcal{F}(\mathcal{M})}(u, v) \leq C \cdot \|u - v\|_1$.

All norms on \mathbb{R}^d are equivalent and thus we could have used any other norm in Definition 2.2.4. Observe, however, that being norm-reducing depends on the norm. With Lemma 2.2.1 in mind, the next remark motivates why we are interested in norm-like set of moves:

Remark 2.2.6. Let \mathfrak{F} be a collection of finite subsets of \mathbb{Z}^d and $\mathcal{M} \subset \mathbb{Z}^d$ be norm-like for \mathfrak{F} . It follows from the definition that there is a constant $C \in \mathbb{N}_{>0}$ such that for all $\mathcal{F} \in \mathfrak{F}$

$$\text{diam}(\mathcal{F}(\mathcal{M})) \leq C \cdot \max\{\|u - v\|_1 : u, v \in \mathcal{F}\}.$$

However, the converse is not true as demonstrated in Example 2.2.7.

Norm-reducing sets are always norm-like, and norm-like sets are in turn always Markov bases, but the reverse of both statements is false in general (see Example 2.2.7 and Example 2.2.8). For collections \mathfrak{F}_A however, every Markov basis is norm-like (Proposition 2.2.9).

Example 2.2.7. For any $i \in \mathbb{N}$, consider the saturated set $\mathcal{F}_i := ([2] \times [i] \times \{0\}) \cup \{(2, i, 1)\}$ with the Markov basis $\{(0, 1, 0), (0, 0, 1), (-1, 0, -1)\}$ (see Figure 2.3). The distance between $(1, 1, 0)$ and $(2, 1, 0)$ in $\mathcal{F}_i(\mathcal{M})$ is $2i$ and thus \mathcal{M} is not norm-like for the collection $\{\mathcal{F}_i : i \in \mathbb{N}\}$. Observe that the diameter of $\mathcal{F}_i(\mathcal{M})$ is bounded from above by $2 \cdot \max\{\|u - v\|_1 : u, v \in \mathcal{F}_i\}$.

Example 2.2.8. Let $d \in \mathbb{N}$ and consider A_d and \mathcal{M}_d from Example 1.3.3. For any $d \geq 3$, the only move from \mathcal{M}_d that can be applied on e_2 in $\mathcal{F}_{A_d,1}(\mathcal{M}_d)$ is the move $e_1 - e_2$. But since $\|(e_2 + e_1 - e_2) - e_3\|_1 = \|e_1 - e_3\|_1$, \mathcal{M}_d is not norm-reducing for \mathfrak{F}_{A_d} . On the other hand, when we cannot find a move that reduces the 1-norm distance of two nodes $u, v \in \mathcal{F}_{A_d,b}$, we instead find two moves $m_1, m_2 \in \mathcal{M}_d$ such that $u + m_1, u + m_1 + m_2 \in \mathcal{F}_{A_d,b}$ and $\|u + m_1 + m_2 - v\|_1 = \|u - v\|_1 - 2$. Thus, the graph-distance of any two elements u and v in $\mathcal{F}_{A_d,b}(\mathcal{M}_d)$ is at most $\|u - v\|_1$ and hence \mathcal{M}_d is norm-like for \mathfrak{F}_{A_d} .



Figure 2.3: The graphs $\mathcal{F}_4(\mathcal{M})$ and $\mathcal{F}_7(\mathcal{M})$ from Example 2.2.7.

The next proposition states that when the sets in a collection come from the same integer matrix, effects as in Example 2.2.7 cannot occur.

Proposition 2.2.9. Let $A \in \mathbb{Z}^{m \times d}$, then any Markov basis of \mathfrak{F}_A is norm-like.

Proof. Let \mathcal{M} be a Markov basis for \mathfrak{F}_A and define $C := \max\{\text{diam}(\mathcal{F}_{A,Ag^+}(\mathcal{M})) : g \in \mathcal{G}_A\}$, which is well since the Graver basis \mathcal{G}_A is a finite set. For $u, v \in \mathcal{F}_{A,b}$, let $v = u + \sum_{i=1}^r g_i$ be a walk from u to v in $\mathcal{F}_{A,b}(\mathcal{G}_A)$ of minimal length. Since the Graver basis is norm-reducing for $\mathcal{F}_{A,b}$ [6, Proposition 6.4], there always exists a path of length at most $\|u - v\|_1$ and hence $r \leq \|u - v\|_1$. Every g_i can be replaced by a path in $\mathcal{F}_{A,Ag_i^+}(\mathcal{M})$ of length at most C and these paths stay in $\mathcal{F}_{A,b}$. This gives a path of length $C \cdot r$, hence $\text{dist}_{\mathcal{F}_{A,b}(\mathcal{M})}(u, v) \leq C\|u - v\|_1$. \square

Proposition 2.2.10. Let $\mathcal{P} \subset \mathbb{Q}^d$ be a polytope with $\dim(\mathcal{P} \cap \mathbb{Z}^d) > 0$ and let \mathcal{M} be a Markov basis for $\mathcal{F}_i := (i \cdot \mathcal{P}) \cap \mathbb{Z}^d$ for all $i \in \mathbb{N}$. There exists a constant $C' \in \mathbb{Q}_{>0}$ such that for all $i \in \mathbb{N}$, $C' \cdot i \leq \text{diam}(\mathcal{F}_i(\mathcal{M}))$. If \mathcal{M} is norm-like for $\{\mathcal{F}_i : i \in \mathbb{N}\}$, then there exists a constant $C \in \mathbb{Q}_{>0}$ such that $\text{diam}(\mathcal{F}_i(\mathcal{M})) \leq C \cdot i$ for all $i \in \mathbb{N}$.

Proof. For the lower bound on the diameter, it suffices to show the existence of C' such that $C' \cdot i \leq \max\{\|u - v\|_1 : u, v \in \mathcal{F}_i\}$ for all $i \in \mathbb{N}$ due to Lemma 2.2.1. Since $\dim(\mathcal{P} \cap \mathbb{Z}^d) > 0$, we can pick distinct $w, w' \in \mathcal{P} \cap \mathbb{Z}^d$. For all $i \in \mathbb{N}$, the vectors $i \cdot w$ and $i \cdot w'$ are in \mathcal{F}_i and hence $i \cdot \|w - w'\|_1 \leq \max\{\|u - v\|_1 : u, v \in \mathcal{F}_i\}$. For the upper bound, assume that the Markov basis \mathcal{M} is norm-like. By Remark 2.2.6, it suffices to show the existence of $C \in \mathbb{Q}_{\geq 0}$ such that $\max\{\|u - v\|_1 : u, v \in \mathcal{F}_i\} \leq i \cdot C$. Now, let $v_1, \dots, v_r \in \mathbb{Q}^d$ such that $\mathcal{P} = \text{conv}_{\mathbb{Q}}(v_1, \dots, v_r)$ and define $C := \max\{\|v_s - v_t\|_1 : s \neq t\}$. Since $\mathcal{F}_i = (i \cdot \mathcal{P}) \cap \mathbb{Z}^d \subset \text{conv}_{\mathbb{Q}}(iv_1, \dots, iv_r)$ for all $i \in \mathbb{N}$, $\max\{\|u - v\|_1 : u, v \in \mathcal{F}_i\} \leq \max\{\|u - v\|_1 : u, v \in i \cdot \mathcal{P}\} \leq \max\{\|iv_s - iv_t\|_1 : s \neq t\} \leq C \cdot i$. \square

[97, Proposition 2.10] shows that the diameter of fiber graphs of $A_{n,n}$ along a certain ray in $\mathbb{N}A_{n,n}$ grows linearly. The following result generalizes this to all matrices:

Corollary 2.2.11. *Let $A \in \mathbb{Z}^{m \times n}$, \mathcal{M} be a Markov basis for \mathfrak{F}_A , and $b \in \mathbb{N}A$ with $\dim(\mathcal{F}_{A,b}) > 0$. Then there exist $C, C' \in \mathbb{Q}_{>0}$ such that $i \cdot C' \leq \text{diam}(\mathcal{F}_{A,i \cdot b}(\mathcal{M})) \leq i \cdot C$ for all $i \in \mathbb{N}$.*

Proof. This follows from Proposition 2.2.10 since \mathcal{M} is norm-like by Proposition 2.2.9. \square

One question remains. We have seen that the diameter of fiber graphs with norm-like Markov bases is bounded from above and below by the 1-norm distance of their elements linearly. Example 2.2.3 shows a collection and a Markov basis where this is not the case, but with one additional move, the linear bound hold for this example as well. Is this the general case?

Question 2.2.12. *Let \mathfrak{F} be a collection of saturated sets in \mathbb{Z}^d that has a Markov basis. Does there exist a Markov basis \mathcal{M} for \mathfrak{F} such that there is a constant $C \in \mathbb{Q}_{>0}$ with*

$$\text{diam}(\mathcal{F}(\mathcal{M})) \leq C \cdot \max\{\|u - v\|_1 : u, v \in \mathcal{F}\}$$

for all $\mathcal{F} \in \mathfrak{F}$?

We now turn our attention to the diameter of compressed fiber graphs. Since the edge set of a graph is contained in the edge set of its compressed version, paths may become shorter after the compression. But, in general, compression may do not change the graph at all:

Example 2.2.13. For any $i \in \mathbb{N}$, let $\mathcal{F}_i := \{(0, 0), (0, 1), (1, 1), (1, 2), \dots, (i, i)\} \subset \mathbb{Z}^2$. The unit vectors $\mathcal{M} = \{e_1, e_2\}$ are a Markov basis for $\{\mathcal{F}_i : i \in \mathbb{N}\}$. However, $\mathcal{F}_i^c(\mathcal{M}) = \mathcal{F}_i(\mathcal{M})$ and thus $\text{diam}(\mathcal{F}_i^c(\mathcal{M})) = \text{diam}(\mathcal{F}_i(\mathcal{M})) = 2i$ is unbounded.

Recall that \sqsubseteq denotes the partial ordering on \mathbb{Z}^d by sign-compatibility (see Section 1).

Lemma 2.2.14. *Let $A \in \mathbb{Z}^{m \times d}$ and $z \in \ker_{\mathbb{Z}}(A)$. There exists $r \in [2d - 2]$, distinct elements $g_1, \dots, g_r \in \mathcal{G}_A$, and $\lambda_1, \dots, \lambda_r \in \mathbb{N}_{>0}$ such that $z = \sum_{i=1}^r \lambda_i g_i$ and $g_i \sqsubseteq z$ for all $i \in [r]$*

Proof. This is [31, Lemma 3.2.3], although it only becomes clear from the original proof in [104, Theorem 2.1] that the appearing elements are indeed all distinct. \square

Proposition 2.2.15. *Let $A \in \mathbb{Z}^{m \times d}$ and $\mathfrak{F} := \{\{x \in \mathbb{Z}^d : Ax = b, l \leq x \leq u\} : l, u \in \mathbb{Z}^d, b \in \mathbb{Z}^m\}$. Then for all $\mathcal{F} \in \mathfrak{F}$, $\text{diam}(\mathcal{F}^c(\mathcal{G}_A)) \leq 2d - 2$.*

Proof. Let $s, t \in \{x \in \mathbb{Z}^d : Ax = b, l \leq x \leq u\}$, then $s - t \in \ker_{\mathbb{Z}}(A)$ and thus $s = t + \sum_{i=1}^r \lambda_i g_i$ with $r \leq 2d - 2$, $\lambda_1, \dots, \lambda_r \in \mathbb{N}_{>0}$, and distinct $g_1, \dots, g_r \in \mathcal{G}_A$ such that $g_i \sqsubseteq s - t$ according to Lemma 2.2.14. By [31, Lemma 3.2.4], all intermediate points $t + \sum_{i=1}^l \lambda_i g_i$, $l \in [r]$, are in $\{x \in \mathbb{Z}^d : Ax = b, l \leq x \leq u\}$ and the claim follows. \square

Lemma 2.2.16. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite and let $\mathcal{F}_i := (i \cdot \text{conv}_{\mathbb{Q}}(\mathcal{F})) \cap \mathbb{Z}^d$ for $i \in \mathbb{N}$. For all $u, v \in \mathcal{F}$, $\text{dist}_{\mathcal{F}_i^c(\mathcal{M})}(iu, iv) \leq \text{dist}_{\mathcal{F}(\mathcal{M})}(u, v)$ for all $i \in \mathbb{N}$.*

Proof. The statement is true if u and v are disconnected in $\mathcal{F}(\mathcal{M})$. Thus, let $u = v + \sum_{j=1}^k m_j$ with $m_j \in \mathcal{M}$ be a path in $\mathcal{F}(\mathcal{M})$ of length $k = \text{dist}_{\mathcal{F}(\mathcal{M})}(u, v)$ and let $i \in \mathbb{N}$. Clearly, $i \cdot u = i \cdot v + i \cdot \sum_{j=1}^k m_j = i \cdot v + \sum_{l=1}^k i \cdot m_j$, so it is left to prove that the elements traversed by this paths are in \mathcal{F}_i . Let $l \in [k]$, since $v + \sum_{j=1}^l m_j \in \mathcal{F}$, we have $i \cdot v + \sum_{j=1}^l i \cdot m_j \in i \cdot \mathcal{F} \subseteq \mathcal{F}_i$. Hence, this is a path in $\mathcal{F}_i^c(\mathcal{M})$ of length $k = \text{dist}_{\mathcal{F}(\mathcal{M})}(u, v)$. \square

We are ready to prove that the diameter of compressed fiber graphs that come from the same integer matrix is bounded from above for all right-hand sides simultaneously.

Theorem 2.2.17. *Let $A \in \mathbb{Z}^{m \times d}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$ and let \mathcal{M} be a Markov basis for \mathfrak{F}_A . There exists a constant $C \in \mathbb{N}$ such that $\text{diam}(\mathcal{F}^c(\mathcal{M})) \leq C$ for all $\mathcal{F} \in \mathfrak{F}_A$.*

Proof. The proof relies on basic properties of the Graver basis \mathcal{G}_A of A . For any $g \in \mathcal{G}_A$, let $\mathcal{F}_g := \mathcal{F}_{A, Ag^+}$ and let $K := \max\{\text{dist}_{\mathcal{F}_g(\mathcal{M})}(g^+, g^-) : g \in \mathcal{G}_A\}$. We show that the diameter of any compressed fiber graph of A is bounded from above by $(2d - 2) \cdot K$. For any $b \in \mathbb{N}A$, let $u, v \in \mathcal{F}_{A,b}$ be arbitrary. According to Proposition 2.2.15, there exists $r \in [2d - 2]$, Graver moves $g_1, \dots, g_r \in \mathcal{G}_A$, and coefficients $\lambda_1, \dots, \lambda_r \in \mathbb{Z}$ such that $u = v + \sum_{i=1}^r \lambda_i g_i$, and $v + \sum_{i=1}^l \lambda_i g_i \in \mathbb{N}^d$ for all $l \in [r]$. According to Lemma 2.2.16, for any $i \in [r]$ there are $m_1^i, \dots, m_{k_i}^i \in \mathcal{M}$ and $\alpha_1, \dots, \alpha_{k_i} \in \mathbb{Z}$ such that $\lambda_i g_i^+ = \lambda_i g_i^- + \sum_{j=1}^{k_i} \alpha_j m_j^i$ is a path in the compression of the fiber graph $\mathcal{F}_{A, A\lambda_i g_i^+}(\mathcal{M})$ of length $k_i \leq K$. Lifting these paths yields a path $u = v + \sum_{i=1}^r \sum_{j=1}^{k_i} \alpha_j m_j^i$ in $\mathcal{F}_{A,b}^c(\mathcal{M})$ of length $r \cdot K \leq (2d - 2) \cdot K$. \square

Remark 2.2.18. Proposition 2.2.9 and Theorem 2.2.17 extend trivially to collections of the form $\mathfrak{F} = \mathfrak{F}_{A_1} \cup \dots \cup \mathfrak{F}_{A_r}$ for integer matrices $A_i \in \mathbb{Z}^{m_i \times d}$.

2.3 Graph degrees

The *minimal degree* of a graph G is $\delta(G) := \min\{\deg_G(v) : v \in V(G)\}$. In this section, miscellaneous observations on the minimal degree of fiber graphs are discussed. The minimal degree of $\mathcal{F}(\mathcal{M})$ is trivially bounded from above by $|\pm \mathcal{M}|$ and its exact value is the solution to the optimization problem

$$\min\{|\{m \in \mathcal{M} : v + m \in \mathcal{F}\}| : v \in \mathcal{F}\},$$

which seems too generic without additional conditions on \mathcal{F} and \mathcal{M} . We thus focus in the following on fibers $\mathcal{F}_{A,b}$ with $A \in \mathbb{Z}^{m \times d}$ and $b \in \mathbb{N}A$. What makes this situation special is that for all finite sets of moves $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$, the graph-degree of any node u in $\mathcal{F}_{A,b}(\mathcal{M})$ can be read off u and \mathcal{M} directly by checking which moves $m \in \mathcal{M}$ satisfy $u + m \in \mathbb{N}^d$. Hence, no information about the surrounding faces of the polytope $\text{conv}_{\mathbb{Q}}(\mathcal{F}_{A,b})$ is needed. When $\mathcal{M} \subseteq \{-1, 0, 1\}^d$, then the graph-degree of u depends only on the support $\text{supp}(u)$ and we can locate the nodes with the smallest degree in $\mathcal{F}_{A,b}(\mathcal{M})$ for any b :

Proposition 2.3.1. *Let $A \in \mathbb{Z}^{m \times d}$ and $\mathcal{M} \subseteq \{-1, 0, 1\}^d$. For any $b \in \mathbb{N}A$, the minimal degree of $\mathcal{F}_{A,b}(\mathcal{M})$ is attained at a vertex of the polytope $\text{conv}_{\mathbb{Q}}(\mathcal{F}_{A,b})$.*

Proof. Let $\deg : \mathcal{F}_{A,b} \rightarrow \mathbb{N}$ be the function that maps a node of $\mathcal{F}_{A,b}(\mathcal{M})$ to its graph-degree. Let $W \subseteq \mathcal{F}_{A,b}$ be the set of vertices of $\text{conv}_{\mathbb{Q}}(\mathcal{F}_{A,b})$ and choose $v \in \mathcal{F}_{A,b} \setminus W$. Then there is $k > 1$ and vertices $w_1, \dots, w_k \in W$ together with coefficients $\lambda_1, \dots, \lambda_k \in \mathbb{Q}_{>0}$ that satisfy $\sum_{i=1}^k \lambda_i = 1$ such that $v = \sum_{i=1}^k \lambda_i w_i$. Since all involved vectors are non-negative, no cancellation appears and hence $\text{supp}(\lambda_i w_i) \subseteq \text{supp}(v)$ for all $i \in [k]$. Since $\lambda_i > 0$, we deduce that $\text{supp}(w_i) \subseteq \text{supp}(v)$. All moves are in $\{-1, 0, 1\}^d$ and hence the graph degree only depends on the support of the respective node. More precisely, the containment $\text{supp}(w_i) \subseteq \text{supp}(v)$ implies $\deg(w_i) \leq \deg(v)$ for all $i \in [k]$ and this finishes the proof. \square

In the proof of Proposition 2.3.1, we can further conclude that

$$\deg(v) = \sum_{i=1}^k \lambda_i \deg(v) \geq \sum_{i=1}^k \lambda_i \deg(w_i).$$

Since $\sum_{i=1}^k \lambda_i w_i = v$, this shows concavity of the graph-degree of $\mathcal{F}_{A,b}(\mathcal{M})$ when $\mathcal{M} \subseteq \{-1, 0, 1\}^d$. Moreover, for any $u \in \mathcal{F}_{A,b}$ and $l \in \ker_{\mathbb{Z}}(A)$, let $\beta \in \mathbb{Z}_{\geq 0}$ be maximal such that $u + \beta l \in \mathcal{F}_{A,b}$ and let $\mu \in \mathbb{Z}_{\leq 0}$ be minimal such that $u + \mu l \in \mathcal{F}_{A,b}$. For any $\mu < t \leq t' < \beta$, $\text{supp}(u + tl) = \text{supp}(u + t'l)$ and since $\mathcal{M} \subseteq \{-1, 0, 1\}^d$, the nodes $u + tl$ and $u + t'l$ have the same graph-degree in $\mathcal{F}_{A,b}(\mathcal{M})$. Moreover, the supports of $u + \beta l$ and $u + \mu l$ are strictly contained in $\text{supp}(u + tl)$ for any $\mu < t < \beta$ and hence among all points on the discrete ray $(u + \mathbb{Z} \cdot l) \cap \mathcal{F}_{A,b}$, the nodes $u + \mu l$ and $u + \beta l$ have the smallest degree.



Figure 2.4: Fiber graphs whose minimal degree is attained at an interior point.

It is easy to see (especially after seeing it) that Proposition 2.3.1 is false when one of the assumptions is violated: On the left-hand side of Figure 2.4, the fiber graph on $\mathcal{F}_{A_3,3}$ that uses the moves $\{2 \cdot m : m \in \mathcal{M}_3\}$ with A_3 and \mathcal{M}_3 as in Example 1.3.3 is shown. The graph on the right-hand side is the fiber graph on the saturated set

$$\left\{ \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ -1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ +1 \\ 0 \end{bmatrix}, \begin{bmatrix} +1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} +2 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} +1 \\ 0 \\ +1 \end{bmatrix}, \begin{bmatrix} +1 \\ 0 \\ -1 \end{bmatrix} \right\}$$

that uses the moves

$$\left\{ \begin{bmatrix} +1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ +1 \\ 0 \end{bmatrix}, \begin{bmatrix} +1 \\ +1 \\ -1 \end{bmatrix}, \begin{bmatrix} +1 \\ +1 \\ +1 \end{bmatrix}, \begin{bmatrix} +1 \\ -1 \\ +1 \end{bmatrix}, \begin{bmatrix} +1 \\ -1 \\ -1 \end{bmatrix}, \begin{bmatrix} +1 \\ 0 \\ -1 \end{bmatrix}, \begin{bmatrix} +1 \\ 0 \\ +1 \end{bmatrix} \right\}.$$

In both graphs, the degree of their unique interior point is strictly smaller than the degree of all the other nodes. Whereas in the first graph, the entries of the moves are not from $\{-1, 0, 1\}$ and the node set is a fiber, the node set of the second graph is not a fiber but the moves are from $\{-1, 0, 1\}^3$, justifying that both assumptions are needed in Proposition 2.3.1. We also obtain an upper bound on the minimal degree that is slightly better than the trivial bound $|\pm \mathcal{M}|$.

Proposition 2.3.2. *Let $A \in \mathbb{Z}^{m \times d}$, $\mathcal{M} \subseteq \{-1, 0, 1\}^d$, and $b \in \mathbb{N}A$. Then $\delta(\mathcal{F}_{A,b}(\mathcal{M})) \leq |\mathcal{M}|$.*

Proof. By Proposition 2.3.1, there exists a vertex v of the polytope $\text{conv}_{\mathbb{Q}}(\mathcal{F}_{A,b})$ whose degree equals the minimal degree of $\mathcal{F}_{A,b}(\mathcal{M})$. Thus, for any $m \in \mathcal{M}$, either $v+m \notin \mathcal{F}_{A,b}$ or $v-m \notin \mathcal{F}_{A,b}$ since v would otherwise be contained in a face whose dimension is at least 1. \square

Recall that an integer matrix is *totally unimodular* if each of its subdeterminants is in $\{-1, 0, 1\}$. For totally unimodular matrices, dilation of their fibers does not change the minimal degree:

Proposition 2.3.3. *Let $\mathcal{M} \subseteq \{-1, 0, 1\}^d$ be a Markov basis for a totally unimodular matrix A . For all $b \in \mathbb{N}A$ and all $i \in \mathbb{N}$, $\delta(\mathcal{F}_{A,b}(\mathcal{M})) = \delta(\mathcal{F}_{A,ib}(\mathcal{M}))$.*

Proof. Total unimodularity yields $\text{conv}_{\mathbb{Q}}(\mathcal{F}_{A,b}) = \{x \in \mathbb{Q}_{\geq 0}^d : Ax = b\}$. Let $v_1, \dots, v_r \in \mathcal{F}_{A,b}$ be the vertices of $\text{conv}_{\mathbb{Q}}(\mathcal{F}_{A,b})$, then $\text{conv}_{\mathbb{Q}}(\mathcal{F}_{A,ib}) = \text{conv}_{\mathbb{Q}}(iv_1, \dots, iv_r)$ for all $i \in \mathbb{N}$. According to Proposition 2.3.1, the minimal degree of $\mathcal{F}_{A,b}$ is attained at a vertex of the polytope, say v_1 . Since the graph-degree of any vector with respect to fiber graphs using \mathcal{M} does not change after multiplying it with a scalar, the minimal degree of $\mathcal{F}_{A,ib}$ is attained at $i \cdot v_1$ and coincides with the minimal degree of $\mathcal{F}_{A,b}$. \square

2.4 Edge-expansions

This section is based on the author's work [119]. There are many different invariants of graphs that measure how well connected it is. Beside *connectivity* and *toughness* [11], the *edge-expansion* is one of the most important concepts:

Definition 2.4.1. Let $G = (V, E)$ be a graph. For any $S \subseteq V$, let $E_G(S) \subseteq E$ be the set of edges with endpoints in S and $V \setminus S$. The *edge-expansion* of G is

$$h(G) := \min \left\{ \frac{|E_G(S)|}{|S|} : S \subset V, 0 < 2|S| \leq |V| \right\}.$$

The invariant $h(G)$ has many names in the literature, like *Cheeger constant* [24] or *isoperimetric number* [90]. As we will see in Chapter 4, the edge-expansions connects nicely statistics and graph theory since it yields a bound on the second largest eigenvalue modulus (Proposition 4.1.9). For simple walks, the edge-expansion of the graph is proportional to the conductance:

Remark 2.4.2. Let G be a d -regular graph, then $\Phi(\mathcal{S}_G) \cdot d = h(G)$.

Let $A \in \mathbb{Z}^{m \times d}$ and let $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis for A . The central question of this section is the following: What happens to the edge-expansion of fiber graphs asymptotically when the right-hand side varies? The following example depicts how a possible answer looks like:

Example 2.4.3. Let A_2 and \mathcal{M}_2 be as in Example 1.3.3. Then $\mathcal{F}_{A_2,i}(\mathcal{M}_2)$ is a path on $i+1$ nodes and hence its edge-expansion is $\frac{2}{i+1}$ if i is odd and $\frac{2}{i}$ when i is even [90, Section 2].

Definition 2.4.4. Let $A \in \mathbb{Z}^{m \times d}$ and let $(b_i)_{i \in \mathbb{N}}$ be a sequence in $\mathbb{N}A$. The sequence $(b_i)_{i \in \mathbb{N}}$ is a *ray* in $\mathbb{N}A$ if there is $b \in \mathbb{N}A$ such that $(b_i)_{i \in \mathbb{N}} = (i \cdot b)_{i \in \mathbb{N}}$.

We need the following terminology for our next definition: For $b \in \mathbb{N}A$, the \mathbb{Q} -relaxation of the fiber $\mathcal{F}_{A,b}$ is the polytope $\mathcal{R}_{A,b} := \{x \in \mathbb{Q}_{\geq 0}^d : Ax = b\}$.

Definition 2.4.5. Let $A \in \mathbb{Z}^{m \times d}$. A sequence $(b_i)_{i \in \mathbb{N}}$ is *dominated* if there exists $b \in \mathbb{N}A$ with $\dim(\mathcal{R}_{A,b}) > 0$ such that $b_i - i \cdot b \in \mathbb{N}A$ for all $i \in \mathbb{N}$ and if there is $u \in \mathcal{F}_{A,b}$ and $w_i \in \mathcal{F}_{A,b_i - i \cdot b}$ with $\text{supp}(w_i) \subseteq \text{supp}(u)$ for all $i \in \mathbb{N}$.

On the one hand, being dominated is a sufficient, though technical, condition on $(b_i)_{i \in \mathbb{N}}$ that is crucial in our proof of the decline of the edge-expansion of $(\mathcal{F}_{A,b_i}(\mathcal{M}))_{i \in \mathbb{N}}$ (Theorem 2.4.16). The prime example of a dominated sequence the reader should have in mind is a ray in the semigroup $\mathbb{N}A$ as in Example 2.4.3. We think it is an interesting task to further relax the conditions from Definition 2.4.5 so that they still prevent the following effects:

Remark 2.4.6. It was shown in Proposition 2.1.7 that for all $i \in \mathbb{N}$ and $c \in \mathbb{N}^m$, the fiber graphs of $A \times I_m$ for $b_i := (b + i \cdot c, i \cdot c)^T$ are all isomorphic to $\mathcal{F}_{A,b}$. In particular, their edge-expansion is constant along the sequence $(b_i)_{i \in \mathbb{N}}$. Assume that the sequence $(b_i)_{i \in \mathbb{N}}$ is dominated. Then there is $f = f_1 \times f_2 \in \mathbb{N}(A \times I_m)$ with $b_i - i \cdot f \in \mathbb{N}(A \times I_m)$ for all $i \in \mathbb{N}$. Thus, for every $i \in \mathbb{N}$ there exists $u_i \times v_i \in \mathbb{N}^{d+m}$ such that

$$\begin{bmatrix} A & I_m \\ 0 & I_m \end{bmatrix} \cdot \begin{bmatrix} u_i \\ v_i \end{bmatrix} = \begin{bmatrix} b + i \cdot c - i \cdot f_1 \\ i \cdot c - i \cdot f_2 \end{bmatrix}. \quad (2.4.1)$$

Since $f \in \mathbb{N}(A \times I_m)$, there is $u \times v \in \mathbb{N}^{d+m}$ such that $A \times I_m \cdot (u, v)^T = f$. This implies $Au = f_1 - f_2$ and, plugged into equation (2.4.1), we obtain $Au_i = b + i \cdot A(-u)$. This yields $u_i + i \cdot u \in \mathcal{F}_{A,b}$ for all $i \in \mathbb{N}$ which is due to $u_i, u \in \mathbb{N}^d$ only possible when $u = 0$. Hence, $f_1 = f_2$ and $\mathcal{F}_{A \times I_m, f} = \{0 \times f_2\}$. In particular, there is no element w_i in $\mathcal{F}_{A \times I_m, b_i - i \cdot f} = \mathcal{F}_{A,b} \times \{i \cdot (c - f_2)\}$ such that $\text{supp}(w_i) \subseteq \text{supp}(0 \times f_2)$ and consequently $(b_i)_{i \in \mathbb{N}}$ is not dominated.

Dominated sequences appear, for instance, as subsequence of sequences whose distance to the facets of $\mathbb{N}A$ becomes arbitrarily large. Let $H_A(b) := \min\{\text{dist}(b, F) : F \text{ facet of } \mathbb{N}A\}$, where $\text{dist}(b, F) \in \mathbb{Q}_{\geq 0}$ denotes the distance between b and the facet $F \subseteq \mathbb{N}A$.

Proposition 2.4.7. Let $A \in \mathbb{Z}^{m \times d}$ and let $(b_i)_{i \in \mathbb{N}}$ be from $\mathbb{N}A$ with $\limsup_{i \in \mathbb{N}} H_A(b_i) = \infty$, then $(b_i)_{i \in \mathbb{N}}$ has a dominated subsequence.

Proof. Let $a_1, \dots, a_d \in \mathbb{Z}^m$ be the columns of A and let $c := a_1 + \dots + a_d$. First, we show the following: For every $k \in \mathbb{N}$, there exists $m_k \in \mathbb{N}$ such that any $b \in \mathbb{N}A$ with $H_A(b) \geq m_k$ is contained in $k \cdot c + \mathbb{N}A$. The set $\mathbb{N}A \setminus (k \cdot c + \mathbb{N}A)$ is contained in finitely many hyperplanes parallel to the facets of $\mathbb{N}A$. Hence, choosing $m_k \in \mathbb{N}$ large enough, every $b \in \mathbb{N}A$ with $H_A(b) \geq m_k$ cannot be in $\mathbb{N}A \setminus (k \cdot c + \mathbb{N}A)$ and hence must be in $k \cdot c + \mathbb{N}A$. The statement of the lemma follows immediately because $\limsup_{i \in \mathbb{N}} H_A(b_i) = \infty$ implies that there is $i_k \in \mathbb{N}$ such that $H_A(b_{i_k}) \geq m_k$. In particular, $(b_{i_k})_{k \in \mathbb{N}}$ is dominated since $\mathcal{F}_{A,c}$ has an element with full support and since $\dim(\mathcal{R}_{A,c}) = \dim(\ker_{\mathbb{Z}}(A)) > 0$. \square

Remark 2.4.8. The reverse of Proposition 2.4.7 is not true. For instance, take the matrix

$$A = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

and the right-hand side $b = (2, 0)^T \in \mathbb{Z}^2$. The ray $(i \cdot b)_{i \in \mathbb{N}}$ is dominated since $\dim(\mathcal{R}_{A,b}) > 0$. However, since $\{i \cdot b : i \in \mathbb{N}\}$ is contained in a facet of $\mathbb{N}A$, $H_A(i \cdot b) = 0$ for all $i \in \mathbb{N}$.

Almost needless to say, being dominated is preserved under taking subsequences:

Lemma 2.4.9. *Every subsequence of a dominated sequence $(b_i)_{i \in \mathbb{N}}$ in $\mathbb{N}A$ is dominated.*

Proof. Let $(w_i)_{i \in \mathbb{N}}$ and $u \in \mathcal{F}_{A,b}$ be as in Definition 2.4.5 and let $(b_{i_j})_{j \in \mathbb{N}}$ be a subsequence of $(b_i)_{i \in \mathbb{N}}$. Consider $w'_j := w_{i_j} + (i_j - j) \cdot u$ for $j \in \mathbb{N}$, then $Aw'_j = b_{i_j} - j \cdot b$ and since $i_j \geq j$, we have $b_{i_j} - j \cdot b \in \mathbb{N}A$. The claim follows then from $\text{supp}(w'_j) \subseteq \text{supp}(u) \cup \text{supp}(w_{i_j}) = \text{supp}(u)$. \square

The edge-expansion of a graph is bounded from above by the number of edges that leave a certain set of nodes divided by the cardinality of this particular set. In a fiber graph, the edges come from a fixed set of moves and hence have a limited outreach. The next construction helps to find the nodes that are incident to outgoing edges:

Definition 2.4.10. Let $A \in \mathbb{Z}^{m \times d}$, $b \in \mathbb{N}A$, and $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$. For $u \in \mathbb{N}^d$, the u -boundary of $\mathcal{F}_{A,b}$ is $\partial_{\mathcal{M}}^u(\mathcal{F}_{A,b}) := \{v \in u + \mathcal{F}_{A,b} : \exists m \in \pm \mathcal{M} : v + m \in \mathbb{N}^d \setminus (u + \mathcal{F}_{A,b})\}$.




Figure 2.5: The sets $\partial_{\mathcal{M}_3}^{(3,0,0)^T}(\mathcal{F}_{A_3,3})$, $\partial_{\mathcal{M}_3 \cup 2 \cdot \mathcal{M}_3}^{(3,0,0)^T}(\mathcal{F}_{A_3,3})$, and $\partial_{\mathcal{M}_3}^{(1,1,1)^T}(\mathcal{F}_{A_3,3})$ in $\mathcal{F}_{A_3,6}$ (white points).

Figure 2.5 justifies that $\partial_{\mathcal{M}}^u(\mathcal{F}_{A,b})$ can indeed be regarded as boundary. With this, the number of outgoing edges in a translated fiber $u + \mathcal{F}_{A,b}$ within a larger fiber can be bounded from above:

Lemma 2.4.11. *Let $A \in \mathbb{Z}^{m \times d}$ and $b, b' \in \mathbb{N}A$ with $2|\mathcal{F}_{A,b}| \leq |\mathcal{F}_{A,b'+b}|$. Then for any finite set $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ and all $u \in \mathcal{F}_{A,b'}$,*

$$h(\mathcal{F}_{A,b'+b}(\mathcal{M})) \leq \frac{2|\mathcal{M}| \cdot |\partial_{\mathcal{M}}^u(\mathcal{F}_{A,b})|}{|\mathcal{F}_{A,b}|}.$$

Proof. Let $G := \mathcal{F}_{A,b'+b}(\mathcal{M})$, then $u + \mathcal{F}_{A,b} \subset V(G)$ and since $2|u + \mathcal{F}_{A,b}| = 2|\mathcal{F}_{A,b}| \leq |\mathcal{F}_{A,b'+b}|$, the following upper bound is immediate by the definition of the edge-expansion:

$$h(G) \leq \frac{|E_G(u + \mathcal{F}_{A,b})|}{|u + \mathcal{F}_{A,b}|}.$$

The edges leaving $u + \mathcal{F}_{A,b}$ in $\mathcal{F}_{A,b'+b}(\mathcal{M}) \subset \mathbb{N}^d$ are precisely those with endpoints in $\partial_{\mathcal{M}}^u(\mathcal{F}_{A,b})$. Every node of $\mathcal{F}_{A,b'+b}(\mathcal{M})$ has at most $|\pm \mathcal{M}|$ incident edges and hence $|E_G(u + \mathcal{F}_{A,b})|$ is bounded from above by $2|\mathcal{M}| \cdot |\partial_{\mathcal{M}}^u(\mathcal{F}_{A,b})|$. \square

The size of the entries in a Markov basis is crucial to determine the size of the boundary. The larger those entries are, the more nodes are in the boundary (Lemma 2.4.12) since more nodes in the shifted fiber $u + \mathcal{F}_{A,b}$ are adjacent to nodes outside of $u + \mathcal{F}_{A,b}$.

Lemma 2.4.12. *Let $A \in \mathbb{Z}^{m \times d}$, $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ a finite set, and $b \in \mathbb{N}A$. Then for all $u \in \mathbb{N}^d$,*

$$\partial_{\mathcal{M}}^u(\mathcal{F}_{A,b}) \subseteq u + \bigcup_{j \in \text{supp}(u)} \bigcup_{r=0}^{\|\mathcal{M}\|_{\infty}} \{w \in \mathcal{F}_{A,b} : w_j = r\}.$$

Proof. Let $v \in \partial_{\mathcal{M}}^u(\mathcal{F}_{A,b})$, then there is $m \in \pm\mathcal{M}$ such that $v + m \in \mathbb{N}^d$, but $v + m \notin u + \mathcal{F}_{A,b}$. Since $v \in u + \mathcal{F}_{A,b}$, there is $w \in \mathcal{F}_{A,b}$ such that $v = u + w$. The vector $w + m$ must have a negative entry, since otherwise $w + m \in \mathbb{N}^d$, that is $w + m \in \mathcal{F}_{A,b}$ which in turn implies $v + m = u + w + m \in u + \mathcal{F}_{A,b}$. Hence, there is $j \in [d]$ such that $(w + m)_j < 0$. Suppose $j \notin \text{supp}(u)$. Then $(u + w + m)_j = (w + m)_j < 0$, which contradicts $u + w + m = v + m \in \mathbb{N}^d$. Thus, $j \in \text{supp}(u)$ and $w_j < -m_j$. Since that means $w_r \leq \|\mathcal{M}\|_{\infty}$, the statement follows. \square

Lemma 2.4.11 allows to bound the edge-expansion by essentially comparing the growth of fibers with the growth of their boundary. The idea is to show that the boundary grows asymptotically slower than the fiber itself. Counting the number of integer points in a polytope is the subject of Ehrhart theory [47]. Let $\mathcal{P} \subset \mathbb{Q}^d$ be rational polytope and consider the map $L_{\mathcal{P}} : \mathbb{N} \rightarrow \mathbb{N}$ which counts the integer points in the i -th dilation $i\mathcal{P}$, i.e. $L_{\mathcal{P}}(i) := |i\mathcal{P} \cap \mathbb{Z}^d|$. Ehrhart's theorem (cf. [15, Theorem 3.23]) says that $L_{\mathcal{P}}$ is a *quasi-polynomial* of degree $r := \dim(\mathcal{P})$, that is there exist periodic maps $c_0, \dots, c_r : \mathbb{N} \rightarrow \mathbb{Z}$ with integral periods such that

$$L_{\mathcal{P}}(t) = c_r(t)t^r + c_{r-1}(t)t^{r-1} + \dots + c_0(t)$$

with c_r not identically zero. This applies to rays in affine semigroups: Since for any $i \in \mathbb{N}$, the integer points of $\mathcal{R}_{A,ib}$ are precisely the elements of $\mathcal{F}_{A,ib}$, $L_{\mathcal{R}_{A,b}}(i) = |\mathcal{F}_{A,ib}|$ for all $i \in \mathbb{N}$ and hence $|\mathcal{F}_{A,ib}|$ grows in i (quasi-)polynomial with degree $\dim(\mathcal{R}_{A,b})$.

Remark 2.4.13. For any integer matrix A and $b \in \mathbb{N}A$, $\dim(\mathcal{R}_{A,b}) \geq \dim(\mathcal{F}_{A,b})$. In particular, if $\dim(\mathcal{F}_{A,b}) > 0$, then $(|\mathcal{F}_{A,ib}|)_{i \in \mathbb{N}}$ is unbounded. If A is totally unimodular, then $\mathcal{R}_{A,b}$ equals $\text{conv}_{\mathbb{Q}}(\mathcal{F}_{A,b})$ and hence the dimensions of $\mathcal{R}_{A,b}$ and $\mathcal{F}_{A,b}$ coincide.

We count the lattice points in the boundary of a fiber with Lemma 2.4.12. However, the components of the set appearing there are not precisely dilates of polytopes and Ehrhart theory does not apply directly. Nevertheless, their growth can be bounded in terms of their dimension.

Lemma 2.4.14. *Let $A \in \mathbb{Z}^{m \times d}$, $b \in \mathbb{N}A$, and fix integers $j \in [d]$ and $l \in \mathbb{N}$. If for all $i \in \mathbb{N}_{>0}$, $\mathcal{R}_{A,ib}$ is not completely contained in the hyperplane $H := \{x \in \mathbb{Q}^d : x_j = l\}$, then there is $C \in \mathbb{N}$ such that the number of integer points in $\mathcal{R}_{A,ib} \cap H$ is bounded from above by $C \cdot i^{\dim(\mathcal{R}_{A,b})-1}$.*

Proof. Write $\mathcal{P} := \mathcal{R}_{A,b}$ and $r := \dim(\mathcal{P})$. For i large enough, the dimension of $(i\mathcal{P}) \cap H$ stabilizes, i.e. there are $r', N \in \mathbb{N}$ such that $r' := \dim(i\mathcal{P} \cap H)$ for all $i \geq N$. The affine space of $i\mathcal{P} \cap H$ is completely contained in H whereas the affine space of $i\mathcal{P}$ has elements outside of H . That implies $r' < r$. Let $A = (a_1, \dots, a_d)$ and A' be submatrix of A omitting the j -th column, then the (bijective) projection of $i\mathcal{P} \cap H$ onto all coordinates different from j is

$$Q_i := \{x \in \mathbb{Q}^{d-1} : A'x = i \cdot b - l \cdot a_j\}.$$

By [117, Proposition 1], there exists finitely many sets C_1, \dots, C_k covering \mathbb{N} such that for $i \in C_j$, the number of integer points in Q_i is a quasi-polynomial of degree r' . \square

Lemma 2.4.15. *Let $p(t) = \sum_{s=0}^r c_s(t)t^s$ be a quasi-polynomial with $r > 0$ and let $k \in \mathbb{N}$ such that $c_r(k) > 0$. There is $n \in \mathbb{N}_{>0}$ and $N \in \mathbb{N}$ such that for all $i \in (k + n \cdot \mathbb{N}) \cap \mathbb{N}_{\geq N}$, $2p(i) < p(i + ni)$.*

Proof. Let $n \geq 2$ such that $c_r(i + ni) = c_r(i)$ for all $i \in \mathbb{N}$ (i.e. if c_r is not a constant, let $n \geq 2$ be the period of c_r). For all $i \in k + n \cdot \mathbb{N}$, $c_r(i + ni) = c_r(i) = c_r(k) > 0$ and hence

$$p(i + ni) - 2 \cdot p(i) = c_r(k) ((1 + n)^r - 2) i^r + \sum_{s=0}^{r-1} (c_s(i + ni)(1 + n)^s - 2c_s(i)) i^s.$$

The sum in the term on the right-hand side is a quasi-polynomial of degree at most $r - 1$ and the left term on the right-hand side is a polynomial of degree $r > 0$ whose leading coefficient is positive due to $n \geq 2$ and $r > 0$. Thus, there is $N \in \mathbb{N}$ such that for all $i \in k + n \cdot \mathbb{N}$ with $i \geq N$,

$$c_r(k) ((1 + n)^r - 2) i^r > - \sum_{s=0}^{r-1} (c_s(i + in)(1 + n)^s - 2c_s(k)) i^s,$$

that is $2p(i) < p(i + ni)$. □

We are now ready to state and prove the main theorem of this section.

Theorem 2.4.16. *Let $A \in \mathbb{Z}^{m \times d}$, let $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis for A , and let $(b_i)_{i \in \mathbb{N}}$ be dominated in $\mathbb{N}A$. Then there exist $C, C' \in \mathbb{N}_{>0}$ such that $h(\mathcal{F}_{A, b_i}(\mathcal{M})) \leq \frac{C}{i}$ for all $i \in C' \cdot \mathbb{N}_{\geq 1}$.*

Proof. Since $(b_i)_{i \in \mathbb{N}}$ is dominated, there exists $b \in \mathbb{N}A$ such that $b'_i := b_i - i \cdot b \in \mathbb{N}A$ for all $i \in \mathbb{N}$. Moreover, there is $u \in \mathcal{F}_{A, b}$ and a sequence $(w_i)_{i \in \mathbb{N}}$ in \mathbb{N}^d such that for all $i \in \mathbb{N}$, $w_i \in \mathcal{F}_{A, b'_i}$ and $\text{supp}(w_i) \subseteq \text{supp}(u)$. By Lemma 2.4.9, the subsequence $(b_{(\|\mathcal{M}\|_{\infty} + 1)i})_{i \in \mathbb{N}}$ is dominated as well. Due to the linear re-parametrization, the statement on the edge-expansion is true for the sequence $(b_i)_{i \in \mathbb{N}}$ if it is true for this particular subsequence. Thus, we replace b_i with $b_{(\|\mathcal{M}\|_{\infty} + 1)i}$, b with $(\|\mathcal{M}\|_{\infty} + 1) \cdot b$, and b'_i with $b'_{(\|\mathcal{M}\|_{\infty} + 1)i}$. Additionally, we replace w_i with $w_{(\|\mathcal{M}\|_{\infty} + 1)i}$ and u with $(\|\mathcal{M}\|_{\infty} + 1) \cdot u$, which does not change the support of u . After these changes, we have $u_i > \|\mathcal{M}\|_{\infty}$ for all $i \in \text{supp}(u)$, which is needed later in the proof. The Ehrhart quasi-polynomial $L_{\mathcal{R}_{A, b}}$ has degree $r := \dim(\mathcal{R}_{A, b})$ and by the definition of being dominated, $r > 0$. Write $L_{\mathcal{R}_{A, b}}(i) = \sum_{s=0}^r c_s(i) i^s$ with c_r not identically zero. Since $L_{\mathcal{R}_{A, b}}(i) = |\mathcal{F}_{A, ib}| > 0$, there exists $k \in \mathbb{N}$ such that $c_r(k) > 0$. By Lemma 2.4.15, there exists $n \in \mathbb{N}_{>0}$ and $N \in \mathbb{N}$ such that $2|\mathcal{F}_{A, ib}| \leq |\mathcal{F}_{A, (i+ni)b}|$ for all $i \in (k + n \cdot \mathbb{N}) \cap \mathbb{N}_{\geq N} =: \mathcal{I}$. By the choice of w_i and u , $A \cdot (w_{i+ni} + ni \cdot u) = b'_{i+ni} + ni \cdot b = b_{i+ni} - ib$ for all $i \in \mathcal{I}$ and hence $w_{i+ni} + ni \cdot u + \mathcal{F}_{A, ib} \subsetneq \mathcal{F}_{A, b_{i+ni}}$. In particular, for any $i \in \mathcal{I}$

$$2|\mathcal{F}_{A, ib}| \leq |\mathcal{F}_{A, (i+ni)b}| = |w_{i+ni} + \mathcal{F}_{A, (i+ni)b}| \leq |\mathcal{F}_{A, b_{i+ni}}|.$$

For any $i \in \mathcal{I}$, set $u_i := w_{i+ni} + ni \cdot u$, then Lemma 2.4.12 gives

$$|\partial_{\mathcal{M}}^{u_i}(\mathcal{F}_{A, ib})| \leq \sum_{j \in \text{supp}(u_i)} \sum_{l=0}^{\|\mathcal{M}\|_{\infty}} |\{v \in \mathcal{F}_{A, ib} : v_j = l\}|. \quad (2.4.2)$$

Since $2|\mathcal{F}_{A, ib}| \leq |\mathcal{F}_{A, b_{i+ni}}|$ and $Au_i = b'_{i+ni} + ni \cdot b$ for all $i \in \mathcal{I}$, an application of Lemma 2.4.11 yields the upper bound on the edge-expansion of the graph $\mathcal{F}_{A, b_{i+ni}}(\mathcal{M})$:

$$h(\mathcal{F}_{A, b_{i+ni}}(\mathcal{M})) \leq \frac{2|\mathcal{M}| \cdot |\partial_{\mathcal{M}}^{u_i}(\mathcal{F}_{A, ib})|}{|\mathcal{F}_{A, ib}|} \leq 2|\mathcal{M}| \cdot \frac{\sum_{j \in \text{supp}(u)} \sum_{l=0}^{\|\mathcal{M}\|_{\infty}} |\{v \in \mathcal{F}_{A, ib} : v_j = l\}|}{|\mathcal{F}_{A, ib}|},$$

where (2.4.2) and $\text{supp}(u_i) = \text{supp}(w_{i+ni}) \cup \text{supp}(u) \subseteq \text{supp}(u)$ was used in the first and second inequality respectively. For any $j \in \text{supp}(u)$ and $l \in \{0, \dots, \|\mathcal{M}\|_\infty\}$, consider the hyperplane $H_{j,l} = \{x \in \mathbb{Q}^d : x_j = l\}$ in \mathbb{Q}^d , then for all $i \in \mathcal{I}$, the number of integer points in $(i \cdot \mathcal{R}_{A,b}) \cap H_{j,l}$ is precisely $L_{j,l}(i) := |\{w \in \mathcal{F}_{A,ib} : w_j = l\}|$. Since $u_j > \|\mathcal{M}\|_\infty$ for all $j \in \text{supp}(u)$, the vector $i \cdot u \in \mathcal{R}_{A,ib} = i \cdot \mathcal{R}_{A,b}$ is not contained in $(i \cdot \mathcal{R}_{A,b}) \cap H_{j,l}$ for all $0 \leq l \leq \|\mathcal{M}\|_\infty$ and all $i \in \mathcal{I}$. Lemma 2.4.14 then implies that for all $j \in \text{supp}(u)$ and $l \in \{0, \dots, \|\mathcal{M}\|_\infty\}$, there is a constant $D_{j,l} \in \mathbb{N}$ such that $L_{j,l}(i) \leq D_{j,l} \cdot i^{r-1}$ for all $i \in \mathbb{N}$. Let $D \in \mathbb{N}$ be the maximum of all $D_{j,l}$, then

$$\begin{aligned} h(\mathcal{F}_{A,b_{i+ni}}(\mathcal{M})) &\leq 2|\mathcal{M}| \cdot \frac{\sum_{j \in \text{supp}(u)} \sum_{l=0}^{\|\mathcal{M}\|_\infty} L_{j,l}(i)}{L_{\mathcal{R}_{A,b}}(i)} \\ &= 2|\mathcal{M}| \cdot \frac{|\text{supp}(u)| \cdot (\|\mathcal{M}\|_\infty + 1) \cdot D \cdot i^{r-1}}{c_r(k)i^r + \sum_{s=0}^{r-1} c_s(i)i^s} \end{aligned}$$

for all $i \in \mathcal{I}$. Hence, there exists a constant $C \in \mathbb{Q}_{>0}$ such that $h(\mathcal{F}_{A,b_j}(\mathcal{M})) \leq \frac{C}{j}$ for $j \in (n+1) \cdot \mathcal{I}$. The existence of C' follows immediately from the construction of $(n+1) \cdot \mathcal{I}$. \square

Remark 2.4.17. The constant C' in Theorem 2.4.16 is due to the many boundary effects and the fluctuations in Ehrhart quasi-polynomials. For instance when $(b_i)_{i \in \mathbb{N}}$ is a ray instead of a dominated sequence and when A is totally unimodular, the set $C' \cdot \mathbb{N}$ is a set of the form $\mathbb{N}_{\geq C''}$. However, our pragmatic bounding of the index set suffices to disprove rapid mixing in Chapter 4.

As a consequence, whenever the distance of a sequence to the facets of the semigroup becomes arbitrarily large, the edge-expansion of a subsequence converges to zero.

Corollary 2.4.18. *Let $A \in \mathbb{Z}^{m \times d}$ have non-trivial kernel and $(b_i)_{i \in \mathbb{N}}$ a sequence in $\mathbb{N}A$ with $\limsup_{i \in \mathbb{N}} H_A(b_i) = \infty$, then $\liminf_{i \in \mathbb{N}} h(\mathcal{F}_{A,b_i}(\mathcal{M})) = 0$ for any finite set $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$.*

Proof. This is Proposition 2.4.7 and Theorem 2.4.16. \square

Remark 2.4.19. Let $G = (V, E)$ be a graph with maximal degree d . It is well-known that the diameter and the edge-expansion of G satisfy the inequality

$$\text{diam}(G) \leq \frac{\log |V|}{\log \left(1 + \frac{h(G)}{d}\right)},$$

see for instance [56, Proposition 1.30]. Let $A \in \mathbb{Z}^{m \times d}$, $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ a Markov basis for A , and $b \in \mathbb{N}A$. Together with the lower bound on the diameter from Corollary 2.2.11, we have

$$h(\mathcal{F}_{A,i \cdot b}(\mathcal{M})) \leq |\pm \mathcal{M}| \left(\exp \left(\frac{\log |\mathcal{F}_{A,i \cdot b}|}{C' \cdot i} \right) - 1 \right)$$

for all $i \in \mathbb{N}$. This upper bound on the edge-expansion, however, cannot be used in general to show the $\mathcal{O}(\frac{1}{i})_{i \in \mathbb{N}}$ bound from Theorem 2.4.16.

3 The fiber dimension of a graph

The study of geometric properties of graphs is a key ingredient in understanding their algorithmic behaviour and combinatorial structure [114, 84]. In [52], the dimension of a graph was introduced as the smallest $d \in \mathbb{N}$ such that the graph can be embedded in \mathbb{R}^d with every edge having unit length. Recently, isometric embeddings of graphs into discrete objects like hypercubes or lattices received a lot of attention and led to many new notions of graph dimension, like the isometric dimension [54], the lattice dimension [51, 68], or the Fibonacci dimension [21] of a graph.

The goal of this chapter is to study embedding of graphs onto fiber graphs. As every graph G can be represented as a fiber graph (Proposition 3.1.1), this motivates the question for the smallest dimension in which such a representation exists, the *fiber dimension* of G (Definition 3.1.5). First, we explore general properties of this dimension and state an upper bound in terms of the chromatic number (Theorem 3.1.13) in the spirit of [52]. We then determine the fiber dimension for a variety of graphs. The fiber dimension of a circuit of length n depends on Euler's totient function and we show that it equals one if and only if $n \in \mathbb{N} \setminus \{3, 4, 6\}$. Circuits whose length is one of the exceptional cases in $\{3, 4, 6\}$ have fiber dimension two (Proposition 3.2.7). We also determine the fiber dimension of complete graphs and show that it is logarithmic in the number of nodes (Theorem 3.3.4). In the end, a connection to *distinct pair-sum polytopes* [27] is established and it is shown how the fiber dimension leads to relations between the number of lattice points and the dimension of the ambient space of these polytopes. In Section 3.5, we give an algorithm that decides whether the fiber dimension of a graph is smaller or equal to two. Eventually, we discuss the obstacles that make the computation of the fiber dimension challenging in higher dimensions. All results of this chapter are based on the author's publication [120].

Convention. In this chapter, all graphs are simple.

3.1 Embeddings

The following proposition is the starting point of our investigation:

Proposition 3.1.1. *Every graph is isomorphic to a fiber graph.*

Proof. Let $G = (\{v_1, \dots, v_n\}, E)$ be a graph and let $\mathcal{F} := \{e_1, \dots, e_n\}$. Then \mathcal{F} is saturated since it is the set of lattice points of the $(n - 1)$ -dimensional simplex. Consider $\mathcal{M} := \{e_i - e_j : \{v_i, v_j\} \in E\}$, then $\mathcal{F}(\mathcal{M})$ is isomorphic to G . \square

Remark 3.1.2. Let G , \mathcal{F} , and \mathcal{M} as in the proof of Proposition 3.1.1 and consider the integer matrix $A = (1, \dots, 1) \in \mathbb{Z}^{1 \times n}$. If G is connected, then it is easy to show that \mathcal{M} is not only a Markov basis for $\mathcal{F}_{A,1}$, but also for $\mathcal{F}_{A,b}$ for any $b \in \mathbb{N}A$.

The restriction to graphs without loops is necessary since the canonical way to model loops in fiber graphs yields regular graphs only. It is not clear how to write the graph obtained from K_2 where only one node has a loop as fiber graph. Naively, Proposition 3.1.1 gives rise to the

following alternative definition of the fiber dimension of a graph G : define it as the smallest natural number $d \in \mathbb{N}$ such that there exists a saturated set $\mathcal{F} \subset \mathbb{Z}^d$ and a finite set $\mathcal{M} \subset \mathbb{Z}^d$ so that $G \cong \mathcal{F}(\mathcal{M})$. The complete graph K_n would then be isomorphic to the fiber graph on $[n]$ that uses $[n-1]$ as a set of moves. That is, the complete graph would have the same dimension as a path, that is one, despite the fact that $n-1$ moves are needed to represent the complete graph as fiber graph, whereas the path requires only one, namely $\{1\}$. To capture the structural information stored in the moves, we put more restrictions on the set of moves:

Definition 3.1.3. Let $\mathcal{M} \subset \mathbb{Z}^d$ be a finite set, then \mathcal{M} is a *set of directions* if $\mathcal{M} = -\mathcal{M}$ and if for all $\lambda \in \mathbb{N}$ with $\lambda \geq 2$ and all $m \in \mathcal{M}$, $\lambda \cdot m \notin \mathcal{M}$.

There is no particular mathematical reason why we require a set of directions to be symmetric, but it simplifies the counting and the construction of fiber graphs in subsequent sections. Observe that a set of directions cannot contain the zero vector. Recall that the *dimension* of a saturated set is the dimension of its convex hull. The next lemma states that every fiber graph can be written as a fiber graph on a full dimensional saturated set.

Lemma 3.1.4. For any d -dimensional saturated set $\mathcal{F} \subset \mathbb{Z}^m$ and any set of directions $\mathcal{M} \subset \mathbb{Z}^k$, there exists a full dimensional saturated set $\mathcal{F}' \subset \mathbb{Z}^d$ and a set of directions $\mathcal{M}' \subset \mathbb{Z}^d$ such that $\mathcal{F}(\mathcal{M}) \cong \mathcal{F}'(\mathcal{M}')$.

Proof. We can assume that $d < m$. The affine transformation given in Lemma 1.2.2 preserves the dimension and the transformed set of directions is still a set of directions. Thus, we can assume that $\mathcal{F} = \mathcal{F}_{A,b}$ for an integer matrix $A \in \mathbb{Z}^{n \times k}$ with $k \geq n$, $b \in \mathbb{Z}^n$, and $\dim(\ker_{\mathbb{Z}}(A)) = k - n \geq d$. We can add rows to A and b without changing the identity $\mathcal{F} = \mathcal{F}_{A,b}$ so that $\dim(\ker_{\mathbb{Z}} A) = k - n = d$. First, we transform A into its Hermite normal form, that is, we write $A = (H, 0) \cdot C$ for a unimodular matrix $C \in \mathbb{Z}^{k \times k}$ and a matrix $H \in \mathbb{Z}^{n \times n}$ of full rank. Let $H^{-1} \in \mathbb{Q}^{n \times n}$ and $C^{-1} \in \mathbb{Q}^{k \times k}$ be the inverse matrices of H and C respectively. Since C is unimodular, $C^{-1} \in \mathbb{Z}^{k \times k}$ and thus let $C_1 \in \mathbb{Z}^{k \times n}$ and $C_2 \in \mathbb{Z}^{k \times d}$ such that $C^{-1} = (C_1, C_2)$ and consider the affine map

$$\psi : \mathbb{Q}^d \rightarrow \mathbb{Q}^k, x \mapsto C^{-1} \begin{pmatrix} H^{-1}b \\ -x \end{pmatrix}.$$

Clearly, ψ is injective and it is straightforward to check that the image of the polytope $\mathcal{P}' := \{v \in \mathbb{Q}^d : C_2 \cdot v \leq C_1 H^{-1}b\} \subset \mathbb{Q}^d$ is $\mathcal{P} := \{x \in \mathbb{Q}_{\geq 0}^k : Ax = b\}$. Since $\mathcal{F} = \mathcal{P} \cap \mathbb{Z}^k \neq \emptyset$, the matrix H satisfies $H^{-1}b \in \mathbb{Z}^n$ (see [31, Theorem 2.3.6]) and since C is unimodular, $\mathcal{F}' := \mathcal{P}' \cap \mathbb{Z}^d$ gets mapped to \mathcal{F} . In particular, $\dim(\mathcal{F}') = \dim(\mathcal{F}) = d$. That is, \mathcal{P}' is full dimensional in \mathbb{Q}^d . Now, consider the set of moves

$$\mathcal{M}' := \{\psi^{-1}(v) - \psi^{-1}(u) : v, u \in \mathcal{F}, v - u \in \mathcal{M}\},$$

then $\mathcal{M}' = -\mathcal{M}'$ and \mathcal{M}' cannot contain multiples. Let $\psi(v') = v$ and $\psi(u') = u$ for $v', u' \in \mathcal{F}'$, then $v' - u' \in \mathcal{M}'$ if and only if $v - u \in \mathcal{M}$. Thus, all edges in $\mathcal{F}'(\mathcal{M}')$ are mapped bijectively to edges in $\mathcal{F}(\mathcal{M})$ under ψ , which proves that these graphs are isomorphic. \square

The set of moves constructed in the proof of Proposition 3.1.1 is in fact a set of directions and combined with Lemma 3.1.4, the following definition is well.

Definition 3.1.5. The *fiber dimension* $\text{fdim}(G)$ of a graph G is the smallest $d \in \mathbb{N}$ such that there is a full dimensional saturated set $\mathcal{F} \subset \mathbb{Z}^d$ and a set of directions $\mathcal{M} \subset \mathbb{Z}^d$ with $G \cong \mathcal{F}(\mathcal{M})$.

Remark 3.1.6. In general, the fiber dimension of a graph G is different than its *dimension* $\dim(G)$ as defined in [52]. For example, the complete graph K_5 is realized as fiber graph in \mathbb{Q}^3 (see Theorem 3.3.4 and Figure 3.2), in contrast to $\dim(K_5) = 4$. The *Euclidean dimension* $\text{Edim}(G)$ of a graph G is the smallest $d \in \mathbb{N}$ such that G is isomorphic to a graph with nodes in \mathbb{R}^d where two nodes are adjacent if and only if they have unit distance. Clearly, $\text{Edim}(G) \geq \dim(G)$ for any graph G and hence the fiber dimension is different than the Euclidean dimension, too.

Remark 3.1.7. The generic embedding in Proposition 3.1.1 and Lemma 3.1.4 imply that the fiber dimension of any graph $G = (V, E)$ is at most $|V| - 1$. Trivially, the empty set is a set of directions, and hence all graphs with $|V| \leq 1$ satisfy $\text{fdim}(G) = 0$ and all graphs where $|V| \geq 2$ but $E = \emptyset$ have fiber dimension one (see also Section 3.2).

Remark 3.1.8. Let $\mathcal{F} \subset \mathbb{Z}^d$ be a saturated set and $\mathcal{M} \subset \mathbb{Z}^d$ be a set of directions. Proposition 1.2.7 shows that when \mathcal{M} is a Markov basis for \mathcal{F} such that $2 \cdot \dim(\mathcal{F}) = |\mathcal{M}|$, then $\mathcal{F}(\mathcal{M})$ is bipartite. Needless to say, not all bipartite fiber graphs satisfy $2 \cdot \dim(\mathcal{F}) = |\mathcal{M}|$, as the embedding of the 8-circuit in \mathbb{Q}^1 with a set of four directions shows (Proposition 3.2.7). In general, any symmetric Markov basis $\mathcal{M} \subset \mathbb{Z}^d$ for a saturated set $\mathcal{F} \subset \mathbb{Z}^d$ satisfies $\frac{1}{2} \cdot |\mathcal{M}| \geq \dim(\mathbb{Q} \cdot \mathcal{M}) = \dim(\mathcal{F})$, and hence Proposition 1.2.7 yields a lower bound on the number of directions in an embedding of non-bipartite graphs: If G is a connected graph with $\chi(G) > 2$ and $\text{fdim}(G) = d$, then any fiber graph embedding in \mathbb{Z}^d needs strictly more than $2d$ directions.

We now explore upper bounds on the fiber dimension. Our first observation is that some graph products are compatible with the cartesian product of saturated sets in \mathbb{Z}^d . Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs. The *cartesian product* $G_1 \times G_2$ is the graph on $V_1 \times V_2$ where (v_1, v_2) is adjacent to (u_1, u_2) if either $v_1 = u_1$ and $\{v_2, u_2\} \in E_2$ or if $v_2 = u_2$ and $\{v_1, u_1\} \in E_1$. Another product is the *tensor product* $G_1 \otimes G_2$, which is the graph on $V_1 \times V_2$ where (v_1, v_2) is adjacent to (u_1, u_2) if $\{v_1, u_1\} \in E_1$ and $\{v_2, u_2\} \in E_2$.

Proposition 3.1.9. *Let G_1, \dots, G_n be graphs, then $\text{fdim}(\times_{i=1}^n G_i) \leq \sum_{i=1}^n \text{fdim}(G_i)$.*

Proof. It suffices to prove the inequality for $n = 2$. Let $\mathcal{F}_1, \mathcal{F}_2, \mathcal{M}_1, \mathcal{M}_2$ such that $G_i \cong \mathcal{F}_i(\mathcal{M}_i)$. The cartesian product $\mathcal{F} := \mathcal{F}_1 \times \mathcal{F}_2$ is saturated and has dimension $\dim(\mathcal{F}_1) + \dim(\mathcal{F}_2)$. Additionally, let $\mathcal{M} := \{(m, 0)^T : m \in \mathcal{M}_1\} \cup \{(0, m)^T : m \in \mathcal{M}_2\}$. It is straightforward to check that $\mathcal{F}(\mathcal{M}) = \mathcal{F}_1(\mathcal{M}_1) \times \mathcal{F}_2(\mathcal{M}_2)$. Hence, $\text{fdim}(G_1 \times G_2) \leq \dim(\mathcal{F})$. \square

Proposition 3.1.10. *Let G_1, \dots, G_n be graphs, then $\text{fdim}(\otimes_{i=1}^n G_i) \leq \sum_{i=1}^n \text{fdim}(G_i)$.*

Proof. This proof is similar to the proof of Proposition 3.1.9. With $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$ the set of directions for $\mathcal{F} = \mathcal{F}_1 \times \mathcal{F}_2$, the graph $\mathcal{F}(\mathcal{M})$ equals $\mathcal{F}_1(\mathcal{M}_1) \otimes \mathcal{F}_2(\mathcal{M}_2)$. \square

Remark 3.1.11. The inequality in Proposition 3.1.9 is sharp for $K_2 \times K_2 = C_4$ (see Theorem 3.3.4 and Proposition 3.2.7).

Proposition 3.1.12. *Let G be a graph and $v \in V(G)$, then $\text{fdim}(G) \leq \text{fdim}(G - v) + 1$.*

Proof. Write $V(G) = \{v_0, \dots, v_n\}$ with $v_0 = v$. Let $d := \text{fdim}(G - v)$ and let $\phi : G - v \rightarrow \mathcal{F}(\mathcal{M})$ be a graph isomorphism that embeds $G - v$ in dimension d for a saturated set $\mathcal{F} \subset \mathbb{Z}^d$ and a set of directions $\mathcal{M} \subset \mathbb{Z}^d$. Let $\mathcal{F}' := \{0, (1, \phi(v_1)), \dots, (1, \phi(v_n))\} \subset \mathbb{Z}^{1+d}$, then \mathcal{F}' is saturated and $\text{conv}_{\mathbb{Q}}(\mathcal{F}')$ has dimension $d + 1$. Let $N \subseteq \{v_1, \dots, v_n\}$ be the neighborhood of v in G and consider $\mathcal{M}' = \{(0, m) : m \in \mathcal{M}\} \cup \{\pm(1, \phi(v_i)) : v_i \in N\}$. Then $G \cong \mathcal{F}'(\mathcal{M}')$. \square

As in [52], we obtain an upper bound on the dimension in terms of the *chromatic number* $\chi(G)$ of the graph G , that is the smallest natural number $k \in \mathbb{N}$ such that G has a k -coloring of its nodes [40, Chapter 5]. Our approach works as follows: First, we construct sets of integer points which represent the color classes of the graph in such a way that we can freely assign moves within them. In a second step, we map the nodes of the graph on these sets and construct the set of directions accordingly.

Theorem 3.1.13. *Let G be a graph with a k -coloring in which r color classes have cardinality 1, then $\text{fdim}(G) \leq 2k - r - 1$.*

Proof. Write $V(G) = V_1 \cup \dots \cup V_k$ and set $n_i := |V_i|$. Define for $i \in [k]$

$$\mathcal{F}_i := \{(e_i, j \cdot e_i)^T \in \mathbb{N}^{2k} : j \in [n_i]\} \subset \mathbb{N}^{2k}$$

and let $\mathcal{F} := \cup_{i=1}^k \mathcal{F}_i$ and $\mathcal{P} := \text{conv}_{\mathbb{Q}}(\mathcal{F})$. To show that $\mathcal{P} \cap \mathbb{Z}^{2k} = \mathcal{F}$, let $u \in \mathcal{P} \cap \mathbb{Z}^{2k}$. Since every \mathcal{F}_i is saturated, there exists $w_i \in \mathcal{F}_i$ and $\lambda_1, \dots, \lambda_k \in \mathbb{Q}$ with $0 \leq \lambda_i \leq 1$ for all $i \in [k]$ and $\sum_{i=1}^k \lambda_i = 1$ such that $u = \sum_{i=1}^k \lambda_i w_i$. The projection of \mathcal{F} onto the first k coordinates is the set of integer points of the standard simplex and thus saturated. The projection of u onto its first k coordinates is e_i for some $i \in [k]$. The only way to build an integer vector u is thus $\lambda_j = 0$ for $j \neq i$ and $\lambda_i = 1$, i.e. $u = w_i \in \mathcal{F}_i$.

Let us now construct a graph on $\mathcal{F} = \mathcal{P} \cap \mathbb{Z}^{2k}$ which is isomorphic to G . For that, let $\phi : \cup_{i=1}^k V_i \rightarrow \cup_{i=1}^k \mathcal{F}_i$ be any bijection which maps elements from V_i to \mathcal{F}_i and consider the set of directions $\mathcal{M} = \{\phi(v) - \phi(w) : \{v, w\} \in E(G)\}$. By construction of \mathcal{M} , ϕ is a graph homomorphism from G to $\mathcal{F}(\mathcal{M})$. Since edges in G do only connect nodes from different color classes, the first k coordinates of any element in \mathcal{M} do only contain elements from $\{-1, 0, 1\}$ and thus \mathcal{M} cannot contain multiples. Next, let $s \in [n_i]$ and $t \in [n_j]$ such that $(e_i, se_i)^T - (e_j, te_j)^T = \phi(u) - \phi(v) \in \mathcal{M}$ with $v, w \in V(G)$. It follows immediately that $\phi(v) = (e_j, te_j)^T$ and hence $\phi(u) = (e_i, se_i)^T$. Thus, ϕ maps edges from G to $\mathcal{F}(\mathcal{M})$ bijectively and hence $\text{fdim}(G) \leq \dim(\mathcal{F})$. The vertices of the polytope \mathcal{P} are $\{(e_1, e_1)^T, (e_1, n_1 e_1)^T, \dots, (e_k, e_k)^T, (e_k, n_k e_k)^T\}$ and since $n_i = 1$ for r indices $i \in [k]$, \mathcal{P} has $2k - r$ vertices and thus $\dim(\mathcal{F}) = \dim(\mathcal{P}) \leq 2k - r - 1$. \square

Corollary 3.1.14. *For any graph G , $\text{fdim}(G) \leq 2 \cdot \chi(G) - 1$.*

Remark 3.1.15. The upper bound in Theorem 3.1.13 is sharp for K_3 (Theorem 3.3.4). By the Four color theorem [7], every planar graph can be written as a fiber graph in \mathbb{Z}^7 .

3.2 Fiber dimension one

There are more graphs of fiber dimension one than paths. To work with this class, we first specialize our definition of a fiber graph to the 1-dimensional case. Recall that a finite set of integers $\mathcal{D} \subset \mathbb{N}$ is *primitive* if no of its elements is a divisor of any other element.

Definition 3.2.1. Let $n \in \mathbb{N}_{\geq 1}$ and let $\mathcal{D} \subseteq [n - 1]$ a primitive set. The *difference graph* $\mathcal{G}_{\mathcal{D}}^n$ has nodes $[n]$ and two nodes i and j are adjacent if $|i - j| \in \mathcal{D}$.

Proposition 3.2.2. *A graph has fiber dimension 1 if and only if it is a difference graph.*

Remark 3.2.3. Primitive sets, finite and infinite ones, are attractive objects in number theory [5, 121, 2]. The number of primitive sets in $[n]$ is sequence A051026 in *OEIS* [107] and its first elements are shown in Table 3.1.

# nodes	3	4	5	6	7	8
# primitive sets	5	7	13	17	33	45
# graphs	3	4	7	11	16	28

Table 3.1: Number of different difference graphs on n nodes.

Remark 3.2.4. In a difference graph on $[n]$, the graph-degree of i and $n + 1 - i$ coincides. Let G be a difference graph on n nodes and $n - 1 \geq d_1 \geq d_2 \geq \dots \geq d_n \geq 0$ be the node degrees. When n is even, then $|\{i \in [n] : d_i = j\}|$ is even for any $j \in [n - 1]$. When n is odd, then all but one degree appear an even number of times in G . The d_i that appears an odd number of times must be even since $\sum_{i=1}^n d_i$ is even by the Handshaking lemma. However, not every graph with these properties is a difference graph (Example 3.2.5).

Example 3.2.5. Let G be the 3×3 grid graph, that is the cartesian product of two paths of length 3. The degree sequence of G is $d := (4, 3, 3, 3, 3, 2, 2, 2, 2)$ and it fulfills the second conditions from Remark 3.2.4. A complete enumeration of all possible degree sequences of difference graphs on 9 nodes yield that d cannot be realized as a difference graph. Thus, $\text{fdim}(G) > 1$. Proposition 3.1.9 gives the upper bound $\text{fdim}(G) \leq 2$ and hence equality.

Lemma 3.2.6. *Let $n \in \mathbb{N}$ with $n \geq 2$ and $\mathcal{D} \subseteq [n - 1]$. If $\mathcal{G}_{\mathcal{D}}^n$ is connected, then $\text{gcd}(\mathcal{D}) = 1$.*

Proof. Since $n \geq 2$, there exists a path between 1 and 2 in $\mathcal{G}_{\mathcal{D}}^n$. Let $d_1, \dots, d_k \in \mathcal{D}$ be the distinct integers that appear in that path and write $1 + \sum_{i=1}^k \lambda_i d_i = 2$ for $\lambda_1, \dots, \lambda_k \in \mathbb{Z} \setminus \{0\}$. Then $\text{gcd}(d_1, \dots, d_k)$ divides 1. \square

Proposition 3.2.7. *For any $n \in \mathbb{N}_{\geq 3}$,*

$$\text{fdim}(C_n) = \begin{cases} 1, & \text{if } n \notin \{3, 4, 6\} \\ 2, & \text{if } n \in \{3, 4, 6\} \end{cases}.$$

Proof. Let $n \geq 3$ with $n \in \mathbb{N} \setminus \{3, 4, 6\}$. We first show that there exists an integer $k \in \mathbb{N}$ with $2 \leq k < \frac{n}{2}$ such that $\text{gcd}(k, n) = 1$. Let $\phi : \mathbb{N} \rightarrow \mathbb{N}$ be Euler's totient function. Since $n \in \mathbb{N} \setminus \{3, 4, 6\}$ and $n \geq 3$, $\phi(n) \geq 4$ and we have for all $k \in [n]$, $\text{gcd}(k, n) = 1$ if and only if $\text{gcd}(n - k, k) = 1$. In particular, elements in $[n]$ that are coprime to n come in pairs $(k, n - k)$ with $k < n - k$. Thus, since $\phi(n) \geq 4$, there must exist $k \in [n]$ with $1 < k < \frac{n}{2}$ such that $\text{gcd}(k, n) = 1$. We now show that $\mathcal{G}_{\{k, n-k\}}^n$ is a circuit of length n . Clearly, $n - k$ is not a multiple of k since this would imply that n is a multiple of k as well which in turn would contradict $\text{gcd}(n, k) = 1$ since $k > 1$. Any node in $\mathcal{G}_{\{k, n-k\}}^n$ has degree 2 and hence it suffices to prove that this graph is connected. Since k and n are coprime, $\langle k + n\mathbb{Z} \rangle = \mathbb{Z}_n$. Now, take distinct $i, j \in [n]$, then there exists $s \in \mathbb{N}$ such that $j + n\mathbb{Z} = i + sk + n\mathbb{Z}$ in \mathbb{Z}_n . For any $r \in [s]$, let $i_r \in [n]$ such that $i_r + n\mathbb{Z} = i + rk + n\mathbb{Z}$. Either $i_r - (n - k)$ or $i_r + k$ (but not both) are in $[n]$ and since their congruence classes in \mathbb{Z}_n coincide, i_{r-1} and i_r are adjacent in $\mathcal{G}_{\{k, n-k\}}^n$. Since $i_s = j$, i and j are connected. It follows that $C_n = \mathcal{G}_{\{k, n-k\}}^n$.

Let now $n \in \{3, 4, 6\}$. Removing one node from C_n yields a path and hence $\text{fdim}(C_n) \leq 2$ due to Proposition 3.1.12. It suffices to show that $\text{fdim}(C_n) > 1$ for $n \in \{3, 4, 6\}$. If $n = 3$, then $C \cong K_3$ and the claim follows from Theorem 3.3.4 proven in the next section. If $n \in \{4, 6\}$, assume that there exists $\mathcal{D} \subseteq [n - 1]$ such that $C_n \cong \mathcal{G}_{\mathcal{D}}^n$ is a difference graph. This already implies $|\mathcal{D}| = 2$ since if $|\mathcal{D}| \geq 3$ or $|\mathcal{D}| = 1$, the node 1 has either degree greater than 3 or is a

leaf respectively. Thus, we can write $\mathcal{D} = \{d_1, d_2\}$. Since $\mathcal{G}_{\mathcal{D}}^n$ is connected, $\gcd(d_1, d_2) = 1$ by Lemma 3.2.6. Hence, the only possible choices for $\{d_1, d_2\}$ to be primitive are $\{2, 3\}$ if $n = 4$ and $\{2, 3\}$, $\{2, 5\}$, $\{3, 4\}$, $\{3, 5\}$, and $\{4, 5\}$ if $n = 6$. In none of these cases, $\mathcal{G}_{\{d_1, d_2\}}^n$ is a circuit. \square



Figure 3.1: The difference graphs $\mathcal{G}_{\{2,3\}}^5 = C_5$ and $\mathcal{G}_{\{3,7\}}^{10} = C_{10}$.

3.3 Complete graphs

We now turn our attention to complete graphs and we give fiber graph embeddings that yield upper bounds on their fiber dimension in the logarithm of the number of their nodes. Our first result concerns the *star graph*, which has fiber dimension two:

Proposition 3.3.1. *For any $n \in \mathbb{N}$, $\text{fdim}(K_{1,n}) = 2$ if $n \geq 3$ and $\text{fdim}(K_{1,n}) = 1$ if $n < 3$.*

Proof. The star graph $K_{1,n}$ is a path for $n < 3$ and hence has fiber dimension one. For $n \geq 3$, let $v \in V(K_{1,n})$ be the node with maximal degree n . Removing v from $K_{1,n}$ gives a graph on $n \geq 3$ nodes without edges, i.e. the fiber dimension of this graph is one by Remark 3.1.7. Proposition 3.1.12 says that $\text{fdim}(K_{1,n}) \leq 2$. Assume that $\text{fdim}(K_{1,n}) = 1$ and let $\mathcal{D} \subseteq [n]$ be a primitive set such that $K_{1,n} \cong \mathcal{G}_{\mathcal{D}}^{n+1}$. The graph isomorphism maps v to some $j \in [n+1]$ and since j must be adjacent to all nodes in $[n+1] \setminus \{j\}$, we have $1 \in \mathcal{D}$. Primitiveness of \mathcal{D} gives $\mathcal{D} = \{1\}$ and thus $\mathcal{G}_{\mathcal{D}}^{n+1}$ is a path. \square

Our next intermediate goal is the computation of $\text{fdim}(K_n)$. The following lemma is the key ingredient to prove that $\text{fdim}(K_n)$ is at least logarithmic in n .

Lemma 3.3.2. *Let $\mathcal{L} \subseteq \mathbb{Z}^d$ be a lattice of full rank and let $\mathcal{F} \subset \mathbb{Z}^d$ be a set such that for any distinct $v, w \in \mathcal{F}$, $v - w \notin \mathcal{L}$. Then $|\mathcal{F}| \leq |\mathbb{Z}^d / \mathcal{L}|$.*

Proof. Write $\mathcal{F} = \{v_1, \dots, v_n\}$ and consider the linear map $\phi : \mathbb{Z}^d \rightarrow \mathbb{Z}^d / \mathcal{L}$, $\phi(v) = v + \mathcal{L}$. By assumption, $\phi(v_i - v_j) \neq 0$ in $\mathbb{Z}^d / \mathcal{L}$ for all $i, j \in [n]$ with $i \neq j$. Assume that there are $i, j \in [n-1]$ with $i \neq j$ such that $\phi(v_n - v_i) = \phi(v_n - v_j)$. Then $\phi(v_i - v_j) = \phi(v_i - v_n + v_n - v_j) = \phi(v_i - v_n) - \phi(v_j - v_n) = 0$, a contradiction. Thus, $\phi(v_n - v_i) \neq \phi(v_n - v_j)$ for all distinct $i, j \in [n-1]$. That is, $|\{\phi(v_n - v_i) : i \in [n-1]\}| = n-1$. The proposition follows from $n-1 = |\{\phi(v_n - v_i) : i \in [n-1]\}| \leq |\mathbb{Z}^d / \mathcal{L}| - 1$. \square

Remark 3.3.3. Since $|\mathbb{Z}^n / \mathcal{L}| = \det(\mathcal{L})$, Lemma 3.3.2 can be seen as a discrete analogue of Blichfeldt's theorem [31, Theorem 2.4.1].

Theorem 3.3.4. *For any $n \in \mathbb{N}$, $\text{fdim}(K_n) = \lceil \log_2 n \rceil$.*

Proof. The complete graph is embedded trivially into $\{0, 1\}^m$ with $\log_2 n \leq m$ and hence its fiber dimension is at most $\lceil \log_2 n \rceil$. Conversely, let $d := \text{fdim}(K_n)$ and $\mathcal{F} \subset \mathbb{Z}^d$ be a d -dimensional saturated set and $\mathcal{M} \subset \mathbb{Z}^d$ a set of directions such that $K_n \cong \mathcal{F}(\mathcal{M})$. Assume there are $v, w \in \mathcal{F}$ such that $v - w \in 2 \cdot \mathbb{Z}^d$. Since $(v+w)_i$ is even for all $i \in [d]$, $v+w \in 2\mathbb{Z}^d$. In particular, $\frac{1}{2}(v+w) \in \mathbb{Z}^d$ and since \mathcal{F} is saturated, $\frac{1}{2}(v+w) \in \mathcal{F}$. This implies that $v-w \in \mathcal{M}$ and $\frac{1}{2}(v-w) \in \mathcal{M}$. Thus, $v-w \notin 2\mathbb{Z}^d$. Due to Lemma 3.3.2, $n = |\mathcal{F}| \leq 2^d$ and thus $d \geq \lceil \log_2 n \rceil$. \square




Figure 3.2: Fiber graph embeddings of $K_5, K_6,$ and K_7 in \mathbb{Q}^3 .

When we replace a node v in a graph G by a set of nodes W such that each node $w \in W$ inherits the neighbors of v , then the fiber dimension of G increases at most in $\log |W|$:

Proposition 3.3.5. *Let H be a graph on $\{v_1, \dots, v_r\}$ and let G be the graph obtained from H where every node v_i is replaced by a finite set W_i such that any node in W_i is adjacent to any node in W_j if and only if $\{v_i, v_j\} \in E(H)$. Then*

$$\text{fdim}(G) \leq \text{fdim}(H) + \lceil \log_2 \max\{|W_1|, \dots, |W_r|\} \rceil.$$

Proof. Set $n_i := |W_i|$ for $i \in [r]$, $m := \lceil \log_2 \max\{n_i : i \in [r]\} \rceil$, and $s := \text{fdim}(H)$. We prove the upper bound by writing G as a fiber graph in \mathbb{Z}^{s+m} . First, let $\mathcal{F} = \{w_1, \dots, w_r\} \subset \mathbb{Z}^s$ be a saturated set and $\mathcal{M} \subset \mathbb{Z}^m$ a set of directions such that $H \cong \mathcal{F}(\mathcal{M})$ and such that v_i gets mapped to w_i by the isomorphism. For any $i \in [r]$, choose an arbitrary set $\mathcal{F}_i \subseteq \{0, 1\}^m$ of size n_i . This is possible since $n_i \leq 2^m$. The set

$$\mathcal{F}' := \cup_{i=1}^r \{w_i\} \times \mathcal{F}_i \subseteq \{0, 1\}^{s+m}$$

has cardinality $\sum_{i=1}^r n_i$ and is saturated since all subsets of $\{0, 1\}^{s+m}$ are. Choose a bijective map $\phi : \cup_{i=1}^r W_i \rightarrow \mathcal{F}'$ that maps nodes from W_i to \mathcal{F}_i and let $\mathcal{M}' := \mathcal{M} \times \{-1, 0, 1\}^m$. Note that since \mathcal{M} is a set of directions, \mathcal{M}' is as well. For distinct $i, j \in [r]$, all elements of $\{w_i\} \times \mathcal{F}_i$ are adjacent to all elements of $\{w_j\} \times \mathcal{F}_j$ whenever w_i and w_j are adjacent in $\mathcal{F}(\mathcal{M}) \cong H$. Moreover, since $0 \notin \mathcal{M}$, there are no edges within the sets $\{w_i\} \times \mathcal{F}_i$ and thus $\mathcal{F}'(\mathcal{M}') \cong G$. \square

Corollary 3.3.6. *For any $n_1, \dots, n_r \in \mathbb{N}$, $\text{fdim}(K_{n_1, \dots, n_r}) \leq \lceil \log_2 r \rceil + \lceil \log_2 \max\{n_i : i \in [r]\} \rceil$.*

Proof. This is a direct consequence of Proposition 3.3.5 with $H = K_r$ and Theorem 3.3.4. \square

3.4 Distinct pair-sum polytopes

For the remainder, we investigate a universal upper bound on the fiber dimension by generalizing the simplex embedding in Proposition 3.1.1. A priori, a move in a set of directions give rise to distinguished edges in a fiber graph. This is different for fiber graphs on saturated sets whose convex hull satisfy the following property:

Definition 3.4.1. A lattice polytope $\mathcal{P} \subset \mathbb{Q}^d$ with $n := |\mathcal{P} \cap \mathbb{Z}^d|$ is a *distinct pair-sum polytope* if $|\mathcal{P} \cap \mathbb{Z}^d + \mathcal{P} \cap \mathbb{Z}^d| = \binom{n}{2} + n$.

Let $\mathcal{P} \subset \mathbb{Q}^d$ be a distinct pair-sum polytope and write $\mathcal{P} \cap \mathbb{Z}^d = \{v_1, \dots, v_n\}$. Distinct pair-sum polytopes are well-studied objects in discrete geometry [27, 30, 17] and their name comes from the property that all the possible sums $2v_1, \dots, 2v_n, v_1 + v_2, v_1 + v_3, \dots, v_{n-1} + v_n$ are distinct. The next proposition is the reason why they are interesting from a fiber graph perspective:

Proposition 3.4.2. *Let $\mathcal{P} \subset \mathbb{Q}^d$ be a distinct pair-sum polytope and $\mathcal{F} := \mathcal{P} \cap \mathbb{Z}^d$. For any graph G on $|\mathcal{F}|$ nodes, there exists a set of directions $\mathcal{M} \subset \mathbb{Z}^d$ such that $G \cong \mathcal{F}(\mathcal{M})$.*

Proof. Let $n := |\mathcal{F}|$, pick an arbitrary bijection $\phi : V(G) \rightarrow \mathcal{F}$, and define

$$\mathcal{M} := \{\phi(u) - \phi(v) : u \text{ and } v \text{ adjacent in } G\}.$$

First, we show that \mathcal{M} does not contain multiples. Assume, there are $m, m' \in \mathcal{M}$ and $\lambda \in \mathbb{N}$ with $\lambda \geq 2$ such that $m = \lambda \cdot m'$. Let $v, w \in \mathcal{F}$ with $v - w = m$, then $w + \lambda \cdot m' = v$. The fact that $w, w + m', w + 2m' \in \mathcal{F}$ are distinct elements that fulfill $(w + m') + (w + m') = w + (w + 2m')$ contradicts that \mathcal{P} is a distinct pair-sum polytope and hence \mathcal{M} is a set of directions. We claim that $G \cong \mathcal{F}(\mathcal{M})$. Clearly, every edge in G is mapped to an edge in $\mathcal{F}(\mathcal{M})$. Conversely, let $v, w \in \mathcal{F}$ such that $v - w \in \mathcal{M}$. Then there exists adjacent nodes $v', w' \in V(G)$ with $\phi(v') - \phi(w') = v - w$. We have to prove that $\phi(v') = v$ and $\phi(w') = w$. If not, then $\phi(v') + w = \phi(w') + v$ implies that two different sums yield the same element in $\mathcal{F} + \mathcal{F}$, which again gives a contraction. \square

It may be of interest whether the reverse of Proposition 3.4.2 is true, that is whether all lattice polytopes with n lattice points that admit an embedding of all graphs on n nodes are distinct pair-sum polytopes. In [27], a distinct pair-sum polytope in \mathbb{Q}^n on 2^n lattice points was constructed for any $n \in \mathbb{N}$. This gives rise to the following result.

Proposition 3.4.3. *Let G be a graph on 2^n nodes, then $\text{fdim}(G) \leq n$.*

Proof. This is [27, Theorem 3] together with Proposition 3.4.2. \square

Explicit classifications of distinct-pair sums in small dimensions are discussed in [30, 17]. Generally, lower bounds on the fiber dimension can be translated to relations between the number of lattice points and the dimension of the ambient space of distinct pair-sum polytopes. The next proposition demonstrates this for complete graphs and rediscovers a bound which was already proven in [27, Theorem 2].

Proposition 3.4.4. *Let $\mathcal{P} \subset \mathbb{Q}^d$ be a distinct pair-sum polytope, then $|\mathcal{P} \cap \mathbb{Z}^d| \leq 2^d$.*

Proof. Let $\mathcal{F} := \mathcal{P} \cap \mathbb{Z}^d$ and $n := |\mathcal{F}|$. According to Proposition 3.4.2, there exists a set of directions $\mathcal{M} \subset \mathbb{Z}^d$ such that $K_n \cong \mathcal{F}(\mathcal{M})$. By the definition of the fiber dimension and Theorem 3.3.4, $\lceil \log_2 n \rceil = \text{fdim}(K_n) \leq d$, i.e. $n \leq 2^d$. \square

Remark 3.4.5. For any $n \in \mathbb{N}$, there exists a distinct pair-sum polytope on n lattice points, namely the $(n - 1)$ -dimensional simplex $\text{conv}_{\mathbb{Q}}(0, e_1, \dots, e_{n-1}) \subset \mathbb{Q}^{n-1}$. Thus, for fixed $n \in \mathbb{N}$, we can ask for the smallest natural number $d \in \mathbb{N}$ such that there exists a distinct pair-sum polytope in \mathbb{Q}^d on n lattice points. Clearly, $d \leq n - 1$ and Proposition 3.4.4 on the other hand gives $\lceil \log_2 n \rceil \leq d$. Given such a minimal d for fixed n , Proposition 3.4.2 implies the fiber dimension of any graph on n nodes is bounded from above by d . However, the embedding of graphs into distinct pair-sum polytopes is far from optimal since, for instance, a path on n nodes has fiber dimension 1 and thus the distance to d is made arbitrarily large. For complete graphs on the other hand, this bound is best-possible. Finding further classes of graphs for which the bound induced by the embeddings into distinct pair-sum polytopes is asymptotically tight provides new insights on the structure of these polytopes in the spirit of Proposition 3.4.4.

3.5 Computational aspects

Is there an algorithm that decides whether $\text{fdim}(G) \leq d$ holds for a given graph G and a given natural number $d \in \mathbb{N}$? We do not have an answer to this question, but given that it is NP-hard to decide whether the Euclidean dimension (Remark 3.1.6) of a graph is less than two [103], it is very likely that, if the fiber dimension can be computed algorithmically, its computation is hard as well. In this section, we develop an algorithm that decides $\text{fdim}(G) \leq 2$ and state obstacles that prevent our method to work for higher dimensions. Clearly, the case $\text{fdim}(G) \leq 1$ is easy to decide since there is, up to translation, precisely one saturated set in \mathbb{Z} with n elements, namely $[n]$. Thus, deciding whether a given graph G on n nodes has fiber dimension one, i.e. is a difference graph, is done by enumerating all primitive sets $\mathcal{D} \subseteq [n-1]$ and comparing $\mathcal{G}_{\mathcal{D}}^n$ with G (see also Example 3.2.5). In discrete geometry, two higher dimensional saturated sets are identified with each other when they are isomorphic in the following sense:

Definition 3.5.1. Two saturated sets $\mathcal{F}, \mathcal{F}' \in \mathbb{Z}^d$ are *affinely isomorphic* if there is $\psi : \mathbb{Q}^d \rightarrow \mathbb{Q}^d$ bijective affine with $\psi(\mathcal{F}) = \mathcal{F}'$. They are *unimodularly isomorphic* if additionally $\psi(\mathbb{Z}^d) = \mathbb{Z}^d$.

A version of this definition for lattice polytopes is in [122]. Not only is the number of lattice points and the dimension preserved under affine isomorphism, it is also compatible with fiber graphs in the following sense:

Lemma 3.5.2. *Let $\mathcal{F}, \mathcal{F}' \in \mathbb{Z}^d$ be saturated and affinely isomorphic sets, then for any set of directions $\mathcal{M} \subset \mathbb{Z}^d$ there is a set of directions $\mathcal{M}' \subset \mathbb{Z}^d$ such that $\mathcal{F}(\mathcal{M}) \cong \mathcal{F}'(\mathcal{M}')$.*

Proof. Let $\psi : \mathbb{Q}^d \rightarrow \mathbb{Q}^d$ be the affine and bijective function that maps \mathcal{F} to \mathcal{F}' and let $A \in \mathbb{Q}^{d \times d}$ and $b \in \mathbb{Q}^d$ such that $\psi(x) = Ax + b$ for all $x \in \mathbb{Q}^d$. Then $\mathcal{F}' = \{Au + b : u \in \mathcal{F}\}$. Since $\text{rank}(A) = d$, the set $\mathcal{M}' = \{A \cdot m : m \in \mathcal{M}\}$ is a set of directions, provided $\mathcal{M} \subset \mathbb{Z}^d$ is one. It is then easy to show that $\mathcal{F}(\mathcal{M}) \cong \mathcal{F}'(\mathcal{M}')$. \square

By Lemma 3.5.2, one naive approach to decide whether a given graph G on n nodes has $\text{fdim}(G) \leq d$ is to enumerate all saturated sets in \mathbb{Z}^d with n elements up to unimodular isomorphism, to enumerate then for each of these saturated sets all possible sets of directions, and finally to check whether any of these fiber graphs is graph-isomorphic to G . Beside its computational effort, this method does not terminate already for $d = 3$: For $m \in \mathbb{N}$, the set

$$\mathcal{R}_m := \left\{ \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ m \end{bmatrix} \right\}$$

is the set of all lattice points of the m -th *Reeve tetrahedra* [102] and the volume of its convex hull can be made arbitrarily large when m varies. Since the volume is preserved under unimodular isomorphism, these sets cannot be pairwise unimodular isomorphic. It is hence impossible to write down all saturated set in \mathbb{Z}^3 with 4 elements up to unimodular isomorphism. For $d = 2$, however, this enumeration process works and one way to see it is by Pick's theorem, which says that the volume of the convex hull of any saturated set $\mathcal{F} \subset \mathbb{Z}^2$ is at most $|\mathcal{F}|$. Together with the next theorem by Lagarias and Ziegler, we can determine all saturated sets with a given volume:

Lemma 3.5.3. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be a saturated set and $v := \lceil \text{vol}(\text{conv}_{\mathbb{Q}}(\mathcal{F})) \rceil$, then \mathcal{F} is unimodularly isomorphic to a saturated set contained in $[d \cdot d! \cdot v]^d$.*

Proof. This is [80, Theorem 2] for the equivalent lattice polytope version of Definition 3.5.1. \square

Proposition 3.5.4. *There is an algorithm that decides $\text{fdim}(G) \leq 2$ for any graph G .*

Proof. Let G be a graph on n nodes. By Pick's theorem [96], the volume of $\text{conv}_{\mathbb{Q}}(\mathcal{F})$ for any saturated set $\mathcal{F} \subset \mathbb{Z}^2$ is bounded from above by $|\mathcal{F}|$. Thus, in dimension two, Lemma 3.5.3 says that any saturated set in \mathbb{Z}^2 with n elements is unimodular isomorphic to a saturated set that is contained in $[4 \cdot n]^2$. Since the set of graphs that can be embedded on a saturated set is preserved under unimodular isomorphism by Lemma 3.5.2, G satisfies $\text{fdim}(G) \leq 2$ if and only if there exists a set of directions $\mathcal{M} \subset \mathbb{Z}^2$ and $\mathcal{F} \subseteq [4 \cdot n]^2$ such that $G \cong \mathcal{F}(\mathcal{M})$. There are only finitely many saturated sets in $[4 \cdot n]^2$, and by going through all of them iteratively and by enumerating for all of them all possible sets of directions, we can explicitly enumerate all graphs on n nodes that have fiber dimension at most 2. Comparing then all fiber graphs obtained that way to G , this method either gives a fiber embedding of G and hence $\text{fdim}(G) \leq 2$, or fails to find such. In the latter case, this means that $\text{fdim}(G) > 2$. \square

The correctness of Proposition 3.5.4 stands and falls with Pick's theorem, which is false for $d \geq 3$ due to Reeve tetrahedra. One way out is to study weaker notions of isomorphism, that are still strong enough to make the equivalence classes on saturated sets in \mathbb{Z}^d with n elements finite. Although the Reeve tetrahedra are pairwise affinely isomorphic, this notion of isomorphism also does not suffice to make the equivalence classes finite: For $m \in \mathbb{N}$, the saturated sets $\mathcal{R}_m \cup \{-e_3\} \subset \mathbb{Z}^3$ have five lattice points, but are not pairwise affinely isomorphic [92]. The following equivalence relation is tailored for working with fiber dimensions:

Definition 3.5.5. For saturated $\mathcal{F}, \mathcal{F}' \subset \mathbb{Z}^d$, write $\mathcal{F} \prec \mathcal{F}'$ if there is for any set of directions $\mathcal{M} \subset \mathbb{Z}^d$ a set of directions $\mathcal{M}' \subset \mathbb{Z}^d$ with $\mathcal{F}(\mathcal{M}) \cong \mathcal{F}'(\mathcal{M}')$ and write $\mathcal{F} \sim \mathcal{F}'$ if $\mathcal{F} \prec \mathcal{F}' \prec \mathcal{F}$.

It is not hard to see that \sim is an equivalence relation and that it is coarser than affine isomorphism by Lemma 3.5.2. We do not know whether the equivalence classes of \sim on saturated sets in \mathbb{Z}^d with n elements are finite, but we think that a sufficiently good understanding of \sim is the key ingredient towards an algorithmic investigation of the fiber dimension.

4 Symmetric fiber walks

In this chapter, we discuss the mixing behaviour of fiber walks on sequences $(\mathcal{F}_{A,b_i}(\mathcal{M}))_{i \in \mathbb{N}}$ for a fixed matrix $A \in \mathbb{Z}^{m \times d}$, a fixed Markov basis $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$, and a sequence $(b_i)_{i \in \mathbb{N}}$ of right-hand sides. The main result is that the second largest eigenvalue modulus of the simple fiber walk is essentially bounded from below by $(1 - \frac{C}{i})_{i \in \mathbb{N}}$ when the sequence $(b_i)_{i \in \mathbb{N}}$ is dominated (Definition 2.4.5). This implies that the simple fiber walk cannot mix rapidly when $(b_i)_{i \in \mathbb{N}}$ has additionally a *meaningful parametrization* (Definition 4.1.3). We show a similar asymptotics of the second largest eigenvalue modulus for symmetric fiber walks. The conclusion we draw from these results is that an adaption of the Markov basis has to take place depending on the right-hand side $b \in \mathbb{Z}^m$. In particular, to obtain rapid mixing in this setting, the size of the Markov basis has to grow. However, adding more moves to a Markov basis on the one hand increases the rejection rate and on the other hand makes the local sampling process computationally more expensive. Thus, it is a fine line to find the proper number of moves to add to simultaneously keep sampling from the Markov basis cheap and to improve the mixing time of the fiber walk. In Section 4.2, we adapt any given Markov basis so that the underlying graph is the complete graph with additional loops and we show that this adaption yields an expander in fixed dimension under mild assumptions on the diameter of the fiber graph (Corollary 4.2.3). The results of this chapter, except Remark 4.1.8, are based on the author's work [119].

4.1 Fixed Markov bases

Despite the fact that the computation of Markov bases has received a lot of attention in the last decade, mixing results on fiber graphs are rare. The next statement is from [36] and it is probably one of the first mixing results on fiber graphs. It states that the mixing time of fiber walks in \mathbb{Z}^2 that use the unit vectors grows quadratically in the diameter of the fiber graph:

Theorem 4.1.1. *There exists constants $C_1, C_2, D_1, D_2 \in \mathbb{Q}_{>0}$ such that for every $\epsilon > 0$ and every saturated set $\mathcal{F} \subset \mathbb{Z}^2$ that has $\{e_1, e_2\}$ as Markov basis, the mixing time of the simple fiber walk \mathcal{S} on $\mathcal{F}(\{e_1, e_2\})$ satisfies*

$$\frac{\text{diam}(\mathcal{F}(\{e_1, e_2\}))^2}{C_1} \cdot \log\left(\frac{C_2}{\epsilon}\right) \leq \mathcal{T}_{\mathcal{S}}(\epsilon) \leq \frac{\text{diam}(\mathcal{F}(\{e_1, e_2\}))^2}{D_1} \cdot \log\left(\frac{D_2}{\epsilon}\right).$$

Proof. This is [36, Theorem 1.1] in the language of this thesis. □

Remark 4.1.2. The upper bound from Theorem 4.1.1 was generalized in [118] to irreducible fiber walks in \mathbb{Z}^d that use $\{e_1, \dots, e_d\}$. However, it is the lower bound that disproves rapid mixing in \mathbb{Z}^2 : Let $\mathcal{F} \subset \mathbb{Z}^2$ be a full dimensional saturated set and $\mathcal{F}_i := (i \cdot \text{conv}_{\mathbb{Q}}(\mathcal{F})) \cap \mathbb{Z}^2$ be the lattice points in the i -th dilation of its convex hull. Combining Theorem 4.1.1 with Corollary 2.2.11, the mixing time of the simple fiber walk on \mathcal{F}_i is in $\Omega(i^2 \cdot \log \epsilon^{-1})$ and hence cannot be bounded by a polynomial in $\log |\mathcal{F}_i|$ since $(|\mathcal{F}_i|)_{i \in \mathbb{N}} \in \mathcal{O}(i^2)$.

We show that a similar effect to Remark 4.1.2 occurs in higher dimension as well. To make use of our results from Chapter 2 and Theorem 2.4.16 in particular, a connecting piece between the sequence of right-hand sides $(b_i)_{i \in \mathbb{N}}$ in $\mathbb{N}A$ and its parameter $(i)_{i \in \mathbb{N}}$ is needed:

Definition 4.1.3. A sequence $(b_i)_{i \in \mathbb{N}}$ in $\mathbb{N}A$ has a *meaningful parametrization* if there exists a polynomial $q \in \mathbb{Q}[t]$ such that $|\mathcal{F}_{A,b_i}| \leq q(i)$ for all $i \in \mathbb{N}$.

Example 4.1.4. Let A_2 as in Example 1.3.3, then $|\mathcal{F}_{A_2,i}| = i + 1$. The sequence $(2^i)_{i \in \mathbb{N}}$ in $\mathbb{N}A_2 = \mathbb{N}$ is thus not meaningfully parametrized, whereas the sequence $(i)_{i \in \mathbb{N}}$ trivially is. The computation of the edge-expansion from Example 2.4.3 and Proposition 4.1.9 below show that the second largest eigenvalue modulus λ_i of the simple fiber walk on $\mathcal{F}_{A_2,i}(\mathcal{M}_2)$ satisfies $\lambda_i \geq 1 - \frac{1}{i}$. Since $\log |\mathcal{F}_{A_2,i}| = \log(i + 1)$, the simple fiber walk on $(\mathcal{F}_{A_2,i}(\mathcal{M}_2))_{i \in \mathbb{N}}$ cannot mix rapidly.

Proposition 4.1.5. Let $A \in \mathbb{Z}^{m \times d}$ and let $(b_i)_{i \in \mathbb{N}}$ be a sequence in $\mathbb{N}A$ with $(\|b_i\|)_{i \in \mathbb{N}} \in \mathcal{O}(i^r)_{i \in \mathbb{N}}$ for some $r \in \mathbb{N}$. Then $(b_i)_{i \in \mathbb{N}}$ has a meaningful parametrization.

Proof. Denote by $a_1, \dots, a_m \in \mathbb{Z}^d$ the rows of A . Since $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$, there exist coefficients $\lambda_1, \dots, \lambda_m \in \mathbb{Q}$ such that $w := \sum_{i=1}^m \lambda_i a_i \in \mathbb{Q}_{>0}^d$. In particular, for any $b \in \mathbb{N}A$ and for any $u \in \mathcal{F}_{A,b}$, we have $\|u\|_{\infty} \cdot \min_{i \in [d]} w_i \leq w^T u \leq m \cdot \|\lambda\|_{\infty} \cdot \|b\|_{\infty}$. Thus,

$$|\mathcal{F}_{A,b}| \leq \left(\frac{m \cdot \|\lambda\|_{\infty} \|b\|_{\infty}}{\min_{i \in [d]} w_i} \right)^d.$$

Hence, if $\|b_i\| \leq C \cdot i^r$ for all $i \in \mathbb{N}$, then $(b_i)_{i \in \mathbb{N}}$ has a meaningful parametrization. \square

Remark 4.1.6. By design, constraint matrices $A \in \mathbb{Z}^{m \times d}$ of log-linear models have the vector $(1, \dots, 1) \in \mathbb{Z}^d$ in their row space to ensure that all elements of $\mathcal{F}_{A,b}$ have the same $\|\cdot\|_1$ -norm, denoted by s_b in the following (compare also to [43] and Section 1.4). For all $u \in \mathcal{F}_{A,b}$, we have $\|b\|_1 = \|Au\|_1 \leq \|A\| \cdot \|u\|_1 = \|A\| \cdot s_b$ for any matrix norm $\|\cdot\|$ that is compatible with $\|\cdot\|_1$. Since the invariant s_b is precisely the sample size in goodness-of-fit tests (Example 1.4.6), a sequence of right-hand sides $(b_i)_{i \in \mathbb{N}}$ grows polynomially in i whenever the sample size does, and hence it has a meaningful parametrization by Proposition 4.1.5 in this case.

The next lemma is one of the lemmas whose statement is almost longer than its proof. It essentially says which types of lower bounds on the second largest eigenvalue modulus suffice so that meaningful parametrized sequences cannot mix rapidly.

Lemma 4.1.7. Let $A \in \mathbb{Z}^{m \times d}$, $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a finite set, $(b_i)_{i \in \mathbb{N}}$ a sequence in $\mathbb{N}A$ with a meaningful parametrization, and λ_i the second largest eigenvalue modulus of the simple fiber walk on $\mathcal{F}_{A,b_i}(\mathcal{M})$. If there exists an infinite subset $\mathcal{I} \subseteq \mathbb{N}$ and $q \in \mathbb{Q}[t]$ with $\lambda_i \geq 1 - \frac{\log q(i)}{i}$ for all $i \in \mathcal{I}$, then the simple fiber walk on $(\mathcal{F}_{A,b_i}(\mathcal{M}))_{i \in \mathbb{N}}$ is not rapidly mixing.

Proof. Assume that the simple fiber walk mixes rapidly. Then there exists a polynomial $p \in \mathbb{Q}_{\geq 0}[t]$ such that for all $i \in \mathcal{I}$,

$$1 - \frac{1}{p(\log |\mathcal{F}_{A,b_i}|)} \geq \lambda_i \geq 1 - \frac{\log q(i)}{i}.$$

This implies that for all $i \in \mathcal{I}$, $\frac{1}{i} \cdot p(\log |\mathcal{F}_{A,b_i}|) \cdot \log q(i) \geq 1$. However, since the parametrization is meaningful, there is $f \in \mathbb{Q}[t]$ such that $|\mathcal{F}_{A,b_i}| \leq f(i)$ and thus $p(\log |\mathcal{F}_{A,b_i}|) \leq p(\log f(i))$, which gives a contradiction since \mathcal{I} is unbounded in \mathbb{N} . \square

The next remark is joint work with Caprice Stanley and shows how the results on the diameter of fiber graphs (Section 2.2) can be used to disprove rapid mixing:

Remark 4.1.8. Let $(b_i)_{i \in \mathbb{N}}$ be a sequence such that there exists $b \in \mathbb{N}A$ with $b_i - i \cdot b \in \mathbb{N}A$ for all $i \in \mathbb{N}$ and let \mathcal{W}_i be any irreducible and aperiodic random walk on $\mathcal{F}_{A,b_i}(\mathcal{M})$ that has the uniform distribution as stationary distribution. Then for any $u \in \mathcal{F}_{A,b_i-i,b}$, $u + \mathcal{F}_{A,i,b} \subseteq \mathcal{F}_{A,b_i}$. Thus, similarly as in Lemma 2.2.1 and Proposition 2.2.10, there exists a constant $D \in \mathbb{N}$ such that $\text{diam}(\mathcal{F}_{A,b_i}(\mathcal{M})) \geq D \cdot i$ for all $i \in \mathbb{N}$. With [81, Section 7.1.2], the mixing times satisfies $\mathcal{T}_{\mathcal{W}_i}(0.25) \geq \frac{1}{2} \cdot \text{diam}(\mathcal{F}_{A,b_i}(\mathcal{M})) \geq \frac{1}{2} \cdot i \cdot D$. When $(b_i)_{i \in \mathbb{N}}$ has a meaningful parametrization, then the mixing time cannot be bounded polynomially in $\log |\mathcal{F}_{A,b_i}|$ and hence $(\mathcal{W}_i)_{i \in \mathbb{N}}$ is not rapidly mixing. Moreover, combined with Theorem 1.1.9 and since the stationary distribution is uniform, we also deduce a lower bound on $\lambda(\mathcal{W}_i)$ which is tailored for Lemma 4.1.7:

$$\lambda(\mathcal{W}_i) \geq 1 - \frac{2 \cdot \log(4 \cdot |\mathcal{F}_{A,b_i}|)}{i \cdot D}. \quad (4.1.1)$$

The next result strengthens the lower bound on the second largest eigenvalue modulus from (4.1.1) even further by using the following connection to the edge-expansion of the graph:

Proposition 4.1.9. *For any d -regular and connected graph G , $\lambda(\mathcal{S}_G) \geq 1 - \frac{2}{d} \cdot h(G)$.*

Proof. This is [66, Theorem 4.11], which states a lower bound for the second largest eigenvalue of A_G , that is valid for the second largest eigenvalue modulus of $\mathcal{S}_G = \frac{1}{d}A_G$. \square

The lower bound on the diameter from Section 2.2 yields an upper bound on the edge-expansion (Remark 2.4.19), which in turn gives a lower bound on the second largest eigenvalue modulus by Proposition 4.1.9. This detour, however, brings us to essentially the same lower bound on the second largest eigenvalue modulus as stated in (4.1.1). With our stronger bound on the edge-expansion from Theorem 2.4.16, we get rid of the $\log q(i)$ term:

Theorem 4.1.10. *Let $A \in \mathbb{Z}^{m \times d}$, $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a finite set, $(b_i)_{i \in \mathbb{N}}$ be a dominated sequence in $\mathbb{N}A$, and λ_i be the second largest eigenvalue modulus of the simple fiber walk on $(\mathcal{F}_{A,b_i}(\mathcal{M}))_{i \in \mathbb{N}}$. Then there exist constants $C, C' \in \mathbb{N}_{\geq 1}$ such that $\lambda_i \geq 1 - \frac{C}{i}$ for $i \in C' \cdot \mathbb{N}$.*

Proof. By Proposition 4.1.9, $\lambda_i \geq 1 - \frac{1}{|\mathcal{M}|} \cdot h(\mathcal{F}_{A,b_i}(\mathcal{M}))$ since the simple fiber walk is a random walk on a $|\pm \mathcal{M}|$ -regular graph and since adding loops does not change the edge-expansion. The theorem then follows from Theorem 2.4.16. \square

It follows immediately from Theorem 4.1.10 that the simple fiber walk on $(\mathcal{F}_{A,b_i}(\mathcal{M}))_{i \in \mathbb{N}}$ is no expander when $(b_i)_{i \in \mathbb{N}}$ is a dominated sequence and with Lemma 4.1.7, it cannot mix rapidly when $(b_i)_{i \in \mathbb{N}}$ has a meaningful parametrization additionally. The next corollary is a template for possible mitigations of the assumptions of Theorem 4.1.10:

Corollary 4.1.11. *Let $A \in \mathbb{Z}^{m \times d}$ and let $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis for A . Let $(b_i)_{i \in \mathbb{N}}$ be from $\mathbb{N}A$ with a meaningful parametrization and suppose there is $p \in \mathbb{Q}[t]$ with $p(\mathbb{N}) \subseteq \mathbb{N}$ such that $(b_{p(i)})_{i \in \mathbb{N}}$ is dominated. Then the simple fiber walk on $(\mathcal{F}_{A,b_i}(\mathcal{M}))_{i \in \mathbb{N}}$ is not rapidly mixing.*

Proof. Clearly, there exists $C \in \mathbb{N}_{>0}$ such that $p(C \cdot (i + 1)) > p(C \cdot i)$ for all i sufficiently large. Let $b'_i := b_{p(C \cdot i)}$, then $(b'_i)_{i \in \mathbb{N}}$ is a subsequence of $(b_i)_{i \in \mathbb{N}}$ and hence it suffices to show that the simple fiber walk on $(\mathcal{F}_{A,b'_i}(\mathcal{M}))_{i \in \mathbb{N}}$ is not rapidly mixing. Since $|\mathcal{F}_{A,b'_i}| = |\mathcal{F}_{A,b_{p(C \cdot i)}}| \leq q(p(C \cdot i))$ for a polynomial $q \in \mathbb{Q}[t]$, $(b'_i)_{i \in \mathbb{N}}$ has a meaningful parametrization. Since $(b_{p(i)})_{i \in \mathbb{N}}$ is dominated, the sequence $(b'_i)_{i \in \mathbb{N}}$ is as well by Lemma 2.4.9. Now, Theorem 4.1.10 implies that the simple fiber walk on $(\mathcal{F}_{A,b'_i}(\mathcal{M}))_{i \in \mathbb{N}}$ is not rapidly mixing. \square

Example 4.1.12. Let $A_{n,m}$ be the constraint matrix of the independence model (Example 1.3.2 and Example 1.4.7) and assume that $n \geq m$. By Remark 4.1.6, we obtain sequences $(b_i)_{i \in \mathbb{N}}$ with a meaningful parametrization whenever the sample size grows polynomial in i . If additionally, $b_i \geq \frac{s}{t} \cdot i \cdot (1, \dots, 1)^T$ for fixed $s, t \in \mathbb{N}$, then $b_{i \cdot t \cdot n} - i \cdot s \cdot (n, \dots, n, m, \dots, m)^T \in \mathbb{N}A_{n,m}$ (where n, \dots, n denotes the m column sums and m, \dots, m denotes the n row sums) and it follows that $(b_{i \cdot t \cdot n})_{i \in \mathbb{N}}$ is dominated since the fiber of $(n, \dots, n, m, \dots, m)^T$ contains an element with full support. Corollary 4.1.11 shows that the simple fiber walk on $(\mathcal{F}_{A_{n,m}, b_i}(\mathcal{M}_{n,m}))_{i \in \mathbb{N}}$ is not rapidly mixing. These assumptions hold for instance when $n = m$ and $b_i := (i, \dots, i) \in \mathbb{N}^{2n}$, even though the node-connectivity under the basic moves $\mathcal{M}_{n,n}$ is best-possible due to [97, Theorem 2.9].

Remark 4.1.13. Let $\mathcal{M} = \{m_1, \dots, m_r\} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis for A . Extending \mathcal{M} by adding a finite number of \mathbb{Z} -linear combinations $\sum_{i=1}^k \lambda_i m_i$ may improve the mixing behaviour in one particular fiber, but since the cardinality of the new set of moves is still finite, this cannot lead to rapid mixing asymptotically due to Theorem 4.1.10. This implies, that the Graver basis has the same asymptotic mixing behaviour as any other finite Markov basis.

The asymptotic behaviour of the second largest eigenvalue modulus as shown in Theorem 4.1.10 is not restricted to simple fiber walks. To prove it for symmetric fiber walks, we apply a common scheme of Markov chain theory: We study the mixing time of a random walk \mathcal{W}_1 by comparing it to the mixing time of a related random walk \mathcal{W}_2 [35, 45]. More precisely, we compare fiber walks with the simple fiber walk on a spectral level with the following lemma:

Lemma 4.1.14. *Let $G = (V, E)$ be a graph and let \mathcal{W}_1 and \mathcal{W}_2 be reversible, aperiodic and irreducible random walks on G with stationary distributions π_1 and π_2 respectively. Assume that all eigenvalues of \mathcal{W}_2 are non-negative and let $C, C' > 0$ such that $\pi_1(x)\mathcal{W}_1(x, y) \leq C\pi_2(x)\mathcal{W}_2(x, y)$ for all distinct $x, y \in V$ and $C'\pi_2(x) \leq \pi_1(x)$ for all $x \in V$, then*

$$1 - \lambda(\mathcal{W}_1) \leq \frac{C}{C'}(1 - \lambda(\mathcal{W}_2)).$$

Proof. This is [115, Lemma 2.5]. Notice that $\lambda(\cdot)$ denotes the spectral gap in [115]. □

Ideally, we let \mathcal{W}_2 be the simple fiber walk and compare it directly with any other fiber walk \mathcal{W}_1 . However, non-negativity of the eigenvalues of the simple fiber walk cannot be guaranteed in general. In [115], a version of Lemma 4.1.14 is given without the condition on the eigenvalues of \mathcal{W}_2 , but then the constant C must satisfy the additional inequalities $\pi_1(x)\mathcal{W}_1(x, x) \leq C\pi_2(x)\mathcal{W}_2(x, x)$ for all $x \in V$. Since it might happen that all the moves in a Markov basis can be applied on a node x in the fiber, we possibly have $\mathcal{W}_2(x, x) = 0$ while $\mathcal{W}_1(x, x) > 0$ may be true at the same time. We work around this issue by letting \mathcal{W}_2 be the lazy version (Definition 1.1.15) of the simple fiber walk which has non-negative eigenvalues by construction. Its second largest eigenvalue modulus can be bounded in terms of the edge-expansion of the graph as well:

Lemma 4.1.15. *For any d -regular graph G , $\lambda(\mathcal{L}(\mathcal{S}_G)) \geq 1 - \frac{1}{d} \cdot h(G)$.*

Proof. Let $\mathcal{W} := \mathcal{L}(\mathcal{S}_G)$, then the uniform distribution is the stationary distribution of \mathcal{W} and $\Phi(\mathcal{W}) = \frac{1}{2} \cdot \Phi(\mathcal{S}_G)$. With Remark 2.4.2, $\Phi(\mathcal{W}) = \frac{1}{2d} \cdot h(G)$ and since all eigenvalues of \mathcal{W} are non-negative, $\lambda(\mathcal{W}) = \lambda_2(\mathcal{W}) \geq 1 - 2 \cdot \Phi(\mathcal{W}) = 1 - \frac{1}{d} \cdot h(G)$. □

Proposition 4.1.16. Fix $A \in \mathbb{Z}^{m \times d}$, $b \in \mathbb{N}A$, and a Markov basis \mathcal{M} for $\mathcal{F}_{A,b}$. Let \mathcal{W} be a reversible, aperiodic and irreducible random walk on $\mathcal{F}_{A,b}(\mathcal{M})$ that converges to $\pi : \mathcal{F}_{A,b} \rightarrow (0, 1)$, then

$$1 - \lambda(\mathcal{W}) \leq 4 \cdot \frac{\max\{\pi(x) : x \in \mathcal{F}_{A,b}\}}{\min\{\pi(x) : x \in \mathcal{F}_{A,b}\}} \cdot h(\mathcal{F}_{A,b}(\mathcal{M})).$$

Proof. Let \mathcal{L} be the lazy simple fiber walk on $\mathcal{F}_{A,b}(\mathcal{M})$ and let π be the stationary distribution of \mathcal{W} . Our goal is to compare \mathcal{W} and \mathcal{L} with Lemma 4.1.14. First, we have for any distinct $x, y \in \mathcal{F}_{A,b}$, $\mathcal{W}(x, y) = 0$ whenever $\mathcal{L}(x, y) = 0$ since both maps are random walks on $\mathcal{F}_{A,b}(\mathcal{M})$ and since \mathcal{L} , as the lazy version of the simple fiber walk, has positive transition probabilities on all edges. For adjacent nodes $x, y \in \mathcal{F}_{A,b}$, $\mathcal{L}(x, y) = \frac{1}{2} \cdot (2 \cdot |\mathcal{M}|)^{-1} > 0$ and $\mathcal{W}(x, y) \leq 1$. Since the stationary distribution of \mathcal{L} is the uniform distribution, we have

$$\frac{\pi(x) \cdot \mathcal{W}(x, y)}{(|\mathcal{F}_{A,b}|)^{-1} \cdot \mathcal{L}(x, y)} \leq |\mathcal{F}_{A,b}| \cdot 4 \cdot |\mathcal{M}| \cdot \max\{\pi(x) : x \in \mathcal{F}_{A,b}\} =: C.$$

With $C' := |\mathcal{F}_{A,b}| \cdot \min\{\pi(x) : x \in \mathcal{F}_{A,b}\} > 0$ and Lemma 4.1.14,

$$1 - \lambda(\mathcal{W}) \leq \frac{C}{C'} \cdot (1 - \lambda(\mathcal{L})) = 4 \cdot |\mathcal{M}| \cdot \frac{\max\{\pi(x) : x \in \mathcal{F}_{A,b}\}}{\min\{\pi(x) : x \in \mathcal{F}_{A,b}\}} \cdot (1 - \lambda(\mathcal{L})).$$

By definition, \mathcal{L} is the lazy version of the simple fiber walk, which in turn is the simple walk on the $|\pm \mathcal{M}|$ -regular graph obtained from $\mathcal{F}_{A,b}(\mathcal{M})$ after adding loops. Lemma 4.1.15 then yields $1 - \lambda(\mathcal{L}) \leq \frac{1}{|\mathcal{M}|} h(\mathcal{F}_{A,b}(\mathcal{M}))$ and hence the claim. \square

Corollary 4.1.17. Let $A \in \mathbb{Z}^{m \times d}$, $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis, and $(b_i)_{i \in \mathbb{N}}$ be a dominated sequence in $\mathbb{N}A$. Let \mathcal{W}_i be a reversible, aperiodic, irreducible, and symmetric random walk on $\mathcal{F}_{A,b_i}(\mathcal{M})$, then there exist constants $C, C' \in \mathbb{N}_{\geq 1}$ such that $\lambda(\mathcal{W}_i) \geq 1 - \frac{C}{i}$ for all $i \in C' \cdot \mathbb{N}$. If $(b_i)_{i \in \mathbb{N}}$ has a meaningful parametrization, then $(\mathcal{W}_i)_{i \in \mathbb{N}}$ cannot mix rapidly.

Proof. The assumptions imply that for any $i \in \mathbb{N}$, the uniform distribution on \mathcal{F}_{A,b_i} is the stationary distribution of \mathcal{W}_i . Proposition 4.1.16 yield that $1 - \lambda(\mathcal{W}_i) \leq 4 \cdot h(\mathcal{F}_{A,b_i}(\mathcal{M}))$ and Theorem 4.1.10 together with Lemma 4.1.7 finish the proof. \square

Remark 4.1.18. Consider the problem of sampling uniformly from $\mathcal{F}_{A,b}$. A symmetric fiber walk can be seen as a probabilistic Turing machine that takes $(d, m, A \in \mathbb{Z}^{m \times d}, b \in \mathbb{Z}^m, \epsilon)$ as input and outputs an element from $\mathcal{F}_{A,b}$ almost uniformly (with distance at most ϵ). Neglecting ϵ , the binary encoding length of the problem instance is essentially $m \cdot d \cdot \log(\max_{k,j} A_{k,j}) + m \cdot \log(\|b\|_{\infty})$. Theorem 4.1.10 says that already for a fixed matrix A , a fixed Markov basis \mathcal{M} , and a ray $(i \cdot b) \in \mathbb{N}A$, the mixing time of the simple fiber walk cannot be bounded by a polynomial in $\log(|\mathcal{F}_{A,i \cdot b}|)$. Thus, the number of computations the simple fiber walk needs to converge for inputs $(d, m, A, i \cdot b)$, where $A \in \mathbb{Z}^{m \times d}$ and b are fixed, cannot be bounded polynomially in the binary encoding length of the input, even when we can compute a Markov basis \mathcal{M} for A efficiently. For more background details on the complexity of uniform random generation, we recommend [70].

4.2 Adapted Markov bases

The lesson learned from the previous section is that the moves in a Markov bases do not suffice to provide a good mixing behaviour asymptotically. A possible way out is to adapt the Markov

basis appropriately so that its size grows with the size of the right-hand side entries. This can be achieved by adding a varying number of \mathbb{Z} -linear combinations of the moves in a way that the edge-expansion of the resulting graph can be controlled. However, a growth of the set of allowed moves comes with an increase of the number of loops, i.e. an increase of the rejection rate of the walk. Let $A \in \mathbb{Z}^{m \times d}$ be a matrix, $\mathcal{M} = \{m_1, \dots, m_k\} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis for A , and $b \in \mathbb{N}A$. For $l \in \mathbb{N}$, let

$$\mathcal{M}(l) = \left\{ \sum_{j=1}^k \lambda_j m_j : \lambda_1, \dots, \lambda_k \in \mathbb{Z}, \sum_{j=1}^k |\lambda_j| \leq l \right\},$$

define $d_{A,b}^{\mathcal{M}} := \text{diam}(\mathcal{F}_{A,b}(\mathcal{M}))$, and let $\mathcal{M}^b := \mathcal{M}(d_{A,b}^{\mathcal{M}})$ be the *adapted Markov basis* for $\mathcal{F}_{A,b}$. Clearly, the fiber graph $\mathcal{F}_{A,b}(\mathcal{M}^b)$ is the complete graph and the transition matrix of the simple fiber walk on $\mathcal{F}_{A,b}(\mathcal{M}^b)$ is thus

$$\frac{1}{|\mathcal{M}^b|} \begin{bmatrix} 1 & 1 & \dots & 1 & 1 \\ 1 & \ddots & & & 1 \\ \vdots & & & & \vdots \\ 1 & & & \ddots & 1 \\ 1 & 1 & \dots & 1 & 1 \end{bmatrix} + \frac{1}{|\mathcal{M}^b|} \begin{bmatrix} |\mathcal{M}^b| - |\mathcal{F}_{A,b}| & 0 & \dots & 0 & 0 \\ 0 & \ddots & & & 0 \\ \vdots & & & & \vdots \\ 0 & & & \ddots & 0 \\ 0 & 0 & \dots & 0 & |\mathcal{M}^b| - |\mathcal{F}_{A,b}| \end{bmatrix}.$$

In particular, its second largest eigenvalue modulus is $1 - \frac{|\mathcal{F}_{A,b}|}{|\mathcal{M}^b|}$, which proves the next proposition.

Proposition 4.2.1. *Let $A \in \mathbb{Z}^{m \times d}$, $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis for A and $(b_i)_{i \in \mathbb{N}}$ a sequence in $\mathbb{N}A$. Suppose there is $r \in \mathbb{N}$ such that $(|\mathcal{F}_{A,b_i}|)_{i \in \mathbb{N}} \in \Omega(i^r)$ and $(|\mathcal{M}^{b_i}|)_{i \in \mathbb{N}} \in \mathcal{O}(i^r)_{i \in \mathbb{N}}$, then the simple fiber walk on $(\mathcal{F}_{A,b_i}(\mathcal{M}^{b_i}))_{i \in \mathbb{N}}$ is an expander.*

We discuss in Remark 4.2.5 how to sample moves from the adapted Markov basis. To make use of Proposition 4.2.1, the growth of the fibers and the adapted Markov bases has to be compared. Again, Ehrhart's theory applies to compute the growth of certain fiber sequences. The asymptotic growth of \mathcal{M}^{b_i} depends on the growth of the diameter of $\mathcal{F}_{A,b_i}(\mathcal{M})$. Hence, we first want to understand how the number of elements in $\mathcal{M}(l)$ grows as a function of $l \in \mathbb{N}$.

Lemma 4.2.2. *Let $\mathcal{M} = \{m_1, \dots, m_k\} \subset \mathbb{Z}^d$, then $(|\mathcal{M}(l)|)_{l \in \mathbb{N}} \in \mathcal{O}(l^{\text{rank}(\mathcal{M})})_{l \in \mathbb{N}}$.*

Proof. We identify the finite set \mathcal{M} with the integer matrix $(m_1, \dots, m_k) \in \mathbb{Z}^{d \times k}$. Denote the k -dimensional cross-polytope by $\mathcal{P} := \{x \in \mathbb{Q}^k : \|x\|_1 \leq 1\}$ and let $\mathcal{P}' := \{\mathcal{M} \cdot x : x \in \mathcal{P}\}$ be its image in \mathbb{Q}^d under \mathcal{M} . With this, we can write $\mathcal{M}(l) = \{\mathcal{M} \cdot x : x \in (l \cdot \mathcal{P}) \cap \mathbb{Z}^k\}$ and hence $\mathcal{M}(l) \subseteq (l \cdot \mathcal{P}') \cap \mathbb{Z}^d$. Since \mathcal{P}' is a polytope, Ehrhart's theorem [15, Theorem 3.23] gives $|(l \cdot \mathcal{P}') \cap \mathbb{Z}^d| \leq C \cdot l^{\dim(\mathcal{P}')}$ for some $C \in \mathbb{Q}_{>0}$ and since $\dim(\mathcal{P}') = \text{rank}(\mathcal{M})$, the claim follows. \square

Corollary 4.2.3. *Let $A \in \mathbb{Z}^{m \times d}$ and let $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ be a Markov basis for A . Let $(b_i)_{i \in \mathbb{N}}$ be a sequence in $\mathbb{N}A$ such that $(|\mathcal{F}_{A,b_i}|)_{i \in \mathbb{N}} \in \Omega(i^{d - \text{rank}(A)})$ and $(d_{A,b_i}^{\mathcal{M}})_{i \in \mathbb{N}} \in \mathcal{O}(i)_{i \in \mathbb{N}}$. Then the simple fiber walk on $(\mathcal{F}_{A,b_i}(\mathcal{M}^{b_i}))_{i \in \mathbb{N}}$ is an expander.*

Proof. Let $r := \dim(\ker_{\mathbb{Z}}(A))$. It suffices to show that $|\mathcal{M}^{b_i}| \leq C \cdot i^r$ for a constant $C \in \mathbb{Q}_{>0}$ since the statement then follows from Proposition 4.2.1. Since \mathcal{M} is a Markov basis for A , $\text{rank}(\mathcal{M}) = r$ and thus Lemma 4.2.2 implies that $|\mathcal{M}(l)| \leq C_1 \cdot l^r$ for a constant $C_1 \in \mathbb{Q}_{\geq 0}$. The assumption implies that there exists $C_2 \in \mathbb{Q}_{\geq 0}$ such that $d_{A,b_i}^{\mathcal{M}} \leq C_2 \cdot i$ for all $i \in \mathbb{N}$. Then, $|\mathcal{M}^{b_i}| = |\mathcal{M}(d_{A,b_i}^{\mathcal{M}})| \leq |\mathcal{M}(C_2 \cdot i)| \leq C_1 \cdot C_2^r \cdot i^r$. \square

Expanders are not per se fast, and Corollary 4.2.3 is an asymptotic statement. That means, for a given matrix $A \in \mathbb{Z}^{m \times d}$, a given Markov basis $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$, and a right-hand side $b \in \mathbb{N}A$, we know by Theorem 4.1.10 that the second largest eigenvalue modulus of the simple fiber walk that uses \mathcal{M} can be arbitrarily close to 1. On the other hand, since $(d_{A,i,b}^{\mathcal{M}})_{i \in \mathbb{N}} \in \mathcal{O}(i)_{i \in \mathbb{N}}$ by Proposition 2.2.10, the second largest eigenvalue modulus of the simple fiber walk that uses the adapted Markov basis $\mathcal{M}^{i,b}$ can be bounded away from 1 strictly. Thus, there exists a threshold $i_0 \in \mathbb{N}$ such that the adapted Markov basis is faster than the conventional Markov basis on $\mathcal{F}_{A,i,b}$ for $i \geq i_0$. The exact value of i_0 depends on the hidden constants in the asymptotic formulations of Corollary 4.2.3 and can be quite small, as in Figure 4.1, but also very large so that the advantages of the adapted Markov bases pay off only for large right-hand sides. For practitioners, the expander property may hence be too much to ask for. Instead of using $\mathcal{M}(l)$, one can use the slower growing set $\mathcal{M}(\log l)$. But already in the simplest example, a logarithmic adaption of the Markov basis fails to create rapidly mixing fiber walks:

Example 4.2.4. Let $\mathcal{F}_i := \mathcal{F}_{A_2, i-1}$ and $\mathcal{M}_i = \mathcal{M}_2(p(\log i))$ for a polynomial $p \in \mathbb{Q}[t]$. With Lemma 2.2.1 and Remark 2.4.19, $h(\mathcal{F}_i(\mathcal{M}_i)) \leq |\pm \mathcal{M}_i| \cdot \left(\exp\left(\frac{p(\log i) \cdot \log(i)}{i-1}\right) - 1 \right)$ and Proposition 4.1.9 yields the following lower bound on the second largest eigenvalue modulus λ_i of the simple fiber walk:

$$\lambda_i \geq 1 - \left(\exp\left(\frac{p(\log i) \cdot \log(i)}{i-1}\right) - 1 \right).$$

Thus, there cannot exist a polynomial $q \in \mathbb{Q}[t]$ such that $1 - \frac{1}{q(\log i)}$ bounds λ_i from above.

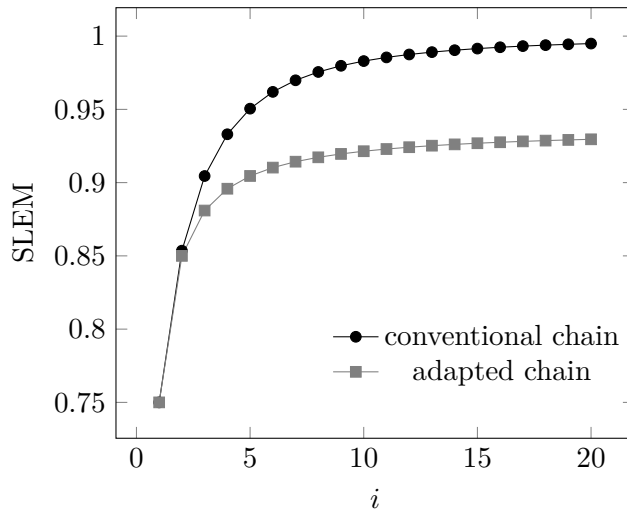


Figure 4.1: The second largest eigenvalue modulus of the simple walk on $\mathcal{F}_{A_3, i}$ using moves from the conventional Markov basis \mathcal{M}_3 and the adapted moves $\mathcal{M}_3(2i)$.

Remark 4.2.5. Running the simple fiber walk on $\mathcal{F}_{A,b}(\mathcal{M}(l))$ for some $l \in \mathbb{N}$ requires to sample from $\mathcal{M}(l)$ uniformly and hence a good understanding of this set is necessary. Basically, we shift the problem of sampling from $\mathcal{F}_{A,b}$ for all $b \in \mathbb{N}A$ where $\mathcal{F}_{A,b}(\mathcal{M})$ has diameter l to the problem of sampling from $\mathcal{M}(l)$, which can be seen as some kind of *rejection sampling* from a larger saturated set $u + \mathcal{M}(l) \supseteq \mathcal{F}_{A,b}$ (Example 4.2.7). For large fibers, one applicable move $m \in \mathcal{M}(l)$ suffices to obtain a sample $u + m \in \mathcal{F}_{A,b}$ that is very close to uniform. Write $\mathcal{M} = \{m_1, \dots, m_k\}$

and $r := \text{rank}(\mathcal{M})$. When $r = k$, then an element λ picked uniformly from $\{u \in \mathbb{Z}^k : \|u\|_1 \leq l\}$ gives rise to an element $\mathcal{M} \cdot \lambda$ that is uniformly generated from $\mathcal{M}(l)$. This is not the case when $r > k$. One approach to sample from $\mathcal{M}(l)$ uniformly in this case is to first compute a lattice basis $\mathcal{B} := \{b_1, \dots, b_r\} \subset \mathbb{Z}^d$ of $\mathcal{M} \cdot \mathbb{Z}^k$ in order to get rid of relations among the moves from \mathcal{M} . Then, we compute for every $i \in [k]$ coefficients $\lambda_1^i, \dots, \lambda_r^i$ such that $m_i = \sum_{j=1}^r \lambda_j^i b_j$. For $C := \sum_{j=1}^r \max_{i \in [k]} |\lambda_j^i|$, we have $\mathcal{M}(l) \subseteq \mathcal{B}(C \cdot l)$. Thus, after sampling coefficients λ from $\{u \in \mathbb{Z}^r : \|u\|_1 \leq C \cdot l\}$ uniformly, we obtain a move $\mathcal{B} \cdot \lambda$ that is sampled uniformly from a superset of $\mathcal{M}(l)$. Since $|\mathcal{B}(C \cdot l)|$ grows as $\mathcal{O}(l^r)_{l \in \mathbb{N}}$, Proposition 4.2.1 remains valid. Sampling from the cross-polytope $\{u \in \mathbb{Z}^r : \|u\|_1 \leq C \cdot l\}$ can be done with the heat-bath method as studied in Chapter 5, which is fast for $l \rightarrow \infty$ (Example 5.2.13).

Remark 4.2.6. There are many heuristics possible to traverse the fiber with the adapted Markov basis. For instance, one can sample from $\mathcal{M}(l)$ where the probability of coefficients $\lambda \in \{u \in \mathbb{Z}^r : u \leq C \cdot l\}$ is inversely proportional to $\|\lambda\|_1$, then shorter moves appear more frequently than longer moves. Basically, every randomized algorithm that generates elements from $\{u \in \mathbb{Z}^k : \|u\|_1 \leq l\}$ according to some mass function ξ gives rise to a sampling scheme on $\mathcal{M}(l)$: Sample first coefficients λ according to ξ and then use the move $\mathcal{M} \cdot \lambda \in \mathcal{M}(l)$. Again, uniform sampling from $\{u \in \mathbb{Z}^k : \|u\|_1 \leq l\}$ can be achieved with the heat-bath walk (Example 5.2.13). To approximate the uniform distribution on the fiber $\mathcal{F}_{A,b}$, it suffices that ξ fulfills $\xi(\lambda) = \xi(-\lambda)$ for all coefficients λ . With an additional Metropolis rejection step (Remark 1.4.4), any mass function on $\mathcal{F}_{A,b}$ with an incomputable normalizing constant can be approximated with $\mathcal{M}(l)$.

Example 4.2.7. Let H_k be as in Example 1.3.1 and consider the sequence $(\mathcal{F}_{H_k, i \cdot e_{2k+1}})_{i \in \mathbb{N}}$. With [62, Section 4], it is easy to show that for any $k \in \mathbb{N}$, the diameter of $\mathcal{F}_{H_k, i \cdot e_{2k+1}}(\mathcal{R}_k)$ is $(2k+1)i$ and hence linear in i . The moves in the reduced Gröbner basis \mathcal{R}_k are linearly independent and hence uniform sampling from the adapted Markov basis $\mathcal{R}_k((2k+1)i)$ is achieved via sampling uniformly from the cross-polytope $\{u \in \mathbb{Z}^{2k+1} : \|u\|_1 \leq (2k+1)i\}$. The node-connectivity of $\mathcal{F}_{H_k, i \cdot e_{2k+1}}(\mathcal{R}_k)$ is 1 (Proposition 2.1.8), but the simple fiber walk in the cross-polytope that uses the unit vectors does not see the bad connectivity. Rejection sampling from $\mathcal{F}_{H_k, i \cdot e_{2k+1}}$ with upper and lower bounds on the entries of its elements yields sampling from a $(2k+1)$ -dimensional hyperrectangle $i \cdot ([C_1] \times \dots \times [C_{2k+1}])$ with constants C_1, \dots, C_{2k+1} , which is easy to sample from, but the rejection rate is larger than sampling from the cross-polytope.

4.3 Varying constraint matrices

Markov bases of constraint matrices coming from statistical problems are often parametrized and they can be explicitly stated for any parameter. For instance, the basic moves $\mathcal{M}_{n,n}$ of the independence model (Example 4.1.12) form a Markov basis for $A_{n,n}$ for every $n \in \mathbb{N}$. Thus, varying the parameter n provides fiber graphs where the set of moves is adapted canonically.

Example 4.3.1. Let $b_n := (1, \dots, 1) \in \mathbb{N}^{2n}$, then the elements of $\mathcal{F}_{A_{n,n}, b_n}$ can be identified with the elements of the symmetric group S_n . Finding a set of generators such that the corresponding Cayley graph on S_n is an expander is an active research field in group theory, see for instance [74]. In [38], it was shown that the simple walk on the Cayley graph of S_n that uses the transpositions mixes rapidly in $\frac{1}{2}n \log n$ many steps. Inspired by shuffling a deck of n cards, a random walk on S_n that uses *riffle shuffles* was studied in [13] and shown to be rapidly mixing as well.

Parametric descriptions of Markov bases can be arbitrarily complicated in general, since by the Universality theorem [33], any integer vector appears as a subvector of a Markov basis element of the three-way no interaction model, when the parameters are large enough. Different than in fixed dimension, where the Markov basis is fixed, the size of the Markov basis is important in the convergence analysis when the dimension varies because the local sampling process of a move can be computationally challenging as the Markov basis becomes larger (Remark 4.3.3). The trade-off between an easily accessible set of moves and a corresponding random walk that has good mixing properties shows the realms of fiber walks in practice. The next proposition illustrates this for H_k from Example 1.3.1, where the overwhelming number of moves in its parametric Graver basis slows the chain down for $k \rightarrow \infty$, despite the fact that the edge-connectivity of these fibers is best-possible by Proposition 2.1.9.

Proposition 4.3.2. *The simple fiber walk on $(\mathcal{F}_{H_k, e_{2k+1}}(\mathcal{G}_{H_k}))_{k \in \mathbb{N}}$ is not rapidly mixing.*

Proof. According to [62, Section 4], $\mathcal{F}_{H_k, e_{2k+1}}(\mathcal{G}_{H_k})$ is isomorphic to the graph on the nodes $\{0, 1\}^{k+1}$ in which two nodes (i_1, \dots, i_{k+1}) and (j_1, \dots, j_{k+1}) are adjacent if either $i_{k+1} = j_{k+1}$ and $\|i - j\|_\infty = 1$, or if $i_{i+1} \neq j_{i+1}$. For any $k \in \mathbb{N}_{>0}$, define

$$S_k := \{(0, i, 0) : i \in \{0, 1\}^{k-1}\} \cup \{(0, i, 1) : i \in \{0, 1\}^{k-1}\},$$

then $|S_k| = \frac{1}{2}|\mathcal{F}_{H_k, e_{2k+1}}|$. Counting the edges that leave S_k yields that for any $(0, i, 0) \in S_k$, there are k with endpoints in $\{(1, i, 0) : i \in \{0, 1\}^{k-1}\}$ and 2^{k-1} with endpoints in $\{(1, i, 1) : i \in \{0, 1\}^{k-1}\}$. The same is true for any $(0, i, 1) \in S_k$. Hence, there are $(k + 2^{k-1}) \cdot 2 \cdot 2^{k-1}$ edges leaving S_k . The edge-expansion of $\mathcal{F}_{H_k, e_{2k+1}}(\mathcal{G}_{H_k})$ is thus bounded from above by $k + 2^{k-1}$. Since $|\mathcal{G}_{H_k}| = 2 \cdot (4^k + 4k)$, Proposition 4.1.9 implies that the second largest eigenvalue of the simple fiber walk is bounded from below by

$$1 - \frac{k + 2^{k-1}}{2 \cdot (4^k + 4 \cdot k)}.$$

This together with $\log |\mathcal{F}_{H_k, e_{2k+1}}| = k + 1$ gives the statement. \square

Remark 4.3.3. Let $A \in \mathbb{Z}^{m \times d}$, $b \in \mathbb{N}^A$ and $\mathcal{X} \subseteq [d]$. Testing the goodness-of-fit of a log-linear model that has structural zeros requires to sample from sets of the form

$$\mathcal{F}_{A,b}^{\mathcal{X}} := \{u \in \mathcal{F}_{A,b} : u_i = 0 \ \forall i \in \mathcal{X}\}.$$

A priori, a Markov basis $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ for A does not make the fiber graphs $\mathcal{F}_{A,b}^{\mathcal{X}}(\mathcal{M})$ connected for all b . The problem when a Markov bases of A still connects sets $\mathcal{F}_{A,b}^{\mathcal{X}}$ with $b_i \geq 1$ was studied in [72] from an algebraic point of view. Interpreting the constraints on the coordinates in \mathcal{X} as linear inequalities, Proposition 2.1.3 tells that the Graver basis \mathcal{G}_A is a Markov basis for these sets. However, as the dimension grows, it can be challenging to sample from the Graver basis uniformly. For instance, the Graver basis of $A_{n,m}$ with additional zero-constraints defined by a set $\mathcal{X} \subset [n] \times [m]$ corresponds to the set of cycles of $K_{n,m}$ where edges from \mathcal{X} are removed. Already for bipartite graphs, sampling a circuit uniformly is challenging [70, Theorem 5.1].

Example 4.3.4. This example of a simple fiber walk where the dimension varies is from [91]. For $a \in \mathbb{R}^d$ and $b \in \mathbb{R}$, consider the set $\mathcal{F}(a, b) := \{u \in \{0, 1\}^d : a^T u \leq b\}$, that is the set of solutions to a knapsack instance. The simple fiber walk on $\mathcal{F}(a, b)$ that samples at each step uniformly from $\pm\{e_1, \dots, e_d\}$ is precisely the random walk studied in [91], and they showed that it converges in polynomially many steps in the dimension, which is rapid mixing.

5 Heat-bath walks

We have seen in the previous chapter that a small diameter is a necessary condition for fast mixing. Since the diameter of all compressed fiber graphs from the same matrix is bounded by a universal constant from above due to Theorem 2.2.17, they are canonical candidates for rapid mixing. *Heat-bath walks*, as studied in [46] recently, are tailored for an execution on compressed fiber graphs. To define them, let $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set and denote for any $u \in \mathcal{F}$ and $m \in \mathbb{Z}^d$, the ray in \mathcal{F} through u along m by $\mathcal{R}_{\mathcal{F},m}(u) := (u + m \cdot \mathbb{Z}) \cap \mathcal{F}$. Given a mass function $\pi : \mathcal{F} \rightarrow [0, 1]$, we define for $x, y \in \mathcal{F}$

$$\mathcal{H}_{\mathcal{F},m}^{\pi}(x, y) := \begin{cases} \frac{\pi(y)}{\pi(\mathcal{R}_{\mathcal{F},m}(x))}, & \text{if } y \in \mathcal{R}_{\mathcal{F},m}(x) \\ 0, & \text{otherwise} \end{cases}.$$

For a finite set $\mathcal{M} \subset \mathbb{Z}^d$ equipped with a mass function $f : \mathcal{M} \rightarrow [0, 1]$, the *heat-bath walk* is

$$\mathcal{H}_{\mathcal{F},\mathcal{M}}^{\pi,f} = \sum_{m \in \mathcal{M}} f(m) \cdot \mathcal{H}_{\mathcal{F},m}^{\pi}. \quad (5.0.1)$$

Observe that \mathcal{F} does not need to be saturated. The heat-bath walk is performed as follows:

Algorithm 2 Heat-bath walk on compressed fiber graphs

Input: $\mathcal{F} \subset \mathbb{Z}^d$, $\mathcal{M} \subset \mathbb{Z}^d$, $v \in \mathcal{F}$, mass functions $f : \mathcal{M} \rightarrow [0, 1]$ and $\pi : \mathcal{F} \rightarrow [0, 1]$, $r \in \mathbb{N}$

1: **procedure** HEATBATH:

2: $v_0 := v$

3: **FOR** $s = 0$; $s = s + 1$, $s < r$

4: Sample $m \in \mathcal{M}$ according to f

5: Sample $v_{s+1} \in \mathcal{R}_{\mathcal{F},m}(v_s)$ according to $\mathcal{R}_{\mathcal{F},m}(v_s) \rightarrow [0, 1]$, $y \mapsto \frac{\pi(y)}{\pi(\mathcal{R}_{\mathcal{F},m}(v_s))}$

6: **RETURN** v_1, \dots, v_r

In other words, the heat-bath walk samples at the current lattice point $u \in \mathcal{F}$ a move $m \in \mathcal{M}$ and walks to a random element in the integer ray $\mathcal{R}_{\mathcal{F},m}(u)$. The authors of [39] discovered that this random walk can be seen as a discrete version of the *hit-and-run* algorithm [82, 116, 83] that has been used frequently to sample from the points of a polytope – not only lattice points. The popularity of the continuous version of the hit-and-run algorithm has not spread to its discrete analogue, and not much is known about its mixing behaviour. One reason is that it is already challenging to guarantee that all points in the underlying set \mathcal{F} can be reached by a random walk that uses moves from \mathcal{M} , whereas for the continuous version, a random sampling from the unit sphere suffices. However, in many situations where a Markov basis is known, the heat-bath walk is evidently fast. For instance, it was shown in [29] that the heat-bath walk on two-way contingency tables under the independence model mixes rapidly when the number of rows is fixed and the basic moves are used. To work around the connectedness issue, a *discrete hit-and-run* algorithm was introduced in [12] for arbitrary finite sets $\mathcal{F} \subset \mathbb{Z}^d$. In each step of this random

walk, a subordinate and unrestricted random walk starts at the current lattice point $u \in \mathcal{F}$ and uses the unit vectors to collect a set of proposals $S \subset \mathbb{Z}^d$. The random walk then moves from u to a random point in $S \cap \mathcal{F}$. Generally speaking, the same methodology is applied by the heat-bath walk, but here, the proposals are on a ray $\mathcal{R}_{\mathcal{F},m}(u)$.




Figure 5.1: Points reached in a simple fiber walk (on the left) and a heat-bath walk (on the right) on a 50×80 grid using the moves $\{(1, 0), (0, 1), (2, 1), (1, 2)\}$ after 1000 steps.

In this chapter, we explore the mixing behaviour of heat-bath walks on lattice points with Markov bases. In Section 5.1, we study in more detail the combinatorial and analytical structure of their transition matrices and prove upper and lower bounds on their second largest eigenvalues. We use the *canonical path approach* from [105] and establish in Theorem 5.2.9 an upper bound on the second largest eigenvalue modulus when the Markov basis is *augmenting* (Definition 5.2.1) and when π is the uniform distribution. From that, we conclude fast mixing results for random walks on lattice points in fixed dimension. In the end, we briefly discuss how the distribution f on the moves \mathcal{M} affects the speed of convergence (Section 5.3). This chapter is based on the joint work [108] with Caprice Stanley.

5.1 Spectral analysis

The underlying graph of the heat-bath walk is the compression $\mathcal{F}^c(\mathcal{M})$ and when the moves in a Markov bases are linearly independent, then the heat-bath walk becomes the *Glauber dynamics*:

Remark 5.1.1. Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite and $\mathcal{M} = \{m_1, \dots, m_d\} \subset \mathbb{Z}^d$ be a linearly independent Markov basis of \mathcal{F} . If the moves are selected uniformly, then the heat-bath walk on \mathcal{F} coincides with the *Glauber dynamics* on \mathcal{F} . To see it, choose $u \in \mathcal{F}$ and let

$$\mathcal{F}' := \{\lambda \in \mathbb{Z}^d : u + \lambda_1 m_1 + \dots + \lambda_d m_d \in \mathcal{F}\}.$$

It is easy to check that \mathcal{F}' is unique up to translation and depends only on \mathcal{F} , u , and \mathcal{M} . Since the vectors in \mathcal{M} are linearly independent, every element of \mathcal{F} can be represented by a unique choice of coefficients in \mathcal{F}' . Thus, the heat-bath walk on \mathcal{F} using \mathcal{M} is equivalent to the heat-bath walk on \mathcal{F}' using the unit vectors as moves. For any unit vector $e_i \in \mathbb{Z}^d$, the ray through an element $v \in \mathcal{F}'$ is $\{w \in \mathcal{F}' : w_j = v_j \forall j \neq i\}$ which is precisely the form in the Glauber dynamics [81, Section 3.3.2].

Although an asymptotically bounded diameter is a necessary condition for good mixing behaviour, it is not sufficient in general: Let G_n be the disjoint union of two complete graphs K_n connected by a single edge, then $\text{diam}(G_n) = 3$, but $h(G_n) \leq \frac{1}{n}$ implies that the simple walk on G_n does not mix rapidly. Thus, heat-bath walks are not per se rapidly mixing and this asks

for a deeper analysis of the spectral decomposition of heat-bath walks. To get started, let us first recall the basic properties of this random walk (compare also to [39, Lemma 2.2]):

Proposition 5.1.2. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets. Let $f : \mathcal{M} \rightarrow [0, 1]$ and $\pi : \mathcal{F} \rightarrow (0, 1)$ be mass functions. Then $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$ is aperiodic, has stationary distribution π , is reversible with respect to π , and all of its eigenvalues are non-negative. The random walk is irreducible if and only if $\{m \in \mathcal{M} : f(m) > 0\}$ is a Markov basis for \mathcal{F} .*

Proof. For any $u \in \mathcal{F}$ and any $m \in \mathcal{M}$, $\mathcal{H}_{\mathcal{F}, m}^{\pi}(u, u) > 0$, there are halting states and thus $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$ is aperiodic. By definition, $\pi(x)\mathcal{H}_{\mathcal{F}, m}^{\pi}(x, y) = \pi(y)\mathcal{H}_{\mathcal{F}, m}^{\pi}(y, x)$ for all $x, y \in \mathcal{F}$ and thus $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$ is reversible with respect to π and π is a stationary distribution. The statement on the eigenvalues is [46, Lemma 1.2]. Let $\mathcal{M}' = \{m \in \mathcal{M} : f(m) > 0\}$ and $f' = f|_{\mathcal{M}'}$, then $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f} = \mathcal{H}_{\mathcal{F}, \mathcal{M}'}^{\pi, f'}$ and thus the heat-bath walk is irreducible if and only if \mathcal{M}' is a Markov basis for \mathcal{F} . \square

Remark 5.1.3. From a computational point of view, the difference of the simple fiber walk and the heat-bath walk is Step 4 of Algorithm 2. More computation is necessary for heat-bath walks at every transition. However, Step 4 can be done efficiently in many cases. As in the Metropolis-Hastings walk (Remark 1.4.4), an incomputable normalizing constant of π cancels. For instance, when π is the uniform distribution, then one needs to sample uniformly from $\mathcal{R}_{\mathcal{F}, m}(v)$ in Step 4. If the input of Algorithm 2 is a saturated set $\mathcal{F} = \{u \in \mathbb{Z}^d : Au \leq b\}$ that is given in \mathcal{H} -representation, then the length of the ray $\mathcal{R}_{\mathcal{F}, m}(v)$ can be computed with a number of operations that is polynomial in the binary encoding length of A and b .

There are situations where the heat-bath walk gives no speed-up compared with the simple fiber walk (Example 5.1.4). Intuitively, adding more moves to the set of allowed moves should improve the mixing time of the fiber walk. Surprisingly, this is not true for heat-bath walks:

Example 5.1.4. Consider the fiber graph sequence from Proposition 4.3.2. The underlying saturated sets are subsets of $\{0, 1\}^{4k+2}$ respectively, and thus every ray along a Graver move has length at most 2. Hence, the transition matrices of the simple fiber walk and the heat-bath walk coincide. Thus, the heat-bath walk is not rapidly mixing for $k \rightarrow \infty$.

Example 5.1.5. Let $\mathcal{F} = [2] \times [5] \subset \mathbb{N}^2$, $\mathcal{M} = \{e_1, e_2, 2e_1 + e_2\}$, and π the uniform distribution on \mathcal{F} . Since $\{e_2, 2e_1 + e_2\}$ is not a Markov basis for \mathcal{F} , any mass function $f : \mathcal{M} \rightarrow [0, 1]$ must have $f(e_1) > 0$ in order to make the corresponding heat-bath walk irreducible. Comparing the second largest eigenvalue modulus of the heat-bath walks that sample uniformly from $\{e_1, e_2\}$ and \mathcal{M} respectively, we obtain

$$\lambda \left(\frac{1}{2} \mathcal{H}_{\mathcal{F}, e_1}^{\pi} + \frac{1}{2} \mathcal{H}_{\mathcal{F}, e_2}^{\pi} \right) = \frac{1}{2} < \frac{2}{3} = \lambda \left(\frac{1}{3} \mathcal{H}_{\mathcal{F}, e_1}^{\pi} + \frac{1}{3} \mathcal{H}_{\mathcal{F}, e_2}^{\pi} + \frac{1}{3} \mathcal{H}_{\mathcal{F}, 2e_1 + e_2}^{\pi} \right).$$

Said in words: Adding $2e_1 + e_2$ to the set of allowed moves slows the walk down. This phenomenon does not appear for the simple fiber walk on \mathcal{F} , where the second largest eigenvalue modulus improves from ≈ 0.905 to ≈ 0.888 when adding the move $2e_1 + e_2$ to the Markov basis.

For the remainder of this section, we primarily focus on heat-bath walks $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$ that converge to the uniform distribution π on a finite, but not necessarily saturated, set \mathcal{F} . In particular, we aim for bounds on its second largest eigenvalue by making use of the decomposition from equation (5.0.1). Our first observations consider its summands $\mathcal{H}_{\mathcal{F}, m}^{\pi}$ that can be well understood analytically (Proposition 5.1.6) and combinatorially (Proposition 5.1.7).




Figure 5.2: Decomposition of the graph in Example 5.1.5.

Proposition 5.1.6. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set, $m \in \mathbb{Z}^d$, and $\pi : \mathcal{F} \rightarrow [0, 1]$ be the uniform distribution. Let $\mathcal{R}_1, \dots, \mathcal{R}_k$ be the disjoint rays through \mathcal{F} along m . Then*

1. $\mathcal{H}_{\mathcal{F}, m}^\pi$ is symmetric and idempotent.
2. $\text{img}(\mathcal{H}_{\mathcal{F}, m}^\pi) = \text{span}_{\mathbb{R}} \left\{ \sum_{x \in \mathcal{R}_1} e_x, \sum_{x \in \mathcal{R}_2} e_x, \dots, \sum_{x \in \mathcal{R}_k} e_x \right\}$.
3. $\ker(\mathcal{H}_{\mathcal{F}, m}^\pi) = \bigoplus_{i=1}^k \text{span}_{\mathbb{R}} \{e_x - e_y : x, y \in \mathcal{R}_i, x \neq y\}$.
4. $\text{rank}(\mathcal{H}_{\mathcal{F}, m}^\pi) = k$ and $\dim \ker(\mathcal{H}_{\mathcal{F}, m}^\pi) = |\mathcal{F}| - k$.
5. The spectrum of $\mathcal{H}_{\mathcal{F}, m}^\pi$ is $\{0, 1\}$.

Proof. Symmetry of $\mathcal{H}_{\mathcal{F}, m}^\pi$ follows from the definition. By assumption, \mathcal{F} is the disjoint union of $\mathcal{R}_1, \dots, \mathcal{R}_k$ and hence there exists a permutation matrix S such that $S\mathcal{H}_{\mathcal{F}, m}^\pi S^T$ is a block matrix whose building blocks are the matrices

$$\frac{1}{|\mathcal{R}_i|} \begin{bmatrix} 1 & \dots & 1 \\ \vdots & & \vdots \\ 1 & \dots & 1 \end{bmatrix} \in \mathbb{Q}^{|\mathcal{R}_i| \times |\mathcal{R}_i|}.$$

Thus, $\mathcal{H}_{\mathcal{F}, m}^\pi$ is idempotent and the rank of $\mathcal{H}_{\mathcal{F}, m}^\pi$ is k . A basis of its image and its kernel can be read off directly and idempotent matrices can only have the eigenvalues 0 and 1. \square

Proposition 5.1.7. *Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets, $\pi : \mathcal{F} \rightarrow [0, 1]$ be the uniform distribution, and let $V_1, \dots, V_c \subseteq \mathcal{F}$ be the nodes of the connected components of $\mathcal{F}(\mathcal{M})$, then*

$$\bigcap_{m \in \mathcal{M}} \text{img}(\mathcal{H}_{\mathcal{F}, m}^\pi) = \text{span}_{\mathbb{R}} \left\{ \sum_{x \in V_1} e_x, \dots, \sum_{x \in V_c} e_x \right\}.$$

Proof. It is clear by Proposition 5.1.6 that the set on the right-hand side is contained in any $\text{img}(\mathcal{H}_{\mathcal{F}, m}^\pi)$ since any V_i decomposes disjointedly into rays along $m \in \mathcal{M}$. To show the other inclusion, write $\mathcal{M} = \{m_1, \dots, m_k\}$ and let for any $i \in [k]$, $\mathcal{R}_1^i, \dots, \mathcal{R}_{n_i}^i$ be the disjoint rays through \mathcal{F} parallel to m_i . In particular, $\{\mathcal{R}_1^i, \dots, \mathcal{R}_{n_i}^i\}$ is a partition of \mathcal{F} for any $i \in [k]$. Let $v \in \bigcap_{i=1}^k \text{img}(\mathcal{H}_{\mathcal{F}, m_i}^\pi)$. By Proposition 5.1.6, for any $i \in [k]$ there exist $\lambda_1^i, \dots, \lambda_{n_i}^i \in \mathbb{Q}$ such that

$$v = \sum_{j=1}^{n_i} \sum_{x \in \mathcal{R}_j^i} \lambda_j^i e_x.$$

If two distinct Markov moves m_i and $m_{i'}$ and two indices $j \in [n_i]$ and $j' \in [n_{i'}]$ satisfy $\mathcal{R}_j^i \cap \mathcal{R}_{j'}^{i'} \neq \emptyset$, then $\lambda_j^i = \lambda_{j'}^{i'}$. We show that for any $i \in [k]$ and any $a \in [c]$, $\lambda_j^i = \lambda_{j'}^a$ when \mathcal{R}_j^i and $\mathcal{R}_{j'}^a$ are a subset of V_a . This implies the proposition. Take distinct $x, x' \in V_a$ and assume that x and x' lie on different rays of m_i and let that be $x \in \mathcal{R}_j^i$ and $x' \in \mathcal{R}_{j'}^i$ with $j \neq j'$. Since x and x' are in the same connected component V_a of $\mathcal{F}(\mathcal{M})$, let $y_{i_0}, \dots, y_{i_r} \in \mathcal{F}$ be the nodes on a minimal path in $\mathcal{F}^c(\mathcal{M})$ with $y_{i_0} = x$ and $y_{i_r} = x'$. For any $s \in [r]$, y_{i_s} and $y_{i_{s-1}}$ are contained in the same ray $\mathcal{R}_{t_s}^{k_s}$ coming from a Markov move m_{k_s} . In particular, $\mathcal{R}_{k_{s-1}}^{t_{s-1}} \cap \mathcal{R}_{t_s}^{k_s} \neq \emptyset$ and due to the observation made above, $\lambda_j^i = \lambda_{t_1}^{k_1} = \lambda_{t_2}^{k_2} = \dots = \lambda_{t_r}^{k_r} = \lambda_{j'}^a$. \square

Definition 5.1.8. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets and $\mathcal{M}' \subseteq \mathcal{M}$. Let \mathcal{V} be the set of connected components of $\mathcal{F}(\mathcal{M} \setminus \mathcal{M}')$ and \mathcal{R} be the set of all rays through \mathcal{F} along all elements of \mathcal{M}' . The *ray matrix* of $\mathcal{F}(\mathcal{M})$ along \mathcal{M}' is $A_{\mathcal{F}}(\mathcal{M}, \mathcal{M}') := (|\mathcal{R} \cap V|)_{R \in \mathcal{R}, V \in \mathcal{V}} \in \mathbb{N}^{\mathcal{R} \times \mathcal{V}}$.

Remark 5.1.9. The ray matrix itself seems to be a very interesting object. For instance, take any saturated set $\mathcal{F} \subset \mathbb{Z}^2$, then, possibly after translation, we can assume that $\mathcal{F} \subseteq [n] \times [m]$ with n and m minimal. Let $\mathcal{M} = \{e_1, e_2\}$ and $\mathcal{M}' = e_1$, then for all $(i, j) \in [n] \times [m]$,

$$A_{\mathcal{F}}(\mathcal{M}, \mathcal{M}')_{ij} = \begin{cases} 1, & \text{if } (i, j) \in \mathcal{F} \\ 0, & \text{if } (i, j) \notin \mathcal{F} \end{cases}.$$

Thus, the ray matrix encodes the integer points of a lattice polytope in its pattern of non-zero entries and we think it is appealing to find properties of the lattice polytope which can be read off the ray matrix and vice-versa.



Figure 5.3: A saturated set $\mathcal{F} \subset [3] \times [7]$ and $A_{\mathcal{F}}(\{e_1, e_2\}, \{e_1\})$.

Proposition 5.1.10. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets, $\pi : \mathcal{F} \rightarrow [0, 1]$ be the uniform distribution, and $\mathcal{M}' \subseteq \mathcal{M}$. Then

$$\ker(A_{\mathcal{F}}(\mathcal{M}, \mathcal{M}')) \cong \bigcap_{m \in \mathcal{M} \setminus \mathcal{M}'} \text{img}(\mathcal{H}_{\mathcal{F}, m}^{\pi}) \cap \bigcap_{m \in \mathcal{M}'} \ker(\mathcal{H}_{\mathcal{F}, m}^{\pi}).$$

Proof. Let V_1, \dots, V_c be the connected components of $\mathcal{F}(\mathcal{M} \setminus \mathcal{M}')$ and $\mathcal{R}_1, \dots, \mathcal{R}_r$ be the rays along elements in \mathcal{M}' . Let $I := \bigcap_{m \in \mathcal{M} \setminus \mathcal{M}'} \text{img}(\mathcal{H}_{\mathcal{F}, m}^{\pi})$ and $K := \bigcap_{m \in \mathcal{M}'} \ker(\mathcal{H}_{\mathcal{F}, m}^{\pi})$. By Proposition 5.1.7, any element of I has the form $v = \sum_{i=1}^c (\lambda_i \sum_{x \in V_i} e_x)$ for $\lambda_1, \dots, \lambda_c \in \mathbb{Q}$. Assume additionally that $v \in \ker(\mathcal{H}_{\mathcal{F}, m}^{\pi})$ for $m \in \mathcal{M}'$ and let $\mathcal{R}_{i_1}, \dots, \mathcal{R}_{i_j}$ be the rays parallel to m , then for any $k \in [j]$, $0 = \sum_{x \in \mathcal{R}_{i_k}} v_x = \sum_{j=1}^c \lambda_j |\mathcal{R}_{i_k} \cap V_j|$. Put differently, a vector $\lambda \in \mathbb{R}^c$ is in the kernel of $(|\mathcal{R}_i \cap V_j|)_{i \in [r], j \in [c]}$ if and only if $\sum_{i=1}^c (\lambda_i \sum_{x \in V_i} e_x) \in I \cap K$. \square

Conditions on the kernel of the ray matrix allow us to give a lower bound on the second largest eigenvalue of the heat-bath walk:

Proposition 5.1.11. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets and π be the uniform distribution. Let $\mathcal{M}' \subseteq \mathcal{M}$ such that $\ker(A_{\mathcal{F}}(\mathcal{M}, \mathcal{M}')) \neq \{0\}$, then $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) \geq 1 - \sum_{m \in \mathcal{M}'} f(m)$ for any mass function $f : \mathcal{M} \rightarrow [0, 1]$.

Proof. Using the isomorphism from Proposition 5.1.10, we can choose a non-zero $v \in \mathbb{Q}^P$ such that $\mathcal{H}_{\mathcal{F}, m}^{\pi} v = v$ for all $m \in \mathcal{M} \setminus \mathcal{M}'$ and $\mathcal{H}_{\mathcal{F}, m}^{\pi} v = 0$ for all $m \in \mathcal{M}'$. In particular

$$\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f} v = \sum_{m \in \mathcal{M}} f(m) \mathcal{H}_{\mathcal{F}, m}^{\pi} v = \sum_{m \in \mathcal{M} \setminus \mathcal{M}'} f(m) \mathcal{H}_{\mathcal{F}, m}^{\pi} v = \sum_{m \in \mathcal{M} \setminus \mathcal{M}'} f(m) v.$$

Since f is a mass function, $1 - \sum_{m \in \mathcal{M}'} f(m)$ is an eigenvalue of $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$. \square

Definition 5.1.12. Let $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set and $m, m' \in \mathbb{Z}^d$ not collinear. The pair (m, m') has the *intersecting ray property in \mathcal{F}* if the following holds: For any pair of rays $\mathcal{R}_1, \mathcal{R}_2$ parallel to m and any pair of rays $\mathcal{R}'_1, \mathcal{R}'_2$ parallel to m' where both $\mathcal{R}_1 \cap \mathcal{R}'_1$ and $\mathcal{R}_2 \cap \mathcal{R}'_2$ are not empty, $\mathcal{R}_1 \cap \mathcal{R}'_2 \neq \emptyset$ implies $\mathcal{R}'_1 \cap \mathcal{R}_2 \neq \emptyset$ and $|\mathcal{R}_1| \cdot |\mathcal{R}'_1|^{-1} = |\mathcal{R}_2| \cdot |\mathcal{R}'_2|^{-1}$. Given a finite set $\mathcal{M} \subset \mathbb{Z}^d$, the graph $\mathcal{F}^c(\mathcal{M})$ has the *intersecting ray property* if all pairs (m, m') with $m, m' \in \mathcal{M}$ have the intersecting ray property in \mathcal{F} .

Example 5.1.13. The compressed fiber graph on $[n_1] \times \cdots \times [n_d] \subset \mathbb{Z}^d$ that uses the unit vectors $\{e_1, \dots, e_d\}$ as moves has the intersecting ray property. On the other hand, consider $\mathcal{F} = \{u \in \mathbb{N}^2 : u_1 + u_2 \leq 1\}$ and take the rays $\mathcal{R}_1 := \{(0, 0), (0, 1)\}$ and $\mathcal{R}_2 := \{(1, 0)\}$ that are parallel to e_2 and the rays $\mathcal{R}'_1 := \{(0, 1)\}$ and $\mathcal{R}'_2 := \{(0, 0), (1, 0)\}$ that are parallel to e_1 . Then $\mathcal{R}_1 \cap \mathcal{R}'_1 = \{(1, 0)\}$ and $\mathcal{R}_2 \cap \mathcal{R}'_2 = \{(0, 1)\}$, but $\mathcal{R}_1 \cap \mathcal{R}'_2 = \{(0, 0)\} \neq \emptyset$ and $\mathcal{R}'_1 \cap \mathcal{R}_2 = \emptyset$.

Proposition 5.1.14. Let $m, m' \in \mathbb{Z}^d$ be not collinear and $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set. The matrices $\mathcal{H}_{\mathcal{F}, m}^\pi$ and $\mathcal{H}_{\mathcal{F}, m'}^\pi$ commute if and only if (m, m') has the intersecting ray property in \mathcal{F} .

Proof. Let $u_1, u_2 \in \mathcal{F}$. Then

$$(\mathcal{H}_{\mathcal{F}, m}^\pi \cdot \mathcal{H}_{\mathcal{F}, m'}^\pi)_{u_1, u_2} = \begin{cases} |\mathcal{R}_{\mathcal{F}, m}(u_1)|^{-1} \cdot |\mathcal{R}_{\mathcal{F}, m'}(u_2)|^{-1}, & \text{if } \mathcal{R}_{\mathcal{F}, m}(u_1) \cap \mathcal{R}_{\mathcal{F}, m'}(u_2) \neq \emptyset \\ 0, & \text{otherwise} \end{cases}.$$

Let $\mathcal{R}_1 := \mathcal{R}_{\mathcal{F}, m}(u_1)$, $\mathcal{R}'_1 := \mathcal{R}_{\mathcal{F}, m'}(u_1)$, $\mathcal{R}_2 := \mathcal{R}_{\mathcal{F}, m}(u_2)$, and $\mathcal{R}'_2 := \mathcal{R}_{\mathcal{F}, m'}(u_2)$. Thus, it is straightforward to check that $(\mathcal{H}_{\mathcal{F}, m}^\pi \cdot \mathcal{H}_{\mathcal{F}, m'}^\pi)_{u_1, u_2} = (\mathcal{H}_{\mathcal{F}, m'}^\pi \cdot \mathcal{H}_{\mathcal{F}, m}^\pi)_{u_1, u_2}$ for all $u_1, u_2 \in \mathcal{F}$ if and only if the pair (m, m') has the intersecting ray property. \square

Lemma 5.1.15. Let $H_1, \dots, H_n \in \mathbb{R}^{n \times n}$ be pairwise commuting matrices. Then any eigenvalue of $\sum_{i=1}^n H_i$ has the form $\lambda_1 + \dots + \lambda_n$ where λ_i is an eigenvalue of H_i .

Proof. This is a straightforward extension of the case $n = 2$ in [67, Theorem 2.4.8.1] and relies on the fact that commuting matrices are simultaneously triangularizable. \square

Proposition 5.1.16. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets and suppose there exists $m \in \mathcal{M}$ such that (m, m') has the intersecting ray property in \mathcal{F} for all $m' \in \mathcal{M}' := \mathcal{M} \setminus \{m\}$. Let V_1, \dots, V_c be the connected components of $\mathcal{F}(\mathcal{M}')$, $\pi_i : V_i \rightarrow [0, 1]$ the uniform distribution, and $f' = (1 - f(m))^{-1} \cdot f|_{\mathcal{M}'}$, then

$$\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f'}) \leq f(m) + (1 - f(m)) \cdot \max\{\lambda(\mathcal{H}_{V_i, \mathcal{M}'}^{\pi_i, f'}) : i \in [c]\}.$$

Proof. Let $\mathcal{H} := \mathcal{H}_{\mathcal{F}, \mathcal{M}'}^{\pi, f'}$ be the heat-bath walk on $\mathcal{F}(\mathcal{M}')$ that samples moves from \mathcal{M}' according to f' , then $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f} = f(m) \cdot \mathcal{H}_{\mathcal{F}, m}^\pi + (1 - f(m)) \cdot \mathcal{H}$. By assumption, all pairs (m, m') with $m' \in \mathcal{M}'$ have the intersecting ray property and thus the matrices $\mathcal{H}_{\mathcal{F}, m}^\pi$ and \mathcal{H} commute according to Proposition 5.1.14. The eigenvalues of all involved matrices are non-negative and thus Lemma 5.1.15 implies that the second largest eigenvalue of $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$ has the form $\lambda + \lambda'$ where $\lambda \in \{0, f(m)\}$ by Proposition 5.1.6 and where λ' is an eigenvalue of $(1 - f(m)) \cdot \mathcal{H}$. The matrix \mathcal{H} is a block matrix whose blocks are the matrices $\mathcal{H}_{V_i, \mathcal{M}'}^{\pi_i, f'} = \mathcal{H}_{V_i, \mathcal{M}'}^{\pi_i, f'}$ and the statement follows. \square

Proposition 5.1.17. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^k$ be finite sets. If $\mathcal{F}^c(\mathcal{M})$ has the intersecting ray property, then $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) \leq 1 - \min(f)$.

Proof. Let $\mathcal{M} = \{m_1, \dots, m_k\}$. The intersecting ray property and Proposition 5.1.14 give that the matrices $f(m_1) \cdot \mathcal{H}_{\mathcal{F}, m_1}^\pi, \dots, f(m_k) \cdot \mathcal{H}_{\mathcal{F}, m_k}^\pi$ commute pairwise. According to Proposition 5.1.6, the spectrum of $f(m_i) \cdot \mathcal{H}_{\mathcal{F}, m_i}^\pi$ is $\{0, f(m_i)\}$. Lemma 5.1.15 yields that the second largest eigenvalue of $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$, which equals the second largest eigenvalue modulus since all of its eigenvalues are non-negative, fulfills $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) = \sum_{i \in I} f(m_i)$ for a subset $I \subseteq [k]$. Since $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) < 1$ and $\sum_{i=1}^k f(m_i) = 1$, we have $I \neq [k]$ and the claim follows. \square

Proposition 5.1.18. *Let $n_1, \dots, n_d \in \mathbb{N}_{>1}$, $\mathcal{F} = [n_1] \times \dots \times [n_d]$, and $\mathcal{M} = \{e_1, \dots, e_d\}$. Then for any positive mass function $f : \mathcal{M} \rightarrow [0, 1]$, $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) = 1 - \min(f)$.*

Proof. Since $\mathcal{F}^c(\mathcal{M})$ has the intersecting ray property, Proposition 5.1.17 shows $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) \leq 1 - \min(f)$. Assume that $\min(f) = f(e_i)$. The connected components of $\mathcal{F}^c(\{e_1, \dots, e_d\} \setminus \{e_i\})$ are the sets $V_j := \{u \in \mathcal{F} : u_i = j\}$ for any $j \in [n_i]$ and the rays through \mathcal{F} along e_i are $\mathcal{R}_k := \{k + s \cdot e_i : s \in [n_i]\}$ for $k \in [n_1] \times \dots \times [n_{i-1}] \times \{0\} \times [n_{i+1}] \times \dots \times [n_d]$. In particular, any ray intersects any connected component exactly once. Thus, the matrix $(|\mathcal{R}_k \cap V_j|)_{k,j}$ is the all-ones matrix, which has a non-trivial kernel. Proposition 5.1.11 implies $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) \geq 1 - f(e_i)$. \square

Remark 5.1.19. In the special case $n := n_1 = \dots = n_d$ and $f : \{e_1, \dots, e_d\} \rightarrow [0, 1]$ the uniform distribution in Proposition 5.1.18, the heat-bath walk on $[n]^d$ is known as *Rook's walk* in the literature. In this case, Proposition 5.1.18 appears as [77, Proposition 2.3]. In [88], upper bounds on the mixing time of the Rook's walk were obtained with *path-coupling*.

By the variational characterization of the eigenvalues of a hermite matrix, the second largest eigenvalue is also the optimal value of a maximization problem:

Proposition 5.1.20. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set, $\mathcal{M} = \{m_1, \dots, m_k\} \subset \mathbb{Z}^d$ be a Markov basis for \mathcal{F} , and π be the uniform distribution on \mathcal{F} . For any $i \in [k]$, let $\mathcal{R}_1^i, \dots, \mathcal{R}_{n_i}^i$ be the disjoint rays through \mathcal{F} along m_i . Then*

$$\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) = \max \left\{ \sum_{i=1}^k f(m_i) \cdot \sum_{j=1}^{n_i} \left(\frac{1}{|\mathcal{R}_j^i|} \sum_{x \in \mathcal{R}_j^i} \sum_{y \in \mathcal{R}_j^i} w_x w_y \right) : \sum_{u \in \mathcal{F}} w_u^2 = 1, \sum_{u \in \mathcal{F}} w_u = 0 \right\}. \quad (5.1.1)$$

Proof. Let $n := |\mathcal{F}|$ and let $\mathcal{B} := \{w \in \mathbb{R}^n : \sum_{i=1}^n w_i^2 = 1, \sum_{i=1}^n w_i = 0\}$. Denote the eigenvalues of $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$ by $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq -1$. All eigenvalues of $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}$ are non-negative, and hence $\lambda_n \geq 0$. In particular, $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) = \lambda_2$. The *second Stiefel manifold* is the set

$$\mathcal{V}_2(\mathbb{R}^n) := \left\{ X \in \mathbb{R}^{n \times 2} : X^T X = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right\} \subset \mathbb{R}^{n \times 2}.$$

Fan's theorem [53, Theorem 1] shows that $\max_{X \in \mathcal{V}_2(\mathbb{R}^n)} \text{trace}(X^T \mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f} X) = \lambda_1 + \lambda_2$. It is straightforward to check that

$$\text{trace}([v \ w]^T \mathcal{H}_{\mathcal{F}, m_i}^\pi [v \ w]) = \sum_{j=1}^{n_i} \left(\frac{1}{|\mathcal{R}_j^i|} \sum_{x \in \mathcal{R}_j^i} \sum_{y \in \mathcal{R}_j^i} (v_x v_y + w_x w_y) \right)$$

for any $v, w \in \mathbb{R}^n$ and $i \in [k]$. Since $\text{trace}(\cdot)$ is multiplicative and since $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f} = \sum_{i=1}^k f(m_i) \mathcal{H}_{\mathcal{F}, m_i}^{\pi}$,

$$\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) = \max_{(v, w) \in \mathcal{V}_2(\mathbb{R}^n)} \left\{ \sum_{i=1}^k \left(f(m_i) \cdot \sum_{j=1}^{n_i} \left(\frac{1}{|\mathcal{R}_j^i|} \sum_{x \in \mathcal{R}_j^i} \sum_{y \in \mathcal{R}_j^i} (v_x v_y + w_x w_y) \right) \right) - 1 \right\}.$$

Let $v := (1, \dots, 1)^T \frac{1}{\sqrt{n}} \in \mathbb{R}^n$ and $w \in \mathcal{B}$. By the definition of \mathcal{B} , $(v, w) \in \mathcal{V}_2(\mathbb{R}^n)$ and thus $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f})$ is greater than the term on the right-hand side of equation (5.1.1). Let w' be a normalized eigenvector of λ_2 , then $w'^T v = 0$ since v is an eigenvector of $\lambda_1 = 1$ by the spectral theorem of symmetric matrices. On the other hand, $\text{trace}([v \ w']^T \mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f} [v \ w']) = \lambda_1 + \lambda_2$ and thus $\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f})$ equals the term on the right-hand side of equation (5.1.1). \square

5.2 Augmenting Markov bases

It follows from our investigation in Section 2.2 that the diameter of all compressed fiber graphs coming from a fixed integer matrix $A \in \mathbb{Z}^{m \times d}$ can be bounded from above by a constant. However, Markov moves can be used twice in a minimal path which makes, a priori, the diameter of the compressed fiber graph larger than the size of the Markov basis. The next definition puts more constraints on the Markov basis and postulates the existence of a path that uses every move from the Markov basis at most once.

Definition 5.2.1. Let $\mathcal{F} \subset \mathbb{Z}^d$ be a finite set and $\mathcal{M} = \{m_1, \dots, m_k\} \subset \mathbb{Z}^d$. An *augmenting path* between distinct $u, v \in \mathcal{F}$ of length $r \in \mathbb{N}$ is a path in $\mathcal{F}^c(\mathcal{M})$ of the form

$$\left(u, u + \lambda_{i_1} m_{i_1}, u + \lambda_{i_1} m_{i_1} + \lambda_{i_2} m_{i_2}, \dots, u + \sum_{k=1}^r \lambda_{i_k} m_{i_k} = v \right) \in \mathcal{F}^{r+1}$$

with distinct $i_1, \dots, i_r \in [k]$. An augmenting path is *minimal* for $u, v \in \mathcal{F}$ if there exists no shorter augmenting path between u and v in $\mathcal{F}^c(\mathcal{M})$. A Markov basis \mathcal{M} for \mathcal{F} is *augmenting* if there is an augmenting path between any distinct nodes. The *augmentation length* $\mathcal{A}_{\mathcal{M}}(\mathcal{F})$ of an augmenting Markov basis \mathcal{M} is the maximum length of all minimal augmenting paths in $\mathcal{F}^c(\mathcal{M})$.

The diameter of compressed fiber graphs that use an augmenting Markov basis is at most the size of the Markov basis. Not every Markov basis is augmenting (Example 2.2.13), but we show that many natural sets of moves have this property. For fiber graphs coming from an integer matrix, an augmenting Markov basis for all of its fibers can be computed (Remark 5.2.3).

Example 5.2.2. Let A_d and \mathcal{M}_d as in Example 1.3.3. We show that \mathcal{M}_d is an augmenting Markov basis for $\mathcal{F}_{A_d, b}$ for any $b \in \mathbb{N}$. Let $u, v \in \mathcal{F}_{A_d, b}$ be distinct, then there exists $i \in [d]$ such that $u_i > v_i$ or $u_i < v_i$, thus, we can walk from u to $u' := u + (u_i - v_i)(e_1 - e_i)$ or from v to $v' := v + (v_i - u_i)(e_1 - e_i)$. In any case, after that augmentation, the pairs (u', v) and (v', u) coincide in the i -th coordinate and thus we find an augmenting path by induction on the dimension d . Since these paths use at most $d - 1$ edges, $\mathcal{A}_{\mathcal{M}_d}(\mathcal{F}_{A_d, b}) \leq d - 1$ for all $b \in \mathbb{N}$.

Remark 5.2.3. Let $A \in \mathbb{Z}^{m \times d}$ with $\ker_{\mathbb{Z}}(A) \cap \mathbb{N}^d = \{0\}$ and let $b \in \mathbb{N}A$. The Graver basis is clearly an augmenting Markov basis for $\mathcal{F}_{A, b}$ for any $b \in \mathbb{N}A$. We claim that when A is totally unimodular, then $\mathcal{A}_{\mathcal{G}_A}(\mathcal{F}_{A, b}) \leq d^2(\text{rank}(A) + 1)$. In particular, the augmentation length

is independent of the right-hand side b . Let $u, v \in \mathcal{F}_{A,b}$ be arbitrary and let $l_i := \min\{u_i, v_i\}$, $w_i := \max\{u_i, v_i\}$, and $c_i := \text{sign}(u_i - v_i) \in \{-1, 0, 1\}$ for $i \in \mathbb{N}$. Then v is the unique optimal value of the linear integer optimization problem

$$\min\{c^T x : Ax = b, l \leq x \leq w, x \in \mathbb{Z}^d\}.$$

A *discrete steepest decent* as defined in [32, Definition 3] using Graver moves needs at most $\|c\|_1 \cdot d \cdot (\text{rank}(A) + 1) \leq d^2 \cdot (\text{rank}(A) + 1)$ many augmentations from u to reach the optimal value v and [32, Corollary 8] ensures that every Graver move is used at most once. Notice that in [32], the variable x is constrained to $x \geq 0$ instead to $x \geq l$, but their argument works in fact for any lower bound.

The bound on the augmentation length of Graver bases from Remark 5.2.3 can be improved in situations where the entries of the Graver elements are from $\{-1, 0, 1\}$:

Proposition 5.2.4. *Let $A \in \mathbb{Z}^{m \times d}$ with $\mathcal{G}_A \subseteq \{-1, 0, 1\}^d$ and let $r \in \mathbb{N}$ such that every choice of r columns of A is linearly independent. Then for all $\mathcal{F} \in \mathfrak{F}_A$, we have $\mathcal{A}_{\mathcal{G}_A}(\mathcal{F}) \leq d - r$.*

Proof. Let $k := d - r$, $b \in \mathbb{N}A$, and pick any distinct $u, v \in \mathcal{F}_{A,b}$. There is an element $g_1 \in \mathcal{G}_A$ such that $u + g_1 \sqsubseteq v$ and $u + g_1 \in \mathcal{F}_{A,b}$. Let $\lambda_1 \in \mathbb{N}$ be maximal such that $u_1 := u + \lambda_1 g_1 \sqsubseteq v$, in particular $(u + \lambda_1 g_1)_{i_1} = v_{i_1}$ for some $i_1 \in [d]$ since $g_1 \in \{-1, 0, 1\}^d$. Again, there exists $g_2 \in \mathcal{G}_A$ such that $u_1 + g_2 \sqsubseteq v$, that is $(g_2)_{i_1} = 0$ and $u_1 + g_2 \in \mathcal{F}_{A,b}$. Choose $\lambda_2 \in \mathbb{N}$ maximal such that $u_1 + \lambda_2 g_2 \sqsubseteq v$, then $(u_1 + \lambda_2 g_2)_{i_2} = v_{i_2}$ for some $i_2 \in [d] \setminus \{i_1\}$. Iterating this procedure, we either arrive at v with fewer than k repetitions yielding an augmenting path of length at most k , or we have constructed at the k -th repetition distinct $i_1, \dots, i_k \in [d]$ and $\lambda_1, \dots, \lambda_k \in \mathbb{N}$ such that the entries of $u' := u + \sum_{i=1}^k \lambda_i g_i$ and v corresponding to the coordinates $\{i_1, \dots, i_k\}$ coincide. By construction, (u, u_1, \dots, u_k) is an augmenting path from u to $u' = u_k$ in the compression of $\mathcal{F}_{A,b}(\mathcal{G}_A)$. Let $a_1, \dots, a_d \in \mathbb{Z}^m$ be the columns of A and set $\mathcal{I} := [d] \setminus \{i_1, \dots, i_k\}$, then $\sum_{j \in \mathcal{I}} a_j u'_j = \sum_{j \in \mathcal{I}} a_j v_j$. Since $|\mathcal{I}| = d - k = r$, the assumption on the columns of A gives that $u'_j = v_j$ for all $j \in \mathcal{I}$ and thus $u' = v$. \square

We now show that the lower bound on the augmentation length observed in Example 5.2.2 cannot be improved. We first need the following easy statement:

Lemma 5.2.5. *Let $v_1, \dots, v_k \in \mathbb{Q}^d$ such that any $v \in \text{span}_{\mathbb{Q}}\{v_1, \dots, v_k\}$ can be represented by a linear combination of r vectors. Then $\dim(\text{span}_{\mathbb{Q}}\{v_1, \dots, v_k\}) \leq r$.*

Proof. Let \mathfrak{B} be the collection of all subsets of $\{v_1, \dots, v_k\}$ with cardinality r . By the assumption, $\cup_{B \in \mathfrak{B}} \text{span}_{\mathbb{Q}}\{B\} = \text{span}_{\mathbb{Q}}\{v_1, \dots, v_k\}$. Since $\dim(\text{span}_{\mathbb{Q}}\{B\}) \leq r$ for all $B \in \mathfrak{B}$ and since \mathfrak{B} is finite, the claim follows. \square

Proposition 5.2.6. *Let $\mathcal{P} \subset \mathbb{Q}^d$ be polytope and let $\mathcal{M} \subset \mathbb{Z}^d$ be an augmenting Markov basis for $\mathcal{F}_i := (i \cdot \mathcal{P}) \cap \mathbb{Z}^d$ for all $i \in \mathbb{N}$. Then $\dim(\mathcal{P}) \leq \sup_{i \in \mathbb{N}} \mathcal{A}_{\mathcal{M}}(\mathcal{F}_i)$.*

Proof. Without restricting generality, we can assume that $0 \in \mathcal{P}$. Let $V := \text{span}_{\mathbb{Q}}\{\mathcal{P}\}$ be the \mathbb{Q} -span of \mathcal{P} , then $\dim(\mathcal{P}) = \dim(V)$. We must have $\dim(\text{span}_{\mathbb{Q}}\{\mathcal{M}\}) = \dim(V)$ since $\dim(\mathcal{P}) = \dim(\text{conv}_{\mathbb{Q}}(\mathcal{F}_i))$ for i sufficiently large and since \mathcal{M} is a Markov basis for \mathcal{F}_i . Define $r := \sup_{i \in \mathbb{N}} \mathcal{A}_{\mathcal{M}}(\mathcal{F}_i)$. The statement is trivially true if $r = \infty$. Otherwise, choose a non-zero $v \in V$ and $u \in \text{relint}(\mathcal{P}) \subset \mathbb{Q}^d$ arbitrarily. Then there exists $\delta \in \mathbb{Q}_{>0}$ such that $u + \delta v \in \mathcal{P}$. Thus, $\frac{1}{\delta}u + v \in \frac{1}{\delta}\mathcal{P}$. Let $c \in \mathbb{N}_{\geq 1}$ such that $i := \frac{c}{\delta} \in \mathbb{N}$ and $w := \frac{c}{\delta}u \in \mathbb{Z}^d$. Then $w + cv = c(\frac{1}{\delta}u + v) \in$

$(i \cdot \mathcal{P}) \cap \mathbb{Z}^d = \mathcal{F}_i$. By assumption, there exists an augmenting path from w to $w + cv$ that uses r elements from \mathcal{M} . That is, the element cv from V can be written as a linear combination of r vectors from \mathcal{M} . Since v was chosen arbitrarily, Lemma 5.2.5 implies $\dim(\mathcal{P}) = \dim(V) \leq r$. \square

Remark 5.2.7. It is a consequence of Proposition 5.2.6 that for any matrix $A \in \mathbb{Z}^{m \times d}$ and an augmenting Markov basis \mathcal{M} , there exists $\mathcal{F} \in \mathcal{P}_A$ such that $\mathcal{A}_{\mathcal{M}}(\mathcal{F}) \geq \dim(\ker_{\mathbb{Z}}(A))$.

Our next result utilizes the techniques from [105] to bound the second largest eigenvalue modulus of random walks. To set up the machinery, $G = (V, E)$ be a graph. For any ordered pair of distinct nodes $(x, y) \in V \times V$, let $p_{x,y} \subseteq E$ be an (x, y) -path in G and let $\Gamma := \{p_{x,y} : (x, y) \in V \times V, x \neq y\}$ be the collection of these paths, then Γ is a *set of canonical paths*. For any edge $e \in E$, let $\Gamma_e := \{p \in \Gamma : e \in p\}$ be the set of paths from Γ that use e . For any symmetric random walk $\mathcal{H} : V \times V \rightarrow [0, 1]$ on G , set

$$\rho(\Gamma, \mathcal{H}) := \frac{\max\{|p| : p \in \Gamma\}}{|V|} \cdot \max_{\{u,v\} \in E} \frac{|\Gamma_{\{u,v\}}|}{\mathcal{H}(u, v)}.$$

Observe that symmetry of \mathcal{H} is needed to make $\rho(\Gamma, \mathcal{H})$ well-defined. The quantity $\rho(\Gamma, \mathcal{H})$ gives rise to an upper bound on the second largest eigenvalue of \mathcal{H} :

Lemma 5.2.8. *Let G be a graph, \mathcal{H} be a symmetric random walk on G , and Γ be a set of canonical paths in G . Then $\lambda_2(\mathcal{H}) \leq 1 - \frac{1}{\rho(\Gamma, \mathcal{H})}$.*

Proof. The stationary distribution of \mathcal{H} is the uniform distribution and thus the statement is a direct consequence of [105, Theorem 5], since $\rho(\Gamma, \mathcal{H})$ is an upper bound on the constant defined in [105, equation 4]. \square

Theorem 5.2.9. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite and let $\mathcal{M} := \{m_1, \dots, m_k\} \subset \mathbb{Z}^d$ be an augmenting Markov basis for \mathcal{F} . Let π be the uniform and f a positive distribution on \mathcal{F} and \mathcal{M} respectively. For $i \in [k]$, let $r_i := \max\{|\mathcal{R}_{\mathcal{F}, m_i}(u)| : u \in \mathcal{F}\}$ and suppose that $r_1 \geq r_2 \geq \dots \geq r_k$. Then*

$$\lambda(\mathcal{H}_{\mathcal{M}, \mathcal{F}}^{\pi, f}) \leq 1 - \frac{|\mathcal{F}| \cdot \min(f)}{\mathcal{A}_{\mathcal{M}}(\mathcal{F}) \cdot \mathcal{A}_{\mathcal{M}}(\mathcal{F})! \cdot 3^{\mathcal{A}_{\mathcal{M}}(\mathcal{F})-1} \cdot 2^{|\mathcal{M}|} \cdot r_1 r_2 \cdots r_{\mathcal{A}_{\mathcal{M}}(\mathcal{F})}}.$$

Proof. Choose for any distinct $u, v \in \mathcal{F}$ an augmenting path $p_{u,v}$ of minimal length in $\mathcal{F}^c(\mathcal{M})$ and let Γ be the collection of all these paths. Let $u + \mu m_k = v$ be an edge in $\mathcal{F}^c(\mathcal{M})$, then our goal is to bound $|\Gamma_{\{u,v\}}|$ from above. Let $\mathcal{S} := \{S \subseteq [r] : |S| \leq \mathcal{A}_{\mathcal{M}}(\mathcal{F}), k \in S\}$ and take any path $p_{x,y} \in \Gamma_{\{u,v\}}$. Then there exists a set $S := \{i_1, \dots, i_s\} \in \mathcal{S}$ of cardinality $s \leq \mathcal{A}_{\mathcal{M}}(\mathcal{F})$ such that $x + \sum_{j=1}^s \lambda_{i_j} m_{i_j} = y$. Since $p_{x,y}$ uses the edge $\{u, v\}$, there is $j \in [s]$ such that $i_j = k$ and $\lambda_{i_j} = \mu$. Since $|\lambda_{i_k}| \leq r_{i_k}$, there are at most

$$s! \cdot (2r_{i_1} + 1) \cdots (2r_{i_{j-1}} + 1) \cdot (2r_{i_{j+1}} + 1) \cdots (2r_{i_s} + 1) \leq s! \cdot 3^{s-1} \prod_{t \in S \setminus \{k\}} r_t$$

paths in $\Gamma_{\{u,v\}}$ that use the moves $m_{i_1}, \dots, m_{i_{j-1}}, m_{i_{j+1}}, \dots, m_{i_s}$. Since every path in $\Gamma_{\{u,v\}}$ uses moves indexed by some set in \mathcal{S} , $|\Gamma_{\{u,v\}}| \leq \sum_{S \in \mathcal{S}} (|S|! \cdot 3^{|S|-1} \prod_{t \in S \setminus \{k\}} r_t)$ and thus we get

$$\frac{|\Gamma_{\{u,v\}}|}{\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}(u, v)} \leq 3^{\mathcal{A}_{\mathcal{M}}(\mathcal{F})-1} \frac{\sum_{S \in \mathcal{S}} (|S|! \prod_{t \in S \setminus \{k\}} r_t)}{f(m_k) \cdot \frac{1}{|\mathcal{R}_{m_k}(u)|}} \leq \frac{3^{\mathcal{A}_{\mathcal{M}}(\mathcal{F})-1} \cdot \mathcal{A}_{\mathcal{M}}(\mathcal{F})! \cdot |\mathcal{S}| \cdot r_1 r_2 \cdots r_{\mathcal{A}_{\mathcal{M}}(\mathcal{F})}}{f(m_k)},$$

where we have used the assumption $r_1 \geq r_2 \geq \dots \geq r_k$. Bounding $|\mathcal{S}|$ from above by $2^{|\mathcal{M}|}$, the claim follows from Lemma 5.2.8 and Proposition 5.1.2, since all eigenvalues are non-negative. \square

The constants involved in the very general bound in Theorem 5.2.9 can be vastly improved in situations where one has more control over the set of moves and the structure of the paths (see Proposition 5.1.18). When the augmentation length of Markov bases, or their size, grows, then the upper bound in Theorem 5.2.9 is not very informative to decide whether a sequence mixes rapidly. However, for a fixed Markov basis in fixed dimension, all bad terms become constants. In this situation, we are left with the asymptotic comparison of the size of the saturated set and the length of the rays. Thus, the growth of the following quantity is essential:

Definition 5.2.10. Let $\mathcal{F} \subset \mathbb{Z}^d$ and $\mathcal{M} \subset \mathbb{Z}^d$ be finite sets. The longest ray through \mathcal{F} along vectors of \mathcal{M} is $\mathcal{R}_{\mathcal{F},\mathcal{M}} := \arg \max\{|\mathcal{R}_{\mathcal{F},m}(u)| : m \in \mathcal{M}, u \in \mathcal{F}\}$.

Corollary 5.2.11. Let $(\mathcal{F}_i)_{i \in \mathbb{N}}$ be a sequence of finite sets in \mathbb{Z}^d and let π_i be the uniform distribution on \mathcal{F}_i . Let $\mathcal{M} \subset \mathbb{Z}^d$ be an augmenting Markov basis for \mathcal{F}_i with $\mathcal{A}_{\mathcal{M}}(\mathcal{F}_i) \leq \dim(\mathcal{F}_i)$ and suppose that $(|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}|^{\dim(\mathcal{F}_i)})_{i \in \mathbb{N}} \in \mathcal{O}(|\mathcal{F}_i|)_{i \in \mathbb{N}}$. Then for any mass function $f : \mathcal{M} \rightarrow (0, 1]$, $(\mathcal{H}_{\mathcal{F}_i,\mathcal{M}}^{\pi_i,f})_{i \in \mathbb{N}}$ is an expander.

Proof. This is a straightforward application of Theorem 5.2.9. \square

Corollary 5.2.12. Let $\mathcal{P} \subset \mathbb{Q}^d$ be a polytope, $\mathcal{F}_i := (i \cdot \mathcal{P}) \cap \mathbb{Z}^d$ for $i \in \mathbb{N}$, and let π_i be the uniform distribution on \mathcal{F}_i . Let $\mathcal{M} \subset \mathbb{Z}^d$ be an augmenting Markov basis for $\{\mathcal{F}_i : i \in \mathbb{N}\}$ such that $\mathcal{A}_{\mathcal{M}}(\mathcal{F}_i) \leq \dim(\mathcal{P})$ for all $i \in \mathbb{N}$. Then for any mass function $f : \mathcal{M} \rightarrow (0, 1]$, the sequence $(\mathcal{H}_{\mathcal{F}_i,\mathcal{M}}^{\pi_i,f})_{i \in \mathbb{N}}$ is an expander.

Proof. Let $r := \dim(\mathcal{P})$. We first show that $(|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}|)_{i \in \mathbb{N}} \in \mathcal{O}(i)_{i \in \mathbb{N}}$. Write $\mathcal{M} = \{m_1, \dots, m_k\}$ and denote by $l_i := \max\{|(u + m_i \cdot \mathbb{Z}) \cap \mathcal{P}| : u \in \mathcal{P}\}$ the length of the longest ray through the polytope \mathcal{P} along m_i . It suffices to prove that $i \cdot (l_k + 1)$ is an upper bound on the length of any ray along m_k through \mathcal{F}_i . For that, let $u \in \mathcal{F}_i$ such that $u + \lambda m_k \in \mathcal{F}_i$ for some $\lambda \in \mathbb{N}$, then $\frac{1}{i}u + \frac{\lambda}{i}m_k \in \mathcal{P}$ and thus $\lfloor \frac{\lambda}{i} \rfloor \leq l_k$, which gives $\lambda \leq i \cdot (l_k + 1)$. With $C := \max\{l_1, \dots, l_k\} + 1$ we have $|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}| \leq C \cdot i$. Ehrhart's theorem [15, Theorem 3.23] gives $(|\mathcal{F}_i|)_{i \in \mathbb{N}} \in \Omega(i^r)_{i \in \mathbb{N}}$ and since $|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}| \leq C \cdot i$, we have $(|\mathcal{R}_{\mathcal{F}_i,\mathcal{M}}|^r)_{i \in \mathbb{N}} \in \mathcal{O}(|\mathcal{F}_i|)_{i \in \mathbb{N}}$. Corollary 5.2.11 gives the claim. \square

Example 5.2.13. Fix $d, r \in \mathbb{N}$ and consider the lattice points of the d -dimensional cross-polytope of radius r : $\mathcal{C}_{d,r} = \{u \in \mathbb{Z}^d : \|u\|_1 \leq r\}$. Then $\mathcal{M}_d = \{e_1, \dots, e_d\}$ is a Markov basis for $\mathcal{C}_{d,r}$ for any $r \in \mathbb{N}$. We show that \mathcal{M}_d is an augmenting Markov basis whose augmentation length is at most d . For that, let $u, v \in \mathcal{C}_{d,r}$ be distinct elements. We claim that there is $i \in [d]$ such that $u_i \neq v_i$ and $u_+ + (v_i - u_i)e_i \in \mathcal{C}_{d,r}$. Let $S \subseteq [d]$ be the set of indices where u and v differ and let $s := r - \|u\|_1$. If $|S| = 1$, the result is clear so assume $|S| \geq 2$. If the result does not hold then for all $i \in S$, $|v_i| - |u_i| > s$. Thus,

$$\|v\|_1 = \sum_{i \notin S} |u_i| + \sum_{i \in S} |v_i| > \sum_{i \notin S} |u_i| + \sum_{i \in S} s + |u_i| = |S_{uv}| \cdot s + \|u\|_1 = (|S| - 1) \cdot s + r.$$

But we assumed that $v \in \mathcal{C}_{d,r}$. It follows that for any pair of points u, v in $\mathcal{C}_{d,r}$, there is a walk, using the unit vectors as moves, that uses each move at most once. Corollary 5.2.11 yields that for any $d \in \mathbb{N}$, the second largest eigenvalue modulus of the heat-bath walk on $\mathcal{C}_{d,r}$ can be strictly bounded away from 1 for $r \rightarrow \infty$.

The next proposition demonstrates how a more careful construction of paths in compressed fiber graphs in the spirit of Theorem 5.2.9 leads to better bounds on the second largest eigenvalue:

Proposition 5.2.14. *Let $a \in \mathbb{N}_{>0}^d$, $b \in \mathbb{N}$, $\mathcal{F} = \{u \in \mathbb{N}^d : a^T \cdot u \leq b\}$, and $\mathcal{M} := \{e_1, \dots, e_d\}$. If π and f are the uniform distributions on \mathcal{F} and \mathcal{M} respectively, then*

$$\lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) \leq 1 - \frac{|\mathcal{F}|}{d^2} \prod_{i=1}^d \frac{a_i}{b}.$$

Proof. Observe that \mathcal{M} is a Markov basis for \mathcal{F} since all nodes are connected with $0 \in \mathcal{F}$. Let $u, v \in \mathcal{F}$ be distinct. We first show that there exists $k \in [d]$ such that $u_k \neq v_k$ and $u + (v_k - u_k)e_k \in \mathcal{F}$. If $u \leq v$, the statement trivially holds. Otherwise, there exists $k \in [d]$ such that $u_k > v_k$ and the vector obtained by replacing the k -th coordinate of u by v_k remains in \mathcal{F} . Now, consider for the following path between u and v : Choose the smallest index $k \in [d]$ such that $u_k \neq v_k$ and such that $u + (v_k - u_k) \cdot e_k \in \mathcal{F}$ and proceed recursively with $u + (v_k - u_k)$ and v . This gives a path $p_{u,v}$ between u and v of length at most d . Let Γ be the collection of all these paths. We want to apply Lemma 5.2.8. Thus, let $x \in \mathcal{F}$ and consider the edge $x \rightarrow x + c \cdot e_s$. Let us count the paths $p_{u,v}$ that use that edge. Let $u, v \in \mathcal{F}$ and let $k_1, \dots, k_r \in [d]$ be distinct indices such that

$$u \rightarrow u + (v_{k_1} - u_{k_1})e_{k_1} \rightarrow u + (v_{k_1} - u_{k_1})e_{k_1} + (v_{k_2} - u_{k_2})e_{k_2} \rightarrow \dots \rightarrow v$$

represents the path $p_{u,v}$ constructed as explained above. Assume that $p_{u,v}$ uses the edge $\{x, x + ce_s\}$ and let $k_l = s$ and $(v_{k_l} - u_{k_l}) = c$. In particular,

$$\begin{aligned} u + (v_{k_1} - u_{k_1})e_{k_1} + \dots + (v_{k_{l-1}} - u_{k_{l-1}})e_{k_{l-1}} &= x \\ x + (v_{k_l} - u_{k_l})e_{k_l} + \dots + (v_{k_r} - u_{k_r})e_{k_r} &= v. \end{aligned}$$

We see that $v_{k_t} = x_{k_t}$ for all $t < l$ and that $u_{k_t} = x_{k_t}$ for all $t \geq l$. In particular, $v_{k_l} = u_{k_l} + c = x_{k_l} + c$ is also fixed. The coordinates u_{k_t} and v_{k_t} are bounded from above by $\frac{b}{a_{k_t}}$ for all $t \in [r]$, and hence there can be at most

$$\left(\prod_{t=1}^{l-1} \frac{b}{a_{k_t}} \right) \cdot \left(\prod_{t=l+1}^r \frac{b}{a_{k_t}} \right).$$

Since k_1, \dots, k_t are distinct coordinate indices, we have

$$\frac{|\Gamma_{x, x+c \cdot e_s}|}{\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}(x, x + c \cdot e_s)} \leq d \cdot \prod_{i=1}^d \frac{b}{a_i}.$$

Lemma 5.2.8 finishes the proof. \square

The heat-bath walk mixes rapidly when an augmenting Markov basis with a small augmentation length is used. We think that it is interesting to ask how an augmenting Markov basis can be obtained and how the augmentation length can be improved.

Question 5.2.15. *Let $\mathcal{M} \subset \mathbb{Z}^d$ be an augmenting Markov basis of $A \in \mathbb{Z}^{m \times d}$. Are there finitely many moves $m_1, \dots, m_k \in \mathbb{Z}^d$ such that the augmentation length of $\mathcal{M} \cup \{m_1, \dots, m_k\}$ on $\mathcal{F}_{A,b}$ is at most $\dim(\ker_{\mathbb{Z}}(A))$ for all $b \in \mathbb{N}$?*

We believe that sampling with the heat-bath walk is always at least as fast as sampling with the simple fiber walk. To prove this, the following question has to be answered negatively:

Question 5.2.16. *Are there a matrix $A \in \mathbb{Z}^{m \times d}$, a right-hand side $b \in \mathbb{N}$, and a Markov basis $\mathcal{M} \subset \ker_{\mathbb{Z}}(A)$ such that $\frac{1}{2^{|\mathcal{M}|}} h(\mathcal{F}_{A,b}(\mathcal{M})) > \Phi(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f})$?*

5.3 Best move selection

The stationary distribution of the heat-bath walk on $\mathcal{F} \subset \mathbb{Z}^d$ is independent of the actual mass function on the Markov moves $\mathcal{M} \subset \mathbb{Z}^d$. The problem of finding the mass function which leads to the fastest mixing heat-bath walk can be formulated as the following optimization problem:

$$\min \left\{ \lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f}) : f : \mathcal{M} \rightarrow (0, 1], \sum_{m \in \mathcal{M}} f(m) = 1 \right\}. \quad (5.3.1)$$

Assume that $\mathcal{M} = \{m_1, \dots, m_k\}$ is a minimal Markov basis for \mathcal{F} , that is $\mathcal{F}(\mathcal{M} \setminus \{m_i\})$ is disconnected for all $i \in [k]$. The map $\lambda : \mathbb{R}^{n \times n} \rightarrow [0, 1]$, which maps an $n \times n$ matrix to its second largest eigenvalue, is a continuous and convex function [18, Section 2.1]. Every $\mu \in (0, 1)^k$ with $\sum_{i=1}^k \mu_i$ gives rise to a probability mass function on \mathcal{M} and the task is then to compute the second largest eigenvalue of

$$\mathcal{H}(\mu_1, \dots, \mu_k) := \sum_{i=1}^k \mu_i \cdot \mathcal{H}_{\mathcal{F}, m_i}^{\pi}.$$

Let $\mu \in [0, 1]^k$ with $\|\mu\|_1 = 1$ such that $\mu_i = 0$ for some $i \in [k]$. Then the random walk $\mathcal{H}(\mu)$ is reducible (Remark 1.1.8) and thus $\lambda(\mathcal{H}(\mu)) = 1$. On the other hand, assume that $\mu_i = 1$ and $\mu_j = 0$ for all $j \in [k]$ with $j \neq i$. Then $\mathcal{H}(\mu) = \mathcal{H}_{\mathcal{F}, m_i}^{\pi}$ thus $\lambda(\mathcal{H}(\mu)) = 1$ due to Proposition 5.1.6. Since none of these particular $\mu \in [0, 1]^k$ yields an optimal solution to (5.3.1), we can extend the set of distributions on \mathcal{M} in the minimization problem, which then becomes

$$\min \{ \lambda(\mathcal{H}(\mu)) : \mu \in \text{conv}_{\mathbb{R}}(e_1, \dots, e_k) \}.$$

For instance, it follows from Proposition 5.1.18 that the optimal value of (5.3.1) for $\mathcal{F} = [n_1] \times \dots \times [n_d]$, $\mathcal{M} = \{e_1, \dots, e_d\}$, and the uniform distribution π on \mathcal{F} is the uniform distribution on \mathcal{M} . Another example where the uniform distribution is the optimal solution to (5.3.1), but where the verification is more involved, is presented in Example 5.3.2. The next statement implies that, in fixed dimension, the asymptotic mixing behaviour of heat-bath walks does not depend on how the moves are selected:

Proposition 5.3.1. *Let $\mathcal{F} \subset \mathbb{Z}^d$ be finite, $\mathcal{M} \subset \mathbb{Z}^d$ be a Markov basis of \mathcal{F} and $\pi : \mathcal{F} \rightarrow (0, 1]$ be a positive mass function. Then for any mass functions $f_1, f_2 : \mathcal{M} \rightarrow (0, 1]$,*

$$\min \left\{ \frac{f_1(m)}{f_2(m)} : m \in \mathcal{M} \right\} \leq \frac{1 - \lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f_1})}{1 - \lambda(\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f_2})} \leq \max \left\{ \frac{f_1(m)}{f_2(m)} : m \in \mathcal{M} \right\}.$$

Proof. This is a straightforward comparison of $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f_1}$ and $\mathcal{H}_{\mathcal{F}, \mathcal{M}}^{\pi, f_2}$ with Lemma 4.1.14. \square

Example 5.3.2. Let $\mathcal{F} = [2] \times [5]$ as in Example 5.1.5 and consider $\mathcal{M} = \{e_1, 2e_1 + e_2\}$. We investigate for which $\mu \in (0, 1)$, the transition matrix $\mu \mathcal{H}_{\mathcal{F}, e_1}^{\pi} + (1 - \mu) \mathcal{H}_{\mathcal{F}, 2e_1 + e_2}^{\pi}$ has the smallest second largest eigenvalue modulus. The characteristic polynomial in $\mathbb{Q}[\mu, x]$ is

$$-\frac{1}{25}x^4(x-1)(\mu+x-1)^6(-5x^2+5x+2\mu^2-2\mu)(-5x^2+5x+4\mu^2-4\mu)$$

and hence the non-zero eigenvalues are

$$\begin{aligned} x_1(\mu) &:= 1, & x_2(\mu) &:= 1 - \mu, \\ x_3(\mu) &:= \frac{1}{2} \left[1 + \sqrt{1 + \frac{8}{5}(\mu^2 - \mu)} \right], & x_4(\mu) &:= \frac{1}{2} \left[1 - \sqrt{1 + \frac{8}{5}(\mu^2 - \mu)} \right], \\ x_5(\mu) &:= \frac{1}{2} \left[1 + \sqrt{1 + 4(\mu^2 - \mu)} \right], & x_6(\mu) &:= \frac{1}{2} \left[1 - \sqrt{1 + 4(\mu^2 - \mu)} \right]. \end{aligned}$$

It is straightforward to check that $x_5(\mu) > \frac{1}{2} > x_6(\mu)$, $x_3(\mu) > \frac{1}{2} > x_4(\mu)$. Since $\mu^2 - \mu < 0$ for $u \in (0, 1)$ and $x_3(\mu) \geq x_6(\mu)$. We can show that $x_4(\mu) \geq x_2(\mu)$ and thus

$$\lambda(\mu \mathcal{H}_{\mathcal{F}, e_1}^\pi + (1 - \mu) \mathcal{H}_{\mathcal{F}, 2e_1 + e_2}^\pi) = \frac{1}{2} \left[1 + \sqrt{1 + \frac{8}{5}(\mu^2 - \mu)} \right].$$

The fastest heat-bath walk on $\mathcal{F}(\mathcal{M})$ which converges to uniform is thus obtained for $\mu = \frac{1}{2}$, i.e. when the moves are selected uniformly. The second largest eigenvalue in this case is $\frac{1}{10}(5 + \sqrt{15}) \approx 0.887$, which is larger than the second largest eigenvalue of the heat-bath walk that selects uniformly from $\{e_1, e_2\}$ (see Proposition 5.1.18).

In our investigation on heat-bath walks with Markov bases, we have seen many cases where the uniform distribution on the Markov moves yields the fastest mixing behaviour among all mass functions on the moves, which brings us to the following question:

Question 5.3.3. *Is there a saturated set $\mathcal{F} \subset \mathbb{Z}^d$ and a Markov basis $\mathcal{M} \subset \mathbb{Z}^d$, where the uniform distribution on \mathcal{M} is not the optimal value of (5.3.1) with π being the uniform distribution.*

6 Parity binomial edge ideals

The binomial edge ideal of a graph was introduced independently in [63] and [93] and it constitutes an example of the beautiful interplay between algebra and graph theory [49, 87, 23, 10]. This chapter is about a very related class of binomial ideals that arise from a graph G . Let \mathbb{k} be any field and $\mathbb{k}[\mathbf{x}, \mathbf{y}] := \mathbb{k}[x_i, y_i : i \in V(G)]$ be the polynomial ring in $2|V(G)|$ indeterminates, then the *parity binomial edge ideal* of G is

$$\mathcal{I}_G := \langle x_i x_j - y_i y_j : \{i, j\} \in E(G) \rangle \subseteq \mathbb{k}[\mathbf{x}, \mathbf{y}].$$

edge ideals, but the combinatorics is subtler. Various properties related to walks in G depend on whether the walk has even or odd length (and hence the name). If G is bipartite, then everything reduces to binomial edge ideals as follows: Let $V(G) = V_1 \dot{\cup} V_2$ be a decomposition into independent sets and consider the ring automorphism of $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ which exchanges x_i and y_i if $i \in V_1$ and leaves all remaining indeterminates invariant. Under this automorphism, the parity binomial edge ideal is the image of the binomial edge ideal of G .

The definition of parity binomial edge ideals was suggested by Rafael Villarreal at the MOCCA Conference 2014 in Levico Terme. He asked if parity binomial edge ideals are radical. Theorem 6.4.5 combined with Remark 6.4.1 says that this is the case if and only if G is bipartite, or $\text{char}(\mathbb{k}) \neq 2$. We also compute the minimal primes of \mathcal{I}_G in Section 6.3. In Proposition 6.4.4, we write \mathcal{I}_G as an intersection of binomial ideals whose combinatorics is simpler, since then a short induction shows that, under the field assumption, all occurring intersections are radical and hence \mathcal{I}_G is radical. A primary decomposition of \mathcal{I}_G when $\text{char}(\mathbb{k}) = 2$ is given in Theorem 6.4.9.

The paper [64] contains a different analysis of radicality of parity binomial edge ideals. In characteristic two, the parity binomial edge ideal \mathcal{I}_G coincides with the ideal L_G defined there; thus radicality is clarified by their Theorem 1.2 which here appears as Remark 6.4.1. If the characteristic of \mathbb{k} is not two, the linear transformation $x_i \mapsto x_i - y_i$, $y_i \mapsto x_i + y_i$ maps the parity binomial edge ideal to the permanental edge ideal Π_G defined in [64, Section 3]. Radicality of this ideal is clarified in [64, Corollary 3.3] by means of a Gröbner bases calculation. Our approach here is different and was developed completely independently. In particular, our proof of radicality cannot use the Gröbner basis by Remark 6.2.1.

Section 6.1, Section 6.3, and Section 6.4 are based on the joint work [73] with Thomas Kahle and Camilo Sarmiento. The final publication is available at Springer via

<http://dx.doi.org/10.1007/s10801-015-0657-3>.

Inspired by [73, Section 3], the recent article [8] determines the universal Gröbner basis of the parity binomial edge ideal of complete graphs and poses a conjecture for the general case. In Section 6.2, we extend our Gröbner basis calculation of the parity binomial edge ideal from [73, Section 3] to its universal Gröbner basis proving the conjecture of [8] partially.

Convention. Let G be a graph. Throughout we assume that G is connected and in particular has no isolated nodes if $|V(G)| \geq 2$. We freely identify ideals of sub-polynomial rings of $\mathbb{k}[\mathbf{x}, \mathbf{y}]$

with their images in $\mathbb{k}[\mathbf{x}, \mathbf{y}]$. Likewise ideals of $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ that do not use some of the indeterminates are considered ideals of the respective subrings. For a sequence of nodes $P = (i_1, \dots, i_r) \in V(G)^r$, $G[P] := G[\{i_1, \dots, i_r\}]$. A binomial is *pure difference* if it equals the difference of two monomials.

6.1 Markov bases

We first study the lattice ideal $\mathcal{J}_G := \mathcal{I}_G : (\prod_{i \in V(G)} x_i y_i)^\infty$, which is an important factor in the primary decomposition of \mathcal{I}_G . Proposition 2.1.1 says that any of its binomial generating system is a Markov basis for its underlying lattice, whose combinatorics is studied now. Recall that walks, paths, cycles, and circuits in graphs have marked start and end nodes in this thesis.

Definition 6.1.1. Let G be a graph. A walk $P = (v_1, \dots, v_r) \in V(G)^r$ is *odd* (*even*) if its length $r - 1$ is odd (even). The *interior* of P is the set $\text{int}(P) = \{v_1, \dots, v_r\} \setminus \{v_1, v_r\}$.




Figure 6.1: A graph with an even walk, but no even path from 4 to 5. The interior of the walk $(4, 3, 1, 2, 3, 6, 5)$ is $\{1, 2, 3, 6\}$.

Observe that the interior of a cycle depends on the choice of its start and end node. We can associate to every (i, j) -walk P a binomial in $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ in the following way:

$$b(P) := \begin{cases} x_i x_j - y_i y_j, & \text{if } P \text{ is odd} \\ x_i y_j - y_i x_j, & \text{if } P \text{ is even} \end{cases}.$$

Multiplied with an appropriate monomial factor, these binomials become elements of \mathcal{I}_G :

Lemma 6.1.2. Let P be a walk in G and for $k \in \text{int}(P)$, let $t_k \in \{x_k, y_k\}$ arbitrary. Then

$$b(P) \cdot \prod_{k \in \text{int}(P)} t_k \in \mathcal{I}_G.$$

Proof. Let i be the start and j be the end node of P . We prove the statement by induction on the length r of P . If $r = 1$, the statement is true by definition, thus assume that $r > 1$. If $\text{int}(P) = \emptyset$, then P is either even with $i = j$, or P is odd, which implies that i has to be adjacent to j . In both cases, $b(P) \in \mathcal{I}_G$. If $\text{int}(P) \neq \emptyset$, pick a node $s \in \text{int}(P)$. Consider first the case that P is an odd walk. Exchanging the roles of i and j if necessary, we can assume that the (i, s) -subwalk of P is odd and that the (s, j) -subwalk is even. Using the induction hypothesis, the binomials corresponding to these walks multiplied with appropriate monomial factors are in \mathcal{I}_G . Now, if $t_s = x_s$, then

$$x_i x_s x_j \prod_{k \in \text{int}(P) \setminus s} t_k \equiv_{\mathcal{I}_G} y_i y_s x_j \prod_{k \in \text{int}(P) \setminus s} t_k \equiv_{\mathcal{I}_G} y_i x_s y_j \prod_{k \in \text{int}(P) \setminus s} t_k$$

where we have first applied a binomial corresponding to the odd (i, s) -subwalk (which may traverse j) and then a binomial corresponding to the even (s, j) -subwalk of P (which may

traverse i). If $t_s = y_s$, then we first apply the (s, j) -walk and then the (i, s) -walk. The induction step for an even walk is similar and omitted. \square

Remark 6.1.3. Lemma 6.1.2 also holds for odd cycles in which case we get that monomial multiples of $x_i^2 - y_i^2$ are contained in \mathcal{I}_G for any node i that is contained in the same connected component as an odd cycle.

Let $\{i, j\} \in E(G)$ and denote $m_{\{i,j\}} := e_i + e_j \in \mathbb{Z}^{V(G)}$, then the generator $x_i x_j - y_i y_j$ has exponent vector $(m_{\{i,j\}}, -m_{\{i,j\}})^T \in \mathbb{Z}^{2|V(G)|}$. The exponent vectors of generators of \mathcal{I}_G generate the lattice

$$\mathcal{L}_G := \mathbb{Z} \left\{ \begin{bmatrix} m_e \\ -m_e \end{bmatrix} : e \in E(G) \right\} = \text{im}_{\mathbb{Z}} \begin{bmatrix} A_G \\ -A_G \end{bmatrix} \subseteq \mathbb{Z}^{2|V(G)|},$$

where A_G is the incidence matrix of G . Thus, \mathcal{L}_G is the Lawrence lifting of $\text{im}_{\mathbb{Z}}(A_G) \subseteq \mathbb{Z}^{|V(G)|}$. A standard fact about Lawrence liftings is that the Graver basis of $\text{im}_{\mathbb{Z}}(A_G)$ can be lifted to a Graver basis of \mathcal{L}_G , which here equals the universal Gröbner basis and any minimal Markov basis of \mathcal{L}_G [14, Proposition 1.1]. To determine the Graver basis of the lattice $\text{im}_{\mathbb{Z}}(A_G)$, let

$$\begin{aligned} \mathcal{M}_G^{\text{odd}} &:= \{e_i + e_j : \text{there is an odd } (i, j)\text{-walk in } G\} \\ \mathcal{M}_G^{\text{even}} &:= \{e_i - e_j : \text{there is an even } (i, j)\text{-walk in } G\} \setminus \{0\}. \end{aligned}$$

Note in particular that if there is an odd (i, i) -walk, then $2 \cdot e_i \in \mathcal{M}_G^{\text{odd}}$.

Proposition 6.1.4. *The Graver basis of $\text{im}_{\mathbb{Z}}(A_G)$ is $\pm(\mathcal{M}_G^{\text{odd}} \cup \mathcal{M}_G^{\text{even}})$.*

Proof. According to Pottier's termination criterion [31, Algorithm 3.3], it suffices to check that the sum of two elements of $\pm(\mathcal{M}_G^{\text{odd}} \cup \mathcal{M}_G^{\text{even}})$ can be reduced to zero sign-consistently. If there are no cancellations in the sum, for example if the two summands have disjoint support, the sum is reduced by either of the summands. Cancellation among elements $e_{i_1} \pm e_{i_2}$ and $e_{j_1} \pm e_{j_2}$ can only occur if $|\{i_1, i_2, j_1, j_2\}| \leq 3$. Without loss of generality assume $i_2 = j_1$. Thus, if cancellation occurs, the sum of two proposed Graver elements must equal $\pm(e_{i_1} \pm e_{j_2})$ and this is either zero or another element in $\pm(\mathcal{M}_G^{\text{odd}} \cup \mathcal{M}_G^{\text{even}})$ by concatenation of walks. \square

Proposition 6.1.5. *For any graph G , $\mathcal{J}_G = \langle b(P) : P \text{ is a walk in } G \rangle$.*

Proof. This is Proposition 6.1.4 and [14, Proposition 1.1]. \square

Example 6.1.6. Due to the odd cycle in the graph G in Figure 6.1, for all pairs (i, j) of nodes with $i \neq j$, both $x_i x_j - y_i y_j$ and $x_i y_j - x_j y_i$ are contained in \mathcal{J}_G . Hence, the ideal \mathcal{J}_G has 15 generators for odd walks and 15 for even walks with disjoint endpoints. Since G is not bipartite, $x_i^2 - y_i^2 \in \mathcal{J}_G$ for all $i \in [6]$. In total, a minimal Markov basis of \mathcal{J}_G consists of 36 generators.

Remark 6.1.7. If G is bipartite, the reachability of nodes with even or odd walks is determined by membership in the color classes. Consequently, for each spanning tree $T \subseteq G$ we have $\mathcal{J}_T = \mathcal{J}_G$. This is not true if G has an odd cycle.

6.2 Universal Gröbner basis

In this section, we extend the computation of the lexicographic Gröbner basis of parity binomial edge ideals from [73, Section 3] to the computation of their universal Gröbner basis. Recall that the universal Gröbner basis of an ideal is the union of all its reduced Gröbner bases. We show that the binomials in the universal Gröbner bases arise from walks as follows: Let P be an (i, j) -walk in G , then a *walk binomial* of P is a binomial $b(P) \cdot \prod_{k \in \text{int}(P)} t_k$ with $t_k \in \{x_k, y_k\}$ arbitrary. When the graph has an odd cycle, then the universal Gröbner basis cannot be square-free:

Remark 6.2.1. The parity binomial edge ideal \mathcal{I}_{C_3} of the 3-circuit C_3 cannot have a square-free initial ideal with respect to any monomial order. This follows from the fact that \mathcal{I}_{C_3} is not radical in $\mathbb{F}_2[\mathbf{x}, \mathbf{y}]$ (Remark 6.4.1). If \mathcal{I}_{C_3} had a square-free Gröbner basis over some field \mathbb{k} , its binomials must be pure difference (since the generators of \mathcal{I}_{C_3} are pure difference). The pure difference property yields that this Gröbner basis would also be a square-free Gröbner basis over every other field, in particular, over \mathbb{F}_2 .

By Lemma 6.1.2, all walk binomials of P are elements of \mathcal{I}_G . The following condition on a walk guarantees that its walk binomials are elements of the universal Gröbner basis.

Definition 6.2.2. An (i, j) -walk P in G is *minimal* if for no $k \in \text{int}(P)$, there is an (i, j) -walk with the same parity as P in $G[P \setminus \{k\}]$. The set of all walk binomials that come from minimal walks in G is denoted by \mathcal{S}_G .

There can be infinitely many minimal walks between two nodes (Figure 6.2), but since the monomial part of any walk binomial is square-free and depends only on the interior of the walk, the set \mathcal{S}_G is finite. For a given ideal $\mathcal{I} \subset \mathbb{k}[x_1, \dots, x_n]$, a binomial $\prod_{i=1}^n x_i^{u_i} - \prod_{i=1}^n x_i^{v_i} \in \mathcal{I}$ is *primitive in \mathcal{I}* if there exists no binomial $\prod_{i=1}^n x_i^{a_i} - \prod_{i=1}^n x_i^{b_i} \in \mathcal{I}$ such that $a \leq u$ and $b \leq v$. This is a natural extension of the definition of primitive binomials in toric ideals from [109, Chapter 4]. A straightforward generalization of [109, Lemma 4.6] beyond toric ideals shows that every element of the universal Gröbner basis of a pure difference ideal is primitive. In [44, 22], the set of all primitive binomials of an ideal \mathcal{I} is called the *Graver basis* of \mathcal{I} and we follow their notation. It is not hard to see that the Graver basis of a pure difference binomial ideal is a finite set [22, Proposition 4.3]. The following was conjectured in [8]:

Conjecture 6.2.3. *For any graph G , the universal Gröbner basis of \mathcal{I}_G , the set \mathcal{S}_G , and the Graver basis of \mathcal{I}_G coincide.*




Figure 6.2: A graph with infinitely many minimal $(1, 5)$ -walks of even and odd length.

It was shown in [8] that Conjecture 6.2.3 is true for complete graphs. We now prove the first part of this conjecture, i.e. the universal Gröbner basis of \mathcal{I}_G equals \mathcal{S}_G for any graph G . The first step is to reduce walk binomials as in Lemma 6.1.2 to zero by minimal walk binomials:

Lemma 6.2.4. *Let P be a walk in G and $t_k \in \{x_k, y_k\}$ for $k \in \text{int}(P)$. Then $b(P) \cdot \prod_{k \in \text{int}(P)} t_k$ reduces to zero modulo \mathcal{S}_G with respect to any monomial ordering.*

Proof. If P is minimal, then all its walk binomials are in \mathcal{S}_G and the statement holds trivially. Assume differently that P is a non-minimal (i, j) -walk. Then there exists a minimal (i, j) -walk P' of the same parity as P with $\text{int}(P') \subsetneq \text{int}(P)$. Thus, $b(P') = b(P)$ and hence the walk binomial of P is a monomial multiple of the walk binomial $b(P') \cdot \prod_{i \in \text{int}(P')} t_k \in \mathcal{S}_G$. \square

The main theorem of this section relies on Buchberger's criterion [76, Theorem 9.12] and we now recall briefly how it works for pure difference binomial ideals. Let $m, m' \in \mathbb{k}[\mathbf{x}]$ be monomials, $b_1, \dots, b_r \in \mathbb{k}[\mathbf{x}]$ be pure difference binomials, and \prec a monomial ordering. Then m can be *reduced* to m' by $\{b_1, \dots, b_r\}$ with respect to \prec if there exist monomials w_1, \dots, w_{r+1} with $w_1 = m$ and $w_{r+1} = m'$ such that $w_i - w_{i+1}$ is a binomial multiple of b_i and such that the leading monomial of b_i divides w_i for all $i \in [r+1]$. A pure difference binomial $b \in \mathbb{k}[\mathbf{x}]$ can be *reduced to zero* by a set of binomials $\mathcal{B} \subseteq \mathbb{k}[\mathbf{x}]$ with respect to \prec if there exists $b_1, \dots, b_r \in \mathcal{B}$ that reduce the leading monomial of b to its trailing monomial with respect to \prec . Since this implies that $0 \in \mathbb{k}[\mathbf{x}]$ is a \prec -*normal-form* of b (see [76, Definition 9.6]) with respect to \mathcal{B} , the following theorem immediately implies that Buchberger's criterion holds for \mathcal{S}_G .

Theorem 6.2.5. *For any monomial ordering \prec on $\mathbb{k}[\mathbf{x}, \mathbf{y}]$, the s -polynomial of any two elements of \mathcal{S}_G reduces to zero modulo \mathcal{S}_G and \prec .*

Remark 6.2.6. Gröbner bases of binomial edge ideals [63] look similar to \mathcal{S}_G , and also its determination in [100] uses related techniques on the paths of the underlying graph. However, there are also many differences: Every reduced Gröbner basis of a binomial edge ideals is square-free and it suffices to consider paths instead of walks. Both is false for parity binomial edge ideals and this makes the reduction of their s -polynomials more involved and technical.

The proof of Theorem 6.2.6 splits into a couple of lemmas, all of them dealing with reductions of different s -polynomials. We state how every possible s -polynomial of elements from \mathcal{S}_G can be reduced to zero and the following definition keeps the notation a lot simpler: For a set $P \subseteq V(G)$, we abbreviate $\mathbf{x}(P) := \prod_{p \in P} x_p$ and $\mathbf{y}(P) := \prod_{p \in P} y_p$. Before starting with the reduction of the s -polynomial of two walk binomials coming from even walks, we need the following lemma:

Lemma 6.2.7. *Let $P, Q \subseteq V(G)$ and $p, q \in V(G)$ with $p \notin P$ and $q \notin Q$, then*

$$\mathbf{x}(Q) \cdot \mathbf{x}((P \cup p) \setminus (Q \cup q)) = \mathbf{x}(p \setminus q) \cdot \mathbf{x}((P \cup Q) \setminus \{p, q\}).$$

Proof. The proof is immediate by the case distinction $p = q$, $p \in Q$, or $p \notin Q \cup \{q\}$. \square

Lemma 6.2.8. *Let $g_P, g_Q \in \mathcal{S}_G$ be walk binomials of even walks P and Q . Then $\text{spol}(g_P, g_Q)$ reduces to zero with respect to \mathcal{S}_G and any monomial ordering \prec .*

Proof. Let P be an even (p_1, p_2) -walk and Q be an even (q_1, q_2) -walk with $p_1 \neq p_2$ and $q_1 \neq q_2$. By exchanging p_1 with p_2 or q_1 with q_2 if necessary, we can assume that $x_{p_1} y_{p_2} \succ y_{p_1} x_{p_2}$ and $x_{q_1} y_{q_2} \succ y_{q_1} x_{q_2}$. Let $P^x, P^y \subseteq \text{int}(P)$ be the indices of the monomial part of g_P that correspond to indeterminates of $\mathbb{k}[\mathbf{x}]$ and $\mathbb{k}[\mathbf{y}]$ respectively and define $Q^x, Q^y \subseteq \text{int}(Q)$ for g_Q analogously, then $g_P = b(P) \cdot \mathbf{x}(P^x) \cdot \mathbf{y}(P^y)$ and $g_Q = b(Q) \cdot \mathbf{x}(Q^x) \cdot \mathbf{y}(Q^y)$. By our assumption, the leading

monomial of g_P is $\mathbf{x}(P^x \cup p_1) \cdot \mathbf{y}(P^y \cup p_2)$ and the leading monomial of g_Q is $\mathbf{x}(Q^x \cup q_1) \cdot \mathbf{y}(Q^y \cup q_2)$. Hence, their s-polynomial $\text{spol}(g_P, g_Q)$ is the difference of the monomial

$$\begin{aligned} & y_{q_1} x_{q_2} \cdot \mathbf{x}(Q^x) \mathbf{y}(Q^y) \cdot \mathbf{x}((P^x \cup p_1) \setminus (Q^x \cup q_1)) \cdot \mathbf{y}((P^y \cup p_2) \setminus (Q^y \cup q_2)) = \\ & = y_{q_1} x_{q_2} \cdot \mathbf{x}(p_1 \setminus q_1) \cdot \mathbf{y}(p_2 \setminus q_2) \cdot \mathbf{x}((P^x \cup Q^x) \setminus \{p_1, q_1\}) \cdot \mathbf{y}((P^y \cup Q^y) \setminus \{p_2, q_2\}) \end{aligned}$$

and the monomial

$$\begin{aligned} & y_{p_1} x_{p_2} \cdot \mathbf{x}(P^x) \cdot \mathbf{y}(Q^y) \cdot \mathbf{x}((Q^x \cup q_1) \setminus (P^x \cup p_1)) \cdot \mathbf{y}((Q^y \cup q_2) \setminus (P^y \cup p_2)) = \\ & = y_{p_1} x_{p_2} \cdot \mathbf{x}(q_1 \setminus p_1) \cdot \mathbf{y}(q_2 \setminus p_2) \cdot \mathbf{x}((P^x \cup Q^x) \setminus \{p_1, q_1\}) \cdot \mathbf{y}(P^y \cup Q^y \setminus \{p_2, q_2\}), \end{aligned}$$

where we have applied Lemma 6.2.7 two times respectively. Thus, for

$$b := y_{q_1} x_{q_2} \cdot \mathbf{x}(p_1 \setminus q_1) \cdot \mathbf{y}(p_2 \setminus q_2) - y_{p_1} x_{p_2} \cdot \mathbf{x}(q_1 \setminus p_1) \cdot \mathbf{y}(q_2 \setminus p_2)$$

and

$$m := \mathbf{x}((P^x \cup Q^x) \setminus \{p_1, q_1\}) \cdot \mathbf{y}(P^y \cup Q^y \setminus \{p_2, q_2\}),$$

the s-polynomial satisfies $\text{spol}(g_P, g_Q) = (-1)^r b \cdot m$, where the parity of $r \in \{1, 2\}$ depends on which of the monomials $y_{q_1} x_{q_2} \cdot \mathbf{x}(p_1 \setminus q_1) \cdot \mathbf{y}(p_2 \setminus q_2)$ and $y_{p_1} x_{p_2} \cdot \mathbf{x}(q_1 \setminus p_1) \cdot \mathbf{y}(q_2 \setminus p_2)$ is greater with respect to \prec . Now we go through all the cases for the end-points of P and Q .

First, assume that $|\{p_1, p_2, q_1, q_2\}| = 4$. To demonstrate all the subtleties of the reduction, this case is shown in detail. We have $b = x_{p_1} y_{p_2} y_{q_1} x_{q_2} - y_{p_1} x_{p_2} x_{q_1} y_{q_2}$ and since $x_{p_1} y_{p_2} \succ y_{p_1} x_{p_2}$ and $x_{q_1} y_{q_2} \succ y_{q_1} x_{q_2}$ by assumption, any monomial of b can be divided by the leading monomial of either $b(P)$ or $b(Q)$. Assume that $r = 2$, i.e. $y_{q_1} x_{q_2} x_{p_1} y_{p_2} \succ y_{p_1} x_{p_2} x_{q_1} y_{q_2}$. We first want to reduce $\text{spol}(g_P, g_Q)$ by a binomial $b' = (x_{p_1} y_{p_2} - y_{p_1} x_{p_2}) \prod_{k \in \text{int}(P)} t_k$ that corresponds to the minimal (p_1, p_2) -walk P . For all $k \in \text{int}(P) \setminus \{q_1, q_2\}$, we choose t_k as in the monomial factor of g_P . If P traverses q_1 or q_2 , we set $t_{q_1} = y_{q_1}$ and $t_{q_2} = x_{q_2}$. The leading term of b' divides the leading term of $\text{spol}(g_P, g_Q)$ and thus the monomial $x_{p_1} y_{p_2} y_{q_1} x_{q_2} \cdot m$ can be reduced to $y_{p_1} x_{p_2} y_{q_1} x_{q_2} \cdot m$ by b' . That is, $\text{spol}(g_P, g_Q)$ can be reduced to

$$(y_{p_1} x_{p_2} x_{q_1} y_{q_2} - y_{p_1} x_{p_2} y_{q_1} x_{q_2}) \cdot m = b(Q) \cdot y_{p_1} x_{p_2} \cdot m.$$

This binomial is a monomial multiple of a walk binomial that corresponds to the (q_1, q_2) -walk Q with an analogous modification of variables in $\{p_1, p_2\} \setminus \text{int}(Q)$. In particular, $\text{spol}(g_P, g_Q)$ reduces to zero by elements of \mathcal{S}_G with respect to \prec . The case $r = 1$ is similar and omitted.

Assume now that $|\{p_1, p_2, q_1, q_2\}| = 3$. First, suppose that $v := p_1 = q_1$ and $p_2 \neq q_2$, then $b = (y_{p_2} x_{q_2} - x_{p_2} y_{q_2}) \cdot y_v$. Gluing the even walks P and Q along their common start node v , we obtain an even (q_2, p_2) -walk W with $\text{int}(W) \subseteq \text{int}(P) \cup \text{int}(Q) \cup v$. Since $\text{spol}(g_P, g_Q) = -b(W) \cdot y_v \cdot m$, the s-polynomial is a monomial multiple of a suitable walk binomial b' of W that might need y_v in its monomial factor. Since b' can be reduced to zero by Lemma 6.2.4, $\text{spol}(g_P, g_Q)$ can be as well. The case $q_2 = p_2$ and $p_1 \neq q_1$ works similar for a binomial corresponding to a (p_1, q_1) -walk and is thus omitted. The next subcase within the case $|\{p_1, p_2, q_1, q_2\}| = 3$ is $v := p_2 = q_1$ and $p_1 \neq q_2$. Here, $b = y_v^2 x_{q_2} x_{p_1} - x_v^2 y_{p_1} y_{q_2}$ and the reduction works similar by applying suitable walk binomials of P and Q . The case $q_1 = p_2$ and $p_2 \neq q_2$ is similar.

Finally, if $|\{p_1, p_2, q_1, q_2\}| = 2$, then due to $p_1 \neq p_2$ and $q_1 \neq q_2$, the only possible case we have to consider is $p_1 = q_1$ and $p_2 = q_2$, since the case $p_1 = q_2$ and $p_2 = q_1$ contradicts our assumption $y_{q_1} x_{q_2} \succ x_{q_1} y_{q_2}$. But then we have $b = 0$. \square

Due to symmetry, the monomial ordering plays only a minor role in the reduction of s-polynomials coming from two even walks. For odd walks, there are more cases to distinguish, but a similar symmetry argument reduces the number of cases to consider as well:

Remark 6.2.9. Let P be an odd (p_1, p_2) -walk and Q be an odd (q_1, q_2) -walk. Which monomial is the leading monomial in walk binomials of P and Q depends on the following cases:

- $x_{p_1}x_{p_2} \succ y_{p_1}y_{p_2}$ and $x_{q_1}x_{q_2} \succ y_{q_1}y_{q_2}$,
- $x_{p_1}x_{p_2} \prec y_{p_1}y_{p_2}$ and $x_{q_1}x_{q_2} \prec y_{q_1}y_{q_2}$,
- $x_{p_1}x_{p_2} \succ y_{p_1}y_{p_2}$ and $x_{q_1}x_{q_2} \prec y_{q_1}y_{q_2}$, and
- $x_{p_1}x_{p_2} \prec y_{p_1}y_{p_2}$ and $x_{q_1}x_{q_2} \succ y_{q_1}y_{q_2}$.

After exchanging the roles of x_i and y_i symbolically, the reduction of $\text{spol}(g_P, g_Q)$ in the first and the second case works similarly. An exchange of the roles of P and Q shows that the same is true for the third and the fourth case. In the following, we thus only demonstrate the cases:

- $x_{p_1}x_{p_2} \succ y_{p_1}y_{p_2}$ and $x_{q_1}x_{q_2} \succ y_{q_1}y_{q_2}$,
- $x_{p_1}x_{p_2} \succ y_{p_1}y_{p_2}$ and $x_{q_1}x_{q_2} \prec y_{q_1}y_{q_2}$.

Our proof that the s-polynomial of two binomials of odd walks reduces to zero distinguishes the cases that at least one of them is a cycle (Lemma 6.2.11) or none of them is (Lemma 6.2.10).

Lemma 6.2.10. *Let $g_P, g_Q \in \mathcal{S}_G$ be walk binomials of an odd (p_1, p_2) -walk P and an odd (q_1, q_2) -walk Q respectively. If either $p_1 = p_2$ or $q_1 = q_2$, then $\text{spol}(g_P, g_Q)$ reduces to zero modulo \mathcal{S}_G with respect to any monomial ordering \prec .*

Proof. Let us first assume that $p := p_1 = p_2$ and $q := q_1 = q_2$. Clearly we only have to consider the case $p \neq q$. Then

$$\text{spol}(g_P, g_Q) = \begin{cases} (x_p^2 y_q^2 - x_q^2 y_p^2) \cdot \mathbf{x}((P^x \cup Q^x) \setminus \{p, q\}) \cdot \mathbf{y}(P^y \cup Q^y), & \text{if } x_p^2 \succ y_p^2 \text{ and } x_q^2 \succ y_q^2 \\ (y_p^2 y_q^2 - x_p^2 x_q^2) \cdot \mathbf{x}((P^x \cup Q^x) \setminus p) \cdot \mathbf{y}((P^y \cup Q^y) \setminus q), & \text{if } x_p^2 \succ y_p^2 \text{ and } x_q^2 \prec y_q^2 \end{cases}$$

and it suffices to consider these cases by Remark 6.2.9. Let b_P be the walk binomial of P whose monomial factor equals the monomial factor g_P on all nodes from $\text{int}(P) \setminus q$ and which uses the variable y_q in the first (i.e. $x_p^2 \succ y_p^2$ and $x_q^2 \succ y_q^2$) and x_q in the second case (i.e. $x_p^2 \succ y_p^2$ and $x_q^2 \prec y_q^2$), provided that $q \in \text{int}(P)$. Similarly, let b_Q be the walk binomial of Q whose monomial factor equals the monomial factor of g_Q on all nodes from $\text{int}(Q) \setminus p$ and which uses the variable y_p if $p \in \text{int}(Q)$. Then $\text{spol}(g_P, g_Q)$ can be reduced to zero by an application of b_P and b_Q in the first case, and an application of b_Q and b_P in the second case.

For the remainder, assume $p := p_1 = p_2$ and $q_1 \neq q_2$. If we have $p = q_1$, then $\text{spol}(g_P, g_Q)$ is

$$\begin{cases} (y_p x_{q_2} - x_p y_{q_2}) \cdot y_p \cdot \mathbf{y}(P^y \cup Q^y) \cdot \mathbf{x}((P^x \cup Q^x) \setminus q_2), & \text{if } x_p^2 \succ y_p^2 \wedge x_p x_{q_2} \succ y_p y_{q_2} \\ (y_p^3 y_{q_2} - x_p^3 x_{q_2}) \cdot \mathbf{x}((P^x \cup Q^x) \setminus p) \cdot \mathbf{y}((P^y \cup Q^y) \setminus q_2), & \text{if } x_p^2 \succ y_p^2 \wedge x_p x_{q_2} \prec y_p y_{q_2} \end{cases}$$

In the first case, the s-polynomial is a monomial multiple of a walk binomial belonging to the even (p, q_2) -walk that arises from gluing the odd (p, p) -walk with the odd (q_1, q_2) -walk along $p = q_1$. Lemma 6.2.4 gives this case. In the second case, we successively apply walk binomials from P and Q that use appropriate variables. The case $p = q_2$ is similar. The case $p_1 \neq p_2$ and $q_1 = q_2$ follows by symmetry. Finally, assume that $p \notin \{q_1, q_2\}$, then $\text{spol}(g_P, g_Q)$ equals

$$\begin{cases} (y_p^2 x_{q_1} x_{q_2} - x_p^2 y_{q_1} y_{q_2}) \cdot \mathbf{x}((P^x \cup Q^x) \setminus \{q_1, q_2, p\}) \cdot \mathbf{y}(P^y \cup Q^y), & \text{if } x_p^2 \succ y_p^2 \wedge x_{q_1} x_{q_2} \succ y_{q_1} y_{q_2} \\ (y_p^2 y_{q_1} y_{q_2} - x_p^2 x_{q_1} x_{q_2}) \cdot \mathbf{x}((P^x \cup Q^x) \setminus p) \cdot \mathbf{y}((P^y \cup Q^y) \setminus \{q_1, q_2\}), & \text{if } x_p^2 \succ y_p^2 \wedge x_{q_1} x_{q_2} \prec y_{q_1} y_{q_2} \end{cases}$$

and its reduction to zero works similarly. \square

Lemma 6.2.11. *Let $g_P, g_Q \in \mathcal{S}_G$ be walk binomials of an odd (p_1, p_2) -walk P and an odd (q_1, q_2) -walk Q respectively. If $p_1 \neq p_2$ and $q_1 \neq q_2$, then $\text{spol}(g_P, g_Q)$ reduces to zero modulo \mathcal{S}_G with respect to any monomial ordering \prec .*

Proof. By straightforward computations, if $x_{p_1}x_{p_2} \succ y_{p_1}y_{p_2}$ and $x_{q_1}x_{q_2} \succ y_{q_1}y_{q_2}$, then

$$\text{spol}(g_P, g_Q) = (y_{p_1}y_{p_2}x_{q_1}x_{q_2} - y_{q_1}y_{q_2}x_{p_1}x_{p_2}) \cdot \mathbf{x}((P^x \cup Q^x) \setminus \{p_1, p_2, q_1, q_2\}) \cdot \mathbf{y}(P^y \cup Q^y)$$

and if $x_{p_1}x_{p_2} \succ y_{p_1}y_{p_2}$ and $x_{q_1}x_{q_2} \prec y_{q_1}y_{q_2}$, then

$$\text{spol}(g_P, g_Q) = (y_{p_1}y_{p_2}y_{q_1}y_{q_2} - x_{q_1}x_{q_2}x_{p_1}x_{p_2}) \cdot \mathbf{x}((P^x \cup Q^y) \setminus \{p_1, p_2\}) \cdot \mathbf{y}((P^y \cup Q^y) \setminus \{q_1, q_2\}).$$

In both cases, the s-polynomial of g_P and g_Q can be reduced to zero by walk binomials from \mathcal{S}_G corresponding to P and Q , where we choose the variables corresponding to nodes in $\{q_1, q_2\} \setminus \text{int}(P)$ and $\{p_1, p_2\} \setminus \text{int}(Q)$ appropriately. \square

Lemma 6.2.12. *Let $g_P, g_Q \in \mathcal{S}_G$ be binomials of an odd walk P and an even walk Q respectively. Then $\text{spol}(g_P, g_Q)$ reduces to zero modulo \mathcal{S}_G with respect to any monomial ordering \prec .*

Proof. Let P be an odd (p_1, p_2) -walk and Q be an even (q_1, q_2) -walk in G . Let $P^x, P^y, Q^x, Q^y \subseteq V(G)$ such that $g_P = b(P) \cdot \mathbf{x}(P^x) \cdot \mathbf{y}(P^y)$ and $g_Q = b(Q) \cdot \mathbf{x}(Q^x) \cdot \mathbf{y}(Q^y)$. We only demonstrate the case where $p_1 \neq p_2$, $x_{p_1}x_{p_2} \succ y_{p_1}y_{p_2}$, and $x_{q_1}y_{q_2} \succ y_{q_1}x_{q_2}$ since all other cases follow by symmetry or work similarly. The s-polynomial $\text{spol}(g_P, g_Q)$ is then

$$(y_{p_1}y_{p_2}y_{q_2} \cdot \mathbf{x}(q_1 \setminus \{p_1, p_2\}) - y_{q_1}x_{q_2} \cdot \mathbf{x}(\{p_1, p_2\} \setminus q_1)) \cdot \mathbf{x}((P^x \cup Q^x \setminus \{q_1, p_1, p_2\}) \setminus q_2) \cdot \mathbf{y}((P^y \cup Q^y) \setminus q_2).$$

If $q_1 \notin \{p_1, p_2\}$, then suitable walk binomials of P and Q reduce $\text{spol}(g_P, g_Q)$ to zero. If otherwise $q_1 \in \{p_1, p_2\}$, say $q_1 = p_1$, then

$$\text{spol}(g_P, g_Q) = (y_{p_1}y_{p_2}y_{q_2} - y_{q_1}x_{q_2}x_{p_2}) \cdot \mathbf{x}((P^x \cup Q^x \setminus \{p_1, p_2\}) \setminus q_2) \cdot \mathbf{y}((P^y \cup Q^y) \setminus q_2).$$

Since $y_{p_1} = y_{q_1}$, $\text{spol}(g_P, g_Q)$ is a monomial multiple of a binomial that corresponds to the odd (p_2, q_2) -walk that arises from gluing the odd walk P with the even walk Q along $p_1 = q_1$ and which uses the variable $y_{p_1} = y_{q_1}$ in its monomial factor. By Lemma 6.2.4, this binomial can be reduced to zero by elements from \mathcal{S}_G and hence $\text{spol}(g_P, g_Q)$ can be reduced to zero by \mathcal{S}_G . The case $q_1 = p_2$ is similar and omitted. \square

Proof of Theorem 6.2.5. Let $g_P, g_Q \in \mathcal{S}_G$ be walk binomials of a (p_1, p_2) -walk P and a (q_1, q_2) -walk Q respectively. If P and Q are both even, then the statement follows from Lemma 6.2.8. If they have different parity, then the statement is Lemma 6.2.12. If both are odd walks and $p_1 = p_2$ or $q_1 = q_2$, then the statement is Lemma 6.2.10. If both are odd walks and $p_1 \neq p_2$ and $q_2 \neq q_1$, then the statement is Lemma 6.2.11. \square

Theorem 6.2.5 implies that \mathcal{S}_G is a Gröbner basis of \mathcal{I}_G for any monomial ordering and hence any reduced Gröbner basis for \mathcal{I}_G can be extracted from \mathcal{S}_G by successively reducing its element further. To show that \mathcal{S}_G equals the universal Gröbner basis, a little more work has to be done.

Lemma 6.2.13. *For any $g \in \mathcal{S}_G$, there exists a monomial ordering \prec on $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ such that g is an element of the reduced Gröbner basis of \mathcal{I}_G with respect to \prec .*

Proof. Let P be the underlying minimal walk of g and let p_1 and p_2 be the start and end node of P respectively. Let $t_k \in \{x_k, y_k\}$ for $k \in \text{int}(P)$ such that $\prod_{k \in \text{int}(P)} t_k$ is the monomial factor of g . Let $P^x := \{k \in \text{int}(P) : t_k = x_k\}$ and $P^y := \{k \in \text{int}(P) : t_k = y_k\}$ and write $P^x := \{v_1, \dots, v_a\}$ and $P^y := \{w_1, \dots, w_b\}$. According to Theorem 6.2.5, the leading monomials of binomials from \mathcal{S}_G generate the initial ideal of \mathcal{I}_G with respect to any monomial ordering. Thus, it suffices to prove that there is a monomial ordering on $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ so that the leading and trailing monomial of g cannot be divided by the leading monomial of any other element from \mathcal{S}_G .

First, assume that P is odd and let \prec be the lexicographic ordering on $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ induced by

$$y_{v_1} \succ x_{v_1} \succ \dots \succ y_{v_a} \succ x_{v_a} \succ x_{w_1} \succ y_{w_1} \succ \dots \succ x_{w_b} \succ y_{w_b} \succ x_{p_1} \succ x_{p_2} \succ y_{p_1} \succ y_{p_2}.$$

The leading monomial of g is then $m := x_{p_1} x_{p_2} \cdot \mathbf{x}(P^x) \cdot \mathbf{y}(P^y)$ and the trailing monomial of g is $y_{p_1} y_{p_2} \cdot \mathbf{x}(P^x) \cdot \mathbf{y}(P^y)$. Assume that there is a (q_1, q_2) -walk Q and a walk binomial $g_Q \in \mathcal{S}_G \setminus \{g\}$ of Q whose leading monomial divides the leading monomial of g . When that Q is odd, then the leading monomial m' of g_Q is either $x_{q_1} x_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$ or $y_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$ with $t'_k \in \{x_k, y_k\}$ for all $k \in \text{int}(Q)$. In the first case, since $m' | m$, we have $\{q_1, q_2\} \subseteq P^x \cup \{p_1, p_2\}$. Since P is a minimal walk, $\{q_1, q_2\} \neq \{p_1, p_2\}$, thus we can assume that $q_1 \in P^x$. By the construction of \prec , we have $y_{q_1} y_{q_2} \succ x_{q_1} x_{q_2}$ and hence $x_{q_1} x_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$ cannot be the leading monomial of g_Q . If $m' = y_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$, then $\{q_1, q_2\} \subseteq P^y$ and hence $x_{q_1} x_{q_2} \succ y_{q_1} y_{q_2}$, which contradicts $m' = y_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$. Thus, Q cannot be an odd walk. If Q is even, we can assume by symmetry that $m' = x_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$ is the leading monomial of g_Q . Since $m' | m$, we have $q_1 \in P^x \cup \{p_1, p_2\}$ and $q_2 \in P^y$. Since $q_1 \in \{p_1, p_2\}$ implies $y_{q_1} x_{q_2} \succ x_{q_1} y_{q_2}$, which in turn contradicts $m' = x_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$. We thus must have $q_1 \in P^x$. This, however, again implies $y_{q_1} x_{q_2} \succ x_{q_1} y_{q_2}$ and hence m' cannot divide m . In the same way we can show that the trailing monomial $y_{p_1} y_{p_2} \cdot \mathbf{x}(P^x) \cdot \mathbf{y}(P^y)$ cannot be divided by leading terms of elements from $\mathcal{S}_G \setminus \{g\}$. This gives the statement when g corresponds to an odd walk P .

Next, assume that P is even and consider the lexicographic ordering on $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ induced by

$$y_{v_1} \succ x_{v_1} \succ \dots \succ y_{v_a} \succ x_{v_a} \succ x_{w_1} \succ y_{w_1} \succ \dots \succ x_{w_b} \succ y_{w_b} \succ x_{p_1} \succ y_{p_2} \succ y_{p_1} \succ x_{p_2},$$

then the leading monomial of g is $m := x_{p_1} y_{p_2} \cdot \mathbf{x}(P^x) \cdot \mathbf{y}(P^y)$. Let $g_Q \in \mathcal{S}_G$ again be a walk binomial of a minimal (q_1, q_2) -walk Q whose leading monomial m' divides m . First, assume that Q is even. By renumbering, we can assume that $m' = x_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$. Thus, we have $q_1 \in P^x \cup \{p_1\}$ and $q_2 \in P^y \cup \{q_2\}$. If $q_1 \in P^x$, then $y_{q_1} \succ x_{q_1}$ and $y_{q_1} \succ y_{q_2}$, hence $y_{q_1} x_{q_2} \succ x_{q_1} y_{q_2}$. This contradicts $m' = x_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$. Thus, we must have $q_1 = p_1$. Since P is minimal, we cannot have $q_2 = p_2$ and thus $q_2 \in P^y$. Then $x_{q_2} \succ y_{q_2}$ and $x_{q_2} \succ x_{p_1} = x_{q_1}$ which yields a contradiction to $m' = x_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$ since $y_{q_1} x_{q_2} \succ x_{q_1} y_{q_2}$. On the other hand, if Q is odd, then $m' = x_{q_1} x_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$ or $m' = y_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$. In the first case, we have $\{q_1, q_2\} \subset P^x \cup \{p_1\}$. Since $q_1 = q_2 = p_1$ is not possible since P is even and hence g is square-free, $P^x \cap \{q_1, q_2\} \neq \emptyset$. Without restricting generality, we can assume that $q_1 \in P^x$. Since $q_2 \in P^x \cup \{p_1\}$, we have in this case that $y_{q_1} y_{q_2} \succ x_{q_1} x_{q_2}$ contradicting $m' = x_{q_1} x_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$. The case $m' = y_{q_1} y_{q_2} \cdot \prod_{k \in \text{int}(Q)} t'_k$ follows similarly. \square

Putting all the pieces together, we get a positive answer to the first part of Conjecture 6.2.3.

Theorem 6.2.14. *For any graph G , \mathcal{S}_G is the universal Gröbner basis of \mathcal{I}_G .*

Proof. The universal Gröbner basis \mathcal{U}_G of \mathcal{I}_G is by definition the union of all the reduced Gröber bases of \mathcal{I}_G . By Lemma 6.2.13, every element in \mathcal{S}_G is contained in a reduced Gröbner basis and

hence $\mathcal{S}_G \subseteq \mathcal{U}_G$. On the other hand, by Theorem 6.2.5, the elements in \mathcal{S}_G fulfill Buchberger's criterion for all monomial orderings and hence $\mathcal{U}_G = \mathcal{S}_G$. \square

Theorem 6.2.14 implies that parity binomial edge ideals of bipartite graphs are radical, which they must be since they are isomorphic to the binomial edge ideal. If the graph contains an odd cycle, then every reduced Gröbner basis has an element whose leading monomial has a square.

6.3 Minimal primes

Generally, the minimal primes of a binomial ideal come in groups corresponding to the sets of indeterminates they contain. To start, we determine the minimal primes of \mathcal{I}_G that contain no indeterminates, that is, the minimal primes of the saturation \mathcal{J}_G . We need the following notation: For any graph G , let $c(G)$ be the number connected components, $c_0(G)$ the number of bipartite connected components, and $c_1(G)$ the number of connected components which contain an odd cycle. The minimal primes of \mathcal{J}_G follow then quickly from the next lemma, together with the results in [48, Section 2].

Lemma 6.3.1. *Apart from zero rows, the Smith normal form of*

$$\begin{bmatrix} A_G \\ -A_G \end{bmatrix}$$

is the diagonal matrix $\text{diag}(1, \dots, 1, 2, \dots, 2)$ whose number of entries 1 is $|V(G)| - c(G)$ and the number of entries 2 equals $c_1(G)$.

Proof. See [59, Theorem 3.3]. \square

The following ideals are the building blocks for the primary decomposition of \mathcal{J}_G . For any connected graph G with an odd cycle, let

$$\mathfrak{p}^+(G) = \langle x_i + y_i : i \in V(G) \rangle \quad \text{and} \quad \mathfrak{p}^-(G) := \langle x_i - y_i : i \in V(G) \rangle.$$

Proposition 6.3.2. *Let G be a graph whose bipartite connected components are $B_1, \dots, B_{c_0(G)}$ and whose non-bipartite connected components are $N_1, \dots, N_{c_1(G)}$. If $\text{char}(\mathbb{k}) \neq 2$, then \mathcal{J}_G is radical, and its minimal primes are the $2^{c_1(G)}$ ideals*

$$\sum_{i=1}^{c_0(G)} \mathcal{J}_{B_i} + \sum_{i=1}^{c_1(G)} \mathfrak{p}^{\sigma_i}(N_i),$$

where σ ranges over $\{+, -\}^{c_1(G)}$. On the other hand, if $\text{char}(\mathbb{k}) = 2$, then

$$\mathcal{J}_G = \sum_{i=1}^{c_0(G)} \mathcal{J}_{B_i} + \sum_{i=1}^{c_1(G)} \mathcal{J}_{N_i}$$

is primary of multiplicity $2^{c_1(G)}$ over the minimal prime $\sum_{i=1}^{c_0(G)} \mathcal{J}_{B_i} + \sum_{i=1}^{c_1(G)} \mathfrak{p}^+(N_i)$.

Proof. Assume first that \mathbb{k} is algebraically closed. According to [48, Corollary 2.2], the primary decomposition of \mathcal{J}_G is determined by the saturations of the character that defines the lattice ideal \mathcal{J}_G . If a graph is disconnected, then its adjacency matrix has block structure according to the connected components. Therefore it suffices to assume that G is connected. If G is bipartite, then Lemma 6.3.1 and [48, Corollary 2.2] imply that the lattice ideal \mathcal{J}_G is prime. We are thus left with the case that G is connected and not bipartite. Assume first that $\text{char}(\mathbb{k}) \neq 2$. Lemma 6.3.1 and [48, Corollary 2.2] together show that \mathcal{J}_G is radical and has two minimal primes. We show that these are precisely $\mathfrak{p}^+(G)$ and $\mathfrak{p}^-(G)$. The first step is $\mathcal{J}_G \subseteq \mathfrak{p}^+(G)$ using Proposition 6.1.5. Let $i, j \in V(G)$, then $x_i x_j - y_i y_j = x_i \cdot (x_j + y_j) - y_j \cdot (x_i + y_i) \in \mathfrak{p}^+(G)$ and $x_i y_j - x_j y_i = x_i \cdot (x_j + y_j) - x_j \cdot (x_i + y_i) \in \mathfrak{p}^+(G)$. Similarly, $\mathcal{J}_G \subseteq \mathfrak{p}^-(G)$. Now let $\mathfrak{p} \supseteq \mathcal{J}_G$ be a prime ideal. If \mathfrak{p} contains $x_i + y_i$ for all i , then it is either equal to $\mathfrak{p}^+(G)$ or not minimal over \mathcal{J}_G . If there exists a node i such that $x_i + y_i \notin \mathfrak{p}$, then since G has an odd cycle and is connected, for any node j there are both an odd and an even (i, j) -walk in G . Thus,

$$(x_i + y_i) \cdot (x_j - y_j) = x_i x_j - y_i y_j + x_j y_i - x_i y_j \in \mathfrak{p}.$$

Since \mathfrak{p} is prime, it contains $x_j - y_j$ for each j and thus $\mathfrak{p}^-(G) \subseteq \mathfrak{p}$. This shows that $\mathfrak{p}^-(G)$ and $\mathfrak{p}^+(G)$ are the minimal primes of \mathcal{J}_G . If $\text{char}(\mathbb{k}) = 2$, then [48, Corollary 2.2] gives that \mathcal{J}_G is primary of multiplicity two over a minimal prime which equals $\mathfrak{p}^+(G) = \mathfrak{p}^-(G)$ by the above computation. It is now evident that the algebraic closure assumption on \mathbb{k} is irrelevant since all saturations of characters are defined over \mathbb{k} . \square

Remark 6.3.3. The graph G is bipartite if and only if \mathcal{J}_G is prime.

When decomposing a pure difference binomial ideal, all components except those over the saturation \mathcal{J}_G contain monomials (for a combinatorial reason see [71, Example 4.14]). Our next step is to determine the indeterminates in the minimal primes. To this end, for any $S \subseteq V(G)$ let G_S be the induced subgraph of G on $V(G) \setminus S$ and $\mathfrak{m}_S := \langle x_s, y_s : s \in S \rangle$.

Lemma 6.3.4. *Let \mathfrak{p} be a minimal prime of \mathcal{I}_G . Then there exists $S \subseteq V(G)$ and a minimal prime \mathfrak{p}' of \mathcal{J}_{G_S} such that $\mathfrak{p} = \mathfrak{m}_S + \mathfrak{p}'$.*

Proof. Let $S := \{s \in V(G) : x_s \in \mathfrak{p} \text{ and } y_s \in \mathfrak{p}\}$. We first show the inclusions

$$\mathcal{I}_G \subseteq \mathfrak{m}_S + \mathcal{J}_{G_S} \subseteq \mathfrak{p}.$$

The first inclusion is clear, while for the second, it suffices to check that $\mathcal{J}_{G_S} \subseteq \mathfrak{p}$. Generators of \mathcal{J}_{G_S} correspond to (i, j) -walks in G_S according to Proposition 6.1.5. Let b be the binomial corresponding to any such walk, and let $\{k_1, \dots, k_r\} \subseteq V(G) \setminus S$ be its interior. By Lemma 6.1.2, $t_{k_1} \cdots t_{k_r} \cdot b \in \mathcal{I}_G \subseteq \mathfrak{p}$ for any choice of indeterminates $t_{k_l} \in \{x_{k_l}, y_{k_l}\}$, with $1 \leq l \leq r$. By the construction of S , there exists some choice such that $t_{k_1} \cdots t_{k_r} \notin \mathfrak{p}$. Since \mathfrak{p} is prime, $b \in \mathfrak{p}$. The minimal primes of $\mathfrak{m}_S + \mathcal{J}_{G_S}$ arise as sums of \mathfrak{m}_S and minimal primes of \mathcal{J}_{G_S} . By minimality, \mathfrak{p} equals $\mathfrak{m}_S + \mathfrak{p}'$ for some minimal prime \mathfrak{p}' of \mathcal{J}_{G_S} . \square

Not all primes of the form $\mathfrak{m}_S + \mathfrak{p}'$ are minimal over \mathcal{I}_G (see Example 6.3.10). As for binomial edge ideals, cut points play a crucial role in determining the sets S which lead to minimal primes, but for parity binomial edge ideals we count connected components differently:

Definition 6.3.5. For any graph G , let $\mathfrak{s}(G) := c_0(G) + c(G) = 2c_0(G) + c_1(G)$. A set $S \subseteq V(G)$ is a *disconnector* of G if $\mathfrak{s}(G_S) > \mathfrak{s}(G_{S \setminus \{s\}})$ for every $s \in S$.

Remark 6.3.6. The empty set is a disconnecter of any graph, and disconnecters cannot contain isolated nodes. A set which deletes all odd cycles from a graph, that is which makes a non-bipartite graph bipartite, is a disconnecter, even though it may preserves connectedness.

Remark 6.3.7. If a graph G has no isolated nodes, then $\mathfrak{s}(G_{\{s\}}) \geq \mathfrak{s}(G)$ for all $s \in V(G)$ and according to Definition 6.3.5 a node s is a disconnecter of G exactly if the inequality is strict. Moreover, one can conclude from Proposition 6.3.8 that $\{s\}$ is a disconnecter of G if and only if $\mathcal{J}_G \not\subseteq \mathfrak{m}_{\{s\}} + \mathcal{J}_{G_{\{s\}}}$.

Proposition 6.3.8. *Let G be a graph and $S \subseteq V(G)$. Then $\mathcal{J}_G \subseteq \mathfrak{m}_S + \mathcal{J}_{G_S}$ if and only if for all (i, j) -walks in G with $i, j \in V(G_S)$, there is an (i, j) -walk in G_S of the same parity.*

Proof. Let $\mathcal{J}_G \subseteq \mathfrak{m}_S + \mathcal{J}_{G_S}$. Let P be an (i, j) -walk in G with $i, j \notin S$, then $b(P) \in \mathcal{J}_G$. Since $b(P) \in \mathbb{k}[x_i, x_j, y_i, y_j]$, and no polynomial in \mathcal{J}_{G_S} uses indeterminates from S , we find $b(P) \in \mathcal{J}_{G_S}$. It follows that $b(P)$ is the binomial of an element of the Graver basis of \mathcal{J}_{G_S} and thus corresponds to an (i, j) -walk in G_S of the same parity.

On the other hand, let $m \in \mathcal{J}_G$ be a move corresponding to a (i, j) -walk in G . If $i \in S$ or $j \in S$, then $m \in \mathfrak{m}_S$. If otherwise $i, j \in V(G_S)$, then $m \in \mathcal{J}_{G_S}$ by assumption. \square

The next lemma states that the indeterminates contained in a minimal prime correspond to a disconnecter of G , and Theorem 6.3.15 below says when the converse is true as well.

Lemma 6.3.9. *Let \mathfrak{p} be a minimal prime of \mathcal{I}_G . There exists a disconnecter $S \subseteq V(G)$ of G and a minimal prime \mathfrak{p}' of \mathcal{J}_{G_S} such that $\mathfrak{p} = \mathfrak{m}_S + \mathfrak{p}'$.*

Proof. Let S and \mathfrak{p}' be as in Lemma 6.3.4. We prove that S is a disconnecter. Assume the converse, then there exists a node $s \in S$ such that $\mathfrak{s}(G_S) \leq \mathfrak{s}(G_{S \setminus \{s\}})$. In particular, $\{s\}$ is not a disconnecter of $G_{S \setminus \{s\}}$ by Remark 6.3.7. According to Remark 6.3.7 and Proposition 6.3.8,

$$\mathcal{J}_{G_{S \setminus \{s\}}} \subseteq \mathfrak{m}_{\{s\}} + \mathcal{J}_{G_S} \subseteq \mathfrak{m}_{\{s\}} + \mathfrak{p}'.$$

Hence, since the ideal on the right-hand side is prime, choose a minimal prime \mathfrak{p}'' of $\mathcal{J}_{G_{S \setminus \{s\}}}$ such that $\mathcal{J}_{G_{S \setminus \{s\}}} \subseteq \mathfrak{p}'' \subseteq \mathfrak{m}_{\{s\}} + \mathfrak{p}'$. Since $x_s, y_s \notin \mathfrak{p}''$, the containment $\mathfrak{p}'' \subsetneq \mathfrak{m}_{\{s\}} + \mathfrak{p}'$ is strict. Then

$$\mathcal{I}_G \subseteq \mathfrak{m}_{S \setminus \{s\}} + \mathfrak{p}'' \subsetneq \mathfrak{m}_S + \mathfrak{p}' = \mathfrak{p}$$

which contradicts the minimality of \mathfrak{p} . \square

Let $S \subseteq V(G)$ be a disconnecter of G . The induced subgraph G_S splits into bipartite components $B_1, \dots, B_{c_0(G_S)}$ and non-bipartite components $N_1, \dots, N_{c_1(G_S)}$. By Proposition 6.3.2 the minimal primes of \mathcal{J}_{G_S} are

$$\mathfrak{p} = \sum_{i=1}^{c_0(G_S)} \mathcal{J}_{B_i} + \sum_{i=1}^{c_1(G_S)} \mathfrak{p}^{\sigma_i}(N_i), \text{ where } \begin{cases} \sigma_i \in \{+, -\}, & \text{if } \text{char}(\mathbb{k}) \neq 2 \\ \sigma_i = +, & \text{if } \text{char}(\mathbb{k}) = 2 \end{cases}. \quad (6.3.1)$$

Not all of these primes lead to minimal primes of \mathcal{I}_G because of the following effect:

Example 6.3.10. Let G be the graph in Figure 6.3. The node 4 is a disconnecter, and $G_{\{4\}}$ consists of the two triangles $N_1 = \{1, 2, 3\}$ and $N_2 = \{5, 6, 7\}$. Choosing for both triangles the positive sign component, we obtain the prime ideal

$$\mathfrak{m}_{\{4\}} + \mathfrak{p}^+(N_1) + \mathfrak{p}^+(N_2) = \mathfrak{m}_{\{4\}} + \langle x_i + y_i : i \in [7] \setminus \{4\} \rangle$$




Figure 6.3: A graph for which one of the primes in (6.3.1) is not a minimal prime.

which is not minimal over \mathcal{I}_G since it contains the prime ideal $\mathfrak{p}^+(G)$. On the other hand, the ideals $\mathfrak{m}_{\{4\}} + \mathfrak{p}^+(N_1) + \mathfrak{p}^-(N_2)$ and $\mathfrak{m}_{\{4\}} + \mathfrak{p}^-(N_1) + \mathfrak{p}^+(N_2)$, each with different signs on the triangles, are minimal over \mathcal{I}_G .

A combinatorial condition on σ in (6.3.1) guarantees that a minimal prime of \mathcal{J}_{G_S} is the binomial part of a minimal prime of \mathcal{I}_G (the monomial part being \mathfrak{m}_S). To see it, let $s \in S$ be such that $c(G_S) > c(G_S \setminus \{s\})$, i.e. when adding s back to G_S some of its connected components are joined. Denote by $\mathcal{C}_{G_S}(s)$ the set of only those connected components of G_S which are joined when adding s .

Definition 6.3.11. Let $S \subseteq V(G)$ be a disconnecter of G . A minimal prime \mathfrak{p} of \mathcal{J}_{G_S} is *sign-split* if for all $s \in S$ such that $\mathcal{C}_{G_S}(s)$ contains no bipartite graphs, the prime summands of \mathfrak{p} corresponding to connected components in $\mathcal{C}_{G_S}(s)$ are not all equal to \mathfrak{p}^+ or all equal to \mathfrak{p}^- .

Remark 6.3.12. If $\mathcal{C}_{G_S}(s)$ contains at least one bipartite graph, then Definition 6.3.11 imposes no restriction and every choice of prime summands is sign-split.

Remark 6.3.13. If $\text{char}(\mathbb{k}) = 2$, then all signs σ in (6.3.1) are fixed. In this case, Definition 6.3.11 can only be satisfied if $\mathcal{C}_{G_S}(s)$ contains a bipartite component for each $s \in S$.

Example 6.3.14. Not every disconnecter $S \subseteq V(G)$ of G admits a sign-split minimal prime for \mathcal{J}_{G_S} , and thus not every disconnecter contributes minimal primes to \mathcal{I}_G . Consider the graph in




Figure 6.4: A disconnecter whose binomial parts cannot be sign-split.

Figure 6.4. The set of square nodes is a disconnecter that does not contribute minimal primes. Adding one of the squares back yields the requirement that the primes on the two now connected triangles have different signs, but these three requirements cannot be satisfied simultaneously.

Theorem 6.3.15. *The minimal primes of \mathcal{I}_G are the ideals $\mathfrak{m}_S + \mathfrak{p}$, where $S \subseteq V(G)$ is a disconnecter of G and \mathfrak{p} is a sign-split minimal prime of \mathcal{J}_{G_S} .*

Proof. According to Lemma 6.3.9, all minimal primes of \mathcal{I}_G have the form $\mathfrak{m}_S + \mathfrak{p}$, where $S \subseteq V(G)$ is a disconnecter and \mathfrak{p} is a minimal prime of \mathcal{J}_{G_S} . We first show that if \mathfrak{p} is sign-split, this ideal is minimal over \mathcal{I}_G . Assume not, then by Lemma 6.3.4 there exists a set $T \subseteq V(G)$ and a minimal prime $\tilde{\mathfrak{p}}$ of \mathcal{J}_{G_T} such that

$$\mathcal{I}_G \subseteq \mathfrak{m}_T + \tilde{\mathfrak{p}} \subsetneq \mathfrak{m}_S + \mathfrak{p}. \quad (6.3.2)$$

This implies $T \subsetneq S$, since if $T = S$, then by Lemma 6.3.4 also $\tilde{\mathfrak{p}} = \mathfrak{p}$. Let $s' \in S \setminus T$, then $G_S \subsetneq G_{S \setminus \{s'\}} \subseteq G_T$. Since S is a disconnecter of G , $\mathfrak{s}(G_S) > \mathfrak{s}(G_{S \setminus \{s'\}})$. We now make a case distinction on $\mathcal{C}_{G_S}(s')$, that is the set of connected components in G_S that are joined to s' in $G_{S \setminus \{s'\}}$. If $\mathcal{C}_{G_S}(s')$ contains at least one bipartite component, adding s' to G_S either this component becomes non-bipartite in $G_{S \setminus \{s'\}}$ or it is joined to another bipartite component of G_S . In the first case, let B be a bipartite component which becomes non-bipartite. There exists $i \in V(B)$ such that $x_i^2 - y_i^2 \in \mathcal{J}_{G_{S \setminus \{s'\}}} \subseteq \mathcal{J}_{G_T} \subseteq \tilde{\mathfrak{p}}$, but $x_i^2 - y_i^2 \notin \mathcal{J}_B$. Since \mathcal{J}_B is a summand of \mathfrak{p} and since $i \in V(B)$, $x_i^2 - y_i^2 \notin \mathfrak{m}_S + \mathfrak{p}$, in contradiction to (6.3.2). In the second case, let B_1 and B_2 be the bipartite components of G_S which are joined to s' . There are $i_1 \in V(B_1)$ and $i_2 \in V(B_2)$ such that there exists an (i_1, i_2) -walk in $G_{S \setminus \{s'\}}$. Independent of the parity of this walk, the corresponding Markov move is not contained in $\mathcal{J}_{B_1} + \mathcal{J}_{B_2}$ since there is no applicable move from the Graver basis. Since \mathcal{J}_{B_1} and \mathcal{J}_{B_2} are summands of \mathfrak{p} involving the indeterminates i_1 and i_2 , there is a binomial which is not in $\mathfrak{m}_S + \mathfrak{p}$ but in $\mathcal{J}_{G_{S \setminus \{s'\}}} \subseteq \tilde{\mathfrak{p}}$ contradicting (6.3.2). Assume now that all components in $\mathcal{C}_{G_S}(s')$ are non-bipartite (there must be at least two of them since $\{s'\}$ is a disconnecter of $G_{S \setminus \{s'\}}$). By assumption, \mathfrak{p} is sign-split and hence there exist distinct components $N_1, N_2 \in \mathcal{C}_{G_S}(s')$ such that $\mathfrak{p}^+(N_1)$ and $\mathfrak{p}^-(N_2)$ are summands of \mathfrak{p} . There is an odd walk from a node $i_1 \in V(N_1)$ to a node $i_2 \in V(N_2)$ in $G_{S \setminus \{s'\}}$, and therefore, $x_{i_1}x_{i_2} - y_{i_1}y_{i_2} \in \mathcal{J}_{G_{S \setminus \{s'\}}} \subseteq \tilde{\mathfrak{p}}$. However, since

$$x_{i_1}x_{i_2} - y_{i_1}y_{i_2} \notin \mathfrak{p}^+(N_1) + \mathfrak{p}^-(N_2),$$

also $x_{i_1}x_{i_2} - y_{i_1}y_{i_2} \notin \mathfrak{p}$. By construction, $i_1, i_2 \notin S$ and thus

$$x_{i_1}x_{i_2} - y_{i_1}y_{i_2} \notin \mathfrak{m}_S + \mathfrak{p}$$

which contradicts (6.3.2). This shows minimality of $\mathfrak{m}_S + \mathfrak{p}$.

Let now $\mathfrak{m}_S + \mathfrak{p}$ be a minimal prime of \mathcal{I}_G . The set S is a disconnecter by Lemma 6.3.9 and thus it remains to prove that \mathfrak{p} is sign-split. To the contrary, assume there is a node $s \in S$ with $c(G_{S \setminus \{s\}}) > c(G_S)$ such that $\mathcal{C}_{G_S}(s) = \{N_1, \dots, N_k\}$ consists exclusively of non-bipartite components, $k \geq 2$, and all summands of \mathfrak{p} corresponding to N_i have the same sign, say $+$. When adding s back to G_S , the components in $\mathcal{C}_{G_S}(s)$ are joined to a single, non-bipartite connected component H in $G_{S \setminus \{s\}}$, whereas all other components of G_S coincide with connected components of $G_{S \setminus \{s\}}$. Since

$$\mathfrak{p}^+(H) = \mathfrak{m}_{\{s\}} + \sum_{i=1}^k \mathfrak{p}^+(N_i) \subsetneq \mathfrak{m}_{\{s\}} + \sum_{i=1}^k \mathfrak{p}^+(N_i),$$

choosing on all other components of $G_{S \setminus \{s\}}$ the same prime component as in G_S , we obtain a prime ideal that is strictly smaller than $\mathfrak{m}_S + \mathfrak{p}$. \square

Remark 6.3.16. Example 6.3.10 and Definition 6.3.11 are valid independent of $\text{char}(\mathbb{k})$. In the above proof, the case of $\text{char}(\mathbb{k}) = 2$ could be simplified, but everything works in general without the need for a case distinction.

6.4 Radicality

The intersection of the minimal primes of \mathcal{I}_G depends on $\text{char}(\mathbb{k})$ so that we do not attempt to compute it directly. Theorem 6.4.5 below says that \mathcal{I}_G is radical if the characteristic is not two.

Observe that we have to use the Gröbner basis of \mathcal{I}_G to show its *non*-radicality in $\text{char}(\mathbb{k}) = 2$. Here is the principal source of field dependence (see also [63, Theorem 1.2]):

Remark 6.4.1. Fix a field \mathbb{k} with $\text{char}(\mathbb{k}) = 2$. The parity binomial edge ideal \mathcal{I}_G is radical in $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ if and only if G is bipartite. Clearly, if G is bipartite, then \mathcal{I}_G is radical. Conversely, let (i_1, \dots, i_{r+1}) with $i_{r+1} = i_1$ be an odd cycle in G . According to Lemma 6.1.2, $((x_{i_1} - y_{i_1})y_{i_2} \cdots y_{i_r})^2 = (x_{i_1}^2 - y_{i_1}^2)y_{i_2}^2 \cdots y_{i_r}^2 \in \mathcal{I}_G$. Fix now a monomial ordering \prec on $\mathbb{k}[\mathbf{x}, \mathbf{y}]$ with $y_i \prec x_i$ for all $i \in V(G)$. By Theorem 6.2.5, the leading monomials of \mathcal{S}_G generate the initial ideal of \mathcal{I}_G . Under this monomial ordering, all walk binomials corresponding to odd walks have two indeterminates from $\mathbb{k}[\mathbf{x}]$ in their leading monomial. Thus, the only binomials from \mathcal{S}_G whose leading monomials divide $x_{i_1}y_{i_2} \cdots y_{i_r}$ correspond to minimal even (i_1, i_k) -walks in $G[\{i_1, \dots, i_r\}]$ with $k \in \{2, \dots, r\}$. Replacements coming from these binomials lead to monomials where x_{i_1} is replaced by y_{i_1} and y_{i_k} is replaced by x_{i_k} . Thus, $x_{i_1}y_{i_2} \cdots y_{i_r} \not\equiv_{\mathcal{I}_G} y_{i_1}y_{i_2} \cdots y_{i_r}$ and hence \mathcal{I}_G is not radical.

Remark 6.4.2. The ideal \mathcal{I}_G is homogeneous with respect to the multigrading $\deg(x_i) = \deg(y_i) = e_i$, where e_i is the i -th standard basis vector of $\mathbb{R}^{|V(G)|}$.

Lemma 6.4.3. *Let $i \in V(G)$ and $m \in \mathcal{I}_G + \mathfrak{m}_{\{i\}}$ be a monomial. Then $m \in \mathfrak{m}_{\{i\}}$.*

Proof. Since it is generated by pure difference binomials, \mathcal{I}_G does not contain any monomials. Thus, any monomial in $\mathcal{I}_G + \mathfrak{m}_{\{i\}}$ is equivalent to one in $\mathfrak{m}_{\{i\}}$ modulo term replacements using binomials in \mathcal{I}_G , but these do not change membership in $\mathfrak{m}_{\{i\}}$ by Remark 6.4.2. \square

Proposition 6.4.4. *For any graph G , $\mathcal{I}_G = \mathcal{J}_G \cap \bigcap_{i \in V(G)} (\mathcal{I}_G + \mathfrak{m}_{\{i\}})$.*

Proof. According to [48, Corollary 1.5], the intersection is binomial. Let b be any binomial in the intersection. For each $i \in V(G)$, there are three cases: Either no term of b is individually contained in $\mathcal{I}_G + \mathfrak{m}_{\{i\}}$, exactly one is, or both are. In the first case, [48, Proposition 1.10] implies $b \in \mathcal{I}_G$. In the second case, it implies that the other monomial is contained in \mathcal{I}_G , which is impossible. Thus it suffices to consider binomials b both of whose monomials are contained in $\mathcal{I}_G + \mathfrak{m}_{\{i\}}$ for all $i \in V(G)$. By Lemma 6.4.3, both monomials of b are contained in $\mathfrak{m}_{\{i\}}$ for each $i \in V(G)$. Since $b \in \mathcal{J}_G$, there is $r \in \mathbb{N}$ and (s_i, t_i) -walks P_i for $i \in [r]$ such that

$$b = x^{h_1}y^{h'_1}b(P_1) + \cdots + x^{h_r}y^{h'_r}b(P_r)$$

with $h_i, h'_i \in \mathbb{N}^n$. We can assume that one monomial of b equals one of the monomials of $x^{h_1}y^{h'_1}b(P_1)$. Thus both monomials of $x^{h_1}y^{h'_1}b(P_1)$ are divisible by at least one indeterminate for each $i \in V(G)$ and, by Lemma 6.1.2, $x^{h_1}y^{h'_1}b(P_1) \in \mathcal{I}_G$. Replacing b by $b - x^{h_1}y^{h'_1}b(P_1)$ and iterating the argument eventually yields $b \in \mathcal{I}_G$. \square

Theorem 6.4.5. *Let G be a graph. If $\text{char}(\mathbb{k}) \neq 2$, then \mathcal{I}_G is a radical ideal.*

Proof. The proof is by induction on the number of nodes n of G . If G has at most one node, then $\mathcal{I}_G = 0$ and the claim holds. Proposition 6.3.2 shows that $\mathcal{I}_G + \mathfrak{m}_{\{i\}} = \mathcal{I}_{G_{\{i\}}} + \mathfrak{m}_{\{i\}}$ for all $i \in V(G)$. Thus Proposition 6.4.4 reads as $\mathcal{I}_G = \mathcal{J}_G \cap \bigcap_{i=1}^n (\mathcal{I}_{G_{\{i\}}} + \mathfrak{m}_{\{i\}})$. By the induction hypothesis, $\mathcal{I}_{G_{\{i\}}}$ is radical and thus $\mathcal{I}_{G_{\{i\}}} + \mathfrak{m}_{\{i\}}$ is radical. Proposition 6.3.2 says that \mathcal{J}_G is radical if $\text{char}(\mathbb{k}) \neq 2$ which yields the result. \square

Theorem 6.4.9 below contains a primary decomposition of \mathcal{I}_G in the case $\text{char}(\mathbb{k}) = 2$. It uses the following lemma, which allows to transport decompositions between different characteristics. Recall that the combinatorics of any binomial ideal I is encoded in its congruence \sim_I which identifies monomials m_1, m_2 , whenever $m_1 - \lambda m_2 \in I$ for some non-zero $\lambda \in \mathbb{k}$. A binomial ideal is *unital* if it is generated by monomials and pure differences of monomials. Then each congruence is the congruence of a unital binomial ideal, though not uniquely.

Lemma 6.4.6. *If a decomposition $I = J_1 \cap \dots \cap J_s$ of a unital binomial ideal I into unital binomial ideals J_i , $i = 1, \dots, s$ is valid in some characteristic, then it is valid in any characteristic.*

Proof. The congruence \sim_I induced by I is the common refinement of the congruences \sim_{J_i} , induced by the J_i , $i = 1, \dots, s$. Thus, in any characteristic, [71, Theorem 9.12] implies that I and $J_1 \cap \dots \cap J_s$ can only differ if one of them contains monomials, but the other does not. This cannot happen since unital binomial ideals contain monomials if and only if they have monomials among the generators. \square

According to Example 6.3.14, not all disconnectors contribute minimal primes. From Definition 6.3.11 it may seem that this is an arithmetic effect. It is not; the primary decomposition of \mathcal{I}_G in characteristic two also witnesses it. For the following definition, recall that a hypergraph is *k-colorable* if the nodes can be colored with k colors so that no edge is monochromatic.

Definition 6.4.7. Let $S \subseteq V(G)$ be a disconnector, and let $s_1, \dots, s_r \in S$ be the nodes such that $\mathcal{C}_{G_S}(s_i)$ consists exclusively of non-bipartite components of G_S . Let \mathcal{H} be the hypergraph whose node set consists of the connected components $\mathcal{C}_{G_S}(s_1) \cup \dots \cup \mathcal{C}_{G_S}(s_r)$ and with edge set $\{\mathcal{C}_{G_S}(s_1), \dots, \mathcal{C}_{G_S}(s_r)\}$. The disconnector S is *effective* if \mathcal{H} is 2-colorable.

Remark 6.4.8. A disconnector is effective if and only if, in characteristic zero, it admits sign-split minimal primes.

Theorem 6.4.9. *Let \mathcal{S} be the set of effective disconnectors of G . Then*

$$\mathcal{I}_G = \bigcap_{S \in \mathcal{S}} (\mathfrak{m}_S + \mathcal{J}_{G_S}). \quad (6.4.1)$$

If $\text{char}(\mathbb{k}) = 2$, then (6.4.1) is a primary decomposition of \mathcal{I}_G .

Proof. For each disconnector $S \in \mathcal{S}$, let $B_1^S, \dots, B_{c_0(G_S)}^S$ be the bipartite components and $N_1^S, \dots, N_{c_1(G_S)}^S$ the non-bipartite components of G_S . Let $\Sigma^S \subseteq \{+, -\}^{c_1(G_S)}$ denote the set of sign patterns that are sign-split. In characteristic zero, by Theorems 6.3.15 and 6.4.5, \mathcal{I}_G decomposes as

$$\mathcal{I}_G = \bigcap_{S \in \mathcal{S}} \bigcap_{\sigma \in \Sigma^S} \left(\mathfrak{m}_S + \sum_{i=1}^{c_0(G_S)} \mathcal{J}_{B_i^S} + \sum_{i=1}^{c_1(G_S)} \mathfrak{p}^{\sigma_i}(N_i^S) \right).$$

The intersection remains valid when intersecting over additional ideals containing \mathcal{I}_G . In particular, the sign-split requirement can be dropped and Σ^S replaced by $\{+, -\}^{c_1(G_S)}$. Carrying out this inner intersection yields the ideals $\mathfrak{m}_S + \mathcal{J}_{G_S}$ by Proposition 6.3.2, and hence (6.4.1) is valid in characteristic zero. Since all involved ideals are unital, Lemma 6.4.6 yields that (6.4.1) is valid in any characteristic. The ideals under consideration are primary if $\text{char}(\mathbb{k}) = 2$ according to Proposition 6.3.2 and thus the second statement follows. \square

Bibliography

- [1] 4ti2 team. 4ti2 – a software package for algebraic, geometric and combinatorial problems on linear spaces. Available at <http://www.4ti2.de>.
- [2] Rudolf Ahlswede, Levon H. Khachatryan, and András Sárközy. On the counting function of primitive sets of integers. *Journal of Number Theory*, 79(2):330–344, 1999.
- [3] David Aldous. Random walks on finite groups and rapidly mixing Markov chains. *Séminaire de Probabilités XVII 1981/82*, 17:243–297, 1983.
- [4] David F. Anderson, Gheorghe Craciun, Manoj Gopalkrishnan, and Carsten Wiuf. Lyapunov Functions, Stationary Distributions, and Non-equilibrium Potential for Reaction Networks. *Bulletin of Mathematical Biology*, 77(9):1744–1767, 2015.
- [5] Ian Anderson. Primitive sequences whose elements have no large prime factors. *Glasgow Mathematical Journal*, 10(1):10, 1969.
- [6] Satoshi Aoki, Hisayuki Hara, and Akimichi Takemura. *Markov Bases in Algebraic Statistics*, volume 199. Springer New York, 2012.
- [7] Kenneth Appel and Wolfgang Haken. Every planar map is four colorable. Part I: Discharging. *Illinois Journal of Mathematics*, 21(3):429–490, 1977.
- [8] Mourtadha Badiane, Isaac Burke, and Emil Sköldbberg. The universal Gröbner basis of a binomial edge ideal. *preprint, arXiv:1601.04575*, pages 1–11, 2016.
- [9] Velleda Baldoni, Nicole Berline, Jesús A. De Loera, Brandon Durta, Matthias Köppe, Stanislav Moreinis, Gregory Pinto, Michele Vergne, and Jianqiu Wu. A user’s guide for LattE. Available at <http://www.math.ucdavis.edu/~latte/>.
- [10] Arindam Banerjee and Luis Núñez-Betancourt. Graph Connectivity and Binomial Edge Ideals. *preprint, arXiv:1605.00314*, pages 1–12, 2016.
- [11] Douglas Bauer, Hajo Broersma, and Edward Schmeichel. Toughness in Graphs – A Survey. *Graphs and Combinatorics*, 22(1):1–35, 2006.
- [12] Stephen Baumert, Archis Ghate, Seksan Kiatsupaibul, Yanfang Shen, Robert L. Smith, and Zeldá B. Zabinsky. Discrete Hit-and-Run for Sampling Points from Arbitrary Distributions Over Subsets of Integer Hyperrectangles. *Operations Research*, 57(3):727–739, 2009.
- [13] Dave Bayer and Persi Diaconis. Trailing the Dovetail Shuffle to its Lair. *The Annals of Applied Probability*, 2(2):294–313, 1992.
- [14] Dave Bayer, Sorin Popescu, and Bernd Sturmfels. Syzygies of unimodular Lawrence ideals. *Journal für die Reine und Angewandte Mathematik*, 534:169–186, 2001.

- [15] Matthias Beck and Sinai Robins. *Computing the Continuous Discretely*. Springer-Verlag New York, 2007.
- [16] Anna M. Bigatti, Robertola Scala, and Lorenzo Robbiano. Computing Toric Ideals. *Journal of Symbolic Computation*, 27(4):351–365, 1999.
- [17] Mónica Blanco and Francisco Santos. On the enumeration of lattice 3-polytopes. *preprint, arXiv:1601.02577*, pages 1–31, 2016.
- [18] Stephen Boyd, Persi Diaconis, and Lin Xiao. Fastest Mixing Markov Chain on a Graph. *SIAM Review*, 46(4):667–689, 2004.
- [19] Robert Brijder. Dominance and T-Invariants for Petri Nets and Chemical Reaction Networks. In Andrew Phillips and Peng Yin, editors, *DNA Computing and Molecular Programming*, pages 1–15. Springer International Publishing, 2015.
- [20] Winfried Bruns, Bogdan Ichim, Tim Römer, Richard Sieg, and Christof Söger. Normaliz – algorithms for rational cones and affine monoids. Available at <http://normaliz.uos.de>.
- [21] Sergio Cabello, David Eppstein, and Sandi Klavžar. The Fibonacci dimension of a graph. *Electronic Journal of Combinatorics*, 18(1):1–23, 2011.
- [22] Hara Charalambous, Apostolos Thoma, and Marius Vladioiu. Binomial fibers and indispensable binomials. *Journal of Symbolic Computation*, 74:578–591, 2016.
- [23] Faryal Chaudhry, Ahmet Dokuyucu, and Viviana Ene. Binomial edge ideals and rational normal scrolls. *Bulletin of the Iranian Mathematical Society*, 41(4):971–979, 2015.
- [24] Jeff Cheeger. *A lower bound for the smallest eigenvalue of the Laplacian*. Problems in Analysis: A Symposium in Honor of Salomon Bochner. Princeton University Press, New Jersey, 1971.
- [25] Haiyan Chen and Fuji Zhang. The rapid mixing of random walks defined by an n-cube. *Advances in Applied Mathematics*, 33(1):136–145, 2004.
- [26] Yuguo Chen, Ian H. Dinwoodie, and Seth Sullivant. Sequential importance sampling for multiway tables. *The Annals of Statistics*, 34(1):523–545, 2006.
- [27] Man-Duen Choi, Tsit Y. Lam, and Bruce Reznick. Lattice Polytopes with Distinct Pair-Sums. *Discrete & Computational Geometry*, 27(1):65–72, 2002.
- [28] Ronald Christensen. *Log-Linear Models*. Springer New York, 1990.
- [29] Mary Cryan, Martin Dyer, Leslie A. Goldberg, Mark Jerrum, and Russell Martin. Rapidly Mixing Markov Chains for Sampling Contingency Tables with a Constant Number of Rows. *SIAM Journal on Computing*, 36(1):247–278, 2006.
- [30] Milos Curcic. *Lattice polytopes with distinct pair-sums*. Phd thesis, University of Illinois, 2013.
- [31] Jesús A. De Loera, Raymond Hemmecke, and Matthias Köppe. *Algebraic and Geometric Ideas in the Theory of Discrete Optimization*. MPS-SIAM Series on Optimization. SIAM, Philadelphia, 2013.

- [32] Jesús A. De Loera, Raymond Hemmecke, and Jon Lee. On Augmentation Algorithms for Linear and Integer-Linear Programming: From Edmonds–Karp to Bland and Beyond. *SIAM Journal on Optimization*, 25(4):2494–2511, 2015.
- [33] Jesús A. De Loera and Shmuel Onn. Markov bases of three-way tables are arbitrarily complicated. *Journal of Symbolic Computation*, 41(2):173–181, 2006.
- [34] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 7.2)*, 2016. Available at <http://www.sagemath.org>.
- [35] Persi Diaconis and Laurent Saloff-Coste. Comparison theorems for reversible Markov chains. *The Annals of Applied Probability*, 3(3):696–730, 1993.
- [36] Persi Diaconis and Laurent Saloff-Coste. Nash inequalities for finite Markov chains. *Journal of Theoretical Probability*, 9(2):459–510, 1996.
- [37] Persi Diaconis and Laurent Saloff-Coste. What Do We Know about the Metropolis Algorithm? *Journal of Computer and System Sciences*, 57(1):20–36, 1998.
- [38] Persi Diaconis and Mehrdad Shahshahani. Generating a random permutation with random transpositions. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 57(2):159–179, 1981.
- [39] Persi Diaconis and Bernd Sturmfels. Algebraic algorithms for sampling from conditional distributions. *The Annals of Statistics*, 26(1):363–397, 1998.
- [40] Reinhard Diestel. *Graph Theory*. Springer-Verlag Berlin Heidelberg, 2nd edition, 2000.
- [41] Adrian Dobra. Dynamic Markov Bases. *Journal of Computational and Graphical Statistics*, 21(2):496–517, 2012.
- [42] Adrian Dobra and Seth Sullivant. A divide-and-conquer algorithm for generating Markov bases of multi-way tables. *Computational Statistics*, 19:347–366, 2004.
- [43] Mathias Drton, Bernd Sturmfels, and Seth Sullivant. *Lectures on Algebraic Statistics*, volume 39 of *Oberwolfach Seminars*. Birkhäuser Basel, 2009.
- [44] Natalia Dück and Karl-Heinz Zimmermann. Graver bases and Niversal Gröbner bases for linear codes [sic]. *International Journal of Pure and Applied Mathematics*, 98(4):18, 2015.
- [45] Martin Dyer, Leslie A. Goldberg, Mark Jerrum, and Russell Martin. Markov chain comparison. *Probability Surveys*, 3(April):89–111, 2006.
- [46] Martin Dyer, Catherine Greenhill, and Mario Ullrich. Structure and eigenvalues of heat-bath Markov chains. *Linear Algebra and its Applications*, 454:57–71, 2014.
- [47] Eugène Ehrhart. *Polynômes arithmétiques et méthode des polyèdres en combinatoire*, volume 35 of *International Series of Numerical Mathematics*. Birkhäuser Verlag, Basel, 1977.
- [48] David Eisenbud and Bernd Sturmfels. Binomial ideals. *Duke Mathematical Journal*, 84(1):1–45, 1996.

- [49] Viviana Ene, Jürgen Herzog, and Takayuki Hibi. Cohen-Macaulay binomial edge ideals. *Nagoya Mathematical Journal*, 204:57–68, 2011.
- [50] Alexander Engström, Thomas Kahle, and Seth Sullivant. Multigraded commutative algebra of graph decompositions. *Journal of Algebraic Combinatorics*, 39(2):335–372, 2014.
- [51] David Eppstein. The lattice dimension of a graph. *European Journal of Combinatorics*, 26(5):585–592, 2005.
- [52] Paul Erdős, Frank Harary, and William T. Tutte. On the dimension of a graph. *Mathematika*, 12(02):118, 1965.
- [53] Ky Fan. On a Theorem of Weyl Concerning Eigenvalues of Linear Transformations. I. *Proceedings of the National Academy of Sciences of the United States of America*, 35(11):652–655, 1950.
- [54] Shannon L. Fitzpatrick and Richard J. Nowakowski. The strong isometric dimension of finite reflexive graphs. *Discussiones Mathematicae Graph Theory*, 20(1):23–38, 2000.
- [55] Georg Frobenius. Über Matrizen aus nicht negativen Elementen. *Sitzungsberichte Preussische Akademie der Wissenschaft, Berlin*, pages 456–477, 1912.
- [56] Giles Gardam. *Expander Graphs and Kazhdan’s Property (T)*. Bachelor’s thesis, University of Sidney, 2012.
- [57] Ewgenij Gawrilow and Michael Joswig. polymake: a framework for analyzing convex polytopes. In Gil Kalai and Günter M. Ziegler, editors, *Polytopes — Combinatorics and Computation*, pages 43–74. Birkhäuser, 2000.
- [58] Daniel R. Grayson and Michael E. Stillman. Macaulay2 – a software system for research in algebraic geometry. Available at <http://www.math.uiuc.edu/Macaulay2/>.
- [59] Jerrold W. Grossman, Devadatta M. Kulkarni, and Irwin E. Schochetman. On the minors of an incidence matrix and its Smith normal form. *Linear Algebra and its Applications*, 218(1995):213–224, 1995.
- [60] W. K. Hastings. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika*, 57(1):97–109, 1970.
- [61] Raymond Hemmecke and Peter N. Malkin. Computing generating sets of lattice ideals and Markov bases of lattices. *Journal of Symbolic Computation*, 44(10):1463–1476, 2009.
- [62] Raymond Hemmecke and Tobias Windisch. On the Connectivity of Fiber Graphs. *Journal of Algebraic Statistics*, 6(1):24–45, 2015.
- [63] Jürgen Herzog, Takayuki Hibi, Freyja Hreinsdóttir, Thomas Kahle, and Johannes Rauh. Binomial edge ideals and conditional independence statements. *Advances in Applied Mathematics*, 45(3):317–333, 2010.
- [64] Jürgen Herzog, Antonio Macchia, Sara S. Madani, and Volkmar Welker. On the ideal of orthogonal representations of a graph. *Advances in Applied Mathematics*, 71:146–173, 2015.

- [65] Serkan Hoşten and Bernd Sturmfels. GRIN: An implementation of Gröbner bases for integer programming. In Egon Balas and Jens Clausen, editors, *Integer Programming and Combinatorial Optimization*, pages 267–276. Springer Berlin Heidelberg, 1995.
- [66] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–562, 2006.
- [67] Roger A. Horn. *Matrix Analysis*. Cambridge University Press, New York, 2nd edition, 2013.
- [68] Wilfried Imrich and Matjaž Kovše. Lattice embeddings of trees. *European Journal of Combinatorics*, 30(5):1142–1148, 2009.
- [69] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989.
- [70] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.
- [71] Thomas Kahle and Ezra Miller. Decompositions of commutative monoid congruences and binomial ideals. *Algebra and Number Theory*, 8(6):1297–1364, 2014.
- [72] Thomas Kahle, Johannes Rauh, and Seth Sullivant. Positive margins and primary decomposition. *Journal of Commutative Algebra*, 6(2):173–208, 2014.
- [73] Thomas Kahle, Camilo Sarmiento, and Tobias Windisch. Parity binomial edge ideals. *Journal of Algebraic Combinatorics*, 44(1):99–117, 2015.
- [74] Martin Kassabov. Symmetric groups and expander graphs. *Inventiones Mathematicae*, 170(2):327–354, 2007.
- [75] Manuel Kauers, Christoph Koutschan, and Doron Zeilberger. Proof of Ira Gessel’s lattice path conjecture. *Proceedings of the National Academy of Sciences*, 106(28):11502–11505, 2009.
- [76] Gregor Kemper. *A Course in Commutative Algebra*. Springer-Verlag Berlin Heidelberg, 2011.
- [77] Steven S. Kim. Mixing Time of a Rook’s Walk. *Undergraduate certificate paper*, 2012.
- [78] Christian Krattenthaler. Lattice Path Enumeration. In Miklos Bona, editor, *Handbook of enumerative combinatorics*, chapter 10, pages 589–680. CRC Press, 2015.
- [79] Martin Kreuzer and Lorenzo Robbiano. *Computational Commutative Algebra 2*. Springer-Verlag, Berlin, Heidelberg, 2005.
- [80] Jeffrey C. Lagarias and Günter M. Ziegler. Bounds for lattice polytopes containing a fixed number of interior points in a sublattice. *Canadian Journal of Mathematics*, 43(5):1022–1035, 1991.
- [81] David A. Levin, Yuval Peres, and Elisabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009.

- [82] László Lovász. Hit-and-run mixes fast. *Mathematical Programming*, 86(3):443–461, 1999.
- [83] László Lovász and Santosh Vempala. Hit-and-Run from a Corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006.
- [84] László Lovász and Katalin Vesztegombi. Geometric representations of graphs. In Gabor Halász, László Lovász, Miklos Simonovits, and Vera Sós, editors, *Paul Erdős and his Mathematics*, volume 11, pages 471–498. Bolyai Society Mathematical Studies, 2002.
- [85] Jeremie Lumbroso, Marni Mishna, and Yann Ponty. Taming reluctant random walks in the positive quadrant. *preprint, arXiv:1603.06321v1*, pages 1–11.
- [86] Peter N. Malkin. Truncated Markov bases and Gröbner bases for Integer Programming. *Applicable Algebra in Engineering Communication and Computing*, 8(4):27, 2006.
- [87] Kazunori Matsuda and Satoshi Murai. Regularity bounds for binomial edge ideals. *Journal of Commutative Algebra*, 5(1):141–149, 2013.
- [88] Cam McLeman, Peter T. Otto, John Rahmani, and Matthew Sutter. Mixing times for the rook’s walk via path coupling. *to appear in Involve*, pages 1–12, 2016.
- [89] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [90] Bojan Mohar. Isoperimetric numbers of graphs. *Journal of Combinatorial Theory, Series B*, 47(3):274–291, 1989.
- [91] Ben Morris and Alistair Sinclair. Random Walks on Truncated Cubes and Sampling 0-1 Knapsack Solutions. *SIAM Journal on Computing*, 34(1):195–226, 2004.
- [92] Benjamin Nill and Johannes Hofschreier. Personal communication, 2016.
- [93] Masahiro Ohtani. Graphs and Ideals Generated by Some 2-Minors. *Communications in Algebra*, 39(3):905–917, 2011.
- [94] Loïc Paulevé, Gheorghe Craciun, and Heinz Koepl. Dynamical properties of Discrete Reaction Networks. *Journal of Mathematical Biology*, 69(1):55–72, 2014.
- [95] Oskar Perron. Zur Theorie der Matrices. *Mathematische Annalen*, 64(2):248–263, 1907.
- [96] Georg Pick. Geometrisches zur Zahlenlehre. *Sonderabdr. Naturw.-medizin. Verein f. Böhmen “Lotos”*, 19:311–319, 1899.
- [97] Samu Potka. Higher connectivity of fiber graphs of Gröbner bases. *Journal of Algebraic Statistics*, 4(1):93–107, 2013.
- [98] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2015. Available at <http://www.r-project.org/>.
- [99] Fabio Rapallo and Ruriko Yoshida. Markov bases and subbases for bounded contingency tables. *Annals of the Institute of Statistical Mathematics*, 62(4):785–805, 2010.

- [100] Johannes Rauh. Generalized binomial edge ideals. *Advances in Applied Mathematics*, 50(3):409–414, 2013.
- [101] Johannes Rauh and Seth Sullivant. Lifting Markov bases and higher codimension toric fiber products. *Journal of Symbolic Computation*, 74:276–307, 2016.
- [102] John E. Reeve. On the volume of lattice polyhedra. *Proceedings of the London Mathematical Society*, 3(7):378–395, 1957.
- [103] Marcus Schaefer. Realizability of Graphs and Linkages. In *Thirty Essays on Geometric Graph Theory*, pages 461–482. Springer New York, 2013.
- [104] András Sebö. Hilbert Bases, Carathéodory’s Theorem and Combinatorial Optimization. In Ravindran Kannan and William R. Pulleyblank, editors, *Proceedings of the 1st Conference on Integer Programming and Combinatorial Optimization*, pages 431–455. University of Waterloo Press, 1990.
- [105] Alistair Sinclair. Improved Bounds for Mixing Rates of Markov Chains and Multicommodity Flow. *Combinatorics, Probability and Computing*, 1(4):351–370, 1992.
- [106] Alistair Sinclair and Mark Jerrum. Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation*, 82(1):93–133, 1989.
- [107] Neil J. A. Sloane. The on-line encyclopedia of integer sequences. published electronically at <https://oeis.org>.
- [108] Caprice Stanley and Tobias Windisch. Heat-bath random walks with Markov bases. *preprint, arXiv:1605.08386*, pages 1–20, 2016.
- [109] Bernd Sturmfels. *Gröbner bases and convex polytopes*. American Mathematical Society, Providence, RI, 1996.
- [110] Seth Sullivant. Small contingency tables with large gaps. *SIAM Journal on Discrete Mathematics*, 18(4):787–793, 2005.
- [111] Seth Sullivant. Toric fiber products. *Journal of Algebra*, 316(2):560–577, 2007.
- [112] Seth Sullivant and Mike Develin. Markov Bases of Binary Graph Models. *Annals of Combinatorics*, 7(4):441–466, 2003.
- [113] Rekha R. Thomas and Robert Weismantel. Truncated Gröbner Bases for Integer Programming. *Applicable Algebra in Engineering, Communication and Computing*, 8(4):241–256, 1997.
- [114] Carsten Thomassen. Embeddings of graphs. *Discrete Mathematics*, 124(1-3):217–228, 1994.
- [115] Mario Ullrich. Rapid mixing of Swendsen–Wang dynamics in two dimensions. *Dissertationes Mathematicae*, 502:1–64, 2014.
- [116] Santosh S. Vempala. Geometric Random Walks: A Survey. *MSRI Combinatorial and Computational Geometry*, 52:573–612, 2005.

- [117] Sven Verdoolaege, Rachid Seghir, Kristof Beyls, Vincent Loechner, and Maurice Bruynooghe. Counting integer points in parametric polytopes using Barvinok’s rational functions. *Algorithmica*, 48(1):37–66, 2007.
- [118] Bálint Virág. Random walks on finite convex sets of lattice points. *Journal of Theoretical Probability*, 11(4):935–951, 1998.
- [119] Tobias Windisch. Rapid mixing and Markov bases. *SIAM Journal on Discrete Mathematics*, 30(4):2130–2145, 2016.
- [120] Tobias Windisch. The fiber dimension of a graph. *preprint, arXiv:1601.04527*, pages 1–10, 2016.
- [121] Zhenxiang Zhang. On a Problem of Erdős Concerning Primitive Sequences. *Mathematics of Computation*, 60(202):827–834, 1993.
- [122] Günter M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer Verlag, New York, 1994.

Index

- affine semigroup, 6, 27
- algebraic statistics, 1, 6, 16
- binomial
 - primitive, 66
 - pure difference, 64
 - square-free, 66
- binomial edge ideal, *see* graph
- Blichfeldt's theorem, 35
- Buchberger's criterion, 67, 72
- canonical path, *see* random walk
- difference graph, *see* graph
- Discrete reaction network, 14
- discrete steepest decent, 57
- distinct pair-sum polytope, *see* polytope
- Ehrhart's theorem, 27, 45, 59
- eigenfunction, 4
- Euler's totient function, *see* function
- Fan's theorem, 55
- fiber, *see* saturated set
- fiber dimension, *see* graph
- fiber graph, 6
 - compression, 8
- Four color theorem, 33
- function
 - concave, 23
 - eigenvalue, 4
 - Euler's totient function, 34
 - mass function, 3
 - periodic, 27
 - quasi-polynomial, 27, 28
 - variance, 11
- Glauber dynamics, *see* random walk
- Gröbner basis
 - of ideal, 15
 - of matrix, 17
 - reduced, 17
 - universal, 66, 71
- Graph
 - bipartite, 73
- graph
 - k -edge-connected, 17
 - k -node-connected, 17
 - r -partite, 2, 36
 - adjacency matrix, 2
 - binomial edge ideal, 63, 67, 72
 - bipartite, 63
 - cartesian product, 32
 - Cayley graph, 47
 - Cheeger constant, 24
 - chromatic number, 33
 - circuit, 2
 - circuit graph, 2
 - complete graph, 2
 - cycle, 2
 - degree, 2, 22
 - diameter, 18
 - difference graph, 33
 - dimension, 32
 - disconnecter, 73
 - effective, 78
 - distance, 18
 - edge-connectivity, 17, 48
 - edge-expansion, 24
 - Euclidean dimension, 32, 38
 - fiber dimension, 31
 - grid graph, 34
 - implicit, 3
 - induced subgraph, 2
 - isoperimetric number, 24
 - minimal degree, 22
 - node-connectivity, 17, 47
 - parity binomial edge ideal, 63

- path, 2
 - augmenting, 56
- planar, 33
- regular, 5
- star graph, 35
- tensor product, 32
- walk, 2
 - even, 64
 - interior, 64
 - minimal, 66
 - odd, 64
- graph product
 - cartesian product, 17, 32
 - tensor product, 32
- Graver basis
 - ideal, 66
 - matrix, 20, 48, 56, 65
 - of lattice, 8
 - of matrix, 8
- Handshaking lemma, 34
- heat-bath walk, *see* random walk
- hit-and-run algorithm, 49
- ideal
 - Gröbner basis, *see* Gröbner basis
 - primary, 78
 - prime, 73
 - pure difference, 66
 - radical, 77
- intersecting ray property, *see* set of moves
- lattice, 6
 - saturated, 6
- lattice ideal, 15
 - toric ideal, 15
- lattice polytope, *see* polytope
- lattice walk, 9
- log-linear model, 12, 48
 - goodness-of-fit, 12, 48
 - hierarchical, 14
 - independence model, 10, 14, 43, 47
 - sample size, 41
 - structural zeros, 48
 - sufficient statistics, 13
- map, *see* function
- Markov basis, 7
 - adapted, 45
 - for matrix, 7
 - minimal, 10, 61
 - of collection, 7
- Markov chain, 3
 - comparison, *see* random walk
 - Monte Carlo, 11, 12
- matrix
 - Graver basis, *see* Graver basis
 - idempotent, 52
 - ray matrix, *see* set of moves
 - Smith normal form, 72
 - totally unimodular, 24, 27
- Metropolis-Hastings walk, *see* random walk
- mixing time, *see* random walk
- parity binomial edge ideal, *see* graph
- Perron-Frobenius theorem, 4
- Pick's theorem, 38, 39
- polytope, 6
 - cross-polytope, 45, 47, 59
 - dilation, 2
 - distinct pair-sum polytope, 30, 36
 - lattice polytope, 6
 - lattice width, 19
 - Reeve tetrahedra, 38
- Pottier's termination criterion, 65
- primitive set, *see* set of moves
- quasi-polynomial, *see* function
- random walk, 3
 - aperiodic, 3
 - burn-in, 11
 - canonical path, 58
 - comparison, 43
 - conductance, 6
 - expander, 5
 - fiber walk, 8
 - Glauber dynamics, 50
 - heat-bath walk, 49
 - irreducible, 3
 - lazy version, 6, 43, 44
 - Metropolis-Hastings walk, 9, 11, 12, 47
 - mixing time, 4
 - rapidly mixing, 5
 - reducible, 3

- reversible, 3
- Rook's walk, 55
- second largest eigenvalue modulus, 4
- simple fiber walk, 8
- simple walk, 3
- stationary distribution, 3
- symmetric, 3
- ray matrix, *see* set of moves
- Reeve tetrahedron, *see* polytope
- Rook's walk, *see* random walk

- saturated set, 6
 - affinely isomorphic, 38
 - boundary, 26
 - dimension, 31
 - fiber, 6
 - relaxation, 25
 - unimodular isomorphic, 38
- second largest eigenvalue modulus, *see* random walk
- sequence
 - dominated, 25
 - meaningful parametrization, 41
 - ray, 25
 - subsequence, 2
- set of directions, *see* set of moves
- set of moves
 - augmentation length, 56
 - augmenting, 56
 - Gröbner basis, *see* Gröbner basis
 - Graver basis, *see* Graver basis
 - intersecting ray property, 54
 - Markov basis, *see* Markov basis
 - norm-like, 19
 - norm-reducing, 19
 - primitive, 33
 - ray matrix, 53
 - set of directions, 31
- sign-split, 75
- simple fiber walk, *see* random walk
- simple walk, *see* random walk
- SLEM, *see* second largest eigenvalue modulus
- Smith normal form, *see* matrix
- Stiefel manifold, 55

- total variation distance, 4

- Universality theorem, 48