Log-concavity, High-dimensional Expanders, and Lorentzian Polynomials

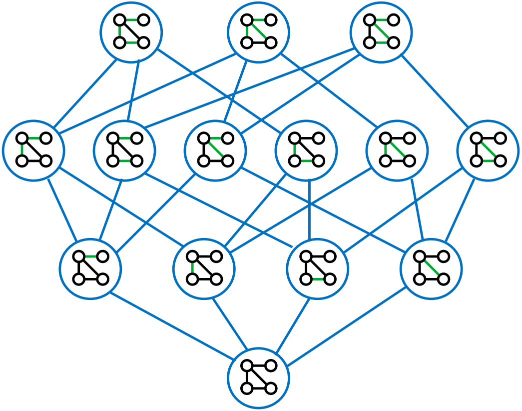

Jonathan Leake (UWaterloo C&O)

CAAC 2026

This Talk

Goal: Log-concavity statements and efficient sampling algorithms

What is **log-concavity?** Given a sequence $(c_k)_{k=0}^d$ of non-negative numbers (often integers), can we prove

$$c_k^2 \geq c_{k-1}c_{k+1} \quad \text{for all } k?$$

What is **sampling?** Given a set $\mathcal{S}$ of mathematical objects, can we efficiently choose one uniformly at random? (Very related to **counting**.)

Two parts:

1. Introduce **Lorentzian polynomials** and **high-dimensional expanders** via examples: independent sets of a matroid and forests of a graph
2. Extend these ideas beyond graphs and matroids to log-concavity and sampling of **generalized graph colorings**

The Connection Between Log-concavity and Sampling

Log-concavity and Sampling

Question: What connects log-concavity and sampling?

Answer: Matrices with bounded second-largest eigenvalue ($\lambda_2 \leq C$)

Log-concavity: Given non-negative $(c_k)_{k=0}^{d}$, we have

$$c_k^2 \geq c_{k-1}c_{k+1} \quad \iff \quad \lambda_2\left(\begin{bmatrix} c_{k+1} & c_k \\ c_k & c_{k-1} \end{bmatrix}\right) \leq 0.$$

Why? One way: $\mathrm{tr} \geq 0 \implies \lambda_1 \geq 0 \implies \det \leq 0$

Sampling: More complicated...

Example: Sampling Spanning Trees

Setup: Given a connected graph $G = (V, E)$, consider the uniform distribution μ on spanning trees of G.

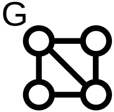

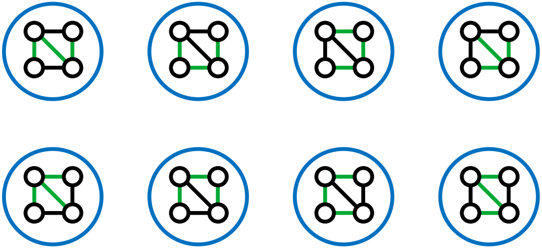

Goal: Sample a uniformly random spanning tree.

- Number of spanning trees can be **exponential** in $|V|, |E|$
- Exact sampling can be hard, so we settle for **approximate sampling**

Sampling Spanning Trees via MCMC

Idea for sampling from μ: Markov chain Monte Carlo (MCMC)

- Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \{\text{spanning trees of } G\}$ and $(T_1, T_2) \in \mathcal{E}$ whenever T_1, T_2 differ by one edge swap
- Do a **random walk** on $\mathcal{G}$, starting at $T \in \mathcal{V}$:
  1. choose random edge (T, T') incident to T
  2. randomly move to T', or else stay at T
  3. repeat
- **Yields uniform sample** as $t \to \infty$

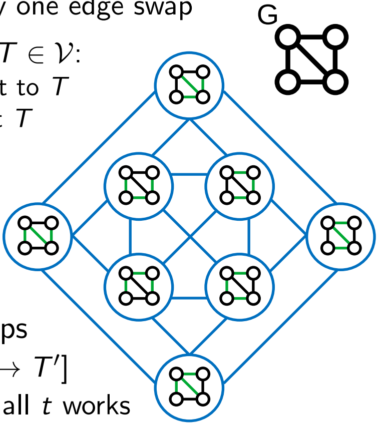

Hope: Approx. sample after small t steps

- **Transition matrix:** $P(T, T') = \mathbb{P}[T \to T']$
- $\lambda_1(P) = 1$, and $\lambda_2(P) \ll 1$ implies small t works
- $\mathcal{G}$ is a λ-**expander** if $\lambda_2(P) \leq \lambda$

Log-concavity and Sampling

Question: What connects log-concavity and sampling?

Answer: Matrices with bounded second-largest eigenvalue ($\lambda_2 \leq C$)

Log-concavity: Given non-negative $(c_k)_{k=0}^d$, we have

$$c_k^2 \geq c_{k-1}c_{k+1} \quad \iff \quad \lambda_2\left(\begin{bmatrix} c_{k+1} & c_k \\ c_k & c_{k-1} \end{bmatrix}\right) \leq 0.$$

Why? One way: $\text{tr} \geq 0 \implies \lambda_1 \geq 0 \implies \det \leq 0$

Sampling: Random walk on $\mathcal{G}$ gives approximate sample

1. Bounded $\lambda_2(P)$ implies **fast mixing**, where P is transition matrix
2. All carries over to **bases of a matroid** in general

Lorentzian Polynomials and Log-concavity

Lorentzian Polynomials

Definition ([Gurvits '09], [Anari-Liu-Oveis Gharan-Vinzant '19], [Brändén-Huh '19],[Brändén-L '21])

Let $\mathcal{C} \subseteq \mathbb{R}^n$ be an open convex cone. A d-homogeneous polynomial $p \in \mathbb{R}[x_1, \ldots, x_n]$ is $\mathcal{C}$-**Lorentzian** if $p \equiv 0$ or

- $\nabla_{\boldsymbol{v}_1} \nabla_{\boldsymbol{v}_2} \cdots \nabla_{\boldsymbol{v}_d} p > 0$ for all $\boldsymbol{v}_i \in \mathcal{C}$, where $\nabla_{\boldsymbol{v}} p := \sum_i v_i \frac{\partial}{\partial x_i} p$
- $\lambda_2(\boldsymbol{H} \nabla_{\boldsymbol{v}_1} \cdots \nabla_{\boldsymbol{v}_{d-2}} p) \leq 0$ for $\boldsymbol{v}_i \in \mathcal{C}$, where $\boldsymbol{H}$ = Hessian

(Original Lorentzian polynomials defined with $\mathcal{C} = \mathbb{R}^n_{>0}$)

Lemma (Lorentzian $\implies$ Log-concavity)

If p is $\mathcal{C}$-Lorentzian and $\boldsymbol{u}, \boldsymbol{v} \in \overline{\mathcal{C}}$ and $p(\boldsymbol{u} \cdot t + \boldsymbol{v} \cdot s) = \sum_{k=0}^{d} \binom{d}{k} c_k t^k s^{d-k}$, then $c_k^2 \geq c_{k-1} c_{k+1}$ for all k.

1. $q(t, s) = p(\boldsymbol{u} \cdot t + \boldsymbol{v} \cdot s)$ is $\mathbb{R}^2_{>0}$-Lorentzian
2. $\boldsymbol{H} \left(\frac{\partial}{\partial t}\right)^{k-1} \left(\frac{\partial}{\partial s}\right)^{d-k-1} q(t, s) \approx \begin{bmatrix} c_{k+1} & c_k \\ c_k & c_{k-1} \end{bmatrix}$ (+ limiting argument)

Examples of Lorentzian Polynomials

For now: Let's consider $\mathcal{C} = \mathbb{R}^n_{>0}$

- **Determinant:** $\det(\sum_i A_i x_i)$ for PSD A_i
- **Volume:** $\mathrm{vol}(\sum_i K_i x_i)$ for compact convex K_i (Alexandrov-Fenchel)
- **Spanning tree:** $\sum_T \prod_{e \in T} x_e$ (or **matroid basis** $T \to B$)
- **Normalized combinatorial polynomials:** Schur, Schubert, conjecturally many others [Huh-Matherne-Mészáros-St. Dizier '22]

How to prove these are Lorentzian?

Theorem (Connectedness + Quadratics [ALOV '19], [BH '19])

A d-homog. p with $\mathbb{R}_{\geq 0}$ coefficients is $\mathbb{R}^n_{>0}$-Lorentzian if and only if

- **Connectedness:** $\boldsymbol{H} \frac{\partial}{\partial x_{i_1}} \cdots \frac{\partial}{\partial x_{i_k}} p$ *is* **irreducible** *for all* $k, i_j \in [n]$
- **Quadratics:** $\lambda_2(\boldsymbol{H} \frac{\partial}{\partial x_{i_1}} \cdots \frac{\partial}{\partial x_{i_{d-2}}} p) \leq 0$ *for all* $i_j \in [n]$

E.g.: for spanning trees, follows from 3-vertex case since $\frac{\partial}{\partial x_e}$ = contraction

Mason's Conjectures

Conjecture: Given connected graph G, let f_k be the number of size-k forests of G (size-k independent sets of a matroid). Then:

$$c_k^2 \geq c_{k-1}c_{k+1}, \quad \text{where} \quad c_k = \frac{f_k}{\binom{|E|}{k}}.$$

- Various strengths of conjecture due to Mason in the '70s
- Proven by [ALOV '19], [BH '19], see also [Adiprasito-Huh-Katz '15], [Huh-Schröter-Wang '18]

Proof.

1. $p(\boldsymbol{x}, y) = \sum_F \left(\prod_{e \in F} x_e\right) y^{|E|-|F|}$ is Lorentzian by C+Q Theorem
2. $q(t,s) = p(\mathbf{1}_{\boldsymbol{x}} \cdot t + \mathbf{1}_y \cdot s) = \sum_{k=0}^{|E|} f_k t^k y^{|E|-k} = \sum_{k=0}^{|E|} \binom{|E|}{k} c_k t^k s^{|E|-k}$
3. $\mathbf{1}_{\boldsymbol{x}}, \mathbf{1}_y \in \mathbb{R}_{\geq 0}^{|E|+1} \implies c_k^2 \geq c_{k-1}c_{k+1}$ for all k by Lemma

□

Summary: From Lorentzian Polynomials to Log-concavity

Main recipe for log-concavity via Lorentzian polynomials:

1. **Encode** the object you want to study as a polynomial p
2. Prove p is $\mathcal{C}$-**Lorentzian**
3. Choose particular vectors $\boldsymbol{u}, \boldsymbol{v} \in \overline{\mathcal{C}}$, such that c_k in

$$p(\boldsymbol{u} \cdot t + \boldsymbol{v} \cdot s) = \sum_{k=0}^{d} \binom{d}{k} c_k t^k s^{d-k}$$

   counts something that you care about

4. Conclude **log-concavity** of the sequence c_k using the Lemma

Sampling via High-dimensional Expanders

Sampling Spanning Trees Revisited

Idea for sampling: Markov chain Monte Carlo (MCMC)

- Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \{\text{spanning trees of } G\}$ and $(T_1, T_2) \in \mathcal{E}$ whenever T_1, T_2 differ by one edge swap
- Do a **random walk** on $\mathcal{G}$, starting at $T \in \mathcal{V}$:
  1. choose random edge (T, T') incident to T
  2. randomly move to T', or else stay at T
  3. repeat
- **Yields uniform sample** as $t \to \infty$

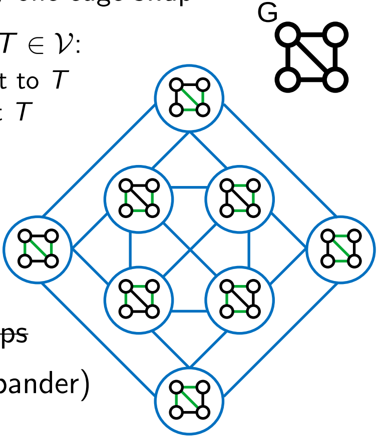

~~**Hope:** Approx. sample after small t steps~~

Goal: Bound $\lambda_2(P) \leq \lambda$ (i.e. $\mathcal{G}$ is λ-expander)

More Structure: Forests

Problem: Unclear how to bound $\lambda_2(P)$ directly

- **Solution:** High-dimensional expanders (HDX)

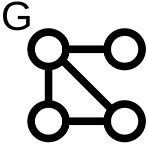

Idea: More structure by including all **forests** of G

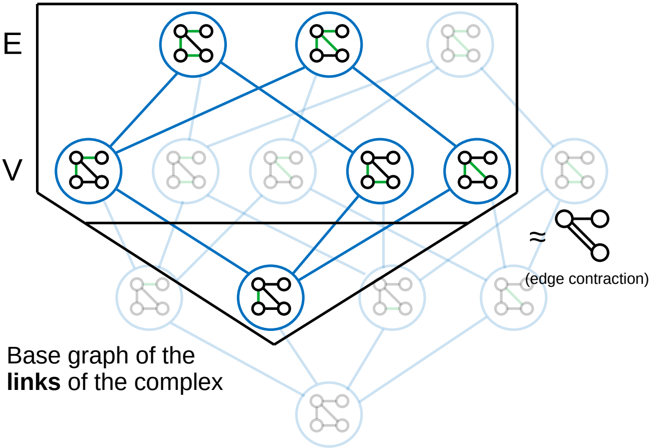

What actually is an HDX?

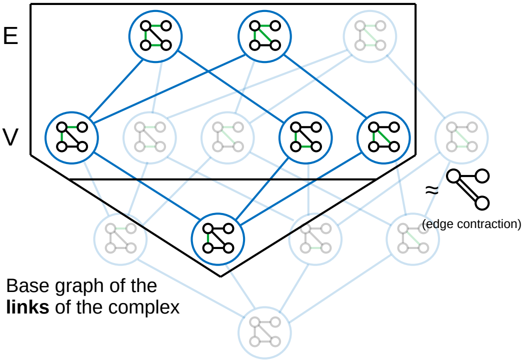

HDX: simplicial complex with the following **expansion** property

Definition (Local spectral expansion [Kaufman-Oppenheim '18]; see also [Dinur-Kaufman '17], [Kaufman-Mass '17], [Oppenheim '18]**)**

We say a complex X is a $(\beta_0, \beta_1, \ldots, \beta_{d-2})$-**expander** if the base graphs of all height-k links of X are β_k-expanders for all k.

From HDX to Approximate Sampling

How to use local spectral expansion?

Main recipe for (approximate) sampling via HDX:

1. **Encode** the objects you want to sample as the facets (top row) of a simplicial complex X
2. Prove **local spectral expansion**: X is a $(\beta_0, \ldots, \beta_{d-2})$-expander
   - trickle-down theorems, correlation decay, coupling, etc.
3. Use **local-to-global theorem**: X is a $(\frac{O(1)}{d}, \frac{O(1)}{d-1}, \frac{O(1)}{d-2}, \ldots)$-expander implies good $\lambda_2(P)$ bound on facet walk; e.g. via [Alev-Lau '20]:

$$(\beta_0, \ldots, \beta_{d-2})\text{-expander} \implies \lambda_2(P) \lesssim 1 - \frac{1}{d}\prod_{i=0}^{d-2}(1-\beta_i)$$

4. Conclude facet walk is **fast mixing**

Local Spectral Expansion via Polynomials

How to prove local spectral expansion?

Theorem (Trickle-down Theorem [ALOV '19], see also [Oppenheim '18]**)**

For complex X, if all base graphs are **connected**, *and all base graphs of* **height-$(d-2)$ links** *are 0-expanders, then X is a $(0, 0, \ldots, 0)$-expander.*

Corollary of C+Q Theorem for Lorentzian polynomials

1. **Encode** facets in polynomial; e.g.: $p(\boldsymbol{x}) = \sum_T \prod_{e \in T} x_e$
2. Partial derivatives correspond to **links (= edge contraction)**
3. Base graph of link of $\sigma \sim \boldsymbol{H}\nabla_{\mathbf{1}}^{d-2-|\sigma|} \left(\prod_{e\in\sigma} \frac{\partial}{\partial x_e}\right) p$
4. Conclude X is a $(0, 0, \ldots, 0)$-**expander**

Theorem ([ALOV '19], [Cryan-Guo-Mousa '20], [Anari-Liu-Oveis Gharan-Vinzant-Vuong '21]; see also [Schild '18]**)**

The "facet walk" for matroid bases (and spanning trees) is fast mixing.

Beyond Forests and Independent Sets of Matroids

What Else?

So far: Always $\mathcal{C} = \mathbb{R}^n_{>0}$ and 0-expansion

- Forces graph forests or independent sets of matroids [BH '19]
- **Can we generalize?** (smaller cone, weaker expansion, etc.?)

Generalized/correspondence graph colorings [Dvořák-Postle '16], e.g.:

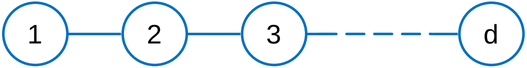

- Each vertex has its own set of colors
- Each edge has its own set of allowed pairs of colors

Set of all partial colorings = **simplicial complex** X

- Faces of complex are sets of **vertex-color pairs**
- Facets of complex $\sim$ complete colorings

Coloring Complex Examples

Obvious example: Proper colorings of a general graph

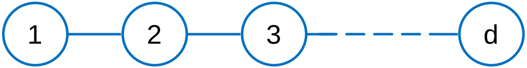

Path complex examples:

1. **Permutations:** vertex k is colored by $S_k \subset [d+1]$ of size k, edges enforce $S_k \subset S_{k+1}$ for all k; **facets = permutations of** $[d+1]$
2. **Linear extensions of poset** P**:** same idea as permutations, but require $x \not\leq_P y$ for $x \in S_{k+1} \setminus S_k$ and $y \in S_k$; **facets = lin. ext.**
3. **Flags of flats:** vertex k is colored by k-dim. subspace W_k of $\mathbb{F}^{d+1}$, edges enforce $W_k \subset W_{k+1}$; **facets = full flags of flats**
4. **Flags of flats of a matroid:** same as flag of flats, but we generalize to **flats of any matroid**, not just representable matroids

Analyzing Path Complexes

Before: Encode spanning trees in $\mathbb{R}^n_{>0}$-Lorentzian polynomial

Problem: Path complex facet-generating polynomial $p(\boldsymbol{x}) = \sum_\sigma \prod_{vc \in \sigma} x_{vc}$ is **not necessarily Lorentzian**.

Solution: New polynomials, smaller cones

- $p_\tau := \left(\prod_{vc \in \tau} \frac{\partial}{\partial x_{vc}}\right) p$, **facet-generating polynomial of link** of τ
- Equivalent **inductive definition** of p_τ via Euler's formula:

$$(d - |\tau|) \cdot p_\tau(\boldsymbol{x}) = \sum_{vc} x_{vc} \cdot p_{\tau \cup \{vc\}}(\boldsymbol{x})$$

- **New polynomials** $\tilde{p}_\tau$ via special linear projections π:

$$(d - |\tau|) \cdot \tilde{p}_\tau(\boldsymbol{x}) = \sum_{vc} x_{vc} \cdot \tilde{p}_{\tau \cup \{vc\}}(\pi_{\tau + vc}(\boldsymbol{x}))$$

Theorem (C+Q Theorem [L-Lindberg-Oveis Gharan '25]; see also [Brändén-L '23]**)**

Given cones $\mathcal{C}_\tau$ s.t. $\pi_{\tau+vc}(\mathcal{C}_\tau) \subseteq \mathcal{C}_{\tau \cup \{vc\}}$ and connected complex X, if $\tilde{p}_\tau$ is $\mathcal{C}_\tau$-Lorentzian for all τ with $|\tau| = d - 2$ then $\tilde{p}_\tau$ is $\mathcal{C}_\tau$-Lorentzian for all τ.

Applications: Sampling

Theorem

For a **connected path complex** X*, if all base graphs of height-*$(d-2)$ *links are* $\frac{1}{2}$*-expanders, then* X *is a* $(\frac{1}{2}, \frac{1}{2}, \ldots, \frac{1}{2})$*-expander.*

Corollaries:

- Sampling of **linear extensions** of a poset [Ferlner-Wernisch '97], [Bubley-Dyer '99]
- Sampling of permutations and linear extensions with respect to **interesting(?) distributions**
- Sampling of **flags of flats** (representable) and **projective** generalizations

Comments:

- Height-$(d-2)$ are $(\frac{1}{2}-\epsilon)$-expanders $\implies$ $(\frac{O(1)}{d}, \frac{O(1)}{d-1}, \ldots)$-expander
- Height-$(d-2)$ are $(\frac{1}{2}+\epsilon)$-expanders $\implies$ bad examples

Applications: Log-concavity

One extra thing: "Increasing" function $\phi(vc) \in \mathbb{R}$ on vertex-color pairs

Theorem

*If X is a ϕ-**expander**, then $(c_k)_{k=0}^d$ is log-concave where*

$$c_k = \sum_{\sigma,\text{facet}} \frac{\prod_{i=0}^d (\phi(\sigma_{i+1}) - \phi(\sigma_i))}{\prod_{i=1}^k (\phi(\sigma_i) - \alpha) \cdot \prod_{i=k+1}^d (\beta - \phi(\sigma_i))},$$

where σ_i is the vertex-color pair for vertex i, and $\alpha \le \phi(\sigma_i) \le \beta$ for all i.

- ϕ-**expander:** C+Q Theorem applies with $\pi_{\tau+vc}$ based on ϕ

Corollaries:

- coefficients of (reduced) **chromatic polynomial** [Huh '12]
- **Heron-Rota-Welsh conjecture** on matroids [Adiprasito-Huh-Katz '18]
- **Stanley's inequality** on linear extensions of a poset [Stanley '81]

Example: Linear Extensions of a Poset

Height-$(d-2)$ links $\iff$ Posets with 3 elements

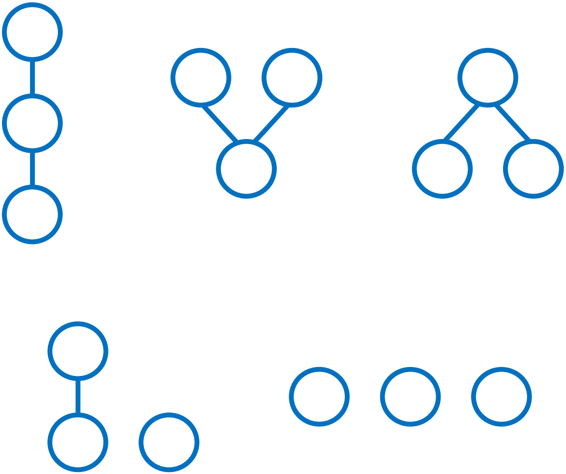

Summary and Future Directions

Part 1:

1. Log-concavity and sampling connected via **second eigenvalue bound**
2. **Lorentzian polynomials** encapsulate second eigenvalue bounds
   - **Log-concavity** statements are natural corollaries
3. Add structure via **high-dimensional expanders** to prove sampling

Part 2:

1. Generalize cones and polynomials
   - Allows applications **beyond spanning trees and matroids**
2. Sampling and log-concavity theorems for **path coloring complexes**

Future directions:

- Coloring complexes for **general graphs**
  - **Conjecture:** sample graph proper colorings
- Other log-concavity statements?
  - Many open questions related to **linear extensions**

Thanks

Thanks!