ISM Summer School: Intersections of Statistical Physics & Combinatorics

Corrine Yap

June 15, 2023

Here are the lecture notes and exercises from the ISM Summer School in Random Trees, Graphs, and Maps, June 5–9, 2023. These are intended as a short introduction to statistical physics for graduate and undergraduate students in combinatorics or discrete probability, so the choice of topics was due to that audience (as well as the personal biases of the author, who is primarily a combinatorialist). The author is indebted to Will Perkins, for teaching her much of the material covered in these notes, and Louigi Addario-Berry, for organizing the summer school.

Contents

1 Lecture 1: Introduction, Phase Transitions and the Ising Model

1.1 Lecture 1 Notes

We started out by looking at a simulation of the Ising model on a subset of $\mathbb{Z}^2$. We observed that at one setting of some parameter β, the configurations looked pretty "random," and at another choice of β, the configurations looked pretty structured or ordered.

1.1.1 Spin Models on Graphs

We set

- $G = (V, E)$ a graph,
- Ω is a set of spins,
- $\sigma : V \to \Omega$ is a spin configuration.

We are interested in describing the "typical" spin configurations, which tells us we should use randomness. So we define a probability distribution.

Examples

(1) Ising model

- $\Omega = \{+1, -1\}$
- $\sigma : V \to \Omega$ is a 2-coloring of V
- $M(G, \sigma) :=$ number of monochromatic edges
- $\beta :=$ inverse temperature (we will only consider $\beta > 0$).
- Gibbs measure $\mu_{G,\beta}(\sigma) \propto e^{\beta M(G,\sigma)}$ i.e.

$$\mu_{G,\beta}(\sigma) = \frac{e^{\beta M(G,\sigma)}}{\sum_{\chi : V \to \Omega} e^{\beta M(G,\chi)}}$$

- partition function $Z_G(\beta) = \sum_{\chi: V \to \Omega} e^{\beta M(G,\chi)}$

Why does the definition of Gibbs measure make sense?

- High temperature: When β is small, e.g. $\beta \to 0$, μ_G approaches the uniform distribution over all spin configurations. This is intuitively what we expect from something "gaseous" or disordered.

- Low temperature: When β is large, e.g. $\beta \to \infty$, μ_G approaches the uniform distribution over the configurations with the *maximum* number of monochromatic edges, i.e. the all-monochromatic configurations. This is again intuitively what we expect from something "solid" or ordered.

(2) Potts model
The Potts model is a generalization of the Ising model to q colors.

- $\Omega = \{1, 2, \ldots, q\} =: [q]$

- $\sigma : V \to \Omega$ is a q-coloring of V

We define the Gibbs measure and partition function to be the same as for the Ising model.
(3) Hardcore model
This is a distribution on independent sets of a graph. An **independent set** $I \subset V(G)$ is a set of vertices such that no two vertices in I are adjacent (meaning no two vertices in I have an edge between them).

- $\mathcal{I}(G)$ is the collection of independent sets in G

- $\lambda > 0$ is the activity or fugacity

- Gibbs measure $\mu_{G,\lambda}(I) \propto \lambda^{|I|}$

- $Z_G(\lambda) = \sum_{J \in \mathcal{I}(G)} \lambda^{|J|}$

Observe: we can rewrite $Z_G(\lambda) = \sum_{k \geq 0} i_k(G)\lambda^k$ where $i_k(G)$ is the number of independent sets of size k in G. In this form, $Z_G(\lambda)$ is commonly known as the *independence polynomial* of the graph G and is well-studied in graph theory.

The hardcore model is also a spin model, under the following framework:

General framework:

- Hamiltonian $H(\sigma) = \sum_{v \in V} f(\sigma_v) + \sum_{uv \in E} g(\sigma_u, \sigma_v)$ for some f, g (σ_v denotes the spin of the vertex v).

- Gibbs measure $\mu_{G,\beta}(\sigma) \propto e^{-\beta H(\sigma)}$

Exercise 1 is to write the hardcore model in this framework. Hint: we can use spins $\Omega = \{0, 1\}$ and let the vertices in I be those with spin 1. Also note that we are allowed to have ∞ in the image of f or g.

1.1.2 Ising Model on $\mathbb{Z}^d$

For more on this section, see Friedli-Velenik textbook in the references.

A main question of interest to statistical physicists is about the notion of phase transitions. Does there exist one? How would we describe it mathematically? There are several ways. One thing to note is that a phase transition is a phenomenon of infinite graphs, or a limiting phenomenon of finite graphs.

Let Λ_n denote the box of sidelength n in $\mathbb{Z}^d$. Let $\mu^\tau_{\Lambda_n}$ be the Gibbs measure on Λ_n with boundary conditions τ.

Fact: The measures $\mu^\tau_{\Lambda_n}$ converge weakly to a measure μ^τ_∞, called an *infinite volume Gibbs measure*. (This result holds for general infinite G as long as the regions Λ_n converge *in the sense of van Hove* to the infinite graph. This is not always the case; for example, the infinite d-regular tree does not have such a sequence of subgraphs.)

We say a **phase transition** occurs at $\beta = \beta_c$ if for

- $\beta < \beta_c$, there exists a unique infinite volume Gibbs measure
- $\beta > \beta_c$, there exist multiple infinite volume Gibbs measures

Theorem 1.1 (Peierls 1936)**.** *The Ising model in $\mathbb{Z}^d$ has a phase transition if $d \geq 2$.*

This is actually an if and only if statement; for $d = 1$, Ising proved in his 1924 PhD thesis that exponential decay of correlations holds for all β, which implies there is no phase transition.

Proof. We will prove the low-temperature statement, i.e. if β is large, there must be more than one infinite volume Gibbs measure. We will do this by showing that the all $+1$ and all -1 boundary conditions give rise to distinct infinite volume measures. To show that they are different, we analyze $\mathbb{E}[\sigma_0^\tau]$, the expected spin at the origin.

Let's first consider the all $+1$ boundary conditions. We will rewrite the Ising model and think about bichromatic edges instead.

- $B(G, \sigma) =$ number of bichromatic edges of σ
- $\mu_{G,\beta}(\sigma) \propto e^{-\beta B(G,\sigma)}$

We define something called *contours* by doing the following:

- Draw a unit box around each -1 spin in the *dual lattice* (place a vertex in the center of each grid square and connect vertices in adjacent grid squares).
- Erase edges separating adjacent -1 spins.
- Round off corners in an unambiguous way e.g. if a northeast and southwest corner meet, merge them, and if a northwest and southeast corner meet, separate them.

Check: this results in a collection of nonintersecting simply connected closed curves. Spins are constant on the regions defined by these curves. Given a collection of contours and the boundary conditions, then, we can recover the original spin configuration.

Main point: The contour edges from this process are in bijective correspondence with the bichromatic edges of σ. So we can rewrite the Ising model as a contour model: if Γ is a collection of contours, then the probability of Γ can be written as

$$\mu^+_{\Lambda_n}(\Gamma) = \frac{\prod_{\gamma\in\Gamma} e^{-\beta|\gamma|}}{\sum_{\Gamma'}\prod_{\gamma\in\Gamma'} e^{-\beta|\gamma|}}$$

where the sum in the denominator is over all collections of contours.

Now we consider the probability that the origin gets spin -1. If we can show this probability is small, then we get that the expected spin at the origin is positive.

Exercise: The origin receives spin -1 if and only if the origin is inside an odd number of contours.

Thus,

$$\begin{aligned}\mu^+_{\Lambda_n}(\sigma_0 = -1) &= \mu^+_{\Lambda_n}(0 \text{ is in an odd number of contours})\\ &\le \mu^+_{\Lambda_n}(0 \text{ is in some contour})\\ &= \mu^+_{\Lambda_n}(\{\Gamma : \text{ there exists } \gamma^\star \in \Gamma \text{ such that } 0 \in Int(\gamma^\star)\})\\ &\le \sum_{\gamma^\star : 0\in Int(\gamma^\star)} \sum_{\Gamma\ni\gamma^\star} \mu^+_{\Lambda_n}(\Gamma)\end{aligned}$$

where $Int(\gamma^\star)$ refers to the interior of the contour, meaning the vertices inside the region whose boundary is $\gamma^\star$. In the last inequality, we use a first moment bound.

Exercise/Claim:

$$\sum_{\Gamma\ni\gamma^\star} \mu^+_{\Lambda_n}(\Gamma) \le e^{-\beta|\gamma^\star|}$$

Hint: To prove this, write out $\mu^+_{\Lambda_n}(\Gamma)$ using the definition of our contour model Gibbs measure. Factor out $e^{-\beta|\gamma^\star|}$ from the numerator. Show that the remaining terms are ≤ 1 by finding an injection from the terms in the numerator to the terms in the denominator.

Assuming the claim holds, we now have

$$\mu^+_{\Lambda_n}(\sigma_0 = -1) \le \sum_{\gamma^\star : 0\in Int(\gamma^\star)} e^{-\beta|\gamma^\star|}$$

To bound this sum, we split up the terms according to the length (number of edges) in the contour. Observe that the shortest contour surrounding 0 must have 4 edges.

$$\mu^+_{\Lambda_n}(\sigma_0 = -1) \le \sum_{k\ge 4} \sum_{\gamma^\star : 0\in Int(\gamma^\star), |\gamma^\star| = k} e^{-\beta k}$$

Now we estimate the number of terms in the inner sum.

Exercise: Show that the number of contours of length k with 0 in its interior is at most $\frac{k}{2} \cdot 4 \cdot 3^{k-1}$ (this type of estimate is often referred to as Peierls' method).

We then have

$$\mu^+_{\Lambda_n}(\sigma_0 = -1) \leq \sum_{k \geq 4} \frac{k}{2} \cdot 4 \cdot 3^{k-1} e^{-\beta k}$$

For β large enough, this series converges and we can conclude that $\mu^+_{\Lambda_n}(\sigma_0 = -1)$ is "small" which tells us that $\mathbb{E}[\sigma_0^+] > \delta$ for some $\delta > 0$. But by symmetry, we can run this proof for the all -1 boundary conditions to conclude that $\mathbb{E}[\sigma_0^-] < -\delta$, which tells us that there are two distinct infinite volume Gibbs measures! □

The Peierls method has since been applied in many different contexts. One significant application is in percolation theory, related to the threshold of the existence of an infinite connected component in $\mathbb{Z}^d$. A quite different application is in recent work on the reconstruction of random binary grids. Links to these two examples are in the references section.

1.2 Lecture 1 References

- phase transitions and Peierls argument: Friedli–Velenik, section 3.2.7; `https://www.unige.ch/math/folks/velenik/smbook/`
- Percolation: lecture notes of Perla Sousi, `http://www.statslab.cam.ac.uk/~ps422/percolation.html`; lecture notes of Hugo Duminil-Copin, `https://www.ihes.fr/~duminil/publi/2017percolation.pdf`
- Reconstructing random pictures (Narayanan, Y.): `https://arxiv.org/abs/2210.09410`

1.3 Lecture 1 Exercises

1. Write the hardcore model as a spin model on a graph by determining the Hamiltonian and Gibbs measure.
2. Compute the hard-core partition function for:
   (a) K_d, the complete graph (clique) on d vertices
   (b) $K_{d,d}$, the complete d-regular bipartite graph (two sets of d vertices L and R with all d^2 edges between L and R present and no others)
   (c) C_6, the cycle on six vertices
3. Let $G = G_1 \cup G_2$, the disjoint union of two graphs G_1, G_2. Prove that

$$Z_G(\lambda) = Z_{G_1}(\lambda) Z_{G_2}(\lambda) .$$

4. Prove the following claim we used in the Peierls argument for phase coexistence of the low-temperature Ising model on $\mathbb{Z}^2$: Fix a contour $\gamma^\star$ containing 0 in its interior. Let $\mu^+_{\Lambda_n}$ be the Gibbs measure defined on collections of contours on the $n \times n$ box with all $+$ boundary conditions. Show that

$$\sum_{\Gamma \ni \gamma^\star} \mu^+_{\Lambda_n}(\Gamma) \leq e^{-\beta|\gamma^\star|}$$

5. Show that the number of contours of length k with 0 in its interior is at most $\frac{k}{2} \cdot 4 \cdot 3^{k-1}$ (this type of estimate is often referred to as Peierls' method).

6. Consider percolation on $\mathbb{Z}^2$ with probability p. Let $\mathcal{C}(0)$ denote the connected component containing 0. Define the threshold

$$p_c = \sup\{p \in [0,1] : \mathbb{P}_p(|\mathcal{C}(0)| = \infty) = 0\}$$

Use Peierls' method (namely, the previous exercise) to show that $p_c < 1$.

2 Lecture 2: Extremal Combinatorics and the Hardcore Model

2.1 Lecture 2 Notes

Today we will focus on the hardcore model and applications to extremal combinatorics. The main questions of interest in extremal combinatorics are about the maximum or minimum of certain quantities related to discrete structures, and what structures achieve these extrema. One of the most classical/foundational examples is the following:
Question: What is the maximum number of edges in an n-vertex triangle-free graph?

Theorem 2.1 (Mantel)**.** *The maximum number of edges in an n-vertex triangle-free graph is $\frac{n^2}{4}$. This is achieved by the complete bipartite graph on parts of size $n/2$.*

We are going to discuss extremal questions about independent sets in d-regular graphs. Throughout, fix n, d, and G a d-regular graph. Let $i(G)$ be the number of independent sets in G.

Recall the partition function of the hardcore model is $Z_G(\lambda) = \sum_{J \in \mathcal{I}(G)} \lambda^{|J|}$. This is in some sense a combinatorial encoding of our model; we saw that $Z_G(\lambda)$ is the same as the independence polynomial of G, and we also have $Z_G(1) = i(G)$.
Question: What is the maximum number of independent sets in a d-regular graph?

Theorem 2.2 (Kahn 2001)**.** *If G is bipartite, then*

$$i(G) \leq i(H_{d,n})$$

where $H_{d,n}$ is the graph consisting of $n/2d$ copies of $K_{d,d}$.

Kahn proved this using the entropy method. Observe that

$$\begin{aligned}
& i(G) \leq i(H_{d,n}) \\
\Leftrightarrow\ & Z_G(1) \leq Z_{H_{d,n}}(1) \\
\Leftarrow\ & Z_G(\lambda) \leq Z_{H_{d,n}}(\lambda) = (Z_{K_{d,d}}(\lambda))^{n/2d} \\
\Leftrightarrow\ & \log Z_G(\lambda) \leq \frac{n}{2d} \log Z_{K_{d,d}}(\lambda)
\end{aligned}$$

Theorem 2.3 (Galvin–Tetali 2004). *If G is bipartite, then*

$$\frac{1}{n} \log Z_G(\lambda) \leq \frac{1}{2d} \log Z_{K_{d,d}}(\lambda)$$

for all $\lambda > 0$.

Theorem 2.4 (Zhao 2010). *The above holds for general d-regular graphs G.*

We discuss the proof of a stronger result which uses the context of the hardcore model.

2.1.1 The Occupancy Method

Let $\alpha_G(\lambda) := \frac{1}{n}\mathbb{E}[|I|]$ be the **occupancy fraction** i.e. the expected fraction of vertices contained in an independent set drawn from the hardcore model. Using the definition of expectation, we can rewrite it as

$$\begin{aligned}
\alpha_G(\lambda) &= \frac{1}{n} \sum_{I \in \mathcal{I}(G)} |I| \mu_{G,\lambda}(I) \\
&= \frac{1}{n} \sum_{I \in \mathcal{I}(G)} |I| \frac{\lambda^{|I|}}{Z_G(\lambda)} \\
&= \frac{1}{n} \frac{\lambda}{Z_G(\lambda)} \sum_{I \in \mathcal{I}(G)} |I| \lambda^{|I|-1} \\
&= \frac{1}{n} \frac{\lambda}{Z_G(\lambda)} (Z_G(\lambda))' \\
&= \frac{\lambda}{n} (\log Z_G(\lambda))'
\end{aligned}$$

Thus, the following result implies the previous three theorems:

Theorem 2.5 (Davies–Jenssen–Perkins–Roberts 2015).

$$\alpha_G(\lambda) \leq \alpha_{K_{d,d}}(\lambda)$$

for all $\lambda > 0$ and d-regular G.

Proof Sketch. To simplify things, we will assume that G is triangle-free.

Let's analyze $\alpha_G(\lambda)$, this time using linearity of expectation. Say a vertex is *blocked* if at least one of its neighbors is in I and *unblocked* otherwise. A simple but key observation is that a blocked vertex cannot be added to I. For notational ease, let $\mathbb{P} = \mu_{G,\lambda}$.

$$\begin{aligned}\alpha_G(\lambda) &= \sum_{v \in V} \mathbb{P}(v \in I) \\ &= \sum_{v \in V} \mathbb{P}(v \in I | v \text{ is unblocked}) \cdot \mathbb{P}(v \text{ is unblocked})\end{aligned}$$

where in the last line, we use the Law of Total Probability.

Now we apply a key property of the hardcore model (which is also true of Gibbs measure in general), which is that it is a *Markov random field.* This means that if we partition $V(G)$ as $A \cup B \cup C$ such that A and C have no edges between them, then the spins in A are independent of the spins in C. In our context, this means that if we condition on $N(v) = \{u : uv \in E\}$, the spin at v is independent of the rest of the graph.

Thus, $\mathbb{P}(v \in I | v \text{ is unblocked})$ is simply the probability that $v \in I$ for I chosen from a single-vertex graph. This probability is

$$\mathbb{P}(v \in I | v \text{ is unblocked}) = \frac{\lambda}{1+\lambda}$$

The remaining term to analyze is $\mathbb{P}(v \text{ is unblocked})$. We again use conditioning to rewrite this

$$\begin{aligned}\mathbb{P}(v \text{ is unblocked}) &= \mathbb{P}(v \text{ is unblocked} | v \text{ has } j \text{ unblocked neighbors}) \cdot \mathbb{P}(v \text{ has } j \text{ unblocked neighbors}) \\ &= \frac{1}{(1+\lambda)^j} \mathbb{P}(v \text{ has } j \text{ unblocked neighbors})\end{aligned}$$

The last line is by the same reasoning as before, using the Markov random field property along with our assumption that G is triangle-free. Putting everything together, we get

$$\alpha_G(\lambda) = \sum_{v \in V} \frac{\lambda}{1+\lambda} \frac{1}{(1+\lambda)^j} \mathbb{P}(v \text{ has } j \text{ unblocked neighbors})$$

Pulling the $\frac{\lambda}{1+\lambda}$ term outside of the summation, here is the key observation: the remaining summation is exactly $\mathbb{E}[\frac{1}{(1+\lambda)^Y}]$ where Y is *the random variable that counts the number of unblocked neighbors of a uniformly random vertex.*

We now rewrite $\alpha_G(\lambda)$ in a second way, to obtain constraints for our random variable Y. Indeed, we can use linearity of expectation again but counting from the perspective of neighborhoods to write

$$\begin{aligned}\alpha_G(\lambda) &= \frac{1}{n}\sum_{v\in V}\frac{1}{d}\sum_{u\in N(v)}\mathbb{P}(u\in I)\\ &= \frac{1}{n}\sum_{v\in V}\frac{1}{d}\sum_{u\in N(v)}\frac{\lambda}{1+\lambda}\mathbb{P}(u \text{unblocked})\end{aligned}$$

where the last line is again by conditioning on u being unblocked. Pulling the $\frac{\lambda}{1+\lambda}$ to the outside, we can see that the inner summation is simply $\mathbb{E}[Y]$.

We thus have the relationship

$$\frac{1}{d}\mathbb{E}[Y] = \mathbb{E}[(1+\lambda)^{-Y}$$

The remainder of the proof uses linear programming. We obtain a set of linear constraints on the random variable Y. We then relax the optimization problem: instead of considering just Y arising from graphs, we consider Y arising from any probability distribution, subject to the constraints and $0 \le Y \le d$. Using LP duality, we find a unique maximizer for α, and because this maximizer actually does correspond to a distribution from a graph (specifically $K_{d,d}$), we can conclude that $K_{d,d}$ is the extremal graph. □

We now briefly introduce a different statistical physics method which allows us to prove an even stronger result, on the level of coefficients.

2.1.2 Cluster Expansion

The cluster expansion is a tool for describing "perturbative" models, such as in settings where we have weakly dependent random variables and want to measure deviations from the independent setting. The main idea is to write $\frac{1}{n}\log Z_G(\lambda)$ (also called the *free energy*) as a (hopefully convergent!) series expansion.

Theorem 2.6 (Davies–Jenssen–Perkins 2021)**.** *For n large enough, all d-regular n-vertex graphs G, and all k,*

$$i_k(G) \le i_k(H_{d,n})$$

They also show the analogous statement for matchings, which is more commonly known as the Upper Matching Conjecture.

The proof idea uses the cluster expansion. Roughly, we want to bound $i_k(H_{d,n}) - i_k(G)$. Although $i_k(G)$ is not the partition function of the hardcore model, it turns out we can write it as the partition function of a different model, called a polymer model. We take the cluster expansion and use the combinatorial description to argue that the series is mainly determined by the density of cycles C_g and C_{g-1} where $\text{girth}(G) \ge g-1$.

We did not have time to go into the details, but here are some formal definitions for your reference:

As an example, consider the "ideal" hard-core model where G is simply n singleton vertices. Then $Z_G(\lambda) = (1+\lambda)^n$. In this case, we know from calculus that

$$\frac{1}{n}\log Z_G(\lambda) = \log(1+\lambda) = \lambda - \frac{\lambda^2}{2} + \frac{\lambda^3}{3} - \cdots$$

and that this converges if $|\lambda| < 1$.

We will see how we can write the free energy as a formal power series even when we are not in this ideal setting, and some nice combinatorial conditions for computation and convergence.
Setup: We will define the expansion for the *multivariate hardcore model.*

- $\vec{\lambda} = (\lambda_v)_{v \in V}$ is a tuple of activities assigned to each vertex,
- $Z_G(\vec{\lambda}) = \sum_{I \in \mathcal{I}(G)} \prod_{v \in I} \lambda_v$ (taking $\lambda_v = \lambda$ for all v gives us our original HC model)
- Γ is a tuple of vertices in G (allowing repeats)
- $H(\Gamma)$ is the **incompatibility graph** of Γ. The vertices are those in the tuple Γ (with multiplicities). Place an edge between two copies of the same vertex and between vertices which form an edge in G.
- a **cluster** is a tuple Γ such that $H(\Gamma)$ is connected
- the **Ursell function** is

$$\phi(H) = \frac{1}{|V(H)|!} \sum_{\substack{A \subset E(H), \\ (V(H),A) \text{ connected}}} (-1)^{|A|}$$

The **cluster expansion** is then

$$\log Z_G(\vec{\lambda}) = \sum_{\text{clusters } \Gamma} \phi(H(\Gamma)) \prod_{v \in \Gamma} \lambda_v$$

For the derivation, see Friedli–Velenik, Chapter 5.
Example: Let G be a singleton vertex v. The possible tuples are $(v), (v,v), (v,v,v), \ldots$. If Γ is a k-tuple of v's, then $H(\Gamma) = K_k$ so each such Γ is a cluster.

Exercise: $\phi(K_k) = \frac{1}{k!}(k-1)!(-1)^{k+1}$

Then $\log Z_G(\lambda) = \sum_k \frac{\lambda^k}{k}(-1)^{k+1}$ as we determined previously.

2.2 Lecture 2 References

- a survey on extremal regular graphs for independent sets by Yufei Zhao, `https://yufeizhao.com/research/extremal-regular-graphs.pdf`

- the entropy method: lecture notes of David Galvin, `https://arxiv.org/abs/1406.7872`
- the occupancy method (Davies, Jenssen, Perkins, and Roberts): `https://arxiv.org/abs/1508.04675`
- the cluster expansion: Friedli-Velenik, chapter 5, `https://www.unige.ch/math/folks/velenik/smbook/`
- the Upper Matching Conjecture (Davies, Jenssen, Perkins): `https://arxiv.org/abs/2004.06695`

2.3 Lecture 2 Exercises

1. Prove that the following probability distribution on independent sets of G is the hard-core model on G at fugacity λ. Pick a subset $S \subseteq V(G)$ by including each vertex independently with probability $\frac{\lambda}{1+\lambda}$ and condition on the event that S is an independent set.

2. Let $i_k(G)$ be the number of independent sets of size k in a graph G.

   (a) Give a probabilistic interpretation (as, say, an expectation) for the quantity $\frac{i_{k+1}(G)}{i_k(G)}$ in terms of the uniform distribution over independent sets of size k in G.

   (b) Prove that for all G of maximum degree Δ on n vertices,

   $$\frac{i_{k+1}(G)}{i_k(G)} \geq \frac{n-(\Delta+1)k}{k+1},$$

   and find a family of graphs for which the inequality is tight.

   (c) Use the above to prove that for all G of maximum degree Δ on n vertices,

   $$\frac{1}{n}\log Z_G(\lambda) \geq \frac{1}{\Delta+1}\log(1+(\Delta+1)\lambda),$$

   and show that the inequality is tight. (Hint: recall that partition functions are multiplicative over disjoint graphs and that $Z_G(\lambda)$ is a polynomial).

3. *If you read about the definition of cluster expansion from the notes:* Write the cluster expansion for $\frac{1}{n}\log Z_G(\lambda)$ (where $Z_G(\lambda)$ is the hard-core partition function) when G is

   (a) a single vertex (show that the Ursell function is what we claimed in the notes)

   (b) a single edge

   (c) a Δ-regular triangle-free graph (compute clusters up to size 3)

3 Lecture 3: Algorithms and the Potts Model

3.1 Lecture 3 Notes

From yesterday, recall that the cluster expansion is a series expansion for $\log Z_G(\lambda)$. We will not utilize the specific definition of the expansion, but simply keep in mind that the cluster expansion is a combinatorial quantity, in that the coefficients are computed according to some connected subgraph structures (clusters and incompatibility graphs).

It turns out that we can define the cluster expansion not only for the (multivariate) hardcore model but also for more general models called "abstract polymer models." As the name suggests, these are more abstract in some sense but the advantage is that we can often rewrite other Gibbs measures in the framework of an abstract polymer model, and this allows us to take advantage of the cluster expansion tools.

3.1.1 Algorithms

Two main algorithmic questions associated with spin models are about approximate counting and sampling. The *counting* problem refers to computing the partition function Z_G. The *sampling* problem refers to producing samples from the Gibbs distribution μ_G. Exact computation of Z_G is #P-hard in general, so we are mainly concerned with approximate counting and sampling.

Given $\delta > 0$, a **δ-relative approximation** of Z is some $\hat{Z}$ such that $e^{-\delta}\hat{Z} \leq Z \leq e^{\delta}\hat{Z}$. An **FPTAS** (a fully polynomial-time approximation scheme) for computing a partition function Z_G is an algorithm whose output is a δ-relative approximation of Z_G and that runs in time polynomial in $\frac{1}{\delta}, n$.

Given $\delta > 0$, a **polynomial-time sampling scheme** for μ_G is a randomized algorithm whose output is a spin configuration σ according to distribution $\hat{\mu}$ such that $\|\mu_G - \hat{\mu}\|_{TV} < \delta$ and that runs in time polynomial in $\frac{1}{\delta}, n$.

For many models that we are interested in (including all three covered in these notes), the problems of approximate counting and sampling are equivalent—that is, given an algorithm for one, you can produce an algorithm for the other.

Exercise: Prove that approximate counting and approximate sampling are equivalent for the hardcore model.

There are three general classes of spin models with respect to efficient approximation algorithms:

1. there exists an efficient algorithm for all graphs. The main question here is to find deterministic algorithms, since many of the known algorithms are randomized (e.g. the ferromagnetic Ising model, the monomer-dimer model)

2. the approximate counting problem is NP-hard in general but there exist subclasses of graphs or regimes of temperatures where efficient algorithms exist (e.g. hardcore model—see reference for Weitz–Sly)

3. not known to be NP-hard but also no efficient algorithms (random or deterministic) are known in general (e.g. ferromagnetic Potts)

The last category includes the complexity class #BIS, which stands for Bipartite Independent Set–the problem of approximating the number of independent sets in a bipartite graph. The complexity of #BIS is open.

There are several avenues to producing efficient algorithms. One which we will not discuss here involves Markov chains, primarily two called the Glabuer dynamics and the Swendsen–Wang dynamics. There has been a lot of work determining fast or slow mixing of these Markov chains on different classes of graphs or regimes of temperatures. Another avenue comes from Weitz and is called the method of "correlation decay."

We will focus on the method of cluster expansion, which came about quite recently with the following landmark result (which I will not state in full detail here):

Theorem 3.1 (Helmuth–Perkins–Regts 2019)**.** *Under some mild conditions, a convergent cluster expansion for a partition function Z_G gives an FPTAS for Z_G.*

The idea is to essentially truncate the series expansion and compute the coefficients in polynomial time. This requires that, for example, the runtime of computing $\log(t)$ many coefficients scales appropriately with t.

We will apply the above theorem as a black box, but this will motivate our study of the Potts model—our goal will be to put the Potts in the framework of an abstract polymer model, and then to show that the cluster expansion converges.

3.1.2 The Potts Model on Expanders

The following will be the main theorem of this section: we say a graph G is an η-expander if for every $A \subset V(G)$ such that $|A| \leq \frac{n}{2}$, we have $|\nabla(A)| \geq \eta|A|$.

Theorem 3.2 (Carlson–Davies–Fraiman–Kolla–Potukuchi–Y. 2022)**.** *For all $\epsilon > 0$, there exists d large enough and an absolute constant C such that for $q \geq d^C$ and for*

- $\beta \leq (1-\epsilon)\beta_0$ *(high-temperature)*
- $\beta \geq (1+\epsilon)\beta_0$ *(low-temperature)*

there exists an FPTAS for the q-color Potts model on d-regular 2-expander graphs.

For comparison to previous results, Helmuth–Jenssen–Perkins also show the existence of an FPTAS for the Potts model but requiring q exponential in d and $\Omega(d)$ expansion of G. However, they obtain algorithms at all temperatures, rather than the gap we have around β_0. Their results are also obtained in the greater generality of the random cluster model.

We will focus on the low-temperature proof, which requires some combinatorial tools.

Low Temperature Proof

Throughout, let G be a d-regular 2-expander. Recall that at low temperatures, we expect the most likely colorings to be the all-monochromatic ones, or small deviations from these. We will show that this is reflected in the partition function.

For $1 \leq i \leq q$, let $S_i = \{\sigma : |\sigma^{-1}(i)| \geq n/2\}$ (the colorings with "majority color" i), and let S_0 be the remaining colorings. We argue that

$$\frac{\sum_{\sigma \in S_0} e^{\beta M(G,\sigma)}}{Z_G(q,\beta)}$$

is small (in fact, less than $q^{-\Omega(\epsilon n/d)}$). Recall that $B(G,\sigma)$ is the number of bichromatic (i.e. non-monochromatic) edges of σ. We break up the sum in the numerator as

$$\sum_{k \geq 1} \sum_{\sigma \in S_0 : B(G,\sigma) = k} e^{\beta M(G,\sigma)}$$

We can bound the number of terms in the inner sum using the following lemma:

Lemma 3.3. *The number of q-colorings with exactly k bichromatic edges is at most $\binom{n}{2k/d} q^{2k/d}$.*

Proof. The proof is an adaptation of Karger's randomized algorithm for computing min-cut of a graph. We sketch the proof here and leave verification of the details to the exercises.

Fix σ with k bichromatic edges. We produce a random coloring σ' and analyze the probability that it is equal to σ:

1. Choose an edge of G uniformly at random and contract it. Delete self-loops. Repeat this step until $|V| \leq \frac{2k}{d}$.

2. Color the remaining vertices by assigning each vertex a color independently and uniformly at random.

Observe that the coloring of the contracted graph maps back to a coloring of the original graph. Call this coloring σ'. To get a lower bound on $\mathbb{P}(\sigma' = \sigma)$, observe that we need the k bichromatic edges to remain uncontracted. The probability of this is at least

$$\left(1 - \frac{k}{nd/2}\right)\left(1 - \frac{k}{(n-1)d/2}\right) \cdots \left(\frac{d}{2k}\right)$$

Here we use our expansion assumption (and actually, we only need a much weaker condition about the size of the min-cut). We also need the vertices to receive the correct colors at the end, and this has probability at least $q^{-2k/d}$. The result follows. □

Plugging in the above bound, we are able to show a sufficiently small upper bound on the weight of colorings in S_0. We omit the details here.

Abstract Polymer Model:

We now put the low-temperature Potts model in the framework of something called an *abstract polymer model.* This will allow us to apply the Helmuth–Perkins–Regts result, assuming we can show that the cluster expansion converges.

To define an abstract polymer model, we need three components: polymers, a compatibility relation, and weights. Our model will capture the "deviations" of colorings from the ground states, which are the q all-monochromatic colorings.

- polymer γ: a connected subgraph on at most $n/2$ vertices (representing the components that do not receive the majority color)
- compatibility: $\gamma \sim \gamma'$ iff $dist_G(\gamma, \gamma') \geq 2$ (i.e. when γ and γ' can be separate components in the same coloring)
- weight: $w_\gamma = e^{-\beta(e_\gamma + \nabla_\gamma)} Z_\gamma(q-1, \beta)$ (the contribution of γ to the total weight of a coloring, where e_γ is the number of edges induced by γ and ∇_γ is the number of boundary edges)

From this, we define the polymer model partition function as

$$\Xi = \sum_{\Gamma \text{ compatible polymers}} \prod_{\gamma \in \Gamma} w_\gamma$$

Exercise: Show that $Z_G(q, \beta) = q^n \Xi$.

We will do cluster expansion on Ξ. We prove that the cluster expansion converges using the following:

Theorem 3.4 (Kotecký–Preiss)**.** *If there exist $f, g : \{polymers\} \to [0, \infty)$ such that for all polymers γ, we have*

$$\sum_{\gamma' \not\sim \gamma} w_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq f(\gamma)$$

then the cluster expansion for Ξ converges.

We again try to analyze the size of the summation. Given a polymer γ, by our definition, $\gamma' \not\sim \gamma$ iff $dist(\gamma, \gamma') \in \{0, 1\}$. In either case, γ' must contain a vertex in $N_G[\gamma] = N_G(\gamma) \cup \gamma$, (the closed neighborhood of γ). Thus, we can rewrite the KP condition sum as

$$\sum_{\gamma' \not\sim \gamma} \text{``} = \sum_{u \in N_G[\gamma]} \sum_{\gamma' \ni u} \text{``} = \sum_{u \in N_G[\gamma]} \sum_{b \geq d} \sum_{\gamma' \ni u, |\nabla(\gamma')| = b} \text{``}$$

where in the last term, we split up the polymers γ' by their boundary size. The reason we do this is to use the following lemma:

Lemma 3.5. *Let G be a d-regular η-expander and let $x \in V$. The number of $A \subset V$ such that A is connected, $|A| \leq \frac{n}{2}$, $|\nabla(A)| = b$, and $x \in A$ is at most $d^{O(1+1/\eta)b/d}$*

We call this a *container-like lemma* since the proof comes from identifying a "container" A_0 associated with each set A, and computing bounds on the number of A corresponding to each A_0 as well as the number of containers A_0.

3.2 Lecture 3 References

- algorithms from the cluster expansion (Helmuth, Perkins, Regts): `https://arxiv.org/abs/1806.11548`
- the Potts model on expanders (Carlson, Davies, Fraiman, Kolla, Potukuchi, Y.): `https://arxiv.org/abs/2204.01923`
- random cluster model on random regular graphs (Helmuth, Jenssen, Perkins): `https://arxiv.org/abs/2006.11580`
- Karger's algorithm: lecture notes of Eric Vigoda, `https://faculty.cc.gatech.edu/~vigoda/7530-Spring10/Kargers-MinCut.pdf`

3.3 Lecture 3 Exercises

1. (a) Fill in the details for the Coloring Lemma proof (mainly, verify the computations).

   (b) Adapt the coloring lemma proof to produce an algorithm for finding a min-cut of a graph (i.e. a minimum-size set of edges whose removal disconnects the graph). This is *Karger's algorithm.*

2. Given a graph G, consider the *random cluster model*, a distribution on $\{0,1\}^{E(G)}$ with parameters $q, \beta \geq 0$ whose Gibbs measure and partition function are defined by

$$\mu_G(A) \propto q^{c(A)}(e^\beta - 1)^{|A|}$$

$$Z_G(q, \beta) = \sum_{A \subset E} q^{c(A)}(e^\beta - 1)^{|A|}$$

   $c(A)$ denotes the number of connected components of (V, A). (Observe that by setting $p := 1 - e^{-\beta}$, this is related to bond percolation with probability p.)

   (a) Conjecture a relationship between the random cluster partition function and the q-color Potts model partition function.

   (b) Prove your conjecture holds by demonstrating a coupling between the Potts and random cluster distributions for integer q (meaning, describe a method for obtaining a random cluster configuration from a Potts configuration with the correct probability, and vice versa)

3. Consider the hard-core model on a family of graphs that is closed under taking subgraphs.

   (a) Define a process to choose an independent set at random by considering the vertices of G one at a time, in order, and use this to show that efficient approximate counting on the family implies approximate sampling.

(b) Write $\mu_G(\emptyset)$ in two ways to show that approximate sampling implies approximate counting.

4 Solutions to Exercises

4.1 Lecture 1 Solutions

1. *Write the hardcore model as a spin model on a graph by determining the Hamiltonian and Gibbs measure.*

   Solution: Let $\Omega = \{0, 1\}$. Let $f(\sigma_v) = -(\log \lambda)\sigma_v$ and let $g(\sigma_u, \sigma_v)$ be ∞ if $\sigma_u = \sigma_v = 1$ and 0 else. Then if σ is a configuration where $\{v : \sigma_v = 1\}$ is not independent, then there exist some $uv \in E$ such that $g(\sigma_u, \sigma_v) = \infty$ so $H(\sigma) = +\infty$ and $\mu_{G,\lambda}(\sigma) \propto 0$. Else, if $I := \{v : \sigma_v = 1\}$ is independent, then $H(\sigma) = |I|(-\log \lambda)$. Taking $\beta = 1$, we get $\mu_{G,\lambda}(\sigma) \propto e^{-\beta|I|(-\log\lambda)} = \lambda^{|I|}$ as desired.

2. *Compute the hard-core partition function for:*

   (a) K_d, the complete graph (clique) on d vertices

   (b) $K_{d,d}$, the complete d-regular bipartite graph (two sets of d vertices L and R with all d^2 edges between L and R present and no others)

   (c) C_6, the cycle on six vertices

   Solution: For (a), observe that the only independent sets are the empty one and singletons, so $Z_{K_d}(\lambda) = 1 + d\lambda$.

   For (b), the independent sets consist of all subsets of L and all subsets of R. Each side has 2^d independent sets, but adding these double-counts the empty set, so $Z_{K_{d,d}}(\lambda) = 2\sum_{k=0}^{d}\binom{d}{k}\lambda^k - 1$.

   For (c), $Z_{C_6}(\lambda) = 1 + 6\lambda + 9\lambda^2 + 2\lambda^3$.

3. *Problem: Let $G = G_1 \cup G_2$, the disjoint union of two graphs G_1, G_2. Prove that*

   $$Z_G(\lambda) = Z_{G_1}(\lambda) Z_{G_2}(\lambda)\,.$$

   Solution: Observe that $I \subset V(G)$ is independent if and only if it can be written as $I_1 \cup I_2$ where $I_1 \in \mathcal{I}(G_1), I_2 \in \mathcal{I}(G_2)$. The statement follows.

4. *Prove the following claim we used in the Peierls argument for phase coexistence of the low-temperature Ising model on $\mathbb{Z}^2$: Fix a contour $\gamma^\star$ containing 0 in its interior. Let $\mu^+_{\Lambda_n}$ be the Gibbs measure defined on collections of contours on the $n \times n$ box with all + boundary conditions. Show that*

   $$\sum_{\Gamma \ni \gamma^\star} \mu^+_{\Lambda_n}(\Gamma) \le e^{-\beta|\gamma^\star|}$$

Solution: The notes give a hint to pull out $e^{-\beta|\gamma^\star|}$ from the numerator, which gives us

$$\mu^+_{\Lambda_n}(\Gamma) = \frac{\prod_{\gamma\in\Gamma} e^{-\beta|\gamma|}}{\sum_{\Gamma'}\prod_{\gamma\in\Gamma'} e^{-\beta|\gamma|}} = e^{-\beta|\gamma^\star|}\frac{\sum_{\Gamma\ni\gamma^\star}\prod_{\gamma\in\Gamma\setminus\gamma^\star} e^{-\beta|\gamma|}}{\sum_{\Gamma'}\prod_{\gamma\in\Gamma'} e^{-\beta|\gamma|}}$$

It is enough to show that the fraction above is at most 1. We can do this by defining an injection from terms in the numerator to terms in the denominator. For each term counted in the numerator i.e. $\Gamma \ni \gamma^\star$, produce a contour collection Γ' by removing $\gamma^\star$ and reversing all the spins in $Int(\gamma^\star)$. This gives an injective correspondence to the terms counted by the denominator.

5. *Show that the number of contours of length k with 0 in its interior is at most $\frac{k}{2}\cdot 4\cdot 3^{k-1}$ (this type of estimate is often referred to as Peierls' method).*

   Solution: Let γ be a contour of length k containing 0. Since γ consists of edges in the dual lattice, there must be some point $(x-1/2, 1/2)$ contained in γ where $x \in \{1, \ldots, [\frac{k}{2}]\}$. Let this be the "starting point" of γ. We then count the number of possible γ by counting the number of ways to "walk" along γ given the starting point. There are 4 directions to proceed in from the first step, and for each subsequent step, there are 3 possible directions, giving 3^{k-1}.

6. *Consider bond percolation on $\mathbb{Z}^2$ with probability p (include each edge independently with probability p; the included edges are sometimes called* open *edges). Let $\mathcal{C}(0)$ denote the connected component containing 0. Define the threshold*

$$p_c = \sup\{p \in [0,1] : \mathbb{P}_p(|\mathcal{C}(0)| = \infty) = 0\}$$

   Use Peierls' method (namely, the previous exercise) to show that $p_c < 1$.

   Solution: See Perla Sousi's lecture notes (linked in the references) for this proof as well as more discussion. The idea is to perform bond percolation on the dual lattice where a dual edge is open iff the edge it crosses from $\mathbb{Z}^2$ is open. Observe that $|\mathcal{C}(0)| < \infty$ if and only if 0 is contained in a cycle of the dual (the "contours" in this context). Now use a first moment bound on the probability that $|\mathcal{C}(0)| < \infty$ and apply an estimate involving the number of dual cycles of length n containing the origin. The conclusion is that for p large enough—say $p > 1-\delta$, then $\mathbb{P}(|\mathcal{C}(0)| < \infty) < 1-\epsilon$ so $p_c \leq 1-\delta$.

4.2 Lecture 2 Solutions

1. *Prove that the following probability distribution on independent sets of G is the hard-core model on G at fugacity λ. Pick a subset $S \subseteq V(G)$ by including each vertex independently with probability $\frac{\lambda}{1+\lambda}$ and condition on the event that S is an independent set.*

Solution: For $S \subset V(G)$, the probability of choosing S in this distribution is

$$\mathbb{P}[S] = \left(\frac{\lambda}{1+\lambda}\right)^{|S|} \left(\frac{1}{1+\lambda}\right)^{n-|S|}$$

The probability of choosing an independent set in this distribution is

$$\sum_{I \in \mathcal{I}(G)} \left(\frac{\lambda}{1+\lambda}\right)^{|I|} \left(\frac{1}{1+\lambda}\right)^{n-|I|} = \frac{Z_G(\lambda)}{(1+\lambda)^n}$$

So the probability of choosing a fixed independent set I is

$$\mathbb{P}[I] = \left(\frac{\lambda}{1+\lambda}\right)^{|I|} \left(\frac{1}{1+\lambda}\right)^{n-|I|} \cdot \frac{(1+\lambda)^n}{Z_G(\lambda)} = \frac{\lambda^{|I|}}{Z_G(\lambda)}$$

as desired.

2. *Let $i_k(G)$ be the number of independent sets of size k in a graph G.*
   (a) *Give a probabilistic interpretation (as, say, an expectation) for the quantity $\frac{i_{k+1}(G)}{i_k(G)}$ in terms of the uniform distribution over independent sets of size k in G.*
   (b) *Prove that for all G of maximum degree Δ on n vertices,*

   $$\frac{i_{k+1}(G)}{i_k(G)} \geq \frac{n-(\Delta+1)k}{k+1},$$

   and find a family of graphs for which the inequality is tight.
   (c) *Use the above to prove that for all G of maximum degree Δ on n vertices,*

   $$\frac{1}{n}\log Z_G(\lambda) \geq \frac{1}{\Delta+1}\log(1+(\Delta+1)\lambda),$$

   and show that the inequality is tight. (Hint: recall that partition functions are multiplicative over disjoint graphs and that $Z_G(\lambda)$ is a polynomial).

Solution: For (a): Let $\mathcal{I}_k(G)$ be the collection of independent sets of size k in G. Define a random variable $X : \mathcal{I}_k(G) \to \mathbb{R}$ where $X(I)$ is the number of independent sets of size $k+1$ containing I. Then

$$\mathbb{E}X = \sum_{I \in \mathcal{I}_k(G)} X(I)Pr(I) = \sum_I X(I)\frac{1}{i_k(G)}$$

Observe that each element of $\mathcal{I}_{k+1}(G)$ contains exactly $k+1$ independent k-sets, so $\sum_I X(I)$ counts each element of $I_{k+1}(G)$ exactly $k+1$ times. Thus,

$$\mathbb{E}X = \frac{(k+1)i_{k+1}(G)}{i_k(G)}$$

For (b): For $I \in \mathcal{I}_k(G)$, let $N(I) = \bigcup_{v \in I} N(v)$. Then

$$|N(I)| \le \sum_{v \in I} |N(v)| \le k\Delta$$

Every vertex in $V \setminus (I \cup N(I))$ can be added to I to make an independent $(k+1)$-set. Thus, the number of independent $(k+1)$-sets containing I is

$$X(I) = n - (k + |N(I)|) \ge n - (\Delta + 1)k$$

This implies $\mathbb{E}X \ge n - (\Delta + 1)k$, and by part (a), we have

$$\frac{i_{k+1}(G)}{i_k(G)} \ge \frac{n - (\Delta + 1)k}{k + 1}$$

This inequality is tight if $|N(I)| = \sum_{v \in I} |N(v)| = k\Delta$ for all $I \in \mathcal{I}_k(G)$, which is achieved when $N(I) = \sqcup_{v \in I} N(v)$ and $|N(v)| = \Delta$ for all v. This holds when G is the disjoint union of $\frac{n}{\Delta+1}$ copies of $K_{\Delta+1}$.

For (c): Let $H = H(\Delta, n)$ be the disjoint union of $\Delta+1$ copies of G, and let $K = K(\Delta, n)$ be the disjoint union of n copies of $K_{\Delta+1}$. Since partition functions are multiplicative over disjoint graphs, we have

$$Z_H(\lambda) = (Z_G(\lambda))^{\Delta+1}, \quad Z_K(\lambda) = (Z_{K_{\Delta+1}}(\lambda))^n$$

From Lecture 1 Exercise 2(a), we know that $Z_{K_{\Delta+1}}(\lambda) = 1 + (\Delta + 1)\lambda$, so the claim is equivalent to showing $Z_H(\lambda) \ge Z_K(\lambda)$ for all λ. In fact, we will show a stronger statement—that the inequality holds on the level of coefficients.

For $k = 1$, we have $i_1(H) = i_1(K) = n(\Delta+1)$. For $k \ge 2$, $i_k(K) = \binom{n}{k}(\Delta+1)^k$, whereas we can inductively/recursively apply the bound from part (b) to get

$$\begin{aligned}
i_k(H) &\ge \frac{n(\Delta+1) - (k-1)(\Delta+1)}{k} i_{k-1}(H) \\
&\ge \cdots \ge \frac{(n-k+1)(\Delta+1)}{k} \cdot \frac{(n-k+2)(\Delta+1)}{k-1} \cdots \frac{n(\Delta+1)}{1} \\
&= \binom{n}{k}(\Delta+1)^k
\end{aligned}$$

Thus, $i_k(H) \ge i_k(K)$ for all k, which implies $Z_H(\lambda) \ge Z_K(\lambda)$, proving the claim. The bound is tight since we can take $G = K_{\Delta+1}$ to get $H = K$.

4.3 Lecture 3 Solutions

1. *(a) Fill in the details for the Coloring Lemma proof (mainly, verify the computations).*

(b) *Adapt the coloring lemma proof to produce an algorithm for finding a min-cut of a graph (i.e. a minimum-size set of edges whose removal disconnects the graph). This is* Karger's algorithm.

Solution: For (a), the main observation is that the mincut size of G is at least d (by d-regularity and 2-expansion) and this does not decrease with edge contractions. So at each step of the edge contractions, if the graph has m vertices, it must have at least $md/2$ edges.

For (b), see notes linked in the references.

2. *Given a graph G, consider the* random cluster model*, a distribution on $\{0,1\}^{E(G)}$ with parameters $q, \beta \geq 0$ whose Gibbs measure is defined by*

$$\mu_G(A) \propto q^{c(A)}(e^\beta - 1)^{|A|}$$

$$Z_G(q,\beta) = \sum_{A \subset E} q^{c(A)}(e^\beta - 1)^{|A|}$$

$c(A)$ denotes the number of connected components of (V, A). (Observe that by setting $p := 1 - e^{-\beta}$, this is related to bond percolation with probability p.)

(a) *Conjecture a relationship between the random cluster partition function and the q-color Potts model partition function.*

(b) *Prove your conjecture by demonstrating a coupling between the Potts and random cluster distributions for integer q (meaning, describe a method for obtaining a random cluster configuration from a Potts configuration with the correct probability, and vice versa)*

Solution: Let Z_G^{RC} be the random cluster partition function and Z_G^P be the Potts partition function. The coupling can be described as follows: Given a random cluster configuration A, produce a q-coloring σ_A by coloring each connected component with one of the q colors uniformly at random. The distribution of this coloring is $\mu_G^{RC}(A)q^{-c(A)} \propto (e^\beta - 1)^{|A|}$

Given a Potts configuration $\sigma : V \to [q]$, we produce a random cluster configuration as follows: for each monochromatic component, do percolation on that component with probability $p = 1 - e^{-\beta}$.

3. *Consider the hard-core model on a family of graphs that is closed under taking subgraphs.*

(a) *Define a process to choose an independent set at random by considering the vertices of G one at a time, in order, and use this to show that efficient approximate counting on the family implies approximate sampling.*

(b) *Write $\mu_G(\emptyset)$ in two ways to show that approximate sampling implies approximate counting.*

Solution: For part (a): Order $V(G)$ as $v_1, v_2, \ldots, v_n$. For each i, we flip an independent coin to decide if v_i is in I. Let I_j be the partial independent set after considering $v_1, \ldots, v_j$.

Let $N[v]$ denote the closed neighborhood of v. We add v_1 with probability

$$\mathbb{P}(v_1 \in I) = \frac{\lambda \cdot Z_{G-N[v_1]}(\lambda)}{Z_G(\lambda)}$$

And in general we add v_k with probability

$$\mathbb{P}(v_k \in I | I \cap \{v_1, \ldots, v_{k-1}\} = I_{k-1}) = \frac{\lambda \cdot Z_{G-\{v_1,\ldots,v_{k-1}-N[I_{k-1}]-N[v_k]}}{Z_{G-\{v_1,\ldots,v_{k-1}\}-N[I_{k-1}]}}$$

Using the FPTAS for computing Z_G (and its subgraphs), we can calculate an $\frac{\epsilon}{n}$ approximation for each of the conditional probabilities so that our final error is ϵ.

For part (b), observe that $\mu_G(\emptyset) = \frac{1}{Z_G(\lambda)}$ so to obtain an ϵ-relative approx for Z, it's enough to find one for $\mu_G(\emptyset)$. We also have

$$\mu_G(\emptyset) = \prod_{i=1}^{n} \mathbb{P}_{G-\{v_1,\ldots,v_{i-1}\}}(v_i \notin I)$$

Each of the marginals can be approximated by repeated sampling from the hard-core model, which we assume we can do efficiently.