

Log-Concave Polynomials II: High-Dimensional Walks and an FPRAS for Counting Bases of a Matroid

Nima Anari¹, Kuikui Liu², Shayan Oveis Gharan², and Cynthia Vinzant³

¹Stanford University, anari@cs.stanford.edu

²University of Washington, liukui17@cs.washington.edu, shayan@cs.washington.edu

³North Carolina State University, clvinzan@ncsu.edu

February 11, 2019

Abstract

We design an FPRAS to count the number of bases of any matroid given by an independent set oracle, and to estimate the partition function of the random cluster model of any matroid in the regime where $0 < q < 1$. Consequently, we can sample random spanning forests in a graph and (approximately) compute the reliability polynomial of any matroid. We also prove the thirty year old conjecture of Mihail and Vazirani that the bases exchange graph of any matroid has expansion at least 1.

Our algorithm and the proof build on the recent results of Dinur, Kaufman, Mass and Oppenheim [KM17; DK17; KO18] who show that high dimensional walks on simplicial complexes mix rapidly if the corresponding localized random walks on 1-skeleton of links of all complexes are strong spectral expanders. One of our key observations is a close connection between pure simplicial complexes and multiaffine homogeneous polynomials. Specifically, if X is a pure simplicial complex with positive weights on its maximal faces, we can associate with X a multiaffine homogeneous polynomial p_X such that the eigenvalues of the localized random walks on X correspond to the eigenvalues of the Hessian of derivatives of p_X .

1 Introduction

Let $\mu : 2^{[n]} \rightarrow \mathbb{R}_{\geq 0}$ be a probability distribution on the subsets of the set $[n] = \{1, 2, \dots, n\}$. We assign a multiaffine polynomial with variables x_1, \dots, x_n to μ ,

$$g_\mu(x) = \sum_{S \subseteq [n]} \mu(S) \cdot \prod_{i \in S} x_i.$$

The polynomial g_μ is also known as the *generating polynomial* of μ . A polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ is d -homogeneous if every monomial of p has degree d . We say μ is d -homogeneous if the polynomial g_μ is d -homogeneous, meaning that $|S| = d$ for any S with $\mu(S) > 0$.

A polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ with nonnegative coefficients is *log-concave* on a subset $K \subseteq \mathbb{R}_{\geq 0}^n$ if $\log p$ is a concave function at any point in K , or equivalently, its Hessian $\nabla^2 \log p$ is negative

semidefinite on K . We say a polynomial p is *strongly log-concave* on K if for any $k \geq 0$, and any sequence of integers $1 \leq i_1, \dots, i_k \leq n$,

$$(\partial_{i_1} \cdots \partial_{i_k} p)(x_1, \dots, x_n)$$

is log-concave on K . In this paper, for convenience and clarity, we only work with (strong) log-concavity with respect to the all-ones vector, $\mathbb{1}$. So, unless otherwise specified, $K = \{\mathbb{1}\}$ in the above definition. We say the distribution μ is *strongly log-concave* at $\mathbb{1}$ if g_μ is strongly log-concave at $\mathbb{1}$. The notion of strong log-concavity was first introduced by Gurvits [Gur09; Gur10] to study approximation algorithms for mixed volume and multivariate generalizations of Newton’s inequalities.

In this paper, we show that the “natural” Monte Carlo Markov Chain (MCMC) method on the support of a d -homogeneous strongly log-concave distribution $\mu : 2^{[n]} \rightarrow \mathbb{R}_{\geq 0}$ mixes rapidly. This chain can be used to generate random samples from a distribution arbitrarily close to μ .

The chain \mathcal{M}_μ is defined as follows. We take the state space of \mathcal{M}_μ to be the support of μ , namely $\text{supp}(\mu) = \{S \subseteq [n] \mid \mu(S) \neq 0\}$. For $\tau \in \text{supp}(\mu)$, first we drop an element $i \in \tau$, chosen uniformly at random from τ . Then, among all sets $\sigma \supset \tau \setminus \{i\}$ in the support of μ we choose one with probability proportional to $\mu(\sigma)$.

It is easy to see that \mathcal{M}_μ is reversible with stationary distribution μ . Furthermore, assuming g_μ is strongly log-concave, we will see that \mathcal{M}_μ is irreducible. We prove that this chain mixes rapidly. More formally, for a state τ of the Markov chain \mathcal{M} , and $\epsilon > 0$, the *total variation mixing time* of \mathcal{M} started at τ with transition probability matrix P and stationary distribution π is defined as follows:

$$t_\tau(\epsilon) = \min\{t \in \mathbb{Z}_{\geq 0} \mid \|P^t(\tau, \cdot) - \pi\|_1 \leq \epsilon\},$$

where $P^t(\tau, \cdot)$ is the distribution of the chain started at τ at time t .

The following theorem is the main result of this paper.

Theorem 1.1. *Let $\mu : 2^{[n]} \rightarrow \mathbb{R}_{\geq 0}$ be a d -homogeneous strongly log-concave probability distribution. If P_μ denotes the transition probability matrix of \mathcal{M}_μ and $X(k)$ denotes the collection of size- k subsets of $[n]$ which are contained in some element of $\text{supp}(\mu)$, then for every $0 \leq k \leq d - 1$, P_μ has at most $|X(k)| \leq \binom{n}{k}$ eigenvalues of value $> 1 - \frac{k+1}{d}$. In particular, \mathcal{M}_μ has spectral gap at least $1/d$, and if τ is in the support of μ and $0 < \epsilon < 1$, the total variation mixing time of the Markov chain \mathcal{M}_μ started at τ is at most*

$$t_\tau(\epsilon) \leq d \log\left(\frac{1}{\epsilon \mu(\tau)}\right).$$

To state the key corollaries of this theorem, we will need the following definition.

Definition 1.2. Given a domain set Ω , which is compactly represented by, say, a membership oracle, and a nonnegative weight function $w : \Omega \rightarrow \mathbb{R}_{\geq 0}$, a *fully polynomial-time randomized approximation scheme (FPRAS)* for computing the partition function $Z = \sum_{x \in \Omega} w(x)$ is a randomized algorithm that, given an error parameter $0 < \epsilon < 1$ and error probability $0 < \delta < 1$, returns a number \tilde{Z} such that $\text{Prob}[(1 - \epsilon)Z \leq \tilde{Z} \leq (1 + \epsilon)Z] \geq 1 - \delta$. The algorithm is required to run in time polynomial in the problem input size, $1/\epsilon$, and $\log(1/\delta)$.

Equipped with this definition, we can now concisely state the main applications of [Theorem 1.1](#). [Theorem 1.1](#) gives us an algorithm to efficiently sample from a distribution which approximates

μ closely in total variation distance. By the equivalence between approximate counting and approximate sampling for self-reducible problems [JVV86], this gives an FPRAS for each of the following:

1. counting the bases of a matroid, and
2. estimating the partition function of the random cluster model for a new range of parameter values

For real linear matroids, we also give an algorithm for estimating the partition function of a generalized version of a k -determinantal point process. Note that these problems are all instantiations of the following: estimate the partition function of some efficiently computable nonnegative weights on bases of a matroid. Furthermore, as the restriction and contraction of a matroid by a subset of the ground set are both (smaller) matroids, problems of this form are indeed self-reducible. In the following sections we discuss these applications in greater depth.

1.1 Counting Problems on Matroids

Let $M = ([n], \mathcal{I})$ be an arbitrary matroid on n elements (see Section 2.3) of rank r . Let μ be the uniform distribution on the bases of the matroid M . It follows that μ is r -homogeneous. Using the Hodge-Riemann relation proved by Adiprasito, Huh, and Katz [AHK18], a subset of the authors proved [AOV18] that for any matroid M , μ is strongly log-concave.¹ This implies that the chain \mathcal{M}_μ converges rapidly to stationary distribution. This gives the first polynomial time algorithm to generate a uniformly random base of a matroid. Note that to run \mathcal{M}_μ we only need an oracle to test whether a given set $S \subseteq [n]$ is an independent set of M . Therefore, with only polynomially many queries (in $n, r, \log(1/\epsilon)$) we can generate a random base of M .

Corollary 1.3. *For any matroid $M = ([n], \mathcal{I})$ of rank r , any basis B of M and $0 < \epsilon < 1$, the mixing time of the Markov chain \mathcal{M}_μ starting at B is at most*

$$t_B(\epsilon) \leq r \log(n^r/\epsilon) \leq r^2 \log(n/\epsilon).$$

To prove this we simply used the fact that a matroid of rank r on n elements has at most $\binom{n}{r} \leq n^r$ bases. There are several immediate consequences of the above corollary. Firstly, by equivalence of approximate counting and approximate sampling for self-reducible problems [JVV86] we can count the number of bases of any matroid given by an independent set oracle up to a $1 + \epsilon$ multiplicative error in polynomial time.

Corollary 1.4. *There is a randomized algorithm that for any matroid M on n elements with rank r given by an independent set oracle, and any $0 < \epsilon < 1$, counts the number of bases of M up to a multiplicative factor of $1 \pm \epsilon$ with probability at least $1 - \delta$ in time polynomial in $n, r, 1/\epsilon, \log(1/\delta)$.*

As an immediate corollary for any $1 \leq k \leq r$ we can count the number of independent sets of M of size k . This is because if we truncate M to independent sets of size at most k it remains

¹Indeed, in [AOV18], it is shown that g_μ satisfies a seemingly stronger property known as “complete log-concavity”, namely that $\partial_{v_1} \cdots \partial_{v_k} g_M$ is log-concave (at $\mathbf{1}$) for any sequence of *directional* derivatives $\partial_{v_1} \cdots \partial_{v_k}$ with nonnegative directions $v_1, \dots, v_k \in \mathbb{R}_{\geq 0}^n$. We will prove in a future companion paper that complete log-concavity is equivalent to strong log-concavity.

a matroid. As a consequence we can generate uniformly random forests in a given graph, and compute the reliability polynomial

$$C_M(p) = \sum_{S \subseteq [n]: \text{rank}(S)=r} (1-p)^{|S|} p^{n-|S|}$$

for any matroid and $0 \leq p \leq 1$, all in polynomial time. Note this latter fact follows from the ability to count the number of independent sets of a fixed size, as the complements of rank- r subsets $S \subseteq [n]$ are precisely the independent sets of the dual of M . Prior to our work, we could only compute the reliability polynomial for graphic matroids due to a recent work of Guo and Jerrum [GJ18a].

One can associate a graph G_M to any matroid M , called *the bases exchange graph*. This graph has a vertex for every basis of M and two bases B, B' are connected by an edge if $|B \Delta B'| = 2$. It follows by the bases exchange property of matroids that this graph is connected. For an unweighted graph $G = (V, E)$, the expansion of a set $S \subset V$ and the graph G are defined as

$$h(S) = \frac{|E(S, \bar{S})|}{|S|} \quad \text{and} \quad h(G) = \min_{S: |S| \leq |\bar{S}|} h(S).$$

Mihail and Vazirani [MV89] conjectured that the bases exchange graph has expansion at least one, i.e., that $h(G_M) \geq 1$, for any matroid M . It turns out that the bases exchange graph is closely related to the Markov chain \mathcal{M}_μ . The following theorem is an immediate consequence of the above corollary.

Theorem 1.5. *For any matroid M , the expansion of the bases exchange graph is at least 1, $h(G_M) \geq 1$.*

1.2 The Random Cluster Model

Another application of this theory is estimating the partition function of the random cluster model. For a matroid $M = ([n], \mathcal{I})$ of rank r and parameters p, q , the partition function of the random cluster model from statistical mechanics due to Fortuin and Kasteleyn [For71; FK72; For72a; For72b] is the following polynomial function associated to M ,

$$Z_M(p, q) = \sum_{S \subseteq [n]} q^{r+1-\text{rank}(S)} p^{|S|}$$

where $\text{rank}(S)$ is the size of the largest independent set contained in S . We note that typically one scales each term by $(1-p)^{n-|S|}$ but up to a normalization factor (and change of variables) the two polynomials are equivalent. We refer interested readers to a recent book of Grimmett [Gri09] for further information. Typically, one considers the special case where M is a graphic matroid, in which case the exponent of q is simply the number of connected components of S . To the best of our knowledge, prior to this work, one could only compute Z_M when $q = 2$ because of the close connection to the Ising model [JS93; GJ17]. Our next result is a polynomial time algorithm that estimates $Z_M(p, q)$ for any $0 < q \leq 1$ and $p \geq 0$.

Theorem 1.6. *For a matroid M with rank function $\text{rank} : 2^{[n]} \rightarrow \mathbb{Z}_{\geq 0}$, parameter $0 < q \leq 1$ and choice of "external field" $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathbb{R}_{>0}^n$, the polynomial*

$$f_{M,k,q}(x_1, \dots, x_n) = \sum_{S \in \binom{[n]}{k}} q^{-\text{rank}(S)} \prod_{i \in S} \lambda_i x_i$$

is strongly log-concave.

Together with [Theorem 1.1](#), this gives an FPRAS for estimating $f_{M,k,q}(\mathbb{1})$ given an independence oracle for the matroid M . Estimating $Z_M(p, q)$ then follows as

$$Z_M(p, q) = q^{r+1} \sum_{k=0}^n p^k f_{M,k,q}(\mathbb{1})$$

and each term is nonnegative. In fact, the polynomial Z_M is closely related to the Tutte polynomial

$$T_M(x, y) = \sum_{S \subseteq [n]} (x-1)^{r-\text{rank}(S)} (y-1)^{|S|-\text{rank}(S)}.$$

Indeed, we can write

$$T_M(x, y) = \frac{1}{(x-1)(y-1)^{r+1}} Z_M(y-1, (x-1)(y-1))$$

Hence, an FPRAS for estimating $Z_M(p, q)$ for $p \geq 0$ and $0 \leq q \leq 1$ gives an FPRAS for estimating $T_M(x, y)$ in the region described by the inequalities $y \geq 1$ and $0 \leq (x-1)(y-1) \leq 1$.

1.3 Determinantal Distributions on Real Linear Matroids

Finally, we show that the class of homogeneous multiaffine strongly log-concave polynomials is closed under raising all coefficients to a fixed exponent less than 1.

Theorem 1.7. *Let $f = \sum_{S \subseteq [n]} c_S \prod_{i \in S} x_i$ be a homogeneous degree- k multiaffine strongly log-concave polynomial. Then $f_\alpha = \sum_{S \subseteq [n]} c_S^\alpha \prod_{i \in S} x_i$ is strongly log-concave for every $0 \leq \alpha \leq 1$.*

We use the above theorem to design a sampling algorithm for determinantal point processes. A *determinantal point process (DPP)* on a set of elements $[n]$ is a probability distribution $\mu : 2^{[n]} \rightarrow \mathbb{R}_{\geq 0}$ identified by a positive semidefinite matrix $L \in \mathbb{R}^{n \times n}$ where for any $S \subseteq [n]$ we have

$$\mu(S) \propto \det(L_S),$$

where L_S is the principal sub-matrix of L indexed by the elements of S . Determinantal point processes are fundamental to the study of a variety of tasks in machine learning, including text summarization, image search, news threading, and diverse feature selection [see, e.g., [KT12](#)]. A *k-determinantal point process (k-DPP)* is a determinantal point process conditioned on the sets S having size k .

Given a positive semidefinite matrix L , let μ be the corresponding k -DPP. We have

$$g_\mu(x) \propto \sum_{S \in \binom{[n]}{k}} \det(L_S) \cdot \prod_{i \in S} x_i.$$

It turns out that the above polynomial is real stable and so it is strongly log-concave over $\mathbb{R}_{\geq 0}^n$ [see, e.g., [AOR16](#)]. Anari, Oveis Gharan, and Rezaei [[AOR16](#)] show that a natural Markov chain with the Metropolis rule mixes rapidly and generates a random sample of μ . The above theorem immediately implies the following log-concavity result.

Corollary 1.8. For every positive semidefinite matrix $L \succeq 0$ and exponent $0 \leq \alpha \leq 1$, the polynomial

$$\sum_{S \in \binom{[n]}{k}} \det(L_S)^\alpha \prod_{i \in S} x_i$$

is strongly log-concave.

It follows from [Theorem 1.1](#) that for any $0 \leq \alpha \leq 1$ we can generate samples from a “smoothed” k -DPP distribution, where for any set S , $\mathbb{P}[S] \propto \det(L_S)^\alpha$, in polynomial time. The weights $\det(L_S)^\alpha$ may be thought of as a way to interpolate between two extremes for selecting diverse data points.

We also note that for $\alpha = 1/2$, it is known that [Corollary 1.8](#) follows from the Brunn-Minkowski theorem applied to appropriately defined zonotopes. For $\alpha = 0$ when the k -DPP has full support, and for $\alpha = 1$ as mentioned earlier, the above polynomial is actually real stable, and hence strongly log-concave. [Theorem 1.7](#) gives a unified proof that all of these polynomials are strongly log-concave.

1.4 Related Works

There is a long line of work on designing approximation algorithms to count the bases of a matroid. Most of these works focus on expansion properties of bases exchange graph. Feder and Mihail [[FM92](#)] showed that for a special class of matroids known as *balanced matroids* [[MS91](#); [FM92](#)], the bases exchange graph has expansion at least 1. A matroid M is balanced if for any minor of M (including M itself), the uniform distribution over its bases satisfies the pairwise negative correlation property. Many of the extensive results in this area [[Gam99](#); [JS02](#); [Jer+04](#); [Jer06](#); [Clo10](#); [CTY15](#); [AOR16](#)] only study approximation algorithms for this limited class of matroids, and not much is known beyond the class of balanced matroids. Unfortunately, many interesting matroids are not balanced. An important example is the matroid of all acyclic subsets of edges of a graph $G = (V, E)$ of size at most k (for some $k < |V| - 1$) [[FM92](#)].

There has been other approaches for counting bases. Guo and Jerrum [[GJ18b](#)] used the popping method to count bases of bicircular matroids. Barvinok and Samorodnitsky [[BS07](#)] designed a *randomized* algorithm that gives, roughly, a $\log(n)^r$ approximation factor to the number of bases of a given matroid with n elements and rank r . In [[AOV18](#)], a subset of the authors gave a deterministic e^r approximation to the number of bases using the fact that $g_\mu(M)$ is log-concave over $\mathbb{R}_{\geq 0}^n$.

There is an extensive literature on hardness of exact computation and inapproximability of the Tutte polynomial and the partition function of the random cluster model. It is known that exact computation of the Tutte polynomial for a graph is #P-hard at all points (x, y) except at $(1, 1)$, $(-1, -1)$, $(0, -1)$, $(-1, 0)$, along the hyperbola $(x - 1)(y - 1) = 1$, and for planar graphs, along the hyperbola $(x - 1)(y - 1) = 2$ [[JVW90](#)], [[Ver91](#)], [[Wel94](#)]. In the realm of inapproximability, it is known that even for planar graphs, there is no FPRAS to approximate the Tutte polynomial for $x > 1, y < -1$ or $y > 1, x < -1$ assuming $\text{NP} \neq \text{RP}$ [[GJ08](#); [GJ12b](#)]. Furthermore, there is no FPRAS for estimating the partition function Z_M of the random cluster model on general graphic matroids when $q > 2$, nor is there an FPRAS for Z_M at $q = 2$ for general binary matroids, unless there is an FPRAS for counting independent sets in a bipartite graph [[GJ12a](#); [GJ13](#); [GJ14](#)].

1.5 Independent Work

In a closely related upcoming work, Brändén and Huh, in a slightly different language, independently prove the strong log-concavity of several of the polynomials that appear in this paper. In upcoming papers, both groups of authors use these techniques to prove the strongest form of Mason’s conjecture and further study closure properties of (strongly) log-concave polynomials.

1.6 Techniques

One of our key observations is a close connection between pure simplicial complexes and multiaffine homogeneous polynomials. Specifically, if X is a pure simplicial complex with positive weights on its maximal faces, we can associate with X a multiaffine homogeneous polynomial p_X such that the eigenvalues of the localized random walks on X correspond to the eigenvalues of the Hessian of derivatives of p_X . Using this correspondence, one can study multiaffine homogeneous

| Weighted Simplicial Complex X | Multiaffine Polynomial p_X |
|---------------------------------|------------------------------|
| Dimension- d | Degree- d |
| Weight of \emptyset | Evaluation at $\mathbf{1}$ |
| Connectivity of Links | Indecomposability |
| Link | Differentiation |
| Local Random Walk | (Normalized) Hessian |

polynomials using techniques from simplicial complexes, and vice versa. To study the walk \mathcal{M}_μ corresponding to a polynomial g_μ , we analyze the simplicial complex corresponding to g_μ . To do this, we leverage recent developments in the area of high-dimensional expanders, which we discuss below.

Given a simplicial complex X (see [Section 2.4](#)) and an ordering of its vertices, one can associate a high dimensional Laplacian matrix to the k -dimensional faces of X . These matrices generalize the classical graph Laplacian and there has been extensive research to study their eigenvalues [see [Lub17](#), and the references therein]. A method known as Garland’s method [[Gar73](#)] relates the eigenvalues of graph Laplacians of 1-skeletons of links of X to eigenvalues of high dimensional Laplacians of X [see [BS97](#); [Opp18](#)].

Recently, Kaufman and Mass [[KM17](#)] studied a high dimensional walk on a simplicial complex, which is closely related to the walk \mathcal{M}_μ that we defined above (see [Section 3](#)). Their goal is to argue that, similar to classical expander graphs, high dimensional walks mix rapidly on a high dimensional expander. Their bounds were improved in a work of Dinur and Kaufman [[DK17](#)], who showed that if all nontrivial eigenvalues of the simple random walk matrix on all 1-skeletons of links of X have absolute value at most λ , then the high dimensional walk on k -faces of X has spectral gap at least $\frac{1}{k+2} - O((k+1)\lambda)$. This was further improved in a recent work of Kaufman and Oppenheim [[KO18](#)]: They showed that if all non-trivial eigenvalues of the simple random walk matrix on all 1-skeleton of links of X are at most λ , then the spectral gap of the high dimensional walk is at least $\frac{1}{k+2} - (k+1)\lambda$. In other words, negative eigenvalues of the random walk matrix do not matter. One only needs positive eigenvalues to be small.

Note that in order to make the spectral gap bounds meaningful one needs $\lambda \ll \frac{1}{k^2}$. In other words, one needs that, except the trivial eigenvalue of 1, all other eigenvalues are either negative or very close to 0. Here is the place where the connection to (strong) log-concavity comes into

the picture. A polynomial p is log-concave at $\mathbb{1}$ if $\nabla^2 p(\mathbb{1})$ has at most one positive eigenvalue. A polynomial p is strongly log-concave if the same holds for all partial derivatives of p . Our main observation is that this property is equivalent to taking $\lambda = 0$ in the corresponding simplicial complex. Namely, we obtain the best possible spectral gap of $\frac{1}{k+2}$ when the simplicial complex comes from a strongly log-concave polynomial.

Our approach has a close connection to the original plan of Feder and Mihail [FM92] who used the negative correlation property of balanced matroids to show that the bases exchange walk mixes rapidly. Unfortunately, most interesting matroids do not satisfy negative correlation. But it was observed [AHK18; HW16; AOV18] that all matroids satisfy a *spectral* negative dependence property. Namely, consider the uniform distribution μ over the bases of a matroid M , and consider the Hessian $\nabla^2 \log(g_\mu)$ of the log of the generating polynomial g_μ at the point $x = \mathbb{1}$. Then M is negatively correlated if and only if all off-diagonal entries of this matrix are non-positive, whereas M being spectrally negatively correlated means that this matrix is negative semidefinite. Spectral negative correlation is precisely what one needs to bound the mixing time of the high dimensional walk on the corresponding simplicial complex.

Structure of the paper. In Section 2 we discuss necessary background on linear algebra, matroids, simplicial complexes and strongly log-concave polynomials. We also provide a useful characterization of strong log-concavity. In Section 3 we discuss and reprove a version of the main theorem of Kaufman and Oppenheim [KO18] on mixing time of high dimensional walks, Theorem 3.3. In Section 4 we use this to prove Theorem 1.1 and the Mihail-Vazirani conjecture, Theorem 1.5. Finally, in Section 5 we first prove our new characterization of strong log-concavity and discuss its applications. Specifically, we give a self-contained proof that the uniform distribution over the bases of a matroid is strongly log-concave and we prove Theorems 1.6 and 1.7.

Acknowledgements. Part of this work was started while the first and last authors were visiting the Simons Institute for the Theory of Computing. It was partially supported by the DIMACS/Simons Collaboration on Bridging Continuous and Discrete Optimization through NSF grant CCF-1740425. Shayan Oveis Gharan and Kuikui Liu are supported by the NSF grant CCF-1552097 and ONR-YIP grant N00014-17-1-2429. Cynthia Vinzant was partially supported by the National Science Foundation grant DMS-1620014.

We thank Lap Chi Lau, Mark Jerrum and Alan Frieze for helpful comments on an earlier version of this manuscript.

2 Preliminaries

First, let us establish some notational conventions. Unless otherwise specified, all logarithms are in base e . All vectors are assumed to be column vectors. For two vectors $\phi, \psi \in \mathbb{R}^n$, we use $\langle \phi, \psi \rangle$ to denote the standard Euclidean inner product between ϕ and ψ . We use $\mathbb{R}_{>0}$ and $\mathbb{R}_{\geq 0}$ to denote the set of positive and nonnegative real numbers, respectively, and $[n]$ to denote $\{1, \dots, n\}$. For a vector $x \in \mathbb{R}^n$ and a set $S \subseteq [n]$, we let x^S denote $\prod_{i \in S} x_i$.

We use ∂_{x_i} or ∂_i to denote the partial differential operator $\partial/\partial x_i$. We denote the gradient of a function or polynomial p by ∇p and the Hessian of p by $\nabla^2 p$.

2.1 Linear Algebra

We say a matrix $A \in \mathbb{R}^{n \times n}$ is *stochastic* if all entries of A are nonnegative and every row adds up to exactly 1. It is well-known that the largest eigenvalue in magnitude of any stochastic matrix is 1 and its corresponding eigenvector is the all-ones vector, $\mathbf{1}$. If the eigenvalues of a matrix $A \in \mathbb{R}^{n \times n}$ are all real, then we order them as

$$\lambda_n(A) \leq \dots \leq \lambda_1(A).$$

A symmetric matrix $A \in \mathbb{R}^{n \times n}$ is positive semidefinite (PSD), denoted $A \succcurlyeq 0$, if all its eigenvalues $\lambda_k(A)$ are nonnegative, or equivalently if for all $v \in \mathbb{R}^n$,

$$v^\top A v \geq 0.$$

Similarly, A is negative semidefinite (NSD), denoted $A \preccurlyeq 0$, if $v^\top A v \leq 0$ for all $v \in \mathbb{R}^n$. Equivalently, a real symmetric matrix is PSD (NSD) if its eigenvalues are nonnegative (nonpositive), respectively.

Theorem 2.1 (Schur Product Theorem [HJ13, Thm 7.5.3]). *If $A, B \in \mathbb{R}^{n \times n}$ are positive semidefinite, then their Hadamard product $A \circ B$, whose entries are $(A \circ B)_{i,j} = A_{i,j}B_{i,j}$, is positive semidefinite.*

Theorem 2.2 (Perron-Frobenius Theorem [HJ13, Ch. 8]). *Let $A \in \mathbb{R}^{n \times n}$ be symmetric and have strictly positive entries. Then A has an eigenvalue λ which is strictly positive. Furthermore, it has multiplicity one and its corresponding eigenvector v has strictly positive entries.*

Theorem 2.3 (Cauchy's Interlacing Theorem [HJ13, Corollary 4.3.9]). *For a symmetric matrix $A \in \mathbb{R}^{n \times n}$ and vector $v \in \mathbb{R}^n$, the eigenvalues of A interlace the eigenvalues of $A + vv^\top$. That is, for $B = A + vv^\top$,*

$$\lambda_n(A) \leq \lambda_n(B) \leq \lambda_{n-1}(A) \leq \dots \leq \lambda_2(B) \leq \lambda_1(A) \leq \lambda_1(B).$$

The following is an immediate consequence:

Lemma 2.4. *Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix and let $P \in \mathbb{R}^{m \times n}$. If A has at most one positive eigenvalue, then PAP^\top has at most one positive eigenvalue.*

Proof. Since A has at most one positive eigenvalue, we can write $A = B + vv^\top$ for some vector $v \in \mathbb{R}^n$ and some negative semidefinite matrix B . Then $PAP^\top = PBP^\top + Pvv^\top P^\top$. First, observe that $PBP^\top \preccurlyeq 0$, since for $x \in \mathbb{R}^m$, $x^\top PBP^\top x = (P^\top x)^\top B(P^\top x) \leq 0$. Second, let $w = Pv \in \mathbb{R}^m$. Then $Pvv^\top P^\top = ww^\top$ and by **Theorem 2.3**, the eigenvalues of PBP^\top interlace the eigenvalues of $PBP^\top + (Pv)(Pv)^\top$. Since all eigenvalues of PBP^\top are nonpositive, $PAP^\top = PBP^\top + ww^\top$ has at most one positive eigenvalue. \square

The following fact is well-known.

Fact 2.5. *Let $A \in \mathbb{R}^{n \times k}$ and $B \in \mathbb{R}^{k \times n}$ be arbitrary matrices. Then, non-zero eigenvalues of AB are equal to non-zero eigenvalues of BA with the same multiplicity.*

Lemma 2.6. *Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix with at most one positive eigenvalue. Then, for any PSD matrix $B \in \mathbb{R}^{n \times n}$, BA has at most one positive eigenvalue.*

Proof. Since $B \succcurlyeq 0$, we can write $B = C^\top C$ for some $C \in \mathbb{R}^{n \times n}$. By [Fact 2.5](#), $BA = C^\top CA$ has the same nonzero eigenvalues as the matrix CAC^\top . Since A has at most one positive eigenvalue, by [Lemma 2.4](#), CAC^\top has at most one positive eigenvalue and so does BA . \square

Lemma 2.7. *Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix with nonnegative entries and at most one positive eigenvalue, and let $w(i) = \sum_{j=1}^n A_{i,j}$. Then,*

$$A \preceq \frac{ww^\top}{\sum_i w(i)}.$$

Proof. Let $W = \text{diag}(w)$. Then, by [Lemma 2.4](#), $\mathcal{A} = W^{-1/2}AW^{-1/2}$ has at most one positive eigenvalue. Observe that the top eigenvector of \mathcal{A} is the \sqrt{w} vector, where $\sqrt{w}(i) = \sqrt{w(i)}$, for all i . In particular, $\mathcal{A}\sqrt{w} = \sqrt{w}$. So, \sqrt{w} is the only eigenvector of \mathcal{A} with positive eigenvalue and we have

$$\mathcal{A} \preceq \frac{\sqrt{w}\sqrt{w}^\top}{\|\sqrt{w}\|^2} = \frac{\sqrt{w}\sqrt{w}^\top}{\sum_i w(i)}.$$

Multiplying both sides of the inequality on the left and right by $W^{1/2}$ proves the lemma. \square

In this paper, we will often switch between different inner products. As such, we highlight the following variational characterization of eigenvalues of a linear operator that is self-adjoint with respect to an arbitrary inner product. In particular, the matrix of the linear operator need not be symmetric.

Theorem 2.8 (Courant-Fischer Theorem). *Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear operator that is self-adjoint with respect to some inner product $\langle \cdot, \cdot \rangle$ (not necessarily Euclidean). If $\lambda_n \leq \dots \leq \lambda_1$ are the eigenvalues of T , then*

$$\lambda_k = \min_U \max_v \langle v, Tv \rangle,$$

where the minimum is taken over all $(n - k)$ -dimensional subspaces $U \subseteq \mathbb{R}^n$ and the maximum is taken over all vectors $v \in U$ with $\langle v, v \rangle = 1$.

When the inner product $\langle \cdot, \cdot \rangle$ is clear, we call a matrix A self-adjoint when $\langle u, Av \rangle = \langle Au, v \rangle$, for all u, v . Similarly we call a self-adjoint A positive semidefinite when for all v

$$\langle v, Av \rangle \geq 0.$$

By [Theorem 2.8](#), this is equivalent to A having nonnegative eigenvalues.

2.2 Markov Chains and Random Walks

For this paper, we consider a Markov chain as a triple (Ω, P, π) where Ω denotes the (finite) state space, $P \in \mathbb{R}_{\geq 0}^{\Omega \times \Omega}$ denotes the transition probability matrix and $\pi \in \mathbb{R}_{\geq 0}^\Omega$ denotes a stationary distribution of the chain (which will be unique for all chains we consider). For $\tau, \sigma \in \Omega$, we use $P(\tau, \sigma)$ to denote the corresponding entry of P , which is the probability of moving from τ to σ . We say a Markov chain is ϵ -lazy if for any state $\tau \in \Omega$, $P(\tau, \tau) \geq \epsilon$. A chain (Ω, P, π) is *reversible* if there is a nonzero nonnegative function $f : \Omega \rightarrow \mathbb{R}_{\geq 0}$ such that for any pair of states $\tau, \sigma \in \Omega$,

$$f(\tau)P(\tau, \sigma) = f(\sigma)P(\sigma, \tau).$$

If this condition is satisfied, then f is proportional to a stationary distribution of the chain. In this paper we only work with reversible Markov chains. Note that being reversible means that the transition matrix P is self-adjoint w.r.t. the following $\langle \cdot, \cdot \rangle$ defined for $\phi, \psi \in \mathbb{R}^\Omega$:

$$\langle \phi, \psi \rangle = \sum_{x \in \Omega} f(x) \phi(x) \psi(x).$$

Reversible Markov chains can be realized as random walks on weighted graphs. Given a weighted graph $G = (V, E, w)$ where every edge $e \in E$ has weight $w(e)$, the non-lazy *simple* random walk on G is the Markov chain that from any vertex $u \in V$ chooses an edge $e = \{u, v\}$ with probability proportional to $w(e)$ and jumps to v . We can make this walk ϵ -lazy by staying at every vertex with probability ϵ . It turns out that if G is connected, then the walk has a unique stationary distribution where $\pi(u) \propto w(u)$, where $w(u) = \sum_{v \sim u} w(\{u, v\})$ is the weighted degree of u .

For any reversible Markov chain (Ω, P, π) , the largest eigenvalue of P is 1. We let $\lambda^*(P)$ denote the second largest eigenvalue of P in absolute value. That is, if $-1 \leq \lambda_n \leq \dots \leq \lambda_1 = 1$ are the eigenvalues of P , then $\lambda^*(P)$ equals $\max\{|\lambda_2|, |\lambda_n|\}$.

Theorem 2.9 ([DS91, Prop 3]). *For any reversible irreducible Markov chain (Ω, P, π) , $\epsilon > 0$, and any starting state $\tau \in \Omega$,*

$$t_\tau(\epsilon) \leq \frac{1}{1 - \lambda^*(P)} \cdot \log\left(\frac{1}{\epsilon \cdot \pi(\tau)}\right).$$

For our results, it will be enough to look at the second largest eigenvalue $\lambda_2(P)$, which we can bound using the conductance of a weighted graph. Consider a weighted graph $G = (V, E, w)$ and a subset $S \subseteq V$ of vertices. We let \bar{S} denote the complement $V \setminus S$. Then the *conductance* of S , denoted by $\text{cond}(S)$, is defined as

$$\text{cond}(S) = \frac{w(E(S, \bar{S}))}{\text{vol}(S)} = \frac{\sum_{e \in E(S, \bar{S})} w(e)}{\sum_{v \in S} w(v)},$$

where $E(S, \bar{S}) = \{\{u, v\} \in E \mid u \in S, v \notin S\}$ is the set of edges between S and \bar{S} , $w(E(S, \bar{S}))$ is the sum of weights of these edges, and the volume $\text{vol}(S)$ is the sum of the weighted degrees of the vertices in S . The conductance of G is then

$$\text{cond}(G) = \min_S \text{cond}(S),$$

where the minimum is taken over subsets $\emptyset \subsetneq S \subsetneq V$ for which $w(S) \leq w(\bar{S})$.

We say G is d -regular if $w(v) = d$ for all $v \in V$.

Theorem 2.10 (Cheeger's Inequalities [AM85; Alo86]). *For any d -regular weighted graph $G = (V, E, w)$,*

$$\frac{d - \lambda_2(A_G)}{2} \leq \text{cond}(G) \leq \sqrt{2(d - \lambda_2(A_G))},$$

where A_G is the weighted adjacency matrix of G given by $(A_G)_{ij} = w(\{i, j\})$.

A direct consequence of the above theorem is that if the (weighted) graph G is connected, i.e., for all proper nonempty subsets S of vertices, $w(E(S, \bar{S})) > 0$, then $\lambda_2(A_G) < d$. If the matrix A_G is stochastic, then the graph is 1-regular, which gives the following.

Corollary 2.11. *If $A \in \mathbb{R}^{n \times n}$ is a stochastic matrix corresponding to a reversible Markov chain with the property that $\sum_{\{i, j\} \cap S = \emptyset} A_{ij} > 0$ for all subsets $\emptyset \subsetneq S \subsetneq [n]$, then $\lambda_2(A) < 1$.*

2.3 Matroids

A *matroid* $M = ([n], \mathcal{I})$ is a combinatorial structure consisting of a ground set $[n]$ of elements and a nonempty collection \mathcal{I} of *independent* subsets of $[n]$ satisfying:

- i) If $S \subseteq T$ and $T \in \mathcal{I}$, then $S \in \mathcal{I}$ (hereditary property).
- ii) If $S, T \in \mathcal{I}$ and $|T| > |S|$, then there exists an element $i \in T \setminus S$ such that $S \cup \{i\} \in \mathcal{I}$ (exchange axiom).

The *rank*, denoted by $\text{rank}(S)$, of a subset $S \subset [n]$ is the size of any maximal independent set of M contained in S . Thus, the independent sets of M are precisely those subsets $S \subset [n]$ for which $\text{rank}(S) = |S|$. We call $\text{rank}([n])$ the rank of M , and if M has rank r , any set $S \in \mathcal{I}$ of size r is called a *basis* of M .

An element $i \in [n]$ is a *loop* if $\{i\} \notin \mathcal{I}$, that is, $\{i\}$ is dependent. Two non-loops $i, j \in [n]$ are *parallel* if $\{i, j\} \notin \mathcal{I}$, that is, $\{i, j\}$ is dependent.

Definition 2.12 (Contraction). Let $M = ([n], \mathcal{I})$ be a matroid and $S \in \mathcal{I}$. Then the *contraction* M/S is the matroid with ground set $[n] \setminus S$ and independent sets $\{T \subseteq [n] \setminus S \mid T \cup S \in \mathcal{I}\}$.

We will use a key property of matroids called the *matroid partition property*. For any matroid $M = ([n], \mathcal{I})$, the non-loops of M can be partitioned into sets S_1, S_2, \dots, S_k for some $1 \leq k \leq n$ with the property that non-loops $j, k \in [n]$ are parallel if and only if they belong to the same set S_i . Indeed, one can check from the axioms for a matroid that being parallel defines an equivalence relation on the non-loop elements of $[n]$ and S_1, \dots, S_k are then the corresponding equivalence classes.

2.4 Simplicial Complexes

A *simplicial complex* X on the ground set $[n]$ is a nonempty collection of subsets of $[n]$ that is downward closed, namely if $\tau \subset \sigma$ and $\sigma \in X$, then $\tau \in X$. The elements of X are called *faces/simplices*, and the *dimension* of a face $\tau \in X$ is defined as $\dim(\tau) = |\tau|$. Note that for convenience and clarity of notation, our definition deviates from the standard definition of $\dim(\tau) = |\tau| - 1$ used by topologists.

The empty set \emptyset has dimension-0. A face of dimension 1 is a *vertex* of X and a face of dimension 2 is called an *edge*. More generally, we write

$$X(k) = \{\tau \in X \mid \dim(\tau) = k\}$$

for the collection of dimension- k faces, or *k-faces/k-simplices*, of X . The dimension of X is the largest k for which $X(k)$ is nonempty, and we say that X is *pure* of dimension d if all maximal faces of X have the dimension d . In this paper we will only consider pure simplicial complexes.

The *link* of a face $\tau \in X$ denoted by X_τ is the simplicial complex on $[n] \setminus \tau$ obtained by taking all faces in X that contain τ and removing τ from them,

$$X_\tau = \{\sigma \setminus \tau \mid \sigma \in X, \sigma \supset \tau\}.$$

Note that if X is pure of dimension d and $\tau \in X(k)$, then X_τ is pure and has dimension $(d - k)$.

For any matroid $M = ([n], \mathcal{I})$ of rank r , the independent sets \mathcal{I} form a pure r -dimensional simplicial complex on $[n]$ called its *independence (or matroid) complex*. Furthermore, for any $S \in \mathcal{I}$, the link \mathcal{I}_S of the independence complex consists precisely of the independent sets of the contraction M/S . There are many other beautiful simplicial complexes associated to matroids, but here we will be mainly interested in the independence complex.

We can equip a simplicial complex with a weight function: $w : X \rightarrow \mathbb{R}_{>0}$ which assigns a positive weight to each face of X . We say that w is *balanced* if for every non-maximal face $\tau \in X$ of dimension k ,

$$w(\tau) = \sum_{\sigma \in X(k+1): \sigma \supset \tau} w(\sigma) \quad (1)$$

For a pure simplicial complex X we can define a balanced weight function by assigning arbitrary positive weights to maximal faces and defining the weight of each lower dimensional face recursively. Indeed, if X is a pure simplicial complex of dimension d and w is a balanced weight function, then, for any $\tau \in X(k)$,

$$w(\tau) = (d - k)! \sum_{\sigma \in X(d): \sigma \supset \tau} w(\sigma).$$

One natural choice is the function which assigns a weight of one to each maximal face, but there are many other interesting choices.

Any balanced weight function on X induces a weighted graph on the vertices of X as follows. The *1-skeleton* of X is the graph on vertices $X(1)$ with edges $X(2)$. Then, restricting w to $X(1)$ and $X(2)$ determines a weighted graph, where $w(v)$ gives the weighted degree of each $v \in X(1)$. The weighted graphs coming from both X and its links X_τ will be useful in later sections.

2.5 Log-Concave Polynomials

We say a polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ is d -homogeneous if every monomial of p has degree d ; equivalently, p is d -homogeneous if $p(\lambda x_1, \dots, \lambda x_n) = \lambda^d p(x_1, \dots, x_n)$ for every $\lambda \in \mathbb{R}$. For a d -homogeneous polynomial p , the following identity, known as Euler's identity, holds:

$$d \cdot p(x) = \sum_{k=1}^n x_k \partial_k p(x). \quad (2)$$

Note that if p is homogeneous then all directional derivatives of p are also homogeneous, so one can apply this to $\partial_i p(x)$ and $\partial_i \partial_j p(x)$ to find that

$$(d - 1) \cdot (\nabla p) = \sum_{k=1}^n x_k \cdot (\nabla(\partial_k p)) = (\nabla^2 p) \cdot x \quad \text{and} \quad (d - 2) \cdot (\nabla^2 p) = \sum_{k=1}^n x_k \cdot (\nabla^2(\partial_k p)). \quad (3)$$

A polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ with nonnegative coefficients is log-concave if $\log p$ is a concave function over $\mathbb{R}_{>0}^n$. For simplicity we also consider the zero polynomial to be log-concave. Equivalently, p is log-concave if the Hessian of $\log p$

$$\nabla^2 \log p = \frac{p \cdot (\nabla^2 p) - (\nabla p)(\nabla p)^\top}{p^2}$$

is negative semi-definite at any point $x \in \mathbb{R}_{>0}^n$, where ∇p is the gradient of p . Since $(\nabla p)(\nabla p)^\top$ is a rank-1 matrix, by Cauchy's interlacing theorem, $p \cdot (\nabla^2 p)$ has at most one positive eigenvalue at any $x \in \mathbb{R}_{>0}^n$. Since p has nonnegative coefficients and x has strictly positive entries, $p(x) > 0$ so $\nabla^2 \log p$ being negative semidefinite is equivalent to $\nabla^2 p \preceq \frac{(\nabla p)(\nabla p)^\top}{p}$, where the right-hand side is a rank-1 positive semidefinite matrix. In particular, $\nabla^2 p$ has at most one positive eigenvalue at x . In [AOV18] it is shown that for homogeneous polynomials p , the converse of this is also true, i.e., if $\nabla^2 p$ has at most one positive eigenvalue at all $x > 0$ then p is log-concave.

Proposition 2.13 ([AOV18]). *A degree- d homogeneous polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ with nonnegative coefficients is log-concave over $\mathbb{R}_{>0}^n$ iff $(\nabla^2 p)(x)$ has at most one positive eigenvalue at all $x \in \mathbb{R}_{>0}^n$.*

Remark 2.14. Whenever p has degree at least 2 and nonnegative coefficients, $(\nabla^2 p)(x)$ has at least one positive entry for any $x \in \mathbb{R}_{>0}^n$. Hence, one can see (via, for example, the variational characterization of eigenvalues [Theorem 2.8](#) and using a test vector with positive entries) that $(\nabla^2 p)(x)$ must have at least one strictly positive eigenvalue. Thus, whenever we write “at most one positive eigenvalue”, we also mean it has “exactly one positive eigenvalue”.

We say a polynomial $p \in \mathbb{R}[x_1, \dots, x_n]$ is *decomposable* if it can be written as a sum of polynomials in disjoint subset of the variables, that is, if there exists a nonempty subset $I \subsetneq [n]$ and nonzero polynomials $g \in \mathbb{R}[x_i : i \in I]$, $h \in \mathbb{R}[x_i : i \notin I]$ for which $f = g + h$. We call f *indecomposable* otherwise.

Lemma 2.15. *If $p \in \mathbb{R}[x_1, \dots, x_n]$ has nonnegative coefficients, is homogeneous of degree at least 2, and log-concave at $\mathbb{1}$, then p is indecomposable.*

Proof. Suppose that p has nonnegative coefficients, is homogeneous of degree ≥ 2 , and is decomposable, with decomposition $p = g + h$ where $g \in \mathbb{R}[x_i : i \in I]$ and $h \in \mathbb{R}[x_i : i \notin I]$. Both g and h are restrictions of p obtained by setting some variables equal to zero, therefore both g and h are log-concave. Then, at $\mathbb{1}$, the Hessians of g and h each have precisely one positive eigenvalue. However, the Hessian of p at this point is a block diagonal matrix with these two blocks, $\nabla^2 g, \nabla^2 h$,

$$\nabla^2 p = \begin{bmatrix} \nabla^2 g & 0 \\ 0 & \nabla^2 h \end{bmatrix}.$$

So, p has exactly two positive eigenvalues, meaning that p is not log-concave, a contradiction. \square

In order to prove several distributions of interest are strongly log-concave, we will prove an equivalent characterization of strongly log-concave polynomials.

Theorem 2.16. *Let $p \in \mathbb{R}[x_1, \dots, x_n]$ be a d -homogeneous polynomial such that:*

1. *for any $0 \leq k \leq d - 2$ and any $(i_1, \dots, i_k) \in [n]^k$, $\partial_{i_1} \cdots \partial_{i_k} p$ is indecomposable, and*
2. *for any $(i_1, \dots, i_{d-2}) \in [n]^{d-2}$, the quadratic $\partial_{i_1} \cdots \partial_{i_{d-2}} p$ is either identically zero, or log-concave at $\mathbb{1}$.*

Then p is strongly log-concave at $\mathbb{1}$.

In [Lemma 2.15](#) we show that the condition that all partial derivatives are indecomposable is necessary for a polynomial to be (strongly) log-concave.

3 Walks on Simplicial Complexes

Consider a pure d -dimensional complex X with a balanced weight function $w : X \rightarrow \mathbb{R}_{>0}$. We will call (X, w) a *weighted complex*. For $1 \leq k < d$, we define a random walk on $X(k)$ known as the *upper k -walk* based on movement from an face in $X(k)$ to a higher-dimensional face and then returning to $X(k)$. Similarly, for $1 \leq k < d$, we define a random walk on $X(k+1)$ known as the *lower k -walk* based on movement from a face in $X(k+1)$ to a lower-dimensional face and then returning to $X(k+1)$.

To define these walks we construct a bipartite graph G_k with one side corresponding to $X(k)$ and the other side corresponding to $X(k+1)$. We connect $\tau \in X(k)$ to $\sigma \in X(k+1)$ with an edge of weight $w(\sigma)$ iff $\tau \subset \sigma$. Now, consider the simple (weighted) random walk on G_k . Given a vertex we choose a neighbor proportional to the weight of the edge connecting the two vertices.

This is a walk on a bipartite graph and is naturally periodic. We can consider the odd steps and even steps, in order to obtain two random walks; one on $X(k)$ called P_k^\wedge , and the other on $X(k+1)$ called P_{k+1}^\vee , where given $\tau \in X(k)$ you take two steps of the walk in G_k to transition to the next k -face with respect to the P_k^\wedge matrix, and similarly, you take two steps in G_k from $\sigma \in X(k+1)$ to transition with respect to P_{k+1}^\vee .

Now, let us formally write down the entries of P_k^\wedge and P_{k+1}^\vee . Given a simplex $\tau \in X(k)$, first among all $k+1$ dimensional simplices $\sigma \in X(k+1)$ that contain τ we choose one proportional to $w(\sigma)$. Then, we delete one of the $|\sigma| = k+1$ elements of σ uniformly at random to obtain a new state τ' . It follows that the probability of transition to τ' is equal to the probability of choosing $\sigma = \tau \cup \tau'$ in the first step, which is equal to $\frac{w(\tau \cup \tau')}{w(\tau)}$ since w is balanced, times the probability of choosing τ' conditioned on $\sigma = \tau \cup \tau'$, which is $\frac{1}{k+1}$. In summary, for $1 \leq k < d$,

$$P_k^\wedge(\tau, \tau') = \begin{cases} \frac{1}{k+1}, & \text{if } \tau = \tau' \\ \frac{w(\tau \cup \tau')}{(k+1)w(\tau)}, & \text{if } \tau \cup \tau' \in X(k+1) \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

Note that upper walk is not defined for $k = d$, because there is no $(d+1)$ -dimensional simplex in X .

Analogously, given $\sigma \in X(k+1)$, first we remove a uniformly random element of σ to obtain τ . Then, among all all $k+1$ simplices $\sigma' \in X(k+1)$ that contain τ we choose one proportional to $w(\sigma')$. It follows that for $1 \leq k < d$,

$$P_{k+1}^\vee(\sigma, \sigma') = \begin{cases} \frac{\sum_{\tau \in X(k): \tau \subset \sigma} \frac{w(\sigma)}{(k+1)w(\tau)}}{\frac{w(\sigma')}{(k+1)w(\sigma \cap \sigma')}}}, & \text{if } \sigma = \sigma' \\ \frac{w(\sigma')}{(k+1)w(\sigma \cap \sigma')}, & \text{if } \sigma \cap \sigma' \in X(k) \\ 0, & \text{otherwise} \end{cases} \quad (5)$$

Observe the corresponding random walks are reversible with respect to the weight function w , i.e., for all $\tau, \tau' \in X(k)$, we have

$$w(\tau)P_k^\wedge(\tau, \tau') = w(\tau')P_k^\wedge(\tau', \tau) \quad w(\tau)P_{k+1}^\vee(\tau, \tau') = w(\tau')P_{k+1}^\vee(\tau', \tau).$$

This implies that both chains have the same stationarity distribution where the probability of $\tau \in X(k)$ is proportional to $w(\tau)$.

Lemma 3.1. For any $1 \leq k < d$, P_k^\wedge and P_{k+1}^\vee are stochastic, self-adjoint w.r.t. the w -induced inner product, PSD, and have the same (with multiplicity) non-zero eigenvalues.

Proof. Let P_k be the transition probability matrix of the simple random walk on G_k . Since G_k is bipartite and we can write

$$P_k = \begin{bmatrix} 0 & P_k^\downarrow \\ P_k^\uparrow & 0 \end{bmatrix}$$

where $P_k^\downarrow \in \mathbb{R}^{X(k+1) \times X(k)}$ and $P_k^\uparrow \in \mathbb{R}^{X(k) \times X(k+1)}$ are stochastic matrices. Note that P_k is self-adjoint w.r.t. the weight-induced inner product given by weights of the stationary distribution. It follows that P_k is self-adjoint w.r.t. the inner product

$$\langle \phi, \psi \rangle = \sum_{\tau \in X(k)} w(\tau) \phi(\tau) \psi(\tau) + (k+1) \sum_{\sigma \in X(k+1)} w(\sigma) \phi(\sigma) \psi(\sigma)$$

Observe that

$$P_k^2 = \begin{bmatrix} P_k^\downarrow P_k^\uparrow & 0 \\ 0 & P_k^\uparrow P_k^\downarrow \end{bmatrix}$$

So, in particular, P_k^2 is PSD and stochastic. Since P_k^\wedge and P_{k+1}^\vee correspond to two step walks on G_k , indeed we can write

$$\begin{aligned} P_k^\wedge &= P_k^\uparrow P_k^\downarrow \\ P_{k+1}^\vee &= P_k^\downarrow P_k^\uparrow \end{aligned}$$

It follows that both matrices are self-adjoint w.r.t. the w -induced inner product, are PSD, and stochastic, and by [Fact 2.5](#) they have the same eigenvalues. \square

Let us specifically study P_1^\wedge . Observe that P_1^\wedge is the transition probability matrix of the simple $(1/2)$ -lazy random walk on the weighted 1-skeleton of X where the weight of each edge $e \in X(2)$ is $w(e)$. We also need to consider the non-lazy variant of this random walk, given by the transition matrix

$$\tilde{P}_1^\wedge = 2(P_1^\wedge - I/2).$$

Similarly, for any face $\tau \in X(k)$, we define the upper random walk on the faces of the link X_τ . Specifically, let $P_{\tau,1}^\wedge$ denote the transition matrix of the upper walk, as above, on the 1-dimensional faces of X_τ , and

$$\tilde{P}_{\tau,1}^\wedge = 2(P_{\tau,1}^\wedge - I/2)$$

be the transition matrix for the non-lazy version.

Definition 3.2 (Local Spectral Expanders, [\[KO18\]](#)). For $\lambda \geq 0$, a pure d -dimensional weighted complex (X, w) is a λ -local-spectral-expander if for every $0 \leq k < d-1$, and for every $\tau \in X(k)$, we have $\lambda_2(\tilde{P}_{\tau,1}^\wedge) \leq \lambda$.

In other words, X is λ -local spectral expander if the spectral gap of the natural random walk on the 1-skeleton of the link of all simplices of X has a spectral gap of at least $1 - \lambda$. In this section we give a somewhat simpler proof of the following special case of the main theorem of [\[KO18\]](#).

Theorem 3.3 ([KÖ18]). *Let (X, w) be a pure d -dimensional weighted 0-local spectral expander and let $0 \leq k < d$. Then, for all $-1 \leq i \leq k$, P_k^\wedge has at most $|X(i)| \leq \binom{n}{i}$ eigenvalues of value $> 1 - \frac{i+1}{k+1}$, where for convenience, we set $X(-1) = \emptyset$ and $\binom{n}{-1} = 0$. In particular, the second largest eigenvalue of P_k^\wedge is at most $\frac{k}{k+1}$.*

Remark 3.4. In other words, P_k^\wedge has very few “big” eigenvalues. For example, P_k^\wedge has exactly one eigenvalue strictly larger than $\frac{k}{k+1}$ corresponding to the maximum eigenvalue (which has value 1) and at most $n = |X(1)|$ eigenvalues strictly larger than $\frac{k-1}{k+1}$. Hence, P_k^\wedge has at most $n - 1$ eigenvalues between $\frac{k-1}{k+1}$ and $\frac{k}{k+1}$. Note that the significance of this theorem is that we are able to establish an estimate on *all* eigenvalues of P_k^\wedge .

For the proof, we will need the following lemma. We remark that the inner product on the space $\mathbb{R}^{X(k)}$ is given by $\langle \phi, \psi \rangle = \sum_{\tau \in X(k)} w(\tau) \phi(\tau) \psi(\tau)$, and that being self-adjoint, PSD, and the Loewner order are defined w.r.t. this inner product.

Lemma 3.5. $P_k^\wedge \preceq \frac{k}{k+1} P_k^\vee + \frac{1}{k+1} I$ for all $0 \leq k < d$.

Proof. For convenience, let $M = P_k^\wedge - \left(\frac{k}{k+1} P_k^\vee + \frac{1}{k+1} I \right)$. Fix $\eta \in X(k-1)$. We will first consider submatrices M_η whose entries are given by the following:

$$M_\eta(\tau, \sigma) = \begin{cases} M(\tau, \sigma), & \text{if } \tau \neq \sigma, \eta = \tau \cap \sigma \\ -\frac{1}{k+1} \cdot \frac{w(\tau)}{w(\eta)}, & \text{if } \tau = \sigma, \tau \supset \eta \\ 0, & \text{otherwise} \end{cases}$$

Note that $M = \sum_{\eta \in X(k-1)} M_\eta$ and hence, it suffices to prove that $M_\eta \preceq 0$ for every $\eta \in X(k-1)$.

Fix $\eta \in X(k-1)$. Let $\tau, \sigma \in X(k)$ with $\tau \neq \sigma$ and $\tau \cap \sigma = \eta$. Then

$$M_{\tau \cap \sigma}(\tau, \sigma) = M(\tau, \sigma) = \frac{1}{k+1} \left(\frac{w(\tau \cup \sigma)}{w(\tau)} - \frac{w(\sigma)}{w(\tau \cap \sigma)} \right) = \frac{1}{k+1} \left(\frac{w(\tau \cup \sigma)w(\tau \cap \sigma) - w(\tau)w(\sigma)}{w(\tau)w(\tau \cap \sigma)} \right)$$

Furthermore, by definition, if $\tau \in X(k)$ with $\tau \supset \eta$, then $M_\eta(\tau, \tau) = -\frac{1}{k+1} \cdot \frac{w(\tau)}{w(\eta)}$. A matrix calculation reveals that

$$M_\eta = \frac{1}{(k+1)w(\eta)} \text{diag}(w_\eta)^{-1} \cdot (w(\eta) \cdot A_\eta - w_\eta w_\eta^\top)$$

where w_η is the $|X(k)|$ -dimensional vector whose non-zero entries are $w(\tau)$ for $\tau \supset \eta$, and A_η is the $|X(k)| \times |X(k)|$ matrix whose non-zero entries are $w(\tau \cup \sigma)$ for $\tau, \sigma \in X(k)$ satisfying $\tau \cup \sigma \in X(k+1)$ and $\tau \cap \sigma = \eta$. Note that M_η is NSD w.r.t. the inner product defined by w , if and only if $\text{diag}(w_\eta) M_\eta$ is NSD in the usual sense, because for any v

$$\langle v, M_\eta v \rangle = v^\top \text{diag}(w_k) M_\eta v = v^\top \text{diag}(w_\eta) M_\eta v,$$

where w_k is the vector of w values on $X(k)$ and for the last equality we used that w_k is the same as w_η on all $\tau \supset \eta$.

Thus, it suffices to prove that $A_\eta \preceq \frac{w_\eta w_\eta^\top}{w(\eta)}$. We view A_η as the weighted adjacency matrix of the 1-skeleton (which we recall is a graph) of the link X_η . Then $\tilde{P}_{\eta,1}^\wedge = \frac{1}{k+1} \text{diag}(w_\eta)^{-1} A_\eta$ gives its

non-lazy simple random walk matrix. As (X, w) is a 0-local spectral expander, $\tilde{P}_{\eta,1}^\wedge$ has at most one positive eigenvalue, whence $A_\eta = (k+1) \text{diag}(w_\eta) \cdot \tilde{P}_{\eta,1}^\wedge$ has at most one positive eigenvalue by [Lemma 2.6](#).

Finally, observe that the weights being balanced enforces that $w(\tau) = \sum_{\sigma \in X(k): \tau \cup \sigma \in X(k+1)} w(\tau \cup \sigma)$ and $w(\eta) = \sum_{\tau \in X(k): \tau \supset \eta} w(\tau)$. That $A_\eta \preceq \frac{w_\eta w_\eta^\top}{w(\eta)}$ then follows immediately by [Lemma 2.7](#). \square

Proof of [Theorem 3.3](#). We go by induction on k . The case $k = 0$ is trivial, as P_0^\wedge is 1×1 . When $k = 1$, we have $P_1^\wedge = \frac{1}{2}(\tilde{P}_1^\wedge + I)$. As (X, w) is a 0-local spectral expander, \tilde{P}_1^\wedge has exactly one positive eigenvalue, with value 1. Hence, P_1^\wedge has eigenvalue 1 with multiplicity 1. All other eigenvalues of P_1^\wedge are less than or equal to $1/2$, of which, there are $|X(1)| - 1$ many. Thus, the base case holds.

Assume the claim holds for some $d-1 > k \geq 0$. Recall by [Lemma 3.1](#), P_{k+1}^\vee has the same non-zero eigenvalues as P_k^\wedge . By [Lemma 3.5](#),

$$P_{k+1}^\wedge \preceq \frac{k+1}{k+2} P_{k+1}^\vee + \frac{1}{k+2} I$$

For $-1 \leq i \leq k$, P_k^\vee has at most $|X(i)|$ eigenvalues $> 1 - \frac{i+1}{k+1}$. Hence, P_{k+1}^\wedge has at most $|X(i)|$ eigenvalues $> \frac{k+1}{k+2} \cdot (1 - \frac{i+1}{k+1}) + \frac{1}{k+2} = 1 - \frac{i+1}{k+2}$. For $i = k+1$, we trivially have that P_{k+1}^\wedge has at most $|X(k+1)|$ eigenvalues > 0 , as P_{k+1}^\wedge is $|X(k+1)| \times |X(k+1)|$. \square

4 From Strongly Log-Concave Polynomials to Local Spectral Expanders

In this section, we prove [Theorem 1.1](#). Let $p = \sum_S c_S x^S \in \mathbb{R}[x_1, \dots, x_n]$ be a d -homogeneous multiaffine strongly log-concave polynomial with nonnegative coefficients. We can construct a pure d -dimensional complex X^p from p as follows: For every term $c_S x^S$ of p we include the d -dimensional simplex S with weight

$$w(S) = c_S. \tag{6}$$

Note that since p has nonnegative coefficients, the above weight function is nonnegative. We turn X^p into a simplicial complex by including all subsets of the d -dimensional simplices and weighting each lower dimensional simplex inductively according to [Eq. \(1\)](#).

Proposition 4.1. *Let $p \in \mathbb{R}[x_1, \dots, x_n]$ be a multiaffine homogeneous polynomial with nonnegative coefficients. If p is strongly log-concave then (X^p, w) is a 0-local-spectral-expander, where $w(S) = c_S$ for every maximal face $S \in X^p$.*

The converse of the above statement also holds true and we will discuss it in the next section. We build up to the proof of [Proposition 4.1](#) and thereby the proof of [Theorem 1.1](#).

We now fix a multiaffine d -homogeneous strongly log-concave polynomial p with nonnegative coefficients. Fix a simplex $\tau \in X^p(k)$ and let $p_\tau = (\prod_{i \in \tau} \partial_i) p$. Note that p_τ is $(d-k)$ -homogeneous.

Lemma 4.2. *For any $0 \leq k \leq d$, and any simplex $\tau \in X^p(k)$, $w(\tau) = (d-k)! \cdot p_\tau(\mathbb{1})$.*

Proof. We prove this by induction on $d-k$. If $\dim(\tau) = d$ then $p_\tau = c_\tau$ and the statement follows immediately from [Eq. \(6\)](#). So, suppose the statement holds for all simplices $\sigma \in X^p(k+1)$ and fix a simplex $\tau \in X^p(k)$. Then by definition,

$$w(\tau) = \sum_{\sigma \in X^p(k+1): \sigma \supset \tau} w(\sigma) = \sum_{i \in X_\tau^p(1)} w(\tau \cup i).$$

Using the inductive hypothesis and the fact that $\partial_i p_\tau = 0$ for $i \notin X_\tau^p(1)$, we then find that

$$w(\tau) = (d-k-1)! \sum_{i \in X_\tau^p(1)} p_{\tau \cup \{i\}}(\mathbf{1}) = (d-k-1)! \sum_{i=1}^n \partial_i p_\tau(\mathbf{1}) = (d-k)! \cdot p_\tau(\mathbf{1}),$$

where the last equality follows from Euler's identity. \square

Recall that $\tilde{P}_{\tau,1}^\wedge$ is the transition probability matrix of the non-lazy random walk on the 1-skeleton of the link X_τ^p . To prove the [Proposition 4.1](#) it is enough to show that $\lambda_2(\tilde{P}_{\tau,1}^\wedge) \leq 0$, i.e. that $\tilde{P}_{\tau,1}^\wedge$ has at most one positive eigenvalue.

Proof of [Proposition 4.1](#). Since p is strongly log-concave, $\nabla^2 p_\tau(\mathbf{1})$ has at most one positive eigenvalue. Let

$$\tilde{\nabla}^2 p_\tau = \frac{1}{d-k-1} \text{diag}(\nabla p_\tau(\mathbf{1}))^{-1} \nabla^2 p_\tau(\mathbf{1}). \quad (7)$$

We claim that

$$\tilde{\nabla}^2 p_\tau = \tilde{P}_{\tau,1}^\wedge. \quad (8)$$

To see this, note that by [Eq. \(4\)](#), for $i, j \in X_\tau^p$,

$$\tilde{P}_{\tau,1}^\wedge(i, j) = \frac{w_\tau(\{i, j\})}{w_\tau(\{i\})} = \frac{w(\tau \cup \{i, j\})}{w(\tau \cup \{i\})}.$$

On the other hand, by [Eq. \(7\)](#),

$$(\tilde{\nabla}^2 p_\tau)(i, j) = \frac{(\partial_i \partial_j p_\tau)(\mathbf{1})}{(d-k-1) \cdot (\partial_i p_\tau)(\mathbf{1})}.$$

The above two are equal by [Lemma 4.2](#), which proves [Eq. \(8\)](#).

Since p has nonnegative coefficients, the vector $\nabla p_\tau(\mathbf{1})$ has nonnegative entries, which implies $\text{diag}(\nabla p_\tau(\mathbf{1})) \succcurlyeq 0$. Since $\nabla^2 p_\tau(\mathbf{1})$ has at most one positive eigenvalue, $\tilde{\nabla}^2 p_\tau(\mathbf{1})$ has at most one positive eigenvalue by [Lemma 2.6](#) as desired. Therefore by [Eq. \(8\)](#), $\tilde{\nabla}^2 p_\tau = \tilde{P}_{\tau,1}^\wedge$ has at most one positive eigenvalue and $\lambda_2(\tilde{P}_{\tau,1}^\wedge) \leq 0$. \square

With this we are ready to prove our main theorem.

Proof of [Theorem 1.1](#). Let μ be a d -homogeneous strongly log-concave distribution and let P_μ be the transition probability matrix of the chain \mathcal{M}_μ . By [Theorem 2.9](#) it is enough to show that $\lambda^*(P_\mu) \leq 1 - 1/d$. Observe that the chain \mathcal{M}_μ is exactly the same as the chain P_d^\vee for the simplicial complex $X^{\mathcal{G}_\mu}$ defined above. Therefore, $\lambda^*(P_\mu) = \lambda^*(P_d^\vee) = \lambda^*(P_{d-1}^\wedge)$, where the last equality follows by [Lemma 3.1](#). Since g_μ is strongly log-concave, by [Proposition 4.1](#), $X^{\mathcal{G}_\mu}$ is 0-local-spectra-expander. Therefore, by [Theorem 3.3](#),

$$\lambda^*(P_{d-1}^\wedge) \leq 1 - \frac{1}{(d-1)+1} = 1 - \frac{1}{d},$$

as desired. \square

4.1 Bases Exchange Walk

In this part we prove [Theorem 1.5](#). Fix a rank r matroid $M = ([n], \mathcal{I})$, let μ denote the uniform distribution on the bases of M , and consider the simplicial complex $X^{\mathcal{g}_\mu}$. As before, recall that P_r^\vee is the transition matrix for the Markov chain \mathcal{M}_μ defined in the introduction. As discussed in [Section 2.2](#), each reversible Markov chain is equivalent to a random walk in a (weighted) undirected graph. Let H_M be the graph corresponding to P_r^\vee . Then the vertices of H_M correspond to bases of M , and the weight of an edge between two bases τ, τ' is,

$$P_r^\vee(\tau, \tau') = P_r^\vee(\tau', \tau).$$

The symmetry of P_r^\vee follows by the fact that P_r^\vee is reversible and $w(\tau) = w(\tau') = 1$. Observe that P_r^\vee is the adjacency matrix of H_M . Furthermore, note that by [Theorem 5.2](#) (proof in [Section 5](#)), \mathcal{g}_μ is a strongly log-concave polynomial and hence, $X^{\mathcal{g}_\mu}$ is a 0-local spectral expander by [Proposition 4.1](#). By [Theorem 3.3](#), $\lambda_2(P_r^\vee) \leq 1 - \frac{1}{r}$. Note that the weighted degree of each vertex of H_M is 1 since P_r^\vee is a stochastic matrix. Therefore, by Cheeger's inequality, [Theorem 2.10](#),

$$\text{cond}(H_M) \geq \frac{1 - \lambda_2(P_r^\vee)}{2} \geq \frac{1 - (1 - 1/r)}{2} = \frac{1}{2r}. \quad (9)$$

Let $G_M = (\mathcal{B}, E)$ be the (unweighted) bases exchange graph associated to M as defined in the introduction. It follows that G_M is the unweighted base graph of H_M . Fix a nonempty set $S \subset \mathcal{B}$ of bases such that $|S| \leq |\mathcal{B}|/2$. We need to show that the expansion $h(S)$ is ≥ 1 . Note that since the weighted degree of each vertex of H_M is 1, $\text{vol}_H(S) = |S|$. It follows that

$$\text{cond}(H_M) \leq \text{cond}(S) = \frac{\sum_{\tau, \tau': \tau \in S, \tau' \notin S} P_r^\vee(\tau, \tau')}{|S|} \leq \frac{\sum_{\tau, \tau': \tau \in S, \tau' \notin S} \frac{1}{2r}}{|S|} = \frac{\frac{1}{2r} |E(S, \bar{S})|}{|S|} = \frac{h(S)}{2r}$$

The second inequality follows by the following fact. Putting the above together with [Eq. \(9\)](#), implies $h(S) \geq 1$ as desired. This completes the proof of [Theorem 1.5](#).

Fact 4.3. For any pair of bases, $\tau, \tau' \in X^{\mathcal{g}_\mu}(r)$,

$$P_r^\vee(\tau, \tau') \leq \frac{1}{2r}.$$

Proof. Suppose $P_r^\vee(\tau, \tau') > 0$. This means that $|\tau \cap \tau'| = r - 1$. Therefore,

$$w(\tau \cap \tau') \geq w(\tau) + w(\tau') = 2.$$

So, $P_r^\vee(\tau, \tau') = \frac{w(\tau \cap \tau')}{rw(\tau \cap \tau')} \leq \frac{1}{2r}$. □

5 Proof of Strong Log-Concavity and Applications

In this part we prove [Theorem 2.16](#) using connections with high dimensional expanders. [Theorem 2.16](#) can be seen as the converse of [Proposition 4.1](#). In fact, the following statement was proved by Oppenheim [[Opp18](#)].

Theorem 5.1 ([[Opp18](#)]). Let (X, w) be a pure d -dimensional weighted simplicial complex such that:

1. for all $\tau \in X(k)$, for $0 \leq k \leq d-2$, the 1-skeleton of X_τ is a connected (weighted) graph, and
2. for any $\tau \in X(d-2)$, $\lambda_2(\tilde{P}_{\tau,1}^\wedge) \leq 0$.

Then (X, w) is a weighted 0-local-spectral-expander.

Our proof of [Theorem 2.16](#) can be seen as a translation of Oppenheim's result into the language of polynomials.² We actually prove a slightly stronger statement, in the sense that a direct translation of [Theorem 5.1](#) corresponds to testing strong log-concavity of a *multiaffine* homogeneous polynomial. In this case, each derivative $\partial_\tau p$ corresponds to taking the link of τ in X^p . Indecomposability of each derivative then corresponds to connectivity of the 1-skeleton of the corresponding link, and log-concavity of quadratics corresponds to $\lambda_2(\tilde{P}_{\tau,1}^\wedge) \leq 0$.

Proof of Theorem 2.16. We proceed by induction on the degree of p . If the degree of p is at most 2, the claim obviously holds. So, suppose $d \geq 3$. For any $1 \leq i \leq n$, let p_i denote $\partial_i p$. By induction, we can assume that for all i , p_i is strongly log-concave (at $\mathbb{1}$).

First, by [Eq. \(3\)](#), $\nabla^2 p(\mathbb{1}) = \frac{1}{d-2} \sum_{i=1}^n \nabla^2 p_i(\mathbb{1})$. By induction and [Proposition 2.13](#), each matrix $\nabla^2 p_i(\mathbb{1})$ has at most one positive eigenvalue.

Instead, we work with the normalized Hessian matrix, $\tilde{\nabla}^2 p = \frac{1}{d-1} \text{diag}(\nabla p(\mathbb{1}))^{-1} \nabla^2 p(\mathbb{1})$ as defined in [Eq. \(7\)](#). Since the normalized Hessian matrix is stochastic, its top eigenvector is the all-ones vector. When working with the normalized Hessian we need to use the correct inner product operators. For a d -homogeneous polynomial p with nonnegative coefficients and $d > 1$, and vectors $\phi, \psi \in \mathbb{R}^n$, define

$$\langle \phi, \psi \rangle_p = (d-1) \sum_{j=1}^n \phi(j) \psi(j) (\partial_j p(\mathbb{1})),$$

which gives the norm $\|\phi\|_p^2 = \langle \phi, \phi \rangle_p$. The following identity is immediate:

$$\langle \phi, (\tilde{\nabla}^2 p) \psi \rangle_p = \langle \phi, \nabla^2 p(\mathbb{1}) \psi \rangle = \langle (\tilde{\nabla}^2 p) \phi, \psi \rangle_p$$

In particular, $\tilde{\nabla}^2 p$ is self-adjoint with respect to $\langle \cdot, \cdot \rangle_p$. Furthermore, by [Eq. \(3\)](#),

$$\langle \phi, (\tilde{\nabla}^2 p) \psi \rangle_p = \langle \phi, \nabla^2 p(\mathbb{1}) \psi \rangle = \frac{1}{d-2} \sum_{k=1}^n \langle \phi, \nabla^2 p_k(\mathbb{1}) \psi \rangle = \frac{1}{d-2} \sum_{k=1}^n \langle \phi, \tilde{\nabla}^2 p_k \psi \rangle_{p_k}. \quad (10)$$

We highlight that the Hessian $\nabla^2 p(\mathbb{1})$ may be viewed as the weighted adjacency matrix of a graph with edge weights $\partial_i \partial_j p(\mathbb{1})$. Thus, our normalized Hessian may be viewed as the associated random walk matrix, and the inner product $\langle \cdot, \cdot \rangle_p$ may be viewed as a change of basis, which converts the random walk matrix into the normalized adjacency matrix.

Let μ be an eigenvalue of $\tilde{\nabla}^2 p$ with eigenvector ϕ . We prove that $\mu \leq \mu^2$. We claim that this is enough for the induction step: First, since $\tilde{\nabla}^2 p$ is stochastic $\mu \leq 1$. Therefore, we either $\mu = 1$ or $\mu \leq 0$. So, to prove that $\tilde{\nabla}^2 p$ has (exactly) one positive eigenvalue, it is enough to show that $\lambda_2(\tilde{\nabla}^2 p) < 1$. But, since p is indecomposable, the underlying (weighted) graph of $\tilde{\nabla}^2 p$ is connected, so by [Corollary 2.11](#), $\lambda_2(\tilde{\nabla}^2 p) < 1$.

²In a companion paper, we will give an alternative proof based purely on elementary calculus and linear algebra.

It remains to prove that $\mu \leq \mu^2$. From Eq. (10),

$$\mu \|\phi\|_p^2 = \langle \phi, (\tilde{\nabla}^2 p) \phi \rangle_p = \frac{1}{d-2} \sum_{k=1}^n \langle \phi, (\tilde{\nabla}^2 p_k) \phi \rangle_{p_k} \quad (11)$$

Decomposing ϕ orthogonally along $\mathbb{1}$ write

$$\phi = \phi_k^{\perp \mathbb{1}} + \phi_k^{\mathbb{1}}$$

where $\phi_k^{\mathbb{1}} = \frac{\langle \phi, \mathbb{1} \rangle_{p_k}}{\|\mathbb{1}\|_{p_k}^2} \mathbb{1}$ and $\phi_k^{\perp \mathbb{1}}$ is orthogonal to $\mathbb{1}$, i.e., $\langle \phi_k^{\perp \mathbb{1}}, \mathbb{1} \rangle_{p_k} = 0$. It follows that

$$\langle \phi_k^{\perp \mathbb{1}}, (\tilde{\nabla}^2 p_k) \phi_k^{\perp \mathbb{1}} \rangle_{p_k} \leq 0.$$

This is because $\tilde{\nabla}^2 p_k$ has exactly one positive eigenvalue with corresponding eigenvector of $\mathbb{1}$. Therefore,

$$\mu \|\phi\|_p^2 \leq \frac{1}{d-2} \sum_{k=1}^n \langle \phi_k^{\mathbb{1}}, (\tilde{\nabla}^2 p_k) \phi_k^{\mathbb{1}} \rangle_{p_k} = \frac{1}{d-2} \sum_{k=1}^n \langle \phi_k^{\mathbb{1}}, \phi_k^{\mathbb{1}} \rangle_{p_k} = \frac{1}{d-2} \sum_{k=1}^n \frac{\langle \phi, \mathbb{1} \rangle_{p_k}^2}{\langle \mathbb{1}, \mathbb{1} \rangle_{p_k}} \quad (12)$$

Next, we rewrite the numerator and denominator of each ratio in the righthand side. We have

$$\langle \mathbb{1}, \mathbb{1} \rangle_{p_k} = (d-2) \sum_{i=1}^n (\partial_i p_k(\mathbb{1})) = (d-2)(d-1) \cdot p_k(\mathbb{1})$$

where we used Eq. (2) for polynomial p_k . Furthermore,

$$\langle \phi, \mathbb{1} \rangle_{p_k} = (d-2) \sum_{i=1}^n \phi(i) (\partial_i p_k(\mathbb{1})) = (d-2) \cdot ((\nabla^2 p(\mathbb{1})) \phi)(k) \quad (13)$$

So, putting the above identities together, we obtain

$$\frac{\langle \phi, \mathbb{1} \rangle_{p_k}}{\langle \mathbb{1}, \mathbb{1} \rangle_{p_k}} = \frac{1}{(d-1) \cdot p_k} ((\nabla^2 p(\mathbb{1})) \phi)(k) = ((\tilde{\nabla}^2 p(\mathbb{1})) \phi)(k) = \mu \cdot \phi(k)$$

where the last identity we crucially used ϕ is the eigenvector of $\tilde{\nabla}^2 p$ corresponding to μ . Plugging into Eq. (12) we get

$$\begin{aligned} \mu \|\phi\|_p^2 &\leq \frac{1}{d-2} \sum_{k=1}^n \frac{\langle \phi, \mathbb{1} \rangle_{p_k}^2}{\langle \mathbb{1}, \mathbb{1} \rangle_{p_k}} = \frac{\mu}{d-2} \sum_{k=1}^n \phi(k) \langle \phi, \mathbb{1} \rangle_{p_k} \\ &= \mu \sum_{k=1}^n \phi(k) ((\nabla^2 p(\mathbb{1})) \phi)(k) = \mu \langle \phi, (\nabla^2 p(\mathbb{1})) \phi \rangle = \mu \langle \phi, (\tilde{\nabla}^2 p) \phi \rangle_p = \mu^2 \|\phi\|_p^2. \end{aligned}$$

The second equality uses Eq. (13) and the last equality uses definition of ϕ . So, $\mu \leq \mu^2$ as desired. \square

5.1 An Alternative Proof of Strong Log-concavity of Bases Generating Polynomial

Theorem 5.2. *Let $M = ([n], \mathcal{I})$ be a matroid of rank r . Then, for every choice of “external field” $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathbb{R}_{>0}^n$, its (weighted) bases generating polynomial*

$$g_M(x_1, \dots, x_n) = \sum_{B \text{ basis}} \lambda^B x^B$$

is strongly log-concave (at $\mathbf{1}$).

Proof. We verify the indecomposability and log-concavity conditions of [Theorem 2.16](#). If i_1, \dots, i_k contains duplicate elements, then multiaffine-ness of g_M forces $\partial_{i_1} \cdots \partial_{i_k} g_M$ to be identically zero. For any subset $S = \{i_1, \dots, i_k\}$, we use ∂_S as a shorthand notation for $\partial_{i_1} \cdots \partial_{i_k}$. If S is not independent, then no basis contains S , and again, $\partial_S g_M = 0$ identically. Hence, we assume $S \in \mathcal{I}$.

We first argue that $\partial_S g_M$ is indecomposable. Observe that $\partial_S g_M$ equals the weighted basis generating polynomial $g_{M/S}$ of the contraction M/S . As M/S is a matroid of rank ≥ 2 , applying the exchange property immediately tells us that $g_{M/S}$ is indecomposable.

Now, we verify log-concavity of all quadratics. Assume $S \in \mathcal{I}$ and $|S| = r - 2$. As S has rank- $(r - 2)$, M/S has rank two. In particular, $\partial_S g_M = g_{M/S}$ is quadratic and $\nabla^2 g_{M/S}$ has entries

$$(\nabla^2 g_{M/S})_{ij} = \begin{cases} \lambda_i \lambda_j, & \text{if } \{i, j\} \text{ is independent in } M/S \\ 0, & \text{otherwise.} \end{cases}$$

We need to prove that the above matrix has at most one positive eigenvalue.

For any set $T \subseteq [n]$, let λ_T denote the vector with i th entry λ_i for $i \in T$ and 0 for $i \notin T$. The matroid partition property tells us that we can partition the non-loops of M/S into blocks $B = B_1 \cup \cdots \cup B_k$. Then

$$\nabla^2 g_{M/S} = \lambda_B \lambda_B^\top - \sum_{i=1}^k \lambda_{B_i} \lambda_{B_i}^\top \preceq \lambda_B \lambda_B^\top.$$

This proves that $\nabla^2 \partial_S g_M$ has at most one positive eigenvalue and thus $\partial_S g_M$ is log-concave. \square

Corollary 5.3. *Let $M = ([n], \mathcal{I})$ be a matroid of rank r . Then the polynomial*

$$\sum_{S \in \mathcal{I}: |S|=k} \lambda^S x^S$$

is strongly log-concave (at $\mathbf{1}$), for every choice of “external field” $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathbb{R}_{>0}^n$.

Proof. This is the weighted basis generating polynomial of the rank- k truncation of M , which is still a matroid. Thus, the claim follows from [Theorem 5.2](#). \square

5.2 The Random Cluster Model on Matroids

Using the the simplified characterization of strong log-concavity, we can also prove [Theorem 1.6](#).

Proof of Theorem 1.6. We verify the indecomposability and log-concavity conditions of [Theorem 2.16](#). Let f denote $f_{M,k,q}$. As before, for the derivatives $\partial_{i_1} \dots \partial_{i_k} f$, by multiaffine-ness of f , we may assume i_1, \dots, i_k does not contain any duplicate elements. Consider $S = \{i_1, \dots, i_k\}$. Note that $\partial_S f$ has a monomial x^T with nonzero coefficient for every $T \subset [n] \setminus S$ with $|T| \leq k - |S|$. Hence, indecomposability of $\partial_S f$ for every $S \subset [n]$ with $|S| \leq k$ is immediate.

Now, we verify log-concavity of quadratics. Let $S \subset [n]$ with $|S| = k - 2$. First, we calculate that

$$(\partial_S f)(x_1, \dots, x_n) = \sum_{T \in \binom{[n]}{k}: T \supset S} q^{-\text{rank}(T)} \lambda^{T \setminus S} x^{T \setminus S} = \sum_{\{i,j\} \in \binom{[n] \setminus S}{2}} q^{-\text{rank}(S \cup \{i,j\})} \lambda_i \lambda_j x_i x_j.$$

Then for elements $i \neq j$ of $[n] \setminus S$, the (i, j) th entry of $\nabla^2 \partial_S f$ is

$$(\nabla^2 \partial_S f)_{ij} = q^{-\text{rank}(S \cup \{i,j\})} \lambda_i \lambda_j = q^{-\text{rank}(S)} q^{-\text{rank}_{M/S}(\{i,j\})} \lambda_i \lambda_j$$

We will show that the matrix $A = q^{\text{rank}(S)} \nabla^2 \partial_S f$ at most one positive eigenvalue. Note that for $i \neq j$ in $[n] \setminus S$, $A_{ij} = q^{-\text{rank}_{M/S}(\{i,j\})} \lambda_i \lambda_j$. From here, consider the vector $v \in \mathbb{R}^n$ with $v_i = 0$ for $i \in S$, $v_i = \lambda_i$ for loops of M/S and $v_i = q^{-1} \lambda_i$ for non-loops of M/S .

Now consider the matrix $vv^\top - A$, we can check that for $i, j \in [n] \setminus S$,

$$(vv^\top - A)_{ij} = \begin{cases} (q^{-2} - q^{-1}) \lambda_i \lambda_j & \text{if } i, j \text{ are parallel non-loops in } M/S, \text{ and} \\ 0 & \text{otherwise} \end{cases}$$

In particular, by the matroid partition property, if B_1, \dots, B_k denote the equivalence classes of non-loops of M/S which are parallel to each other, then

$$vv^\top - A = (q^{-2} - q^{-1}) \sum_{j=1}^k \lambda_{B_j} \lambda_{B_j}^\top$$

where λ_{B_j} is the vector with entries $\lambda_{B_j}(i) = \lambda_i$ if $i \in B_j$ and $\lambda_{B_j}(i) = 0$ otherwise.

As $0 < q \leq 1$, $q^{-2} - q^{-1} \geq 0$, in which case the right-hand side is positive semidefinite and $A \preceq vv^\top$. We conclude A has at most one positive eigenvalue as desired. \square

Remark 5.4. Observe that

$$q^r f_{M,r,q}(x_1, \dots, x_n)$$

converges to the bases generating polynomial coefficient-wise as $q \rightarrow 0$. Hence, one can view this as a stronger result than strong log-concavity of the bases generating polynomial.

5.3 Geometric Scaling of Coefficients

In this section we prove [Theorem 1.7](#).

Proof of Theorem 1.7. If $k = 0, 1$, the claim is obvious so assume $k \geq 2$. The claim is obvious when $\alpha = 1$, and the case $\alpha = 0$ follows by taking coefficient-wise limits as $\alpha \rightarrow 0$. Hence, we will also assume $0 < \alpha < 1$. Finally, we will assume that all coefficients c_S are strictly positive. The result

for general strongly log-concave polynomials then follows by taking coefficient-wise limits.

Let $T \in \binom{[n]}{k-2}$. We must prove that $\nabla^2 \partial_T f_\alpha$ has at most one positive eigenvalue. Observe that we may concisely write

$$\begin{aligned}\nabla^2 \partial_T f &= \left[c_{T \cup \{i,j\}} \right]_{ij} \\ \nabla^2 \partial_T f_\alpha &= \left[c_{T \cup \{i,j\}}^\alpha \right]_{ij}\end{aligned}$$

As $\nabla^2 \partial_T f$ has at most one positive eigenvalue, and all entries are nonnegative, we may write

$$\nabla^2 \partial_T f = vv^\top - A$$

for some vector $v \in \mathbb{R}^n$ and a positive semidefinite matrix A . Note that since $\nabla^2 \partial_T f$ has strictly positive entries, the Perron-Frobenius Theorem (see [Theorem 2.2](#)) tells us that the entries of v are strictly positive. In particular, $c_{T \cup \{i,j\}} = v_i v_j - A_{ij} > 0$ where $v_i v_j > 0$. Our goal is to write

$$c_{T \cup \{i,j\}}^\alpha = (v_i v_j - A_{ij})^\alpha = v_i^\alpha v_j^\alpha \left(1 - \frac{A_{ij}}{v_i v_j} \right)^\alpha$$

and then Taylor expand $\left(1 - \frac{A_{ij}}{v_i v_j} \right)^\alpha$. Consider the function $\varphi_\alpha(x) = (1-x)^\alpha$, whose Taylor expansion about zero we recall is

$$\sum_{k=0}^{\infty} \frac{\prod_{j=0}^{k-1} (\alpha - j)}{k!} \cdot (-1)^k x^k = \sum_{k=0}^{\infty} \left(\prod_{j=0}^{k-1} \frac{\alpha - j}{1 + j} \right) \cdot (-1)^k x^k = 1 - \sum_{k=1}^{\infty} \left(\prod_{j=0}^{k-1} \left| \frac{\alpha - j}{1 + j} \right| \right) \cdot x^k$$

where for the last equality, we crucially use the fact that $0 < \alpha < 1$. The interval of convergence of this power series contains $(-1, 1)$, as if $a_k = (-1)^k \prod_{j=0}^{k-1} \frac{\alpha - j}{1 + j}$, then

$$\left| \frac{a_{k+1} x^{k+1}}{a_k x^k} \right| = |x| \cdot \frac{\alpha - k}{1 + k} \rightarrow |x| \quad \text{as } k \rightarrow \infty$$

gives a radius of convergence of 1 by the Ratio Test. We verify that this power series also converges at $x \in \{-1, 1\}$ so that we have convergence on the entire closed interval $[-1, 1]$. We can then apply this to $\left(1 - \frac{A_{ij}}{v_i v_j} \right)^\alpha$ since $\nabla^2 \partial_T f$ have nonnegative entries and A being positive semidefinite imply that

$$\det(A_{S,S}) = A_{ii} A_{jj} - A_{ij}^2 \geq 0 \implies |A_{ij}| \leq \sqrt{A_{ii} A_{jj}} = v_i v_j$$

This shows $\frac{A_{ij}}{v_i v_j} \in [-1, 1]$.

To show that the power series converges at $x \in \{-1, 1\}$, it suffices to prove that $\sum_{k=1}^{\infty} \prod_{j=0}^{k-1} \left| \frac{\alpha - j}{1 + j} \right| < \infty$. For this, we first bound each term for $k > 1$ as

$$\prod_{j=0}^{k-1} \left| \frac{\alpha - j}{1 + j} \right| = \frac{\alpha}{k} \cdot \frac{\prod_{j=1}^{k-1} (j - \alpha)}{\prod_{j=1}^{k-1} j} = \frac{\alpha}{k} \prod_{j=1}^{k-1} \left(1 - \frac{\alpha}{j} \right) \leq \frac{\alpha}{k} \cdot e^{-\alpha \sum_{j=1}^{k-1} \frac{1}{j}} \leq \frac{\alpha}{k} \cdot e^{-\alpha \log(k-1)} < \frac{\alpha}{(k-1)^{1+\alpha}}$$

Hence,

$$\sum_{k=1}^{\infty} \prod_{j=0}^{k-1} \left| \frac{\alpha - j}{1 + j} \right| < \alpha + \alpha \sum_{k>1}^{\infty} \frac{1}{(k-1)^{1+\alpha}} < \infty$$

where we crucially use that $\alpha > 0$.

Having verified that the power series is valid for every entry of our matrix, we have

$$\begin{aligned} \nabla^2 \partial_T f_\alpha &= \left[v_i^\alpha v_j^\alpha \right]_{ij} \circ \left(\mathbf{1}\mathbf{1}^\top - \sum_{k=1}^{\infty} \left(\prod_{j=0}^{k-1} \left| \frac{\alpha - j}{1 + j} \right| \right) \cdot \left[\frac{A_{ij}}{v_i v_j} \right]^{\circ k} \right) \\ &= \underbrace{\left[v_i^\alpha v_j^\alpha \right]_{ij}}_{(1)} - \underbrace{\sum_{k=1}^{\infty} \left(\prod_{j=0}^{k-1} \left| \frac{\alpha - j}{1 + j} \right| \right) \cdot \left(\left[v_i^\alpha v_j^\alpha \right]_{ij} \circ \left[\frac{A_{ij}}{v_i v_j} \right]^{\circ k} \right)}_{(2)} \end{aligned}$$

Here, we recall that $A \circ B$ denotes the Hadamard product of A, B , where $(A \circ B)_{ij} = A_{ij} B_{ij}$. Similarly, $A^{\circ k}$ denotes the k -iterated Hadamard product of A with itself.

All we must do is prove that (1) and (2) are both positive semidefinite, and that (1) is rank-1. Observe that

$$\left[v_i^\alpha v_j^\alpha \right]_{ij} = [v_i^\alpha]_i \cdot [v_i^\alpha]_i^\top \quad \left[\frac{1}{v_i v_j} \right]_{ij} = \left[\frac{1}{v_i} \right]_i \cdot \left[\frac{1}{v_i} \right]_i^\top$$

This tells us (1) is positive semidefinite and rank-1. For the second, observe that

$$\left[\frac{A_{ij}}{v_i v_j} \right]_{ij} = A \circ \left[\frac{1}{v_i v_j} \right]_{ij}$$

As $A \succcurlyeq 0$ by assumption, this matrix is positive semidefinite by the Schur Product Theorem (see [Theorem 2.1](#)). Again, inductively applying the Schur Product Theorem, we have $\left[v_i^\alpha v_j^\alpha \right]_{ij} \circ \left[\frac{A_{ij}}{v_i v_j} \right]^{\circ k} \succcurlyeq 0$ for every k . As (2) is a nonnegative linear combination of positive semidefinite matrices, it is positive semidefinite. \square

Remark 5.5. Note that this operation does not preserve complete log-concavity when f is not assumed to be multiaffine. For example, consider the degree-2 bivariate polynomial $f(x, y) = ax^2 + bxy + cy^2$, where $a, b, c > 0$. Here,

$$\nabla^2 f = \begin{bmatrix} 2a & b \\ b & 2c \end{bmatrix}$$

so log-concavity amounts to $\det(\nabla^2 f) = 4ac - b^2 \leq 0$, i.e. $b^2 \geq 4ac$. Now, raise each coefficient to the power α . Then,

$$\nabla^2 f_\alpha = \begin{bmatrix} 2a^\alpha & b^\alpha \\ b^\alpha & 2c^\alpha \end{bmatrix}$$

so log-concavity amounts to $\det(\nabla^2 f_\alpha) = 4a^\alpha c^\alpha - b^{2\alpha} \leq 0$, i.e. $b^2 \geq 4^{1/\alpha} ac$. Clearly, as one decreases α to 0, this inequality gets stronger, which certainly isn't implied by log-concavity of f .

The problem lies in the fact that when you differentiate a monomial that contains variables with multiplicities, you will obtain "factorial coefficients" which are not raised to the power α . The operation must be modified appropriately to take this into account. We defer the discussion of the appropriate generalization of this operation for non-multiaffine polynomials to a future article.

References

- [AHK18] Karim Adiprasito, June Huh, and Eric Katz. "Hodge theory for combinatorial geometries". In: *Annals of Mathematics* 188.2 (2018), pp. 381–452.
- [Alo86] N Alon. "Eigenvalues and expanders". In: *Combinatorica* 6 (2 Jan. 1986), pp. 83–96. issn: 0209-9683.
- [AM85] N. Alon and V. Milman. "Isoperimetric inequalities for graphs, and superconcentrators". In: *Journal of Combinatorial Theory, Series B* 38.1 (Feb. 1985), pp. 73–88.
- [AOR16] Nima Anari, Shayan Oveis Gharan, and Alireza Rezaei. "Monte Carlo Markov Chain Algorithms for Sampling Strongly Rayleigh Distributions and Determinantal Point Processes". In: *COLT*. 2016, pp. 103–115.
- [AOV18] Nima Anari, Shayan Oveis Gharan, and Cynthia Vinzant. "Log-concave polynomials, entropy, and a deterministic approximation algorithm for counting bases of matroids". In: *FOCS*. to appear. 2018.
- [BS07] Alexander Barvinok and Alex Samorodnitsky. "Random weighting, asymptotic counting, and inverse isoperimetry". In: *Israel Journal of Mathematics* 158.1 (Mar. 2007), pp. 159–191. issn: 1565-8511.
- [BS97] W. Ballmann and J. Świątkowski. "On L2-cohomology and property (T) for automorphism groups of polyhedral cell complexes". In: *Geom. Funct. Anal.* 7.4 (1997), pp. 615–645.
- [Clo10] Brian D. Cloteaux. "Approximating the Number of Bases for Almost All Matroids". In: *Congressus Numerantium* 202 (2010), pp. 149–154.
- [CTY15] Emma Cohen, Prasad Tetali, and Damir Yeliussizov. "Lattice path matroids: negative correlation and fast mixing". In: *arXiv preprint arXiv:1505.06710* (2015).
- [DK17] I. Dinur and T. Kaufman. "High Dimensional Expanders Imply Agreement Expanders". In: *FOCS*. 2017, pp. 974–985.
- [DS91] Persi Diaconis and Daniel Stroock. "Geometric bounds for eigenvalues of Markov chains". In: *The Annals of Applied Probability* (1991), pp. 36–61.
- [FK72] Cornelis Marius Fortuin and Pieter Willem Kasteleyn. "On the random-cluster model: I. Introduction and relation to other models". In: *Physica* 57 (4 1972), pp. 536–564.
- [FM92] Tomás Feder and Milena Mihail. "Balanced matroids". In: *STOC*. 1992, pp. 26–38.
- [For71] Cornelis Marius Fortuin. "On the random-cluster model". PhD thesis. Leiden University, 1971.

- [For72a] Cornelis Marius Fortuin. “On the random-cluster model: II. The percolation model”. In: *Physica* 58 (3 1972), pp. 393–418.
- [For72b] Cornelis Marius Fortuin. “On the random-cluster model: III. The simple random-cluster model”. In: *Physica* 59 (4 1972), pp. 545–570.
- [Gam99] Anna Gambin. “On approximating the number of bases of exchange preserving matroids”. In: *International Symposium on Mathematical Foundations of Computer Science*. Springer. 1999, pp. 332–342.
- [Gar73] H. Garland. “p-adic curvature and the cohomology of discrete subgroups of p-adic groups”. In: *Annals of Mathematics* 97.3 (1973), pp. 375–423.
- [GJ08] L. A. Goldberg and M. Jerrum. “Inapproximability of the Tutte polynomial”. In: *Inform. and Comput* 206.7 (2008), pp. 908–929.
- [GJ12a] L. A. Goldberg and M. Jerrum. “Approximating the partition function of the ferromagnetic Potts model”. In: *J. ACM* 59.5 (2012), pp. 1–31.
- [GJ12b] L. A. Goldberg and M. Jerrum. “Inapproximability of the Tutte polynomial of a planar graph”. In: *Computational Complexity* 21 (4 2012), pp. 605–642.
- [GJ13] L. A. Goldberg and M. Jerrum. “Approximating the Tutte polynomial of a binary matroid and other related combinatorial polynomials”. In: *Journal of Computer and System Sciences* 79 (1 2013), pp. 68–78.
- [GJ14] L. A. Goldberg and M. Jerrum. “The complexity of computing the sign of the Tutte polynomial”. In: *SIAM J. Comput.* 43 (2014), pp. 1921–1952.
- [GJ17] Heng Guo and Mark Jerrum. “Random cluster dynamics for the Ising model is rapidly mixing”. In: *SODA*. 2017, pp. 1818–1827.
- [GJ18a] Heng Guo and Mark Jerrum. “A Polynomial-Time Approximation Algorithm for All-Terminal Network Reliability”. In: *ICALP*. 2018, 68:1–68:12.
- [GJ18b] Heng Guo and Mark Jerrum. “Approximately counting bases of bicircular matroids”. 2018.
- [Gri09] G. R Grimmett. *The Random-Cluster Model*. Berlin: Springer-Verlag, 2009.
- [Gur09] Leonid Gurvits. “A polynomial-time algorithm to approximate the mixed volume within a simply exponential factor”. In: *Discrete & Computational Geometry* 41.4 (2009), pp. 533–555.
- [Gur10] Leonid Gurvits. “On multivariate Newton-like inequalities”. In: *Advances in Combinatorial Mathematics*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 61–78.
- [HJ13] Roger A Horn and Charles R Johnson. *Matrix analysis*. 2nd ed. Cambridge university press, 2013.
- [HW16] June Huh and Botong Wang. “Enumeration of points, lines, planes, etc”. In: *arXiv preprint arXiv:1609.05484* (2016).
- [Jer+04] Mark Jerrum, Jung-Bae Son, Prasad Tetali, and Eric Vigoda. “Elementary bounds on Poincaré and log-Sobolev constants for decomposable Markov chains”. In: *Annals of Applied Probability* (2004), pp. 1741–1765.

- [Jer06] Mark Jerrum. “Two Remarks Concerning Balanced Matroids”. In: *Combinatorica* 26.6 (2006), pp. 733–742.
- [JS02] Mark Jerrum and Jung Bae Son. “Spectral gap and log-Sobolev constant for balanced matroids”. In: *FOCS*. 2002, pp. 721–729.
- [JS93] Mark Jerrum and Alistair Sinclair. “Polynomial-time approximation algorithms for the Ising model”. In: *SIAM J. Comput.* 22.5 (1993), pp. 1087–1116.
- [JVV86] Mark Jerrum, Leslie Valiant, and Vijay Vazirani. “Random Generation of Combinatorial Structures from a Uniform Distribution”. In: *Theoretical Computer Science* 43 (1986), pp. 169–188.
- [JVW90] F. Jaeger, D. L. Vertigan, and D. J. A. Welsh. “On the computational complexity of the Jones and Tutte polynomials”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 108 (1 1990), pp. 35–53.
- [KM17] Tali Kaufman and David Mass. “High Dimensional Random Walks and Colorful Expansion”. In: *ITCS*. 2017, 4:1–4:27.
- [KO18] Tali Kaufman and Izhar Oppenheim. “High Order Random Walks: Beyond Spectral Gap”. In: *APPROX/RANDOM*. 2018, 47:1–47:17.
- [KT12] Alex Kulesza and Ben Taskar. “Determinantal point processes for machine learning”. In: *Foundations and Trends in Machine Learning* 5(2-3) (2012), pp. 123–286.
- [Lub17] Alexander Lubotzky. “High Dimensional Expanders”. 2017.
- [MS91] Milena Mihail and Madhu Sudan. *Connectivity Properties of Matroids*. Tech. rep. EECS Department, University of California, Berkeley, Dec. 1991.
- [MV89] M. Mihail and U. Vazirani. “On the expansion of 0/1 polytopes”. In: *Journal of Combinatorial Theory*. B (1989).
- [Opp18] Izhar Oppenheim. “Local spectral expansion approach to high dimensional expanders part I: Descent of spectral gaps”. In: *Discrete and Computational Geometry* 59.2 (2018), pp. 293–330.
- [Ver91] D.L. Vertigan. “On the computational complexity of tutte, jones, homfly and kauffman invariants”. PhD thesis. University of Oxford, 1991.
- [Wel94] D.J.A Welsh. “The computational complexity of knot and matroid polynomials”. In: *Discrete Mathematics* 124 (1-3 1994), pp. 251–269.