Counting independent sets with the cluster expansion

Will Perkins
UIC

joint w/ Matthew Jenssen (Oxford)

How can we use statistical physics tools and intuition in combinatorics?

The cluster expansion is a useful and versatile tool for enumeration

Statistical physics

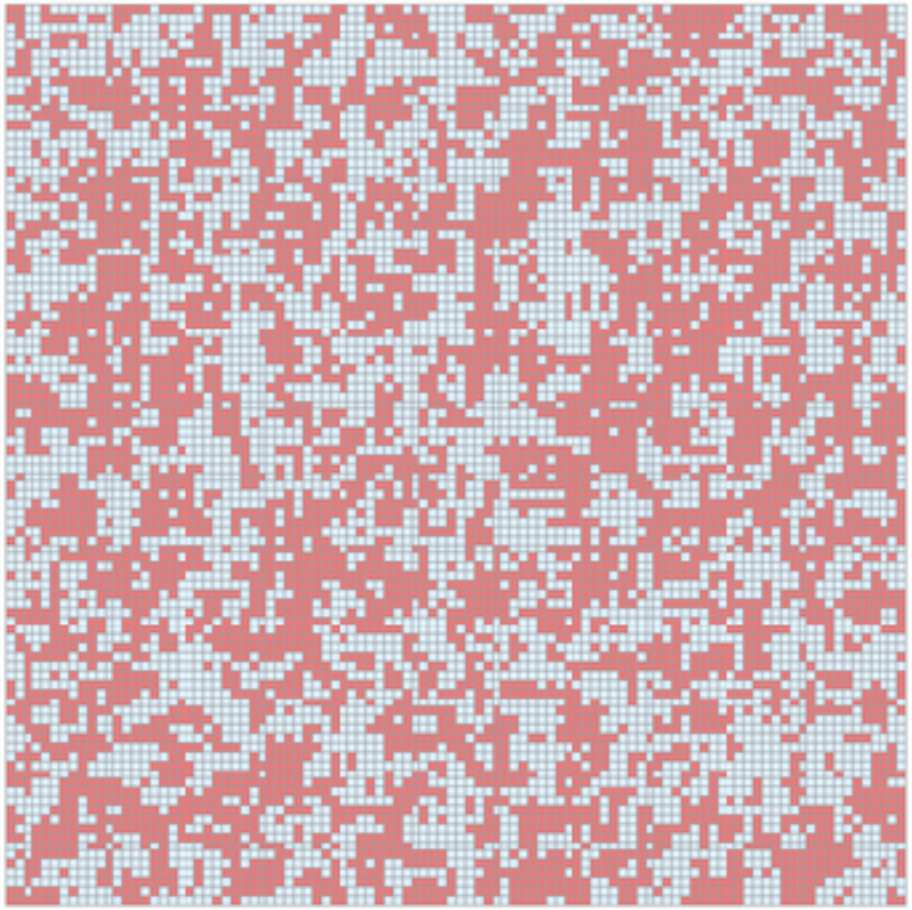

Model a physical material (a gas, magnet) as a random spin configuration on a graph

Probability distribution determined by local interactions.

Statistical physics

Ising model: +1, −1 spins assigned to vertices:

$$\mu(\sigma) = \frac{e^{\beta \sum_{(u,v)\in E} \sigma_u \sigma_v}}{Z_G(\beta)}$$

$$Z_G(\beta) = \sum_{\sigma} e^{\beta \sum_{(u,v)\in E} \sigma_u \sigma_v}$$

Energy, inverse temperature, partition function

High and low temperatures

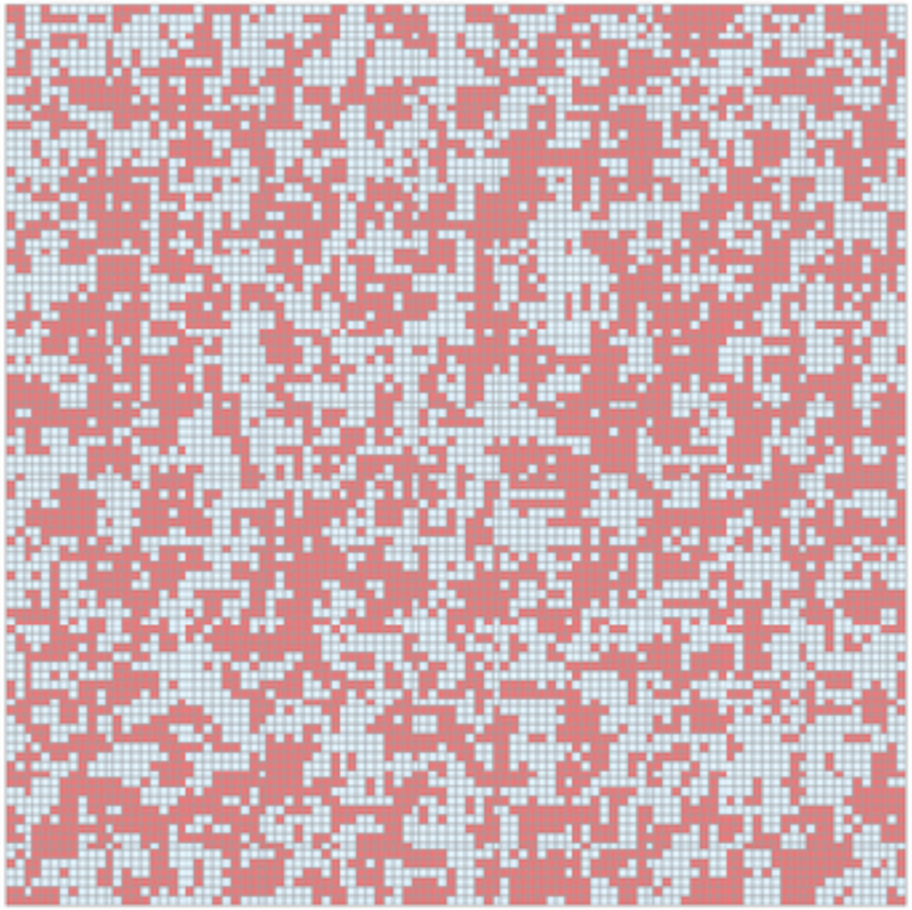

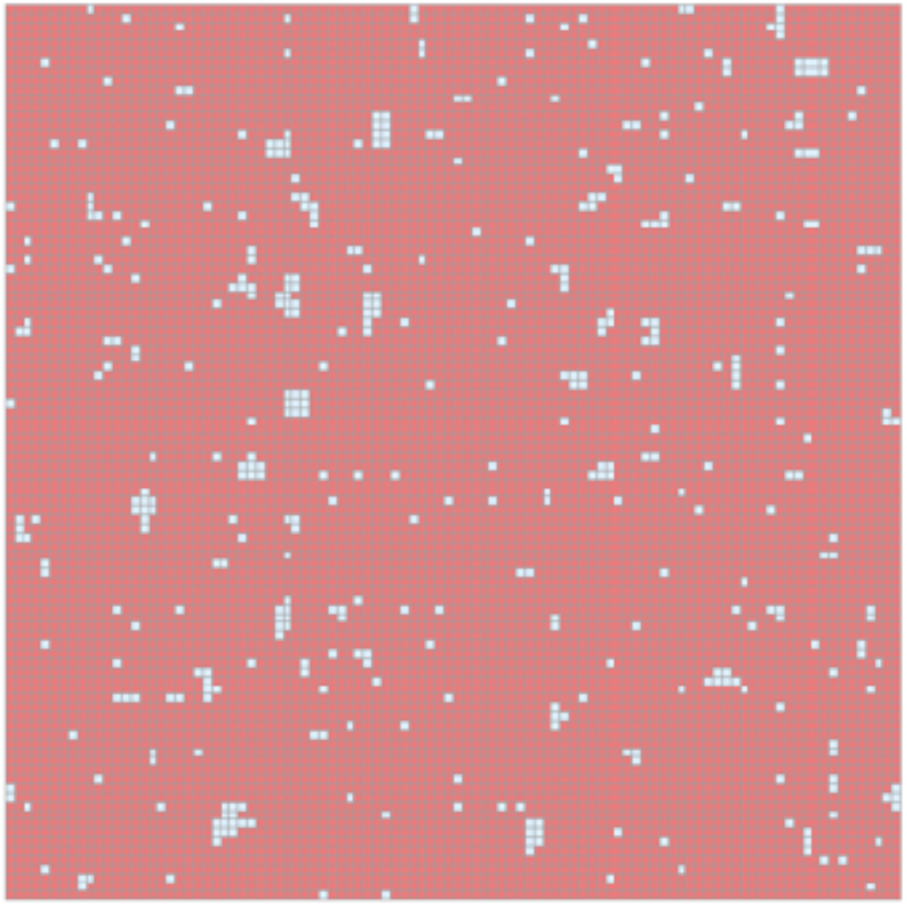

High temperature

- Weak interactions
- Decay of correlations
- Disorder
- Rapid mixing

Low temperature

- Strong interactions
- Strong correlations
- Long-range order
- Slow mixing

Hard-core model

Hard-core model: weighted independent sets

$$\mu(I) = \frac{\lambda^{|I|}}{Z_G(\lambda)}$$

$$Z_G(\lambda) = \sum_I \lambda^{|I|}$$

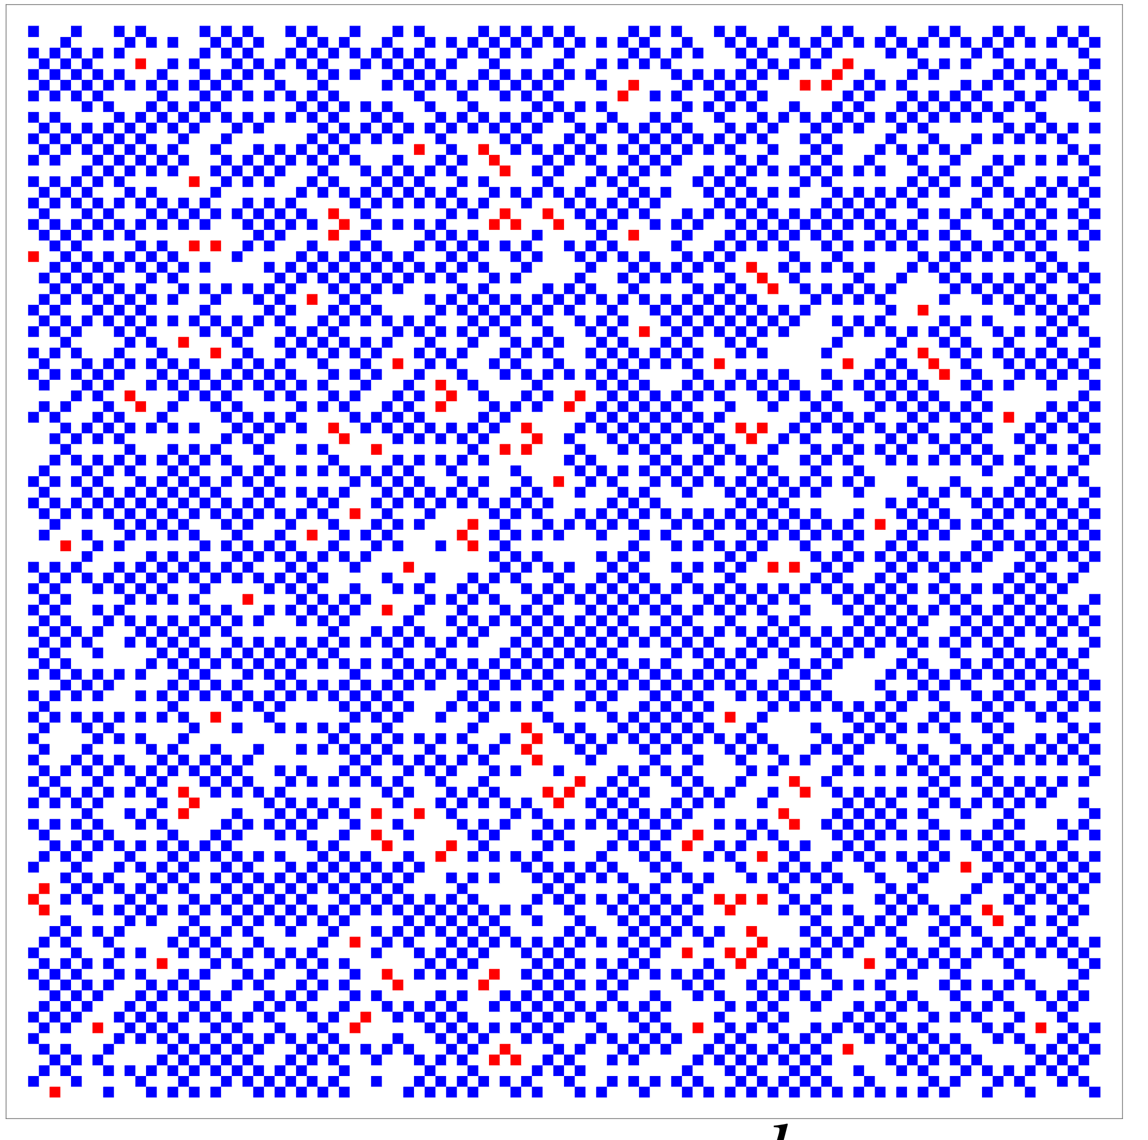

Large λ is akin to low temperature. On $\mathbb{Z}^d$ (bipartite) there are two ground states (even/odd)

Statistical physics perspective

1. Computing (approximately) Z_G is the central task

Weighted counting

2. Understand **correlations** between spins

Properties of a `typical' object

3. Understand changes in correlations as λ, β changes

Extra parameter can help

Combinatorics at low temperature

Combinatorics

- Find an extremal object satisfying a property
- Count the number of such objects
- Describe a typical such object

Physics

- Find a ground state
- Compute the partition function
- Describe the Gibbs measure

Independent sets in the hypercube

Let Q_d be the discrete hypercube and consider its independent sets.

Theorem 1 (Korshunov-Sapozhenko, 1983)

$$i(Q_d) = Z_{Q_d}(1) = (2 + o(1))\sqrt{e}2^{2^{d-1}}$$

Sapozhenko's later simplification of the proof has proved very influential: an example of the method of graph containers

Independent sets in the hypercube

Theorem 2. (Galvin, 2011)

For $\lambda > \sqrt{2} - 1$,

$$Z(\lambda) = (2 + o(1)) \cdot \exp\left[\frac{\lambda}{2}\left(\frac{2}{1+\lambda}\right)^d\right](1+\lambda)^{2^{d-1}}$$

For $\lambda = \Omega\left(\log d \cdot d^{-1/3}\right)$, asymptotics of log Z

Independent sets in the hypercube

Theorem 3. (Jenssen, P.)

1) Asymptotics of $Z_{Q_d}(\lambda)$ for all fixed λ. E.g

$$Z = (2 + o(1)) \cdot \exp\left[\frac{\lambda}{2}\left(\frac{2}{1+\lambda}\right)^d\left(1 + \frac{(2\lambda^2 + \lambda^3)d(d-1) - 2}{4(1+\lambda)^d}\right)\right](1+\lambda)^{2^{d-1}} \text{ for } \lambda > 2^{1/3} - 1$$

2) Expansion to arbitrary precision, e.g.

$$i(Q_d) = 2\sqrt{e} \cdot 2^{2^{d-1}}\left(1 + \frac{3d^2 - 3d - 2}{8 \cdot 2^d} + \frac{243d^4 - 646d^3 - 33d^2 + 436d + 76}{384 \cdot 2^{2d}} + O\left(d^6 \cdot 2^{-3d}\right)\right)$$

3) Essentially **complete probabilistic understanding** of the structure typical (weighted) independent sets (Poisson limits, CLT's, large deviations, correlation decay...)

Perturbative tools

Idea: precisely measure deviations from a simple, known distribution or configuration.

High temperature: expand around iid spins

Low temperature: expand around the constant configurations

Mathematically: control an infinite series expansion of the log partition function: the cluster expansion

Cluster expansion

Infinite series representation for $\log Z$

- How to compute the terms?
- When does the series converge?

Setting: multivariate hard-core model

$$Z_G(\vec{\lambda}) = \sum_{I} \prod_{v \in I} \lambda_v$$

Series for $\log Z_G(\vec{\lambda})$

Cluster expansion

Cluster: multiset of vertices that induce a connected subgraph in G

Ursell function:

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

Example:

Cluster expansion

Formally,

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

Example:

Convergence?

When does the cluster expansion converge?

Classic topic in statistical physics (Mayer, Penrose,...)

Kotecky-Preiss

Given $a : V \to [0,\infty)$, suppose for all v,

$$\sum_{u \in \{v\} \cup N(v)} e^{a(u)} |\lambda_u| \leq a(v).$$

Then the cluster expansion converges absolutely, and

$$\sum_{\Gamma \ni v} \left| \phi(H(\Gamma)) \prod_{u \in \Gamma} \lambda_u \right| \leq a(v).$$

Convergence?

Example: G has max degree Δ and $|\lambda_v| \leq \lambda$

Combinatorics at low temperature

We can use the cluster expansion to account for **deviations** from a ground state.

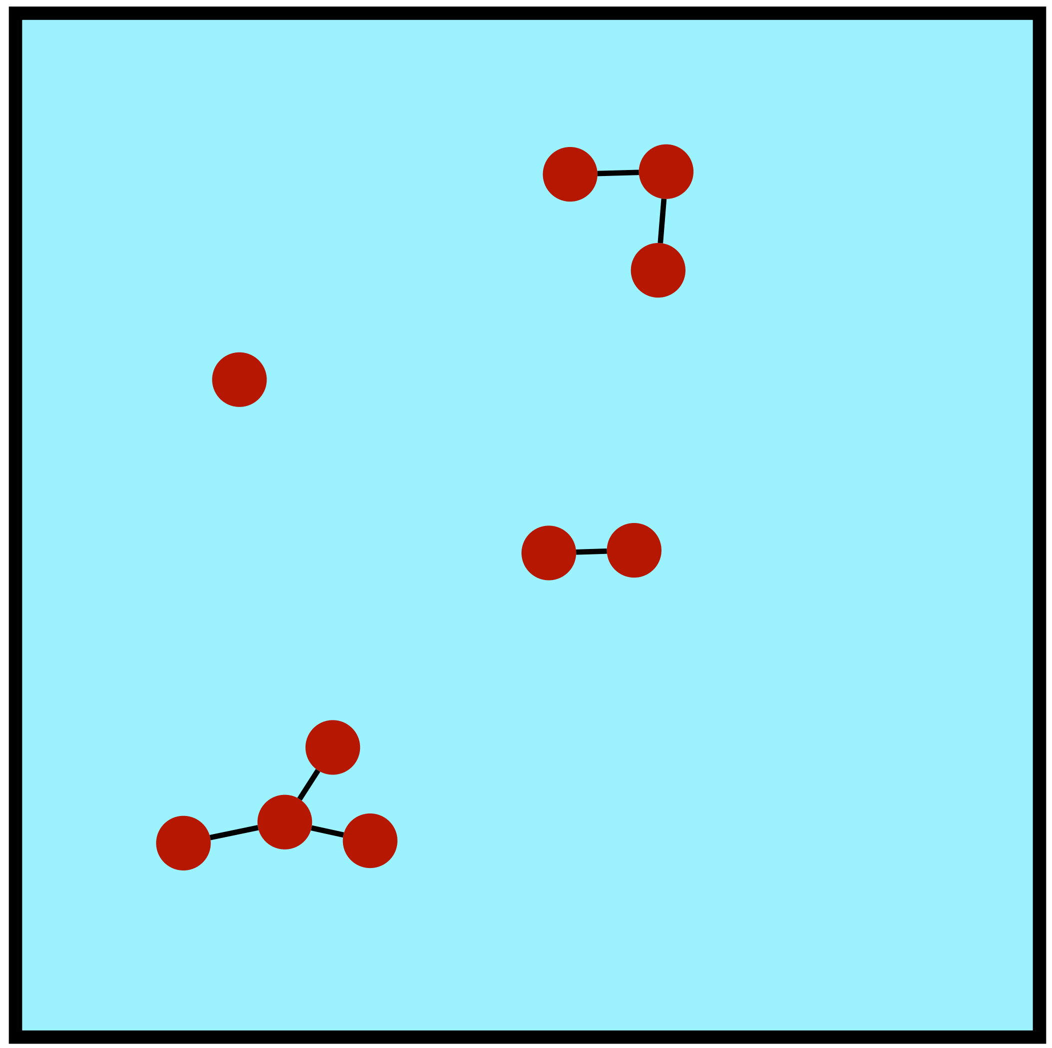

For the Ising model, deviations from the blue ground state are connected components of red vertices

Independent sets in bipartite graphs

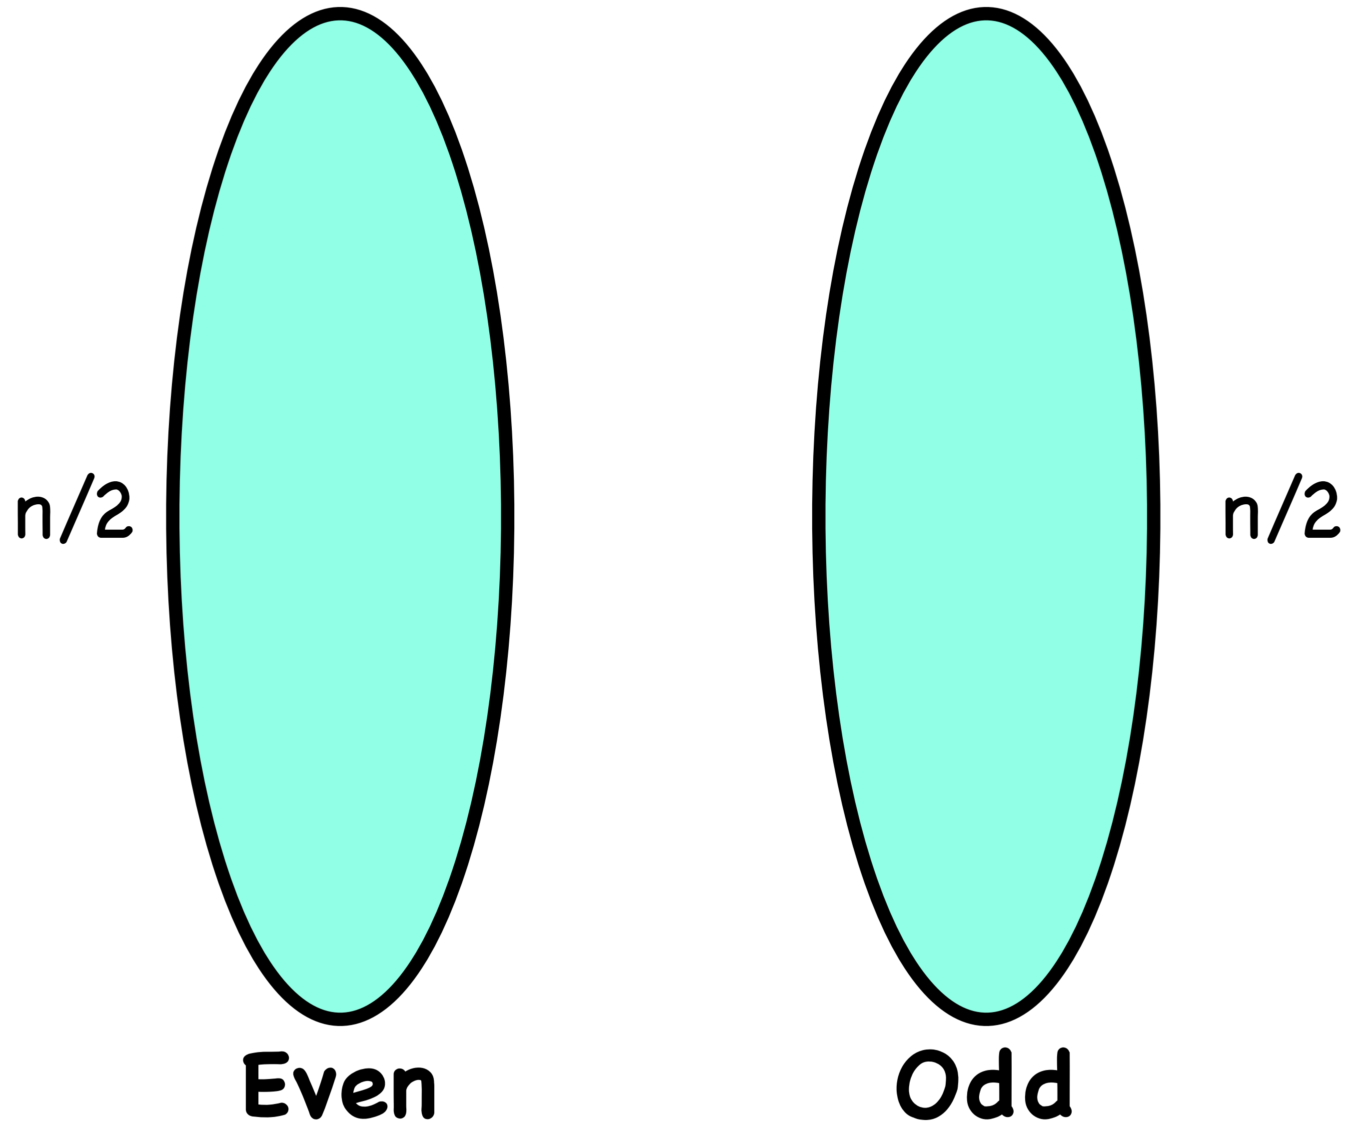

Two **ground states**: Even (no odds), Odd (no evens) each of weight $(1+\lambda)^{n/2}$ (or $2^{n/2}$ for unweighted case)

Independent sets in bipartite graphs

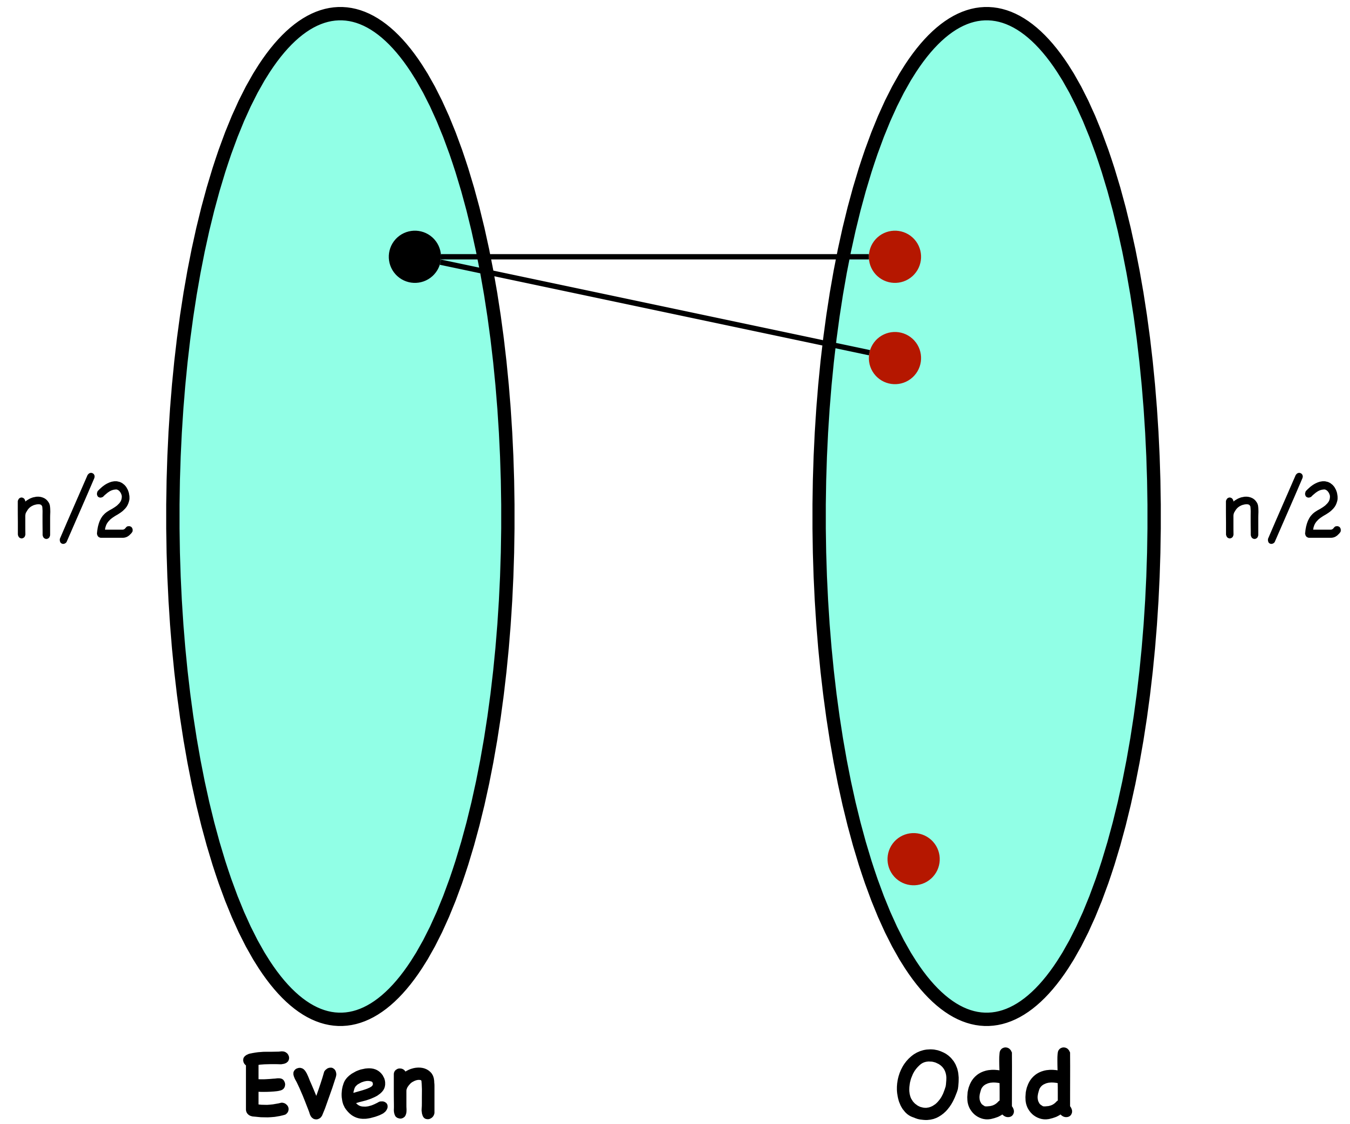

Deviations from Even are 2-linked components of odd vertices.

How much does a deviation **cost** (relative to ground state)?

Independent sets in bipartite graphs

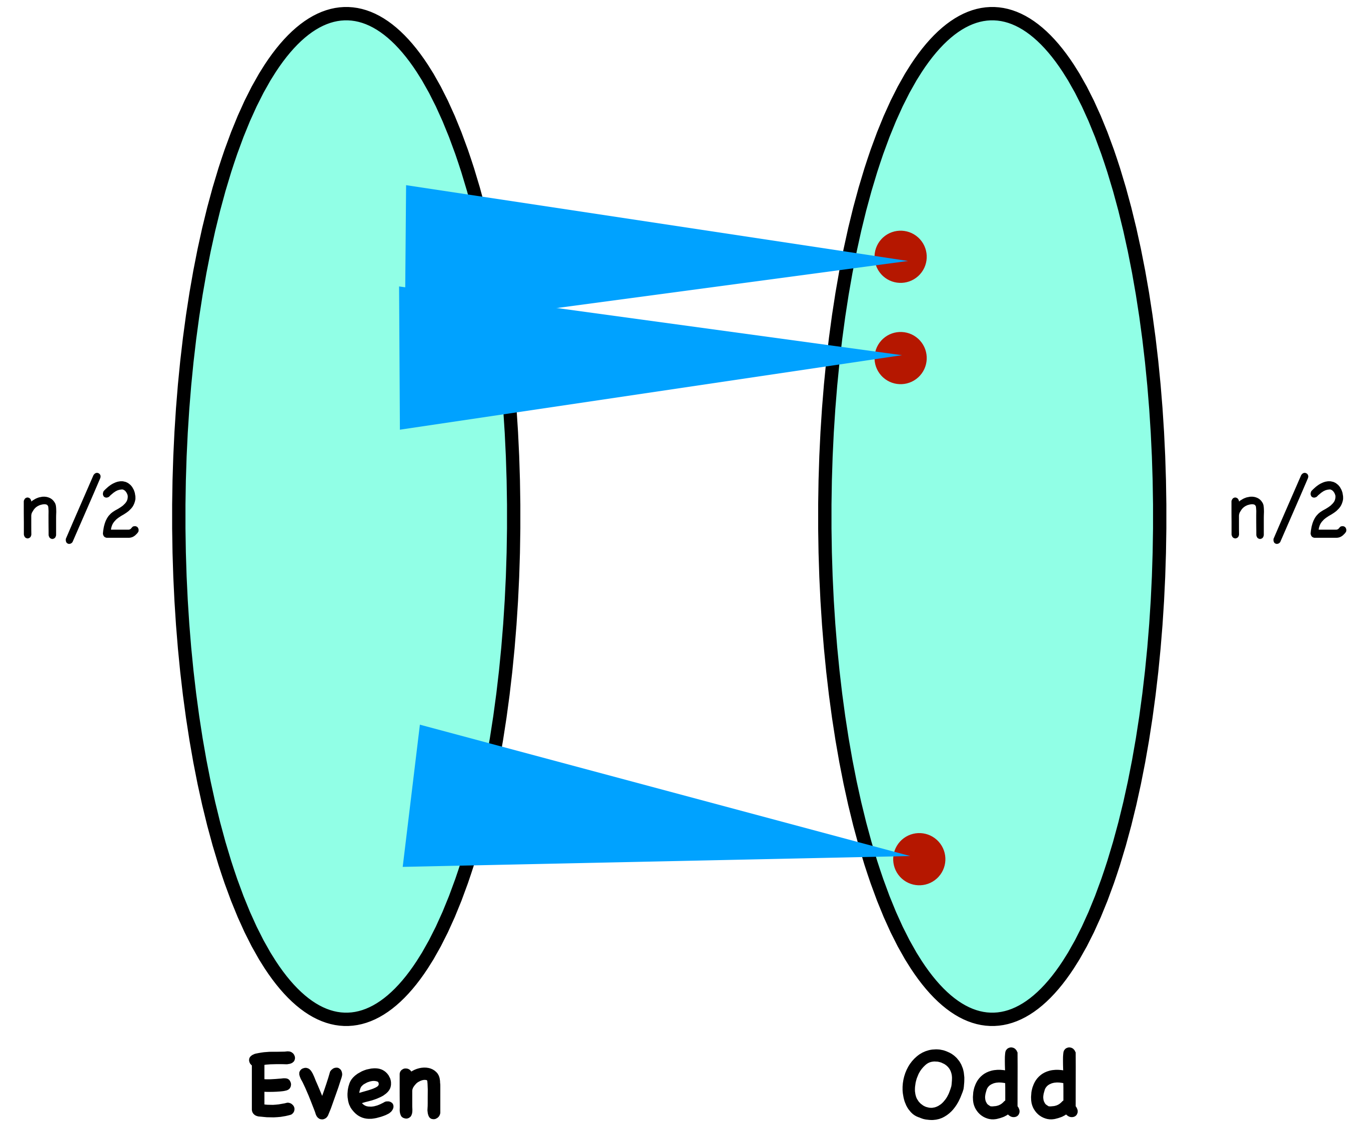

A deviation γ costs $\frac{\lambda^{|\gamma|}}{(1+\lambda)^{|N(\gamma)|}}$. The cost **factorizes** over deviations!

Independent sets in bipartite graphs

Now we rewrite the partition function in terms of deviations (**polymers**):

$$Z = (1+\lambda)^{n/2} \sum_{\Gamma} \prod_{\gamma \in \Gamma} \frac{\lambda^{|\gamma|}}{(1+\lambda)^{|N(\gamma)|}}$$

sum is over all collections of **compatible** polymers (union is not 2-linked)

We want to measure contributions from **mostly even** or **mostly odd** configurations separately.

Independent sets in bipartite graphs

Define a **polymer** as a defect of size at most n/10, say.

Define the **even partition function**:

$$Z_{even} = (1+\lambda)^{n/2} \cdot \Xi_{even} = (1+\lambda)^{n/2} \sum_{\Gamma} \prod_{\gamma \in \Gamma} \frac{\lambda^{|\gamma|}}{(1+\lambda)^{|N(\gamma)|}}$$

sum is over all collections of **compatible** (small) polymers.

Define Z_{odd} analogously.

Independent sets in bipartite graphs

What have we gained by defining Z_{even}, Z_{odd}?

Depends on the graph G (and the value of λ)!

If G is a **sufficiently good expander** and λ **is large** enough then:

1) $Z \approx Z_{even} + Z_{odd}$ (exponentially small rel. error)
2) **Cluster expansion** for Ξ_{even}, Ξ_{odd} converges

Cluster expansion convergence allows us to approximate Z very accurately (algorithmically!) and gives strong correlation decay properties

Independent sets in bipartite graphs

Warm-up example: random Δ-regular bipartite expander graphs (paper w/ Jenssen & Keevash)

Hypercube: not a great expander for large sets but **graph containers** to the rescue

Expander graphs

Let G be a Δ-regular bipartite graph on n vertices so that $|N(S)| \geq (1+\alpha)|S|$ for every $|S| \leq n/4$ on one side of the bipartition.

With $w(\gamma) = \dfrac{\lambda^{|\gamma|}}{(1+\lambda)^{|N(\gamma)|}}$, we have:

$$\Xi_{even} = \sum_{\Gamma} \prod_{\gamma \in \Gamma} w(\gamma)$$

This is a **multivariate hard-core model**!
(with the incompatibility graph on polymers)

The **larger** λ the smaller the fugacities

Expander graphs

Check the Kotecky-Preiss condition:

$$\sum_{\gamma' \nsim \gamma} e^{|\gamma'|} w(\gamma') \leq |\gamma|$$

Expansion gives the bound: $$w(\gamma) \leq \frac{\lambda^{|\gamma|}}{(1+\lambda)^{(1+\alpha)|\gamma|}}$$

Expander graphs

Need an approximation lemma:

$$\frac{Z - Z_{odd} - Z_{even}}{Z} \leq e^{-n}$$

This is a kind of Markov chain **slow-mixing** result: it says there is a bottleneck in the state space

The algorithm: compute cluster expansion up to size log n for Ξ_{even}, Ξ_{odd}, exponentiate, add together, and multiply by $(1+\lambda)^{n/2}$.

Hypercube

Back to Q_d.

Small sets are very good expanders: $|S| \leq d^2$ expand by cd.

Large sets are not: expand by factor $(1 + O(1/\sqrt{d}))$.

Using the same argument we'd need large λ.

Convergence

For larger γ, we use the container lemma of Sapozhenko and Galvin: let $\mathcal{G}(a,b)$ be the family of 2-linked A with $|[A]| = a, |N(A)| = b$.

For all $\lambda \geq \frac{C_0 \log d}{d^{1/3}}, a \leq 2^{d-2}$

$$\sum_{A \in \mathcal{G}(a,b)} \frac{\lambda^{|A|}}{(1+\lambda)^b} \leq 2^d \exp\left(-\frac{C_1(b-a)\log d}{d^{2/3}}\right).$$

This works perfectly to verify the K-P condition for large polymers!

Consequences

With cluster expansion convergence + approximation lemma we can derive many consequences:

$$Z \approx 2(1+\lambda)^{2^{d-1}} \cdot \Xi_{even}$$

Approximation to μ:

pick even/odd at random

sample a configuration X of polymers with probability $\prod_{\gamma \in X} w(\gamma)/\Xi_{even}$

Independently sample unblocked even vertices.

Terms of the cluster expansion

The cluster expansion gives an asymptotic expansion for the log partition function

Theorem 6. (Jenssen, P.)
For $\lambda = \Omega\left(\log d \cdot d^{-1/3}\right)$,

$$Z(\lambda) = 2(1+\lambda)^{2^{d-1}} \cdot \exp\left[\sum_{j=1}^{k} L_j + \epsilon_k\right]$$

The term L_j only dependents on clusters of size j,

and $\epsilon_k = o(L_k)$.

Terms of the cluster expansion

What do the terms L_j look like?

The hard-core model

What does a typical independent set from μ_λ look like?

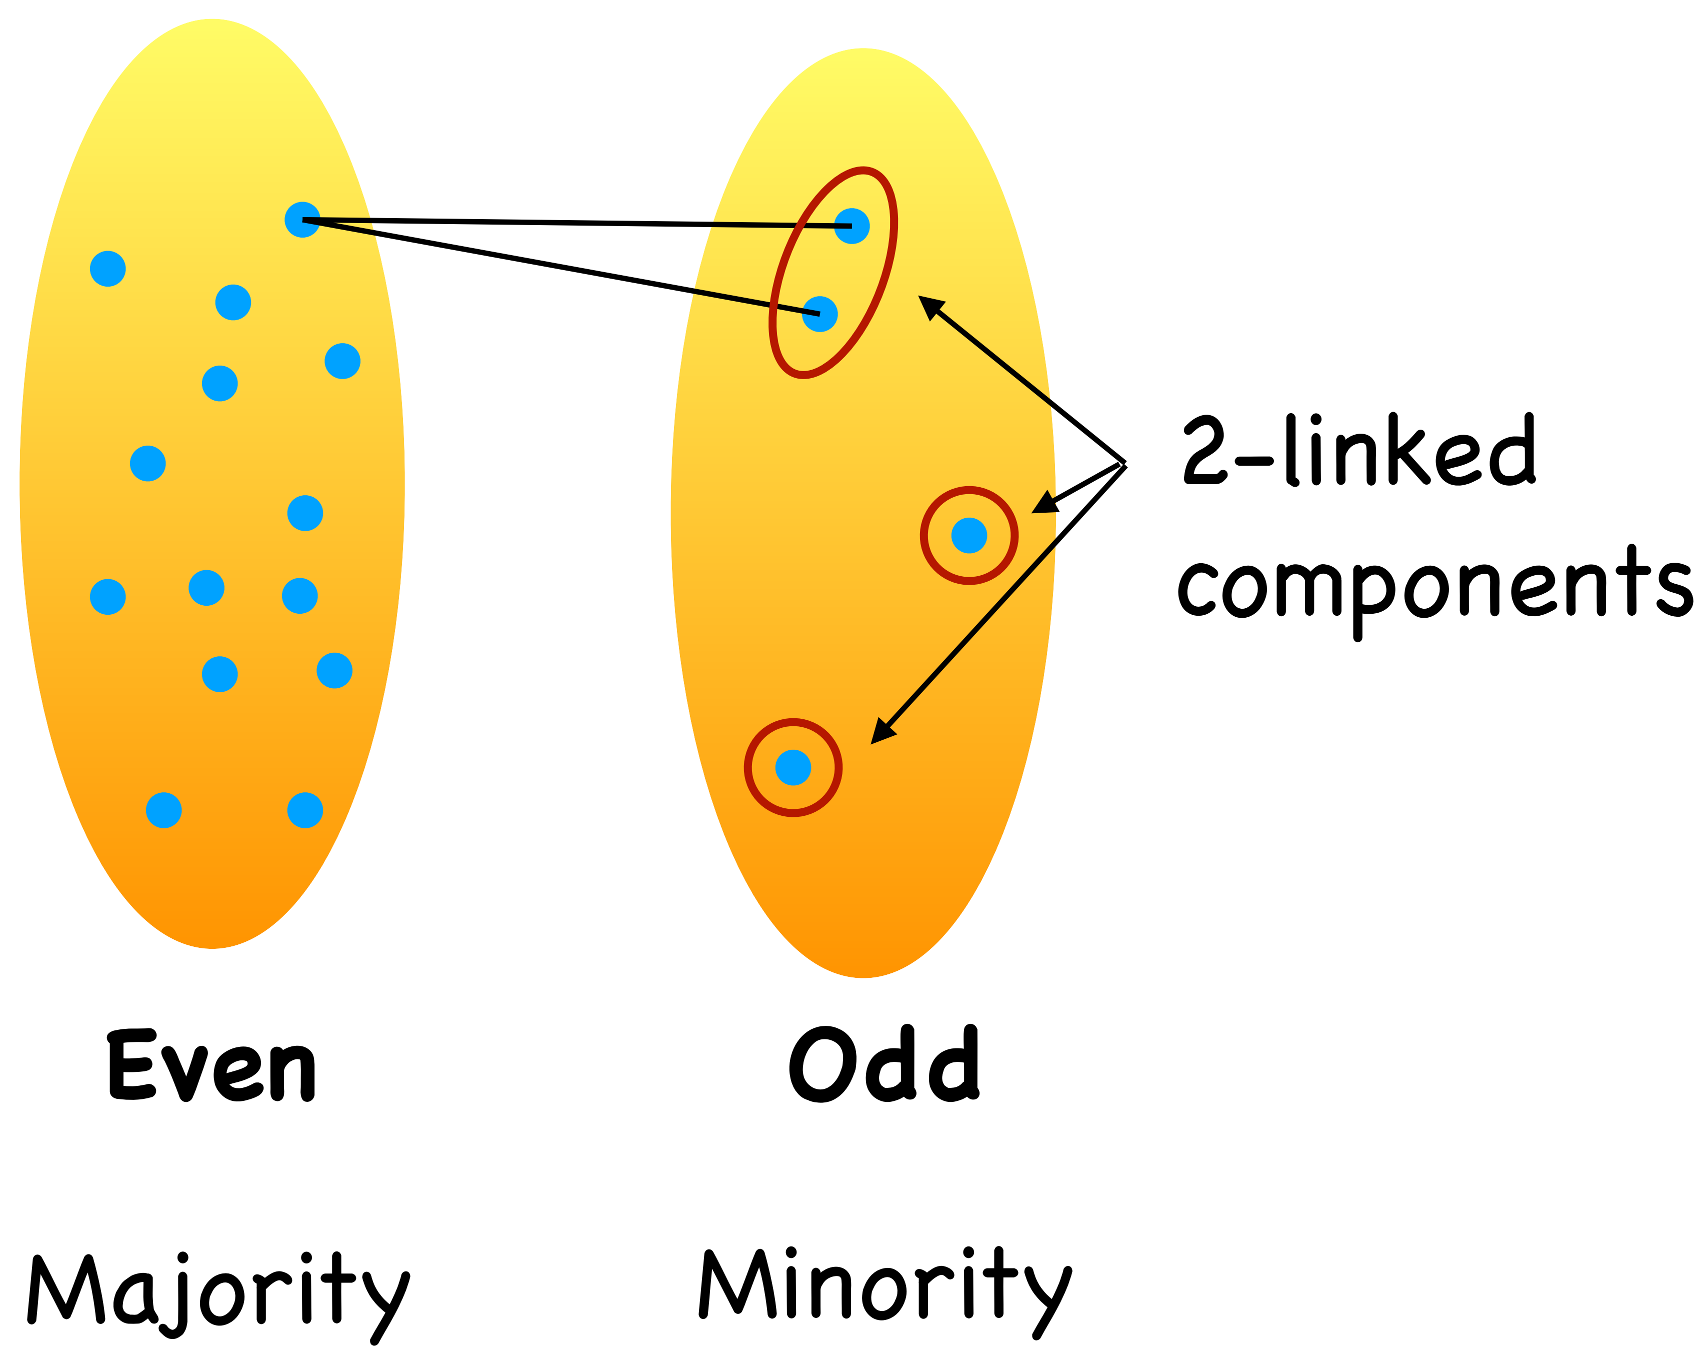

The hard-core model

Theorem 3. (Galvin, 2011)

For $\lambda = 1 + s/d$, the number of occupied vertices on the minority side is asymptotically Poisson with mean $e^{-s/2}/2$

If $\lambda > 2^{1/(m+1)} - 1$ then whp the largest 2-linked occupied component on the minority side is of size at most m

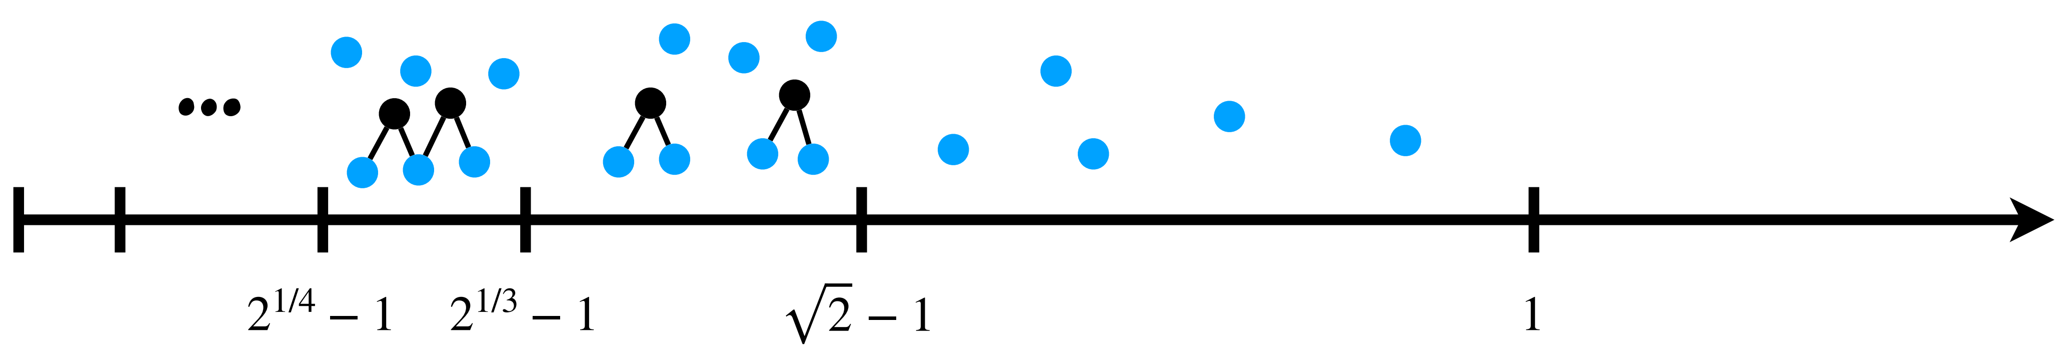

New results

Theorem 4. (Jenssen, P.)

The threshold for emergence of 2-linked components of size m is:

$$\lambda = 2^{1/m} - 1 + \frac{2^{1+1/m}(m-1)\log d}{md} + \frac{s}{d}$$

If $s \to \infty$ then whp there are no components of size m

If $s \to -\infty$ then whp there are components of size m

If s is constant, then the number of m-components is asymptotically Poisson

New results

2-linked components have an isomorphism type as a graph:

m=1

m=2

m=3

m=4

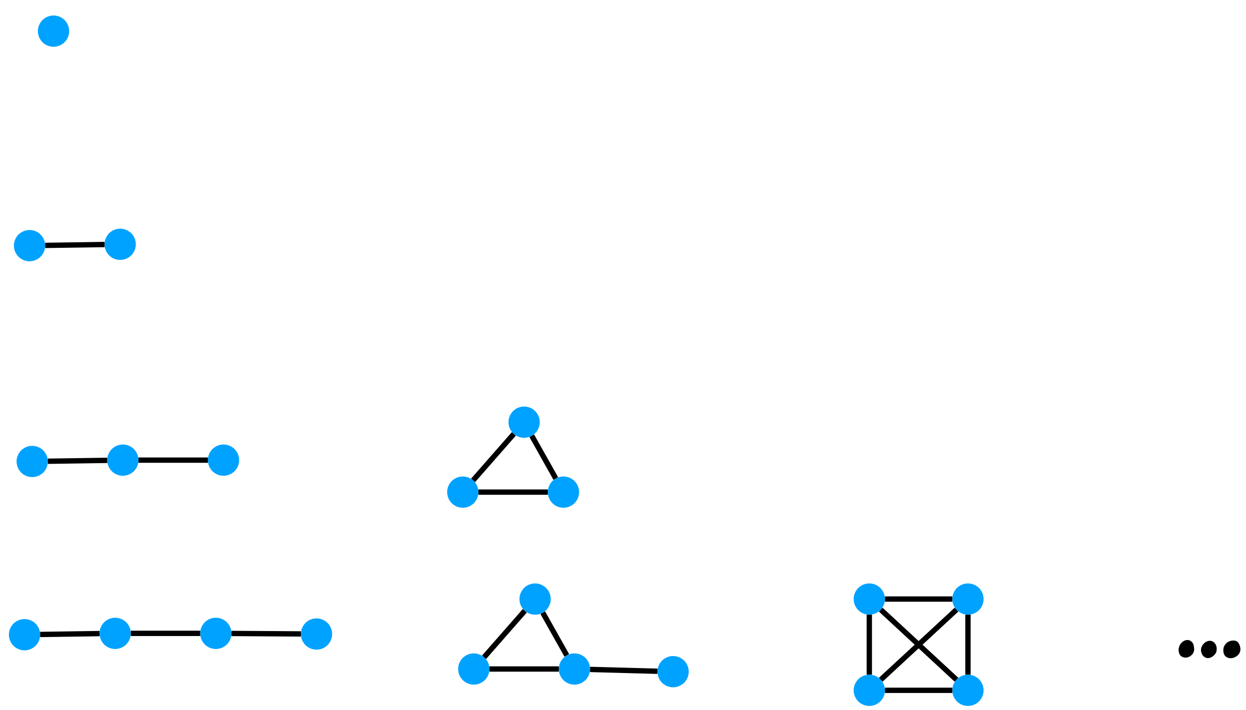

What is the joint distribution of the numbers of each type of 2-linked component on the minority side?

Independent Normals, Poissons!

Further applications of the method

Balogh, Garcia, Li: asymptotics for number of independent sets in middle 2 layers of Q_d

Jenssen, Keevash: asymptotics for number of q-colorings of Q_d for all q. (Galvin q=3, Kahn-Park q=4)

Davies, Jenssen, P.: proof of the Upper Matching Conjecture for large n. Cluster expansion for $i_k(G)$.

Many low-temperature algorithmic results.

Further reading

This paper: https://arxiv.org/abs/1907.00862

Galvin: https://www3.nd.edu/~dgalvin1/pdf/countingindsetsinQd.pdf

Friedli-Vilenik: https://www.unige.ch/math/folks/velenik/smbook/

Scott-Sokal: https://people.maths.ox.ac.uk/scott/Papers/lllshort.pdf

algorithms: https://arxiv.org/abs/1806.11548

Open problems

1. Can we apply the cluster expansion to other applications of graph containers?
2. What about hypergraph containers?
3. Further applications of statistical physics tools?

Thank you!