

The Edwards–Sokal Coupling for the Potts Higher Lattice Gauge Theory on \mathbb{Z}^d

by
Yakov Shklarov
B.Sc., University of Victoria, 2021

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF SCIENCE

in the Department of Mathematics and Statistics

© Yakov Shklarov, 2023
University of Victoria

All rights reserved. This thesis may not be reproduced in whole or in part, by photocopy or other means, without the permission of the author.

The Edwards–Sokal Coupling for the Potts Higher Lattice Gauge Theory on \mathbb{Z}^d

by
Yakov Shklarov
B.Sc., University of Victoria, 2021

Supervisory Committee

Dr. Anthony Quas (Department of Mathematics and Statistics, University of Victoria)

Co-Supervisor

Dr. Gourab Ray (Department of Mathematics and Statistics, University of Victoria)

Co-Supervisor

Abstract

The Edwards–Sokal coupling of the standard Potts model with the FK–Potts (random-cluster) bond percolation model can be generalized to arbitrary-dimension cells. In particular, the Potts lattice gauge theory on \mathbb{Z}^d has a graphical representation as a plaquette percolation measure. We systematically develop these previously-known results, using the frameworks of cubical (simplicial) homology and discrete Fourier analysis.

We show that, in the finite-volume setting, the Wilson loop expectation of a higher cycle γ is equal to the probability that γ is a homological boundary in the higher FK–Potts model. We also prove the strong FKG property of the higher FK–Potts model. These results culminate in a simple proof for the existence of infinite-volume limits in the higher Potts model and, in certain cases, of their invariance under translations and other symmetries. Additionally, we thoroughly examine the behavior of boundary conditions as they relate to the Edwards–Sokal coupling, for the purpose of understanding the higher Potts Gibbs states. In particular, we discuss spatial Markov properties and conditioning in the higher FK–Potts model, and generalize to more general boundary conditions the FKG property, the aforementioned identity for Wilson loop expectations, and a result about monotonicity in the coupling strength parameter. Also, we prove a theorem regarding the sharpness of thresholds of increasing symmetric events for the higher FK–Potts model with periodic boundary conditions.

In the final section, we describe some matrix-based sampling algorithms. Lastly, we prove a new characterization of the ground states of the random-cluster model, motivated by the problem of understanding the ground states in the higher FK–Potts model.

Contents

Supervisory Committee	ii
Abstract	iii
Contents	iv
List of Tables	vi
List of Figures	vii
Acknowledgments	viii
Introduction	1
1 Prelude: The Ising and Potts models and the random-cluster model	3
2 Preliminaries	9
2.1 Finite abelian groups and their duals	9
2.2 Cubical homology	21
2.3 Probabilistic couplings and the stochastic ordering	33
3 The higher Potts and FK–Potts models and their coupling, in finite volume with free boundary condition	40
3.1 The higher Potts model	40
3.2 The higher FK–Potts model	42
FKG and its applications	46
3.3 The Edwards–Sokal coupling	51
3.4 Using the coupling	55
4 Boundary conditions and the spatial Markov property	59
4.1 Spin conditions	59
4.2 Spatial Markov properties	65
4.3 Conditioning in the higher FK–Potts model	69
4.4 Further results	76
4.5 Sharp threshold	79

5	Infinite volume	82
5.1	Gibbs states of the higher Potts model	82
5.2	Thermodynamic limits of the higher Potts model	87
6	Odds and ends	94
6.1	Examples of pathological surfaces	94
6.2	Algorithms	97
	Counting cocycles	97
	Conditional sampling	97
	Coupling from the past	98
6.3	Ground states of the random-cluster model	98
	References	108
	Appendix: Code listings	113

List of Tables

1	Notation for cubical homology	34
2	Discrete exterior calculus vs. cubical homology	35

List of Figures

1	The discrete 2-torus	28
2	Chain complex and cochain complex	31
3	Spin conditions	61
4	Example set X_Λ	85
5	Surface with cross-handle	94
6	Surface with double cross-handle	95
7	Previous surfaces glued together	96
8	The square lattice S_4	103
9	The set S for various graphs	104
10	The set S for the square graph S_4	105
11	A union-of-plaquettes configuration in S_6	106
12	Southeast edges in S_5	106

Acknowledgments

This thesis owes its existence to Gourab Ray, who suggested the idea for the project, spent many hours explaining to me various aspects of probability, helped keep me motivated, and made many suggestions and corrections. Thank you also to Anthony Quas, who taught me the theory of Gibbs states and the thermodynamic formalism (and, over the years, lots more besides), and also gave numerous useful suggestions for improvement.

I am grateful to Mathew Lewis and Nathaniel Butler for some early conversations, to Mikayla Holmes for keeping me on track, and to my parents for their support.

This work was funded in part by NSERC (CGS-M).

Introduction

Lattice spin models have been studied via graphical representations since at least the 1970s [Gri06, p. 341]. The idea behind a graphical representation is to express quantities of interest in terms of other quantities derived from a configuration space of a graph (see [Dum16] for a general introduction.) One of the most famous graphical representations, the *random-cluster model*, is a family of bond percolation measures that can be probabilistically coupled to the Ising model and to its n -spin generalization, the Potts model. Results from the theory of percolation carry across the coupling to prove statements about the Ising model. The random-cluster representation is well-studied and even has its own textbook [Gri06]. Our aim here is to generalize to a certain class of lattice gauge models, which are essentially like the Potts model except with spins assigned to the edges of the hypercubic lattice rather than to its vertices. In fact, even more generally, we'll describe a class of couplings for the corresponding *higher lattice gauge theories* which assign spins to elements of higher-dimension cells (plaquettes, 3-cubes, etc.)¹

Lattice gauge theories are meant to serve as discretizations of certain quantum field theories [Cha19]. Our higher random-cluster model has been floating around the literature in some form for several decades, but it hasn't been mathematically developed in an entirely rigorous way. The key difficulty is that extending to higher dimension requires homology theory (this has been known for some time [DW82; AF84].) In particular, Aizenman and Fröhlich showed that the formula for the random-cluster model does not extend in an obvious way to a useful plaquette percolation model [AF84]. This thesis consists of a precise development of a model suitable for coupling to the Potts lattice gauge theory, with rigorous proofs of a few initial results.

The purpose of developing the graphical representation is to understand the behavior of certain observables called *Wilson loops* in the lattice gauge theory. This problem has received considerable attention in the last several years (e.g., [Cha20; Cao20; FLV21]) and this, in part, is what spurred the present work.

During the preparation of this thesis, a preprint was released by Duncan and Schweinhart with substantial overlap with this thesis [DS23]. Duncan and Schweinhart write in considerable depth, presenting some results in homological percolation—area and perimeter laws, and a theorem about hypersurfaces in the infinite-volume limit of the torus, with applications to Swendsen–Wang dynamics. Also, they discuss duality. Their preprint requires the parameter q to be prime, and

¹Higher gauge theory can be quite challenging algebraically—see [BH11; Pfe03]—but we won't need such advanced algebraic machinery here.

the proof there of the “probability = expectation” theorem (there Theorem 5, here Theorem 41) is not valid for non-prime q [DS23, p. 11]. However, for the model given here (and also mentioned in [DS23, p. 11]) this result does indeed hold for all integer $q \geq 1$. Also, Duncan and Schweinhart refer to [HS16], which proves FKG for prime q . The proof given herein (theorem 35) is different because it must work for arbitrary integer $q \geq 1$, not only prime q , and so it cannot rely on Betti numbers. This is important because ultimately we’d like to consider more general gauge groups. Finally, the discussion of boundary conditions here (section 4) is more comprehensive, and the discussion of infinite-volume limits (section 5) is very different from that in [DS23, §4.2].

This thesis contains no new theorems regarding phase transitions or the decay of correlations. Instead, its purpose is to develop the graphical representation in full rigor, to establish some basic results as a foundation for further work, and to serve as an expository introduction to the area. Perhaps the key results are the strong FKG property (theorems 35 and 62), which gives rise to many useful properties of the higher FK–Potts (generalized random-cluster) model; and “expectation equals probability” (theorems 41 and 64), which explains why the coupling is useful.

1 Prelude: The Ising and Potts models and the random-cluster model

This section is an informal introduction to our graphical representation, which is described more rigorously and in greater generality in section 3.

Recall that the *Ising model* on a finite box $V = \{-n, \dots, n\}^d \subseteq \mathbb{Z}^d$ is the probability distribution on the spins-on-vertices configuration space $\Sigma := \{-1, +1\}^V$ given by

$$\pi_\beta(\sigma) := \frac{\exp\left(\frac{1}{2}\beta \sum_{v \sim w} \sigma_v \sigma_w\right)}{\sum_{\sigma' \in \Sigma} \exp\left(\frac{1}{2}\beta \sum_{v \sim w} \sigma'_v \sigma'_w\right)}, \quad \beta \in (0, \infty), \sigma \in \Sigma := \{-1, +1\}^V,$$

where the sums are over all pairs (v, w) of adjacent vertices in V . (Actually, the Ising model is more general: Here we're assuming free boundary condition, uniform interaction strength, zero external field, and the hypercubic lattice as the underlying graph.) The parameter β is analogous to a physical system's inverse temperature (reciprocal of temperature as measured from absolute zero.) The factor $H(\sigma) = -\frac{1}{2} \sum_{v \sim w} \sigma_v \sigma_w$, called the *Hamiltonian*, is a kind of generalized energy function. As temperature rises, the measure π_β converges to the uniform distribution; as temperature falls, π_β puts more and more probability mass on the configurations with most spins equal. In the infinite-volume limit, an abrupt phase transition appears. There are many ways to characterize this phase transition—for example, by studying the spatial decay of correlations between two vertices. The study of this phase transition is inspired by the Curie transition in physical ferromagnetic materials (although real-world magnets are much more complicated, and not described very well by the Ising model.)

The slightly more general *Potts model* (introduced in [Pot52]; see also [Wu82]) allows the spins to come from a general finite cyclic group $\mathbb{Z}/q\mathbb{Z}$, (or rather, a set of size q , because the group structure isn't used):

$$\pi_{\beta, q}(\sigma) := \frac{\exp\left(\beta \sum_{v \sim w} \llbracket \sigma_v = \sigma_w \rrbracket\right)}{\sum_{\sigma' \in \Sigma} \exp\left(\beta \sum_{v \sim w} \llbracket \sigma'_v = \sigma'_w \rrbracket\right)}, \quad \beta \in (0, \infty), q \in \mathbb{Z}_{\geq 2}, \sigma \in \Sigma := (\mathbb{Z}/q\mathbb{Z})^V,$$

where $\llbracket \cdot \rrbracket$ is the indicator function.² Different Hamiltonians can be given for this vertex configuration space Σ . Sometimes, instead of $H(\sigma) = -\sum_{v \sim w} \llbracket \sigma_v = \sigma_w \rrbracket$, we take the Hamiltonian $H(\sigma) = -\sum_{v \sim w} \sigma_v \cdot \sigma_w$ where $\sigma_v \cdot \sigma_w := \cos 2\pi(\sigma_v - \sigma_w)$ (now we *are* using the group structure),

²The symbol $\llbracket \cdot \rrbracket$ is also called the *Iverson bracket*; for a predicate P it's defined as $\llbracket P \rrbracket := \begin{cases} 1, & P \text{ true,} \\ 0, & P \text{ false.} \end{cases}$ We'll make heavy use of the indicator function, so the Iverson bracket was chosen over the notation $\mathbf{1}_P$ or $\mathbf{1P}$ in order to keep the notation clean.

which gives the so-called *clock model* or *planar Potts model*. The clock model was introduced in the same paper as the Potts model [Pot52]; see also [Dum16, pp. 3–4] for a more modern exposition. Much more generally, it’s possible to allow the spins to come from an arbitrary compact Lie group—this is especially useful for studying gauge theories [Cao20, §1.1]. However, we’ll work exclusively with $\mathbb{Z}/q\mathbb{Z}$ and the Potts model.

Let E be the set of all nearest-neighbor edges, or “bonds”, between the vertices $V = \{-n, \dots, n\}^d$. Let $\Omega := \{0, 1\}^E$. We’ll write edges as ordered pairs $e = (v, w)$, for $v, w \in \mathbb{Z}^d$, always in the forwards orientation (i.e., the sum of the d components of w is one greater than the sum of the components of v). For any particular configuration $\omega \in \Omega$, an edge e is considered *open* if $\omega_e = 1$ and *closed* if $\omega_e = 0$. Let $o(\omega)$ and $c(\omega)$ be the number of edges that are open and closed, respectively. A *cluster* is a connected component in the graph (V, E) . Let $k(\omega)$ be the number of clusters, including isolated vertices. The *random-cluster model* is the probability distribution on edge configurations

$$\varphi_{p,q}(\omega) := \frac{1}{Z_{RC}(p,q)} (1-p)^{c(\omega)} p^{o(\omega)} q^{k(\omega)},$$

$$p \in (0, 1), \quad q \in (0, \infty), \quad \omega \in \Omega := \{0, 1\}^E,$$

where the random-cluster *partition function* $Z_{RC}(p, q)$ is the normalizing constant $\sum_{\omega \in \Omega} (1-p)^{c(\omega)} p^{o(\omega)} q^{k(\omega)}$. Note that for $q = 1$ this reduces to independent Bernoulli(p) bond percolation.

The random-cluster model with $q = 2$ is sometimes called the *FK–Ising model*; with q restricted to $2, 3, 4, \dots$, it’s called the *FK–Potts model*. For these choices of parameter q there are couplings to the Ising and Potts models, which we’ll now describe.

The *Edwards–Sokal coupling* [ES88] of the Potts and FK–Potts models is

$$\mu_{p,q}(\sigma, \omega) := \frac{1}{Z_{ES}(p,q)} (1-p)^{c(\omega)} p^{o(\omega)} \mathbb{I}[(\sigma, \omega) \in F],$$

$$p \in (0, 1), \quad q \in \{2, 3, 4, \dots\}, \quad (\sigma, \omega) \in \Sigma \times \Omega,$$

where $Z_{ES}(p, q)$ is the normalizing constant like before, and $\mathbb{I}[(\sigma, \omega) \in F]$ is 1 if the endpoints of each ω -open edge have equal spins in σ , and 0 otherwise. In the former case, we say that the configuration (σ, ω) is *valid*, or that the configurations σ and ω are *compatible*, and we let F be the set of all valid configurations $(\sigma, \omega) \in \Sigma \times \Omega$ (the notation F is from [Gri06, p. 8].) So the measure $\mu_{p,q}$ is the product of the iid Bernoulli(p) measure on edges and the uniform measure on Σ , conditioned on the event F .

Note that each open edge imposes a linear constraint on Σ (viewed as a \mathbb{Z} -module): If an edge $e = (v, w)$ is open in ω then, in order for $(\sigma, \omega) \in F$, it's necessary that $\sigma_v - \sigma_w = 0$. The closed edges impose no constraints. This algebraic perspective will be useful when we generalize to higher dimensions, so we'll describe it in some more detail. For every edge configuration $\omega \in \Omega$, let $A_\omega = (\mathbb{Z}/q\mathbb{Z})^{O(\omega)}$ where $O(\omega)$ is the set of all open edges in ω . Endow A_ω with the product group structure, and likewise for $\Sigma = (\mathbb{Z}/q\mathbb{Z})^V$. Define a group homomorphism $f_\omega : \Sigma \rightarrow A_\omega$ as follows. For every ω -open edge $e = (v, w)$ (oriented forwards) let $f_\omega(\sigma)(e) = \sigma_w - \sigma_v$. Then $(\sigma, \omega) \in F$ if and only if σ belongs to the kernel of f_ω . Later, we'll call f_ω the *coboundary map*, and elements of its kernel *cocycles*.

The following result may be found in [ES88; Gri06, §1.4]. We include the proof in full detail so as to make the general case (proposition 38) more approachable.

Proposition 1. *For every $p \in (0, 1)$ and $q \in \{2, 3, 4, \dots\}$, the probability measure $\mu_{p,q}(\sigma, \omega)$ is a coupling of $\pi_{\beta,q}$ and $\varphi_{p,q}$, where $\beta = -\log(1 - p)$ (or, equivalently, $p = 1 - e^{-\beta}$.)*

Proof. The first marginal of $\mu_{p,q}$ is (omitting the normalizing factor for clarity)

$$\begin{aligned}
\sum_{\omega \in \Omega} \mu_{p,q}(\sigma, \omega) &\propto \sum_{\omega \in \Omega} (1-p)^{c(\omega)} p^{o(\omega)} \mathbb{I}[(\sigma, \omega) \in F] \\
&= \sum_{\omega \in \Omega} (1-p)^{c(\omega)} p^{o(\omega)} \prod_{\substack{\omega_e=1 \\ e \equiv (v,w) \in E}} \mathbb{I}[\sigma_v = \sigma_w] \\
&= \sum_{\omega \in \Omega} \left(\prod_{\substack{\omega_e=0 \\ e \equiv (v,w) \in E}} (1-p) \right) \left(\prod_{\substack{\omega_e=1 \\ e \equiv (v,w) \in E}} p \mathbb{I}[\sigma_v = \sigma_w] \right) \\
&= \prod_{(v,w) \in E} ((1-p) + p \mathbb{I}[\sigma_v = \sigma_w]) \quad (\text{via expansion}) \\
&= \prod_{(v,w) \in E} (\mathbb{I}[\sigma_v = \sigma_w] + (1-p) \mathbb{I}[\sigma_v \neq \sigma_w]) \\
&= \mathbb{1}^{\{(v,w) \in E \mid \sigma_v = \sigma_w\}} (1-p)^{\{(v,w) \in E \mid \sigma_v \neq \sigma_w\}} \\
&= \exp \left(-\beta \sum_{(v,w) \in E} \mathbb{I}[\sigma_v \neq \sigma_w] \right) \\
&\propto \exp \left(\beta \sum_{(v,w) \in E} \mathbb{I}[\sigma_v = \sigma_w] \right) \\
&\propto \pi_{\beta,q}(\sigma), \quad \sigma \in \Sigma.
\end{aligned}$$

Call an open edge or open cluster in ω *monochromatic* if all its vertices have equal spin in σ . The second marginal of $\mu_{p,q}$ is

$$\begin{aligned}
& \sum_{\sigma \in \Sigma} \mu_{p,q}(\sigma, \omega) \\
& \propto (1-p)^{c(\omega)} p^{o(\omega)} \sum_{\sigma \in \Sigma} \mathbb{I}[(\sigma, \omega) \in F] \\
& = (1-p)^{c(\omega)} p^{o(\omega)} |\{\sigma \in \Sigma : \text{each } \omega\text{-open edge is monochromatic}\}| \\
& = (1-p)^{c(\omega)} p^{o(\omega)} |\{\sigma \in \Sigma : \text{each open cluster in } \omega \text{ is monochromatic}\}| \\
& = (1-p)^{c(\omega)} p^{o(\omega)} q^{k(\omega)} \quad (\text{making } k(\omega) \text{ independent choices from } q \text{ possible spins}) \\
& \propto \varphi_{p,q}(\omega), \quad \omega \in \Omega. \quad \square
\end{aligned}$$

The second part of the proof shows that $q^{k(\omega)}$ is precisely the number of cocycles, that is, $q^{k(\omega)} = |\ker f_\omega|$. In higher dimensions, we'll have to replace the factor $q^{k(\omega)}$ with a more general expression (actually, we'll just write "number of cocycles"; see eq. (6) and proposition 33.)

The conditional measures of $\mu_{p,q}$ have a simple description. To sample a vertex configuration conditional on a given edge configuration, assign a spin uniformly and independently to each cluster. To sample an edge configuration conditional on a given vertex configuration, open each edge uniformly with probability p wherever two incident vertices have equal spins, and leave all remaining edges closed. See [Gri06, Figure 1.3] for a graphical illustration of these conditional sampling procedures. For the proof, see proposition 40.

If $q = 2$ then the probability that two vertices a and b belong to the same cluster is equal to the expectation of the function $(-1)^{\sigma_b - \sigma_a}$. This last expression has a topological interpretation, which will extend to the higher-dimensional setting. Consider the unitary character $[0] \mapsto 1, [1] \mapsto -1$ of the group $\mathbb{Z}/2\mathbb{Z} = \{[0], [1]\}$. An edge path connecting a to b has an oriented boundary consisting of the two points a and b (actually, orientation doesn't matter when $q = 2$.) The product of characters over this oriented boundary is $(-1)^{\sigma_b - \sigma_a}$. The general result is theorem 41, where the product of characters is denoted by W_γ . But here's a simplified proof of the special case, which is essentially the same result as [Gri06, Theorem 1.16].

Proposition 2. *For every $p \in (0, 1)$, every pair of vertices $a, b \in V$ satisfies³*

$$\pi_{\beta,2}((-1)^{\sigma_b - \sigma_a}) = \varphi_{p,2}(a \leftrightarrow b),$$

³The notation $\pi_{\beta,2}X$ indicates the expectation of the random variable X with respect to the measure $\pi_{\beta,2}$.

where $a \leftrightarrow b$ is the event that a and b belong to the same open edge cluster, and $p = 1 - e^{-\beta}$.

Proof. Recall the preceding results about the marginals and conditionals of $\mu_{p,2}$. Conditioning on the edge configuration gives

$$\begin{aligned} \mu_{p,2}((-1)^{\sigma_b - \sigma_a} \mid \omega) &= \mu_{p,2}(\sigma_b = \sigma_a \mid \omega) - \mu_{p,2}(\sigma_b \neq \sigma_a \mid \omega) \\ &= \begin{cases} 1 - 0 & \text{if } a \leftrightarrow b, \\ \frac{1}{2} - \frac{1}{2} & \text{if } a \not\leftrightarrow b \end{cases} \\ &= \llbracket a \leftrightarrow b \rrbracket, \quad \omega \in \Omega. \end{aligned}$$

The ' $\frac{1}{2}$ ' is because the conditional measure independently assigns to each open cluster a uniform spin from $\mathbb{Z}/2\mathbb{Z}$, so distinct clusters' spins are equal with probability $\frac{1}{2}$.

Now take expectations with respect to $\mu_{p,2}$ on both sides. \square

In the Potts lattice gauge theory, elements of $\mathbb{Z}/q\mathbb{Z}$ are assigned to the edges of the graph (V, E) instead of to the vertices. For now we'll consider only $q = 2$ (the Ising lattice gauge theory.) A *plaquette* is a two-dimensional square of side length 1 embedded in \mathbb{R}^d , all of whose vertices are integer lattice points. Let L be the set of all plaquettes in \mathbb{R}^d that are included in the box $[-n, n]^d$. Then, for every plaquette in L , each of its four edges may be identified with an element of E . In general, we'll also need to consider the orientation of the edges (section 2.2), but for $q = 2$ this may be ignored because $-1 \equiv 1 \pmod{2}$.

In the gauge theory we no longer care about the vertex set V , so in place of the graph (V, E) we'll work with the hypergraph (E, L) . The configuration space for the Ising gauge theory is $\Sigma := (\mathbb{Z}/2\mathbb{Z})^E$. The configuration space for the associated gauge FK-Ising model is $\Omega = \{0, 1\}^L$: each plaquette is either closed (0) or open (1). For $\sigma \in \Sigma$ and $Q \in L$, write $\sigma_Q := (-1)^{\sigma_{e_1} + \sigma_{e_2} + \sigma_{e_3} + \sigma_{e_4}}$ where $e_1, e_2, e_3, e_4 \in E$ are the four edges incident to Q . We'll say that the configurations $\sigma \in \Sigma$ and $\omega \in \Omega$ are compatible if $\omega(Q) = 1 \implies \sigma_Q = 1$ for every $Q \in L$; that is, each open plaquette has even edge sum. Again, let $F \subseteq \Sigma \times \Omega$ be the set of compatible pairs of configurations, and define a probability measure on $\Sigma \times \Omega$ in precisely the same way as before,

$$\mu_{p,2}(\sigma, \omega) = \frac{1}{Z_{ES}(p, 2)} (1 - p)^{c(\omega)} p^{o(\omega)} \llbracket (\sigma, \omega) \in F \rrbracket, \quad p \in (0, 1).$$

Computing marginals using the same technique as above, we see that the first marginal of

$\mu_{p,2}(\sigma, \omega)$ is

$$\pi_{\beta,2}(\sigma) = \frac{1}{Z_P(\beta, 2)} \exp \left(\beta \sum_{Q \in \mathcal{L}} \llbracket \sigma_Q = 1 \rrbracket \right), \quad \beta \in (0, \infty), \sigma \in \Sigma,$$

where $p = 1 - e^{-\beta}$ as before. The second marginal is

$$\varphi_{p,2}(\omega) = \frac{1}{Z_{\text{FKP}}(p, 2)} (1-p)^{c(\omega)} p^{o(\omega)} |\{\sigma \in \Sigma \mid (\sigma, \omega) \in F\}|, \quad p \in (0, 1), \omega \in \Omega.$$

The conditional measures, too, are analogous to before. To sample $\omega \in \Omega$ conditional on $\sigma \in \Sigma$, for each plaquette Q , if $\sigma_Q = 1$ then let Q be open independently with probability p , and if $\sigma_Q = -1$ then take Q closed. To sample $\sigma \in \Sigma$ conditional on $\omega \in \Omega$, pick uniformly an element of $\{\sigma \in \Sigma \mid (\sigma, \omega) \in F\}$.

A *Wilson loop* is a closed walk in the graph (repeated vertices and edges are allowed; also, we allow the trivial walk with one vertex and no edges.) For a Wilson loop γ of length $n \geq 0$ containing edges (e_1, \dots, e_n) , and for edge spin configuration $\sigma \in \Sigma$, define the *Wilson loop variable* $W_\gamma(\sigma) := (-1)^{\sum_i \sigma_{e_i}}$. According to theorem 41, the Wilson loop expectation $\pi_{\beta,2} W_\gamma$ coincides with the probability, with respect to $\varphi_{p,2}$, that γ is a boundary of some homological surface consisting of open plaquettes (here $q = 2$ so a “homological surface” is simply a set of plaquettes; more generally we consider 2-chains over $\mathbb{Z}/q\mathbb{Z}$ as defined in section 2.2.) Note the analogy to the event $a \leftrightarrow b$ from before: a path joining a with b is a one-dimensional surface with boundary $\{a, b\}$. A “Wilson loop” in the classical (non-gauge) Ising model is therefore simply a pair of vertices.

For a thorough introduction to the random-cluster model, see the textbook [Gri06]. For an overview of the Ising and Potts model, see the lecture notes [Dum20]. We’ve taken V to be a fixed-size finite box in the hypercubic lattice, but the essential question is the behavior of Wilson loop expectations (and other observables) in the infinite-volume limit. More on this in section 5.

2 Preliminaries

We review some prerequisites in algebra and topology. References are listed at the end of each subsection.

2.1 Finite abelian groups and their duals

The results in this section are elementary, but some of them are hard to find in the literature. For that reason, they were derived as needed, though of course no claim of originality is made.

Let G be a finite abelian group. A *character* of G is a homomorphism from G into the circle group $\mathbb{T} = \{z \in \mathbb{C} \mid |z| = 1\}$. The *dual* \widehat{G} of G , also denoted by G^\wedge , is the set of all characters of G endowed with pointwise group operation $(\chi + \psi)(g) := \chi(g)\psi(g)$. Note well that the *sum* of two characters is their pointwise *product* (see the discussion on page 40 about conventions.) It can be shown (by invoking the structure theorem for finite abelian groups) that $G \cong \widehat{\widehat{G}}$. The *natural* (or *canonical*) map $\eta : G \rightarrow \widehat{\widehat{G}}$, $g \mapsto (\chi \mapsto \chi(g))$ is an isomorphism of G with its *bidual* $\widehat{\widehat{G}}$.

The situation is similar to that of a finite-dimensional vector space V : The algebraic bidual V^{**} is naturally isomorphic to V . In fact, if we take $G = (\mathbb{Z}/p\mathbb{Z})^d$ for prime p and any $d \geq 0$, then the dual of G as a vector space over the field $\mathbb{Z}/p\mathbb{Z}$ is precisely its dual as a finite abelian group, and the notions of natural isomorphism also coincide. This is because every character $\chi : G \rightarrow \mathbb{T}$ satisfies, for every $g \in G$, the identity $1 = \chi(0) = \chi(pg) = \chi(g)^p$, so every element of $\text{im}(\chi)$ is a p th root of unity. Identify the group of p th roots of unity with the additive group $\mathbb{Z}/p\mathbb{Z}$. Conveniently, scalar multiplication in the vector space $(\mathbb{Z}/p\mathbb{Z})^d$ by elements of the field $\mathbb{Z}/p\mathbb{Z}$ is merely repeated addition, so the character χ is a linear functional. And, conversely, every linear functional is a character.

Now back to the general case. Let A and B be finite abelian groups, and let $\alpha : A \rightarrow B$ be a homomorphism. Its *dual map* $\alpha^* : \widehat{B} \rightarrow \widehat{A}$, $\chi \mapsto \chi \circ \alpha$ is a homomorphism of the dual groups. Note that $(\alpha \circ \beta)^* = \beta^* \circ \alpha^*$ for any two composable homomorphisms α and β . The operation sending α to its *bidual map* $\alpha^{**} := (\alpha^*)^* : \widehat{\widehat{A}} \rightarrow \widehat{\widehat{B}}$ is compatible with the natural maps $\eta_A : A \rightarrow \widehat{\widehat{A}}$ and $\eta_B : B \rightarrow \widehat{\widehat{B}}$ in the sense that

$$\alpha^{**} \circ \eta_A = \eta_B \circ \alpha, \tag{1}$$

because

$$\begin{aligned} (\alpha^{**} \circ \eta_A)(g)(\chi) &= (\alpha^{**}(\eta_A(g)))(\chi) = (\eta_A(g) \circ \alpha^*)(\chi) = \eta_A(g)(\chi \circ \alpha) \\ &= \chi(\alpha(g)) = \eta_B(\alpha(g))(\chi) = (\eta_B \circ \alpha)(g)(\chi) \quad \text{for every } g \in A \text{ and } \chi \in \widehat{B}. \end{aligned}$$

Equation (1) is essentially what justifies the term *natural map*.

The *annihilator* $\text{Ann}_G S$ of a set $S \subseteq G$ is the set of all characters that kill S :

$$\text{Ann}_G S := \left\{ \chi \in \widehat{G} \mid \chi(s) = 1 \text{ for every } s \in S \right\}.$$

Evidently, the annihilator is always a subgroup of \widehat{G} . We will often drop the subscript and write simply $\text{Ann } S$ when the group G is clear from context.

Let H be a subgroup of G . The restriction to H of any character of G is a character of H . On the other hand, every character χ of H can be extended to a character of G (for a proof via recursion on the index of H , see [Pey20, §1.2.1, Lemma 1].) The restriction map $\rho : \widehat{G} \rightarrow \widehat{H}$ is a homomorphism, as is the inflation map $\tau : \widehat{G/H} \rightarrow \widehat{G}$, $\chi \mapsto (g \mapsto \chi(gH))$. It's straightforward to show that the sequence

$$\{1\} \rightarrow \widehat{G/H} \xrightarrow{\tau} \widehat{G} \xrightarrow{\rho} \widehat{H} \rightarrow \{1\} \quad (2)$$

is exact [Pey20, §1.2.1, Lemma 2]. We collect a few consequences for future reference.

Fact 3. *Let A and B be finite abelian groups, and let $\alpha : A \rightarrow B$ be a homomorphism. If α is surjective, then α^* is injective. If α is injective, then α^* is surjective.*

Proof. If α is surjective, then clearly the two maps $\alpha^*(\chi_1) = \chi_1 \circ \alpha \in \widehat{B}$ and $\alpha^*(\chi_2) = \chi_2 \circ \alpha \in \widehat{B}$ are distinct whenever the maps $\chi_1, \chi_2 \in \widehat{A}$ are distinct.

Now assume that α is injective and let $\psi \in \widehat{A}$. Let $H = \text{im } \alpha$ and define the character $\psi' : H \rightarrow \mathbb{T}$, $\alpha(a) \mapsto \psi(a)$. Let χ be any character of G that extends ψ' . Then $\alpha^*(\chi) = \chi \circ \alpha = \psi$. \square

Fact 4. *Let G be a finite abelian group, and H a subgroup of G . Every character of H can be extended to a character of G in $|G|/|H|$ distinct ways. This holds in particular for the trivial character, i.e., $|\text{Ann}_G H| = |G|/|H|$.*

Proof. From eq. (2) we see $|\ker(\rho)| = |\text{im}(\tau)| = |\widehat{G/H}| = |G/H| = |G|/|H|$. Since the restriction ρ is a surjective homomorphism, by the first isomorphism theorem all its fibers have equal size $|G|/|H|$:

That is, every element of \widehat{H} can be extended to an element of \widehat{G} in $|G|/|H|$ ways. In particular, $\text{Ann}_G H = \ker(\rho)$, so $|\text{Ann}_G H| = |G|/|H|$. \square

Fact 5. For every subgroup H of a finite abelian group G ,

$$\eta(H) = \text{Ann}(\text{Ann } H),$$

where $\eta : G \rightarrow \widehat{\widehat{G}}$ is the natural map.

Proof. Applying fact 4 with \widehat{G} in place of G and with $\text{Ann}_G H$ in place of H gives

$$|\text{Ann}_{\widehat{G}}(\text{Ann}_G H)| = \frac{|\widehat{G}|}{|\text{Ann}_G H|} = \frac{|\widehat{G}|}{|G|/|H|} = |H| = |\eta(H)|.$$

But every set $S \subseteq G$ satisfies $\eta(S) \subseteq \text{Ann}(\text{Ann } S)$, so $\eta(H) = \text{Ann}(\text{Ann } H)$. \square

Fact 6. Let G be a finite abelian group. The map $\text{Ann} : G \rightarrow \widehat{G}$ induces a bijection from the set of subgroups of G to the set of subgroups of \widehat{G} . Moreover, for subgroups A, B of G ,

$$\begin{aligned} \text{Ann}(A \cap B) &= \text{Ann } A + \text{Ann } B && \text{and} \\ \text{Ann}(A + B) &= \text{Ann } A \cap \text{Ann } B. \end{aligned}$$

Proof. Let $\text{Sub } G$ denote the family of subgroups of G , ordered by inclusion. By fact 5, the map $\text{Ann} : \text{Sub } G \rightarrow \text{Sub } \widehat{G}$ is injective and the map $\text{Ann} : \text{Sub } \widehat{G} \rightarrow \text{Sub } \widehat{\widehat{G}}$ is surjective. But $G \cong \widehat{\widehat{G}}$ and thus $\text{Sub } G \cong \text{Sub } \widehat{\widehat{G}}$. Therefore, $\text{Ann} : \text{Sub } G \rightarrow \text{Sub } \widehat{G}$ is bijective. It is order-reversing in the sense that $C \subseteq D \implies \text{Ann } C \supseteq \text{Ann } D$ for every $C, D \in \text{Sub } G$. Likewise, the map $\text{Ann} : \text{Sub } \widehat{G} \rightarrow \text{Sub } \widehat{\widehat{G}}$ is order-reversing, so by fact 5 $\text{Ann } C \supseteq \text{Ann } D \implies C \subseteq D$ for every $C, D \in \text{Sub } G$. Thus, Ann is an order anti-isomorphism⁴ between the partially ordered sets $\text{Sub } G$ and $\text{Sub } \widehat{G}$.

Every pair of subgroups $A, B \in \text{Sub } G$ has greatest lower bound $A \cap B$ and least upper bound $A + B$ (i.e., $(\text{Sub } G, +, \cap)$ is a lattice), and likewise for $\text{Sub } \widehat{G}$. Any order anti-isomorphism sends greatest lower bounds to least upper bounds, and sends least upper bounds to greatest lower bounds, as can be seen by unrolling the definitions. \square

As a side note, without proof: the pair of maps $\text{Ann} : \mathcal{P}(G) \rightarrow \mathcal{P}(\widehat{G})$ and $\text{Ann} : \mathcal{P}(\widehat{G}) \rightarrow \mathcal{P}(\widehat{\widehat{G}})$ form an antitone Galois connection between the powerset lattices (after identifying G with $\widehat{\widehat{G}}$),

⁴An anti-isomorphism between partial orders P and Q is a bijection $\alpha : P \rightarrow Q$ such that $p_1 \leq p_2 \iff \alpha(p_1) \geq \alpha(p_2)$.

whose Galois closed elements are the subgroups.

The following result will be used in the proof of proposition 59.

Fact 7. *Let A and B be finite abelian groups, and let $\alpha : A \rightarrow B$ be a homomorphism. Then*

$$\begin{aligned} (\alpha^*)^{-1} \circ \text{Ann}_A &= \text{Ann}_B \circ \alpha \quad \text{and} \\ \alpha^* \circ (\alpha^*)^{-1} \circ \text{Ann}_A \circ \alpha^{-1} &= \alpha^* \circ \text{Ann}_B, \end{aligned}$$

where α , α^* , α^{-1} , and $(\alpha^*)^{-1}$ denote the respective induced maps between powersets (e.g., $\alpha^* : \mathcal{P}(\widehat{B}) \rightarrow \mathcal{P}(\widehat{A})$.)

If α^* is surjective, then

$$\text{Ann}_A \circ \alpha^{-1} = \alpha^* \circ \text{Ann}_B.$$

Proof. For every subset $S \subseteq A$,

$$\begin{aligned} ((\alpha^*)^{-1} \circ \text{Ann}_A)(S) &= \{\chi \in \widehat{B} \mid \alpha^*(\chi) \in \text{Ann}_A S\} \\ &= \{\chi \in \widehat{B} \mid \chi \circ \alpha \in \text{Ann}_A S\} \\ &= \{\chi \in \widehat{B} \mid \chi \in \text{Ann}_B \alpha(S)\} \\ &= (\text{Ann}_B \circ \alpha)(S). \end{aligned}$$

This proves the first identity. To prove the second identity from the first, compose each side with α^* on the left and α^{-1} on the right to get

$$\alpha^* \circ (\alpha^*)^{-1} \circ \text{Ann}_A \circ \alpha^{-1} = \alpha^* \circ \text{Ann}_B \circ \alpha \circ \alpha^{-1}.$$

The right-hand side here is equal to $\alpha^* \circ \text{Ann}_B$ because, for every $T \subseteq B$,

$$\begin{aligned} \alpha^*(\text{Ann}_B(\alpha(\alpha^{-1}(T)))) &= \alpha^*(\text{Ann}_B(T \cap \text{im } \alpha)) \\ &= \alpha^*(\text{Ann}_B T + \text{Ann}_B(\text{im } \alpha)) \\ &= \alpha^*(\text{Ann}_B T) + \alpha^*(\text{Ann}_B(\text{im } \alpha)) \\ &= \alpha^*(\text{Ann}_B T) + \{\chi \circ \alpha \mid \chi \in \text{Ann}_B(\text{im } \alpha)\} \\ &= \alpha^*(\text{Ann}_B T) + \{0\} \\ &= \alpha^*(\text{Ann}_B T). \end{aligned}$$

This proves the second identity. Finally, if α^* is surjective, then $\alpha^* \circ (\alpha^*)^{-1}$ is the identity on $\mathcal{P}(\widehat{A})$. \square

Fact 8. For every homomorphism α of finite abelian groups,

$$\begin{cases} \text{Ann}(\text{im } \alpha) = \ker(\alpha^*) & \text{and} \\ \text{Ann}(\ker \alpha) = \text{im}(\alpha^*). \end{cases}$$

Proof. Let A and B be finite abelian groups, and let $\alpha : A \rightarrow B$ be a homomorphism. Then

$$\begin{aligned} \text{Ann}(\text{im } \alpha) &= \{\chi \in \widehat{B} \mid \chi(b) = 0 \text{ for each } b \in \text{im } \alpha\} \\ &= \{\chi \in \widehat{B} \mid \chi \circ \alpha = 0\} = \ker(\alpha^*), \quad \text{and} \\ \text{Ann}(\ker \alpha^*) &= \text{Ann}(\text{Ann}(\text{im } \alpha)) = \eta_B(\text{im } \alpha) \\ &= \text{im}(\alpha^{**} \circ \eta_B) = \text{im}(\alpha^{**}). \end{aligned}$$

It's not hard to see that every homomorphism has a predual (because the map $\alpha \mapsto \alpha^*$ is a bijection from $\text{hom}(A, B)$ to $\text{hom}(\widehat{B}, \widehat{A})$.) Thus, we may replace α^* with α to obtain $\text{Ann}(\ker \alpha) = \text{im}(\alpha^*)$. \square

Recall the structure theorem for finite abelian groups [Pey20, §1.2.2]; [Rom12, Theorem 5.8], which states that there exists a direct sum decomposition

$$G \cong \bigoplus_{1 \leq i \leq n} \mathbb{Z}/q_i\mathbb{Z}$$

where $n \geq 0$ and $q_1, \dots, q_n \geq 2$. By the Chinese remainder theorem, it's possible to arrange for each q_i to be a prime power.

Definition 9. Given a direct sum of groups $G = \bigoplus_{i \in \mathcal{J}} G_i$ or a direct product of groups $G = \prod_{i \in \mathcal{J}} G_i$,⁵ the *coordinate projections* are the maps $\rho_j : G \rightarrow G_j$, $(g_i)_{i \in \mathcal{J}} \mapsto g_j$, and the *coordinate injections* are the maps $\kappa_j : G_j \rightarrow G$, $g_j \mapsto (g_j \text{ for } i = j, 0 \text{ for } i \neq j)_{i \in \mathcal{J}}$.

Let $\omega_j := \kappa_j \circ \rho_j : G \rightarrow G$. (The map ω_j sets to 0 all coordinates except the j th.)

Wherever it's necessary to be explicit about the group, we'll instead write $\rho_{G,j}$, $\kappa_{G,j}$, $\omega_{G,j}$. \triangle

Let $k \geq 1$ and take $\zeta_k = e^{2\pi i/k}$. It's easy to show that there is an isomorphism $\mathbb{Z}/k\mathbb{Z} \cong \widehat{\mathbb{Z}/k\mathbb{Z}}$ given by $g \mapsto (h \mapsto \zeta_k^{hg})$, where hg is the product of natural numbers h, g [CST18, p. 50]. Moreover,

⁵Recall that the *direct product* is the set of all tuples whereas the *direct sum* is the set of all finitely-supported tuples, both with componentwise group operation [Rom12, p. 152].

each character χ of a direct sum $\bigoplus_{1 \leq i \leq n} G_i$ is specified uniquely by the characters $\chi \circ \kappa_i$ of the components, and this identification gives an isomorphism

$$\widehat{\bigoplus_{1 \leq i \leq n} G_i} \cong \bigoplus_{1 \leq i \leq n} \widehat{G_i}.$$

Combining this with the structure theorem proves that $G \cong \widehat{\widehat{G}}$ holds for every finite abelian group G : If G comes with a specified decomposition into cyclic groups, and a generator is specified for each of those cyclic groups, then there's a corresponding isomorphism $G \cong \widehat{\widehat{G}}$ via "components"⁶

$$(g_1, \dots, g_n) \mapsto ((h_1, \dots, h_n) \mapsto \zeta_{k_1}^{h_1 g_1} \dots \zeta_{k_n}^{h_n g_n}). \quad (3)$$

Let G_1, \dots, G_n and H_1, \dots, H_n be finite abelian groups and let $\alpha_i : G_i \rightarrow H_i$ be a homomorphism for each $i = 1, \dots, n$. Define their direct sum $\alpha = \bigoplus_i \alpha_i : \bigoplus_i G_i \rightarrow \bigoplus_i H_i$ as

$$\left(\bigoplus_{1 \leq i \leq n} \alpha_i \right) (g_1, \dots, g_n) = (\alpha_1(g_1), \dots, \alpha_n(g_n)).$$

Then we may take the dual $\alpha^* : \widehat{\bigoplus_i H_i} \rightarrow \widehat{\bigoplus_i G_i}$ "componentwise" as

$$\left(\bigoplus_{1 \leq i \leq n} \alpha_i \right)^* = \bigoplus_{1 \leq i \leq n} \alpha_i^*$$

(the equality is to be interpreted in the sense of our componentwise identifications $\widehat{\bigoplus G_i} \cong \bigoplus \widehat{G_i}$, $\widehat{\bigoplus H_i} \cong \bigoplus \widehat{H_i}$.) For example, if $G_i = H_i$ and if each $\alpha_i : G_i \rightarrow G_i$ is either the zero map or the identity map, then the same is true of the dual $(\bigoplus_i \alpha_i)^* : \widehat{\bigoplus_i G_i} \rightarrow \widehat{\bigoplus_i G_i}$. In other words, if $\bigoplus_i \alpha_i$ may be expressed as a projection onto some subset of coordinates followed by an injection back into $\bigoplus_i G_i$, then the same holds for $(\bigoplus_i \alpha_i)^*$. This will be useful in section 4, so we'll present two cases formally as facts 10 and 11.

Fact 10. *Let $G = G_1 \oplus \dots \oplus G_n$ (where $n \geq 1$) for finite abelian groups G, G_1, \dots, G_n . Identify $\widehat{\widehat{G}}$ with*

⁶It's often pointed out that this isomorphism is non-canonical, and that is indeed the case if we work in the unadorned category of groups. But it is natural in the category of finite abelian groups decorated with decompositions into cyclic subgroups and with specified generators for each cyclic component (also, the isomorphism gives special status to the first primitive q th root of unity for each q .) Unfortunately, the covariant functor involved with this natural isomorphism is uninteresting, and instead we care about the contravariant functor sending a morphism α to its dual α^* . This is why it's conceptually cleaner to keep separate the two notions of a group and its dual: They are different objects because they play a role in different operations. This will be made more clear in section 2.2, where α and α^* will be the boundary and coboundary map, respectively.

$\widehat{G}_1 \oplus \cdots \oplus \widehat{G}_n$ as above (that is, a character $\chi \in \widehat{G}$ is identified with $(\chi_1, \dots, \chi_n) \in \widehat{G}_1 \oplus \cdots \oplus \widehat{G}_n$ where $\chi((g_1, \dots, g_n)) = \chi_1(g_1)\chi_2(g_2)\cdots\chi_n(g_n)$.)

For every $j = 1, \dots, n$, the dual maps of the coordinate projection $\rho_{G,j} : G \rightarrow G_j$ and the coordinate injection $\kappa_{G,j} : G_j \rightarrow G$ are the coordinate injection $\kappa_{\widehat{G},j} : \widehat{G}_j \rightarrow \widehat{G}$ and the coordinate projection $\rho_{\widehat{G},j} : \widehat{G} \rightarrow \widehat{G}_j$, respectively.

Proof. For readability, we prove the result for $n = 2$ (the proof extends in the obvious way to arbitrary n .)

For all $g_i \in G_i$ and $\chi_i \in \widehat{G}_i$ (where $i = 1, 2$),

$$\begin{aligned} (\rho_{G,1}^*(\chi_1))(g_1, g_2) &= \chi_1(\rho_{G,1}(g_1, g_2)) \\ &= \chi_1(g_1) \\ &= (\chi_1, 0)(g_1, g_2), \end{aligned}$$

which proves $\rho_{G,1}^* = \kappa_{\widehat{G},1}$, and

$$\begin{aligned} \kappa_{G,1}^*((\chi_1, \chi_2))(g_1) &= (\chi_1, \chi_2)(\kappa_{G,1}(g_1)) \\ &= (\chi_1, \chi_2)(g_1, 0) \\ &= \chi_1(g_1), \end{aligned}$$

which proves $\kappa_{G,1}^* = \rho_{\widehat{G},1}$. □

Fact 11. Let $G = G_1 \oplus \cdots \oplus G_n$ (where $n \geq 1$) for finite abelian groups G, G_1, \dots, G_n . Identify \widehat{G} with $\widehat{G}_1 \oplus \cdots \oplus \widehat{G}_n$ as above (that is, a character $\chi \in \widehat{G}$ is identified with $(\chi_1, \dots, \chi_n) \in \widehat{G}_1 \oplus \cdots \oplus \widehat{G}_n$ where $\chi((g_1, \dots, g_n)) = \chi_1(g_1)\chi_2(g_2)\cdots\chi_n(g_n)$.)

The duality relation $\omega_{G,j}^* = \omega_{\widehat{G},j}$ holds for every $j = 1, \dots, n$ (see definition 9.)

Proof. By fact 10, the composition $\omega_{G,j} = \kappa_{G,j} \circ \rho_{G,j}$ has dual $\omega_{G,j}^* = (\kappa_{G,j} \circ \rho_{G,j})^* = \rho_{G,j}^* \circ \kappa_{G,j}^* = \kappa_{\widehat{G},j} \circ \rho_{\widehat{G},j} = \omega_{\widehat{G},j}$. □

The Fourier transform [CST18, §2.4] of a function $f : G \rightarrow \mathbb{C}$ (for a finite abelian group G) is the function $\widehat{f} : \widehat{G} \rightarrow \mathbb{C}$ (also denoted $\mathcal{F}\{f\}$) defined as

$$\widehat{f}(\chi) \equiv \mathcal{F}\{f\}(\chi) = \sum_{g \in G} f(g)\overline{\chi(g)}, \quad \chi \in \widehat{G}. \quad (4)$$

Recall that the characters are linearly independent over \mathbb{C} ; in fact, any two distinct characters are orthogonal with respect to the standard inner product $\langle f_1, f_2 \rangle = \sum_{g \in G} f_1(g) \overline{f_2(g)}$ [CST18, §2.3]. The characters form a basis for the \mathbb{C} -vector space \mathbb{C}^G : Every function $f : G \rightarrow \mathbb{C}$ may be expressed as a linear combination of characters using the Fourier inversion formula [CST18, §2.4]

$$f = \frac{1}{|G|} \sum_{\chi \in \widehat{G}} \widehat{f}(\chi) \chi.$$

In particular, the Fourier transform $\mathcal{F} : \mathbb{C}^G \rightarrow \mathbb{C}^{\widehat{G}}$, $f \mapsto \widehat{f}$ is an isomorphism (between vector spaces over \mathbb{C} .)

Next, we derive a simple result that will be used to prove theorem 41.

Fact 12. *Let G be a finite abelian group and H a subgroup. Then*

$$\frac{1}{|\text{Ann } H|} \sum_{\chi \in \text{Ann } H} \chi(g) = \mathbb{1}[g \in H], \quad g \in G.$$

Proof. Consider the homomorphism $e_g : \text{Ann } H \rightarrow \mathbb{C}$, $\chi \mapsto \chi(g)$ and its fibers $e_g^{-1}(z)$, $z \in \mathbb{C}$. By the first homomorphism theorem for groups, all nonempty fibers have equal cardinality. The image of e_g is a finite subgroup of the circle group $\mathbb{T} \subseteq \mathbb{C}$, so either it coincides with $\{1\}$ or is rotationally symmetric about 0. These two cases are distinguished by $g \in H$ because of fact 5 (we have $\eta(g) = e_g \in \text{Ann Ann } H$ if and only if $g \in H$). \square

Fact 12 is equivalent to the Poisson summation formula—which we won't need directly, but include anyway for its aesthetic value. This formula may be found in [Ter99, p. 199] and [CST18, p. 60].

Fact 13 (Poisson summation formula). *Let G be a finite abelian group and H a subgroup. For every function $f : G \rightarrow \mathbb{C}$, the averages over the cosets of H are*

$$\frac{1}{|H|} \sum_{h \in H} f(gh) = \frac{1}{|G|} \sum_{\chi \in \text{Ann } H} \widehat{f}(\chi) \chi(g), \quad g \in G.$$

Proof. For every $h \in H$, the Fourier transform (eq. (4)) of the function $g \mapsto f(gh)$, evaluated at a

character $\psi \in \widehat{G}$, is

$$\begin{aligned}
\mathcal{F}\{g \mapsto f(gh)\}(\psi) &= \sum_{g \in G} f(gh) \overline{\psi(g)} \\
&= \sum_{g \in G} f(g) \overline{\psi(gh^{-1})} \quad (\text{by change of variables } g := gh) \\
&= \psi(h) \sum_{g \in G} f(g) \overline{\psi(g)} \\
&= \psi(h) \widehat{f}(\psi).
\end{aligned}$$

Accordingly, the Fourier transform of the left-hand side of the Poisson summation formula, as a function of g , is

$$\psi \mapsto \frac{\widehat{f}(\psi)}{|H|} \sum_{h \in H} \psi(h), \quad \psi \in \widehat{G}.$$

The right-hand side of the Poisson summation formula has Fourier transform

$$\begin{aligned}
\psi &\mapsto \frac{1}{|G|} \sum_{\chi \in \text{Ann } H} \widehat{f}(\chi) \langle \chi, \psi \rangle \\
&= \frac{1}{|G|} \llbracket \psi \in \text{Ann } H \rrbracket \widehat{f}(\psi) |G| \\
&= \llbracket \psi \in \text{Ann } H \rrbracket \widehat{f}(\psi), \quad \psi \in \widehat{G}.
\end{aligned}$$

So the Fourier transforms of the two sides coincide if and only if

$$\frac{1}{|H|} \sum_{h \in H} \psi(h) = \llbracket \psi \in \text{Ann } H \rrbracket, \quad \psi \in \widehat{G}.$$

But this follows by duality (fact 5) from fact 12 by taking \widehat{G} in place of G and $\text{Ann } H$ in place of H . □

To justify the earlier statement that fact 12 and fact 13 are equivalent, here is a proof in the other direction, too.

Proof of fact 12 from fact 13. Let f be the indicator of the identity, $f(g) = \llbracket g = 0 \rrbracket$. Then $\widehat{f} \equiv 1$. The Poisson summation formula (fact 13) gives

$$\frac{1}{|H|} \sum_{h \in H} \llbracket g = h^{-1} \rrbracket = \frac{1}{|G|} \sum_{\chi \in \text{Ann } H} \chi(g).$$

The sum on the left-hand side vanishes if $g \notin H$ and evaluates to 1 if $g \in H$, so this formula reduces to

$$[[g \in H]] = \frac{1}{|G|/|H|} \sum_{\chi \in \text{Ann } H} \chi(g).$$

Now put $|\text{Ann } H| = |G|/|H|$ (fact 4.) □

The following is an unrelated inequality that we'll use to prove theorems 35 and 62.

Fact 14. *Let G and G' be finite abelian groups and let $\alpha : G \rightarrow G'$ be a homomorphism. Let A and B be subgroups of G , and let D be a subgroup of G' . Then*

$$|\alpha(A + B) + D| |\alpha(A \cap B) + D| \leq |\alpha(A) + D| |\alpha(B) + D|.$$

In particular,

$$|\alpha(A + B)| |\alpha(A \cap B)| \leq |\alpha(A)| |\alpha(B)|.$$

Proof. The special case follows from the general case by putting $D = \{0\}$. But for clarity we'll prove the special case first, and extend to the general case by passing to the quotient G'/D .

The map α satisfies

$$\alpha(A + B) = \alpha(A) + \alpha(B) \quad \text{and} \quad \alpha(A \cap B) \subseteq \alpha(A) \cap \alpha(B).$$

For any two subgroups $K, N \subseteq G'$ the second isomorphism theorem states $(K + N)/N \cong K/(K \cap N)$ and therefore $|K + N| |K \cap N| = |K| |N|$. Taking $K = \alpha(A)$ and $N = \alpha(B)$, and combining with the preceding identity and inclusion, gives

$$\begin{aligned} |\alpha(A + B)| |\alpha(A \cap B)| &\leq |\alpha(A) + \alpha(B)| |\alpha(A) \cap \alpha(B)| \\ &= |\alpha(A)| |\alpha(B)|. \end{aligned}$$

This proves the special case $D = \{0\}$.

Now let $\bar{\alpha} = \pi \circ \alpha : G \rightarrow G'/D$ where $\pi : G' \rightarrow G'/D, g' \mapsto g' + D$. For every subset $S \subseteq G$,

$$\frac{1}{|D|} |\alpha(S) + D| = |\bar{\alpha}(S)|.$$

Put $\bar{\alpha}$ in place of α to get $|\bar{\alpha}(A + B)| |\bar{\alpha}(A \cap B)| \leq |\bar{\alpha}(A)| |\bar{\alpha}(B)|$, that is,

$$\frac{1}{|D|} |\alpha(A + B) + D| \frac{1}{|D|} |\alpha(A \cap B) + D| \leq \frac{1}{|D|} |\alpha(A) + D| \frac{1}{|D|} |\alpha(B) + D|.$$

Multiplying through by $|D|^2$ completes the proof. \square

Finally, although so far we have been discussing finite groups only, we will need two definitions and a result about infinite direct products and direct sums for section 5.

Definition 15. Let G be a finite abelian group and \mathcal{J} a set (not necessarily finite.) We will write

$$G^{\mathcal{J}} := \bigotimes_{i \in \mathcal{J}} G = \{f : \mathcal{J} \rightarrow G\} \quad \text{and}$$

$$G^{(\mathcal{J})} := \bigoplus_{i \in \mathcal{J}} G = \{f : \mathcal{J} \rightarrow G \mid f(i) = 0 \text{ for all but finitely many } i \in \mathcal{J}\},$$

(The former is the direct product and the latter is the direct sum. Recall that they are both abelian groups with componentwise group operation, $(f + f')(i) = f(i) + f'(i)$ for $i \in \mathcal{J}$.) \triangle

For infinite abelian groups, the definitions of *character* and *dual group* are unchanged from before. But in general it's no longer the case that $G \cong \widehat{\widehat{G}}$ for infinite G . However, we will have use for the following result, which may be compared to the discussion on page 13 about the finite case.

Fact 16. Let G be a finite abelian group, and let \mathcal{J} be a set (not necessarily finite.) There is a group isomorphism

$$\psi : \widehat{G^{(\mathcal{J})}} \rightarrow \widehat{G}^{\mathcal{J}}, \quad \chi \mapsto (i \mapsto \chi_i)$$

where

$$\chi_i : G \rightarrow \mathbb{C}, \quad g \mapsto \chi(\kappa_i(g)),$$

$$\kappa_i(g) : \mathcal{J} \rightarrow G, \quad j \mapsto \begin{cases} g, & i = j; \\ 0, & i \neq j. \end{cases}$$

Proof. To check that ψ is a group homomorphism, we verify

$$\begin{aligned}
\psi(\chi - \chi') &= (i \mapsto (\chi - \chi')_i) \\
&= (i \mapsto (g \mapsto (\chi - \chi')(\kappa_i(g)))) \\
&= (i \mapsto (g \mapsto \chi(\kappa_i(g))\chi'(\kappa_i(g))^{-1})) \\
&= (i \mapsto (g \mapsto \chi_i(g)\chi'_i(g)^{-1})) \\
&= (i \mapsto (\chi_i - \chi'_i)) \\
&= (i \mapsto \chi_i) - (i \mapsto \chi'_i) \\
&= \psi(\chi) - \psi(\chi').
\end{aligned}$$

To check that ψ is a bijection, we claim that it has inverse

$$\psi^{-1} : \widehat{G}^J \rightarrow \widehat{G}^{(J)}, \quad (\zeta_i)_{i \in J} \mapsto \left((g_i)_{i \in J} \mapsto \prod_{i \in J} \zeta_i(g_i) \right),$$

where the sum is finite because the element $(g_i)_{i \in J} \in G^{(J)}$ has finite support. We verify

$$\begin{aligned}
(\psi \circ \psi^{-1})((\zeta_i)_{i \in J}) &= \psi \left((g_i)_{i \in J} \mapsto \prod_{i \in J} \zeta_i(g_i) \right) \\
&= i \mapsto \left(g \mapsto \left((g_i)_{i \in J} \mapsto \prod_{i \in J} \zeta_i(g_i) \right) (\kappa_i(g)) \right) \\
&= i \mapsto (g \mapsto \zeta_i(g)) \\
&= (\zeta_i)_{i \in J}, \quad (\zeta_i)_{i \in J} \in \widehat{G}^J
\end{aligned}$$

and

$$\begin{aligned}
(\psi^{-1} \circ \psi)(\chi) &= \psi^{-1}(i \mapsto \chi_i) \\
&= \left((g_i)_{i \in J} \mapsto \prod_{i \in J} \chi_i(g_i) \right) \\
&= \left((g_i)_{i \in J} \mapsto \prod_{i \in J} \chi(\kappa_i(g_i)) \right) \\
&= ((g_i)_{i \in J} \mapsto \chi((g_i)_{i \in J})) \quad (\text{because } \chi \text{ is a homomorphism}) \\
&= \chi, \quad \chi \in \widehat{G}^{(J)}.
\end{aligned}$$

□

Throughout, we deal only with abelian groups. Non-abelian groups are more challenging. It becomes necessary to keep track of higher-dimension representations, because there aren't enough characters. In general, for a finite group G , its group of characters \widehat{G} is isomorphic to $G/[G, G]$, where $[G, G]$ is the commutator subgroup [Pey20, §1.3.2].

For a more thorough introduction to these ideas, refer to [CST18, Chapter 1], [Pey20, Chapters 1, 2], and [Ter99, Part I].

2.2 Cubical homology

We describe the homology theory of certain subsets of \mathbb{R}^d called “cubical sets”. This is essentially a special case of simplicial homology, where the simplicial complexes comprise axis-aligned cubes in \mathbb{R}^d .

We adopt the formalism of computational homology. This is for pragmatic and aesthetic reasons. First, homology has numerous practical applications to data analysis [PR15], and as a result there are many software packages available for computing homology groups and related invariants. It will be easier to set up simulations later if we don't have to translate between different conventions. Second, the notation is conceptually crisp, in that it can be quickly defined in full rigor without any knowledge of differential forms. This makes it (perhaps) more accessible to probability theorists.

A competing formalism, the discrete exterior calculus, is ubiquitous in the existing literature on lattice gauge theory. It might be said to enjoy the benefit of greater geometrical clarity—for instance, the idea of orientation is made explicit. Readers familiar with the exterior calculus of differential forms may wish to browse the “dictionary”, table 2. It should be stressed that there is no fundamental mathematical difference between the two formalisms.

Let G be an abelian group (not necessarily finite), and fix an integer dimension $d \geq 1$.

An *elementary cube of dimension* $k \geq 0$ (or a *k-cube*) is a unit cube $[0, 1]^k$ embedded in \mathbb{R}^d with vertices lying on the integer lattice \mathbb{Z}^d . In other words, an elementary cube is a cartesian product $Q = I_1 \times \cdots \times I_d$ where for each $1 \leq i \leq d$ either $I_i = [n_i, n_i + 1]$ or $I_i = \{n_i\}$ for some $n_i \in \mathbb{Z}$, and its *dimension* $\dim Q$ is the dimension of its affine hull or, equivalently, the number of non-degenerate factors in the cartesian product (here *degenerate* means a singleton, $\{n_i\}$.) For elementary cubes Q' and Q , we say that Q' is a *face* of Q if $Q' \subseteq Q$, and a *primary face* or *facet* if it is a face with $\dim Q' = \dim Q - 1$.

The *elementary cell* $\overset{\circ}{Q}$ associated with an elementary cube Q is the relative interior of Q (i.e.,

the interior of Q considered as a subset of its affine hull). It's geometrically intuitive that every elementary cube is the disjoint union of the relative interiors of its faces (this can be proved by induction on the dimension of the elementary cube, and holds more generally for all convex polytopes [Zie94, p. 61].)

Although we've defined elementary cubes as subsets of \mathbb{R}^d , we'll employ them only as combinatorial elements—what's relevant is only their dimensions and inclusions between them.

The collection of all elementary cubes is denoted by \mathcal{K} ; the collection of all elementary cubes of dimension $k \in \mathbb{Z}$ is denoted by \mathcal{K}_k . Thus, $\mathcal{K} = \bigcup_{k=0}^d \mathcal{K}_k$, and $\mathcal{K}_k = \emptyset$ for $k < 0$ and $k > d$.

A *cubical set* is a subset $X \subseteq \mathbb{R}^d$ that can be expressed as a finite (possibly empty) union of elementary cubes, not necessarily all of the same dimension. For a cubical set X we define $\mathcal{K}_k(X) := \{Q \in \mathcal{K}_k \mid Q \subseteq X\}$ and $\mathcal{K}(X) := \{Q \in \mathcal{K} \mid Q \subseteq X\}$. The collection of all elementary cubes in X together with their inclusions and dimensions, $(\mathcal{K}(X), \subseteq, \dim)$, is called the *cubical complex* associated with X .

As an example, consider the cubical set $X = [0, 1]^d \subseteq \mathbb{R}^d$. Then $\mathcal{K}_d(X) = \{[0, 1]^d\}$ and $\mathcal{K}_0(X)$ consists of 2^d singletons (the vertices of X .) It's straightforward to show that $|\mathcal{K}_k(X)| = \binom{d}{k} 2^{d-k}$ for $0 \leq k \leq d$.

Note that every $(k-1)$ -cube (i.e., elementary cube of dimension $k-1$) is a face of some k -cube, that is, $\bigcup \mathcal{K}_{k-1} \subseteq \bigcup \mathcal{K}_k$ for $1 \leq k \leq d$, but it is not the case that $\bigcup \mathcal{K}_{k-1}(X) \subseteq \bigcup \mathcal{K}_k(X)$ holds for every cubical set X (for example, if X is the singleton $\{0\}$ then $\bigcup \mathcal{K}_0(X) = \{0\} \not\subseteq \emptyset = \bigcup \mathcal{K}_1(X)$.)

If X is a cubical set with $\mathcal{K}_k(X) = \emptyset$ for all $k \geq 2$, then the cubical complex $\mathcal{K}(X)$ may be viewed as a finite graph $G = (V, E) = (\mathcal{K}_0(X), \mathcal{K}_1(X))$ embedded in \mathbb{R}^d .

A (*cubical*) *chain* of dimension $k \in \mathbb{Z}$ (or simply a *k-chain*) with coefficients in G is a finitely-supported map $\mathcal{K}_k \rightarrow G$. More formally, define $C_k(G)$ to be the direct sum $\bigoplus_{\mathcal{K}_k} G$; the elements of $C_k(G)$ are called *k-chains*. For a chain $c \in C_k(G)$ and for $Q \in \mathcal{K}_k$, the group element $c(Q) \in G$ is called the *coefficient* of Q in c .

Now fix some cubical set $X \subseteq \mathbb{R}^d$. For every $k \in \mathbb{Z}$ we define $C_k(X, G) := \bigoplus_{\mathcal{K}_k(X)} G$. An element $c \in C_k(X, G)$ will be called a *k-chain in X*. So a *k-chain in X* is an assignment of an element of G to each k -cube in X . In particular, $C_k(X, G)$ is the trivial group whenever $k < 0$ or $k > d$. We identify $C_k(X, G)$ with a subgroup of $C_k(G)$ in the obvious way, by putting coefficient 0 on every k -cube outside X .

If $Q \in \mathcal{K}_k(X)$ and $g \in G$, then g_Q will denote the k -chain

$$g_Q(P) := \begin{cases} g, & Q = P; \\ 0, & Q \neq P. \end{cases}$$

In particular, if G is the additive group of a ring with additive identity 0 and multiplicative identity 1 , then the element $1_Q \in C_k(X, G)$ is the indicator of Q .

Next, we define a boundary operator, which sends a k -chain to a $(k-1)$ -chain. For a k -cube $Q \in \mathcal{K}_k$ and for $g \in G$, the following definition specifies the boundary of g_Q to be a chain supported on the facets of Q whose nonzero coefficients have values $\pm g$, with opposing signs on each pair of opposing facets, and extended by additivity to all of $C_k(G)$. To make this precise, write $Q = I_1 \times \cdots \times I_d \subseteq \mathbb{R}^d$ where each factor I_j is either degenerate (a singleton) or a unit interval, and label the nondegenerate factors as I_{i_1}, \dots, I_{i_k} with $i_1 < \dots < i_k$, assuming for now that $1 \leq k \leq d$. For $1 \leq j \leq k$, denote the j th pair of facets of Q as

$$\begin{aligned} Q_j^- &:= I_1 \times \cdots \times I_{i_{j-1}} \times \{m_j\} \times I_{i_{j+1}} \times \cdots \times I_d, \\ Q_j^+ &:= I_1 \times \cdots \times I_{i_{j-1}} \times \{m_j + 1\} \times I_{i_{j+1}} \times \cdots \times I_d \quad \text{where } I_{i_j} = [m_j, m_j + 1]. \end{aligned}$$

For $1 \leq k \leq d$, the *boundary operator* $\partial_k : C_k(G) \rightarrow C_{k-1}(G)$ is

$$\partial_k g_Q := \sum_{1 \leq j \leq k} (-1)^{j-1} (g_{Q_j^+} - g_{Q_j^-}),$$

extended by additivity to all of $C_k(G)$. Note that each elementary cube comes with an ordering on its coordinates, and the boundary operator depends on this ordering: some of the signs of the boundary will change if coordinates are permuted, because of the $(-1)^{j-1}$ factor. This defines $\partial_k : C_k(G) \rightarrow C_{k-1}(G)$ for $1 \leq k \leq d$. Observe that the operators ∂_k are group homomorphisms. For $k \leq 0$ and $k > d$, define $\partial_k : C_k(G) \rightarrow C_{k-1}(G)$ to be the zero homomorphism (this is the only option because $C_k(G)$ is trivial for $k < 0$ and $k > d$.) An alternative definition for ∂_k is given in [KMM04], where the definition above is listed instead as a result [KMM04, Proposition 2.36].

We also define $\partial_k : C_k(X, G) \rightarrow C_{k-1}(X, G)$ to be the restriction of ∂_k to $C_k(X, G)$. To keep the notation compact, we'll avoid indicating the domain explicitly (by writing something like $\partial_k(X, G)$.) The domain for ∂_k will always be clear from context. Furthermore, in most contexts it won't cause confusion to write ∂ in place of ∂_k .

A helpful special case to consider is $G = \mathbb{Z}/2\mathbb{Z}$, where $+1 = -1$ so all signs on coefficients may be ignored. For more general G , negating a coefficient on an elementary cube may be thought of as reversing the orientation of the cube (though we do not explicitly define orientation.) The alternating-sign convention in the definition of ∂_k is arranged for the sake of the following boundary-of-boundary result (fact 17). This result is essential to every homology theory. There is a proof by induction on d in [KMM04, Proposition 2.37], which is written for $G = \mathbb{Z}$ but carries over to the general case with only trivial modifications. The proof presented below is much simpler.⁷

Fact 17 (Boundary relation). *For every $k \in \mathbb{Z}$,*

$$\partial_{k-1} \circ \partial_k = 0.$$

Proof. For $k \leq 1$ and $k > d$ the result is trivial, so assume $2 \leq k \leq d$. With Q as in the definition of ∂_k , for $1 \leq j < \ell \leq k$ let

$$\begin{aligned} Q_{j,\ell}^- &= I_1 \times \cdots \times I_{i_{j-1}} \times \{m_j\} \times I_{i_{j+1}} \times \cdots \times I_{i_{\ell-1}} \times \{m_\ell\} \times I_{i_{\ell+1}} \times \cdots \times I_d, \\ Q_{j,\ell}^{+-} &:= I_1 \times \cdots \times I_{i_{j-1}} \times \{m_j\} \times I_{i_{j+1}} \times \cdots \times I_{i_{\ell-1}} \times \{m_\ell + 1\} \times I_{i_{\ell+1}} \times \cdots \times I_d, \\ Q_{j,\ell}^{+ -} &:= I_1 \times \cdots \times I_{i_{j-1}} \times \{m_j + 1\} \times I_{i_{j+1}} \times \cdots \times I_{i_{\ell-1}} \times \{m_\ell\} \times I_{i_{\ell+1}} \times \cdots \times I_d, \\ Q_{j,\ell}^{++} &:= I_1 \times \cdots \times I_{i_{j-1}} \times \{m_j + 1\} \times I_{i_{j+1}} \times \cdots \times I_{i_{\ell-1}} \times \{m_\ell + 1\} \times I_{i_{\ell+1}} \times \cdots \times I_d, \end{aligned}$$

where $I_{i_j} = [m_j, m_j + 1]$ and $I_{i_\ell} = [m_\ell, m_\ell + 1]$. Applying the definition of the boundary operator twice shows

$$\begin{aligned} (\partial_{k-1} \circ \partial_k)(g_Q) &:= \sum_{1 \leq j' < j \leq k} (-1)^{j-1} (-1)^{j'-1} (g_{Q_{j',j}^{++}} - g_{Q_{j',j}^{+-}} - g_{Q_{j',j}^{+ -}} + g_{Q_{j',j}^{--}}) \\ &\quad + \sum_{1 \leq j \leq j' < k} (-1)^{j-1} (-1)^{j'-1} (g_{Q_{j,j'+1}^{++}} - g_{Q_{j,j'+1}^{+-}} - g_{Q_{j,j'+1}^{+ -}} + g_{Q_{j,j'+1}^{--}}) \\ &= \sum_{1 \leq j' < j \leq k} (-1)^{j-1} (-1)^{j'-1} (g_{Q_{j',j}^{++}} - g_{Q_{j',j}^{+-}} - g_{Q_{j',j}^{+ -}} + g_{Q_{j',j}^{--}}) \\ &\quad + \sum_{1 \leq j < j'' \leq k} (-1)^{j-1} (-1)^{j''-2} (g_{Q_{j,j''}^{++}} - g_{Q_{j,j''}^{+-}} - g_{Q_{j,j''}^{+ -}} + g_{Q_{j,j''}^{--}}) \\ &= 0. \end{aligned}$$

□

Example 18. Let Q be a 2-cube (a *plaquette*). Then $\partial_2 g_Q$ is supported on its four edges. The signs

⁷Thank you to Anthony Quas for pointing out this “two-line proof”.

on the opposing sides of the plaquette may be thought of as orientations for the edges. If the edges are visualized as directed arrows, then the four arrows all sit head-to-tail. The boundary of each edge is supported on its endpoints, and the net contribution to each of the four vertices is 0.

More explicitly, let $d = 2$ and $Q = [0, 1] \times [0, 1] \subseteq \mathbb{R}^2$. In the notation above, $I_1 = I_2 = [0, 1]$ so $i_1 = 1$ and $i_2 = 2$. We have

$$\begin{aligned} Q_1^- &= \{0\} \times [0, 1], & Q_2^- &= [0, 1] \times \{0\}, \\ Q_1^+ &= \{1\} \times [0, 1], & Q_2^+ &= [0, 1] \times \{1\}. \end{aligned}$$

So for every $g \in G$

$$\partial_2 g_Q = g_{Q_1^+} - g_{Q_1^-} - g_{Q_2^+} + g_{Q_2^-}.$$

Now consider $Q_1^- = \{0\} \times [0, 1]$; this has one nondegenerate interval, $i_1 = 2$. So (and similarly for the other three 1-cubes)

$$\begin{aligned} (Q_1^-)_1^- &= \{0\} \times \{0\}, & (Q_2^-)_1^- &= \{0\} \times \{0\}, \\ (Q_1^-)_1^+ &= \{0\} \times \{1\}, & (Q_2^-)_1^+ &= \{1\} \times \{0\}, \\ (Q_1^+)_1^- &= \{1\} \times \{0\}, & (Q_2^+)_1^- &= \{0\} \times \{1\}, \\ (Q_1^+)_1^+ &= \{1\} \times \{1\}, & (Q_2^+)_1^+ &= \{1\} \times \{1\}. \end{aligned}$$

Therefore, writing $(i, j) := \{i\} \times \{j\}$ as shorthand,

$$\begin{aligned} \partial_1(\partial_2 g_Q) &= \partial_1 g_{Q_1^+} - \partial_1 g_{Q_1^-} - \partial_1 g_{Q_2^+} + \partial_1 g_{Q_2^-} \\ &= (g_{(Q_1^+)_1^+} - g_{(Q_1^+)_1^-}) - (g_{(Q_1^-)_1^+} - g_{(Q_1^-)_1^-}) - (g_{(Q_2^+)_1^-} - g_{(Q_2^+)_1^+}) + (g_{(Q_2^-)_1^+} - g_{(Q_2^-)_1^-}) \\ &= (g_{(1,1)} - g_{(1,0)}) - (g_{(0,1)} - g_{(0,0)}) - (g_{(1,1)} - g_{(0,1)}) + (g_{(1,0)} - g_{(0,0)}) \\ &= 0. \end{aligned} \quad \triangle$$

The collection $(C_k(G), \partial_k)_{k \in \mathbb{Z}}$ is called the *cubical chain complex for \mathbb{R}^d with coefficients in G* . The collection $(C_k(X, G), \partial_k)_{k \in \mathbb{Z}}$ is called the *cubical chain complex for X with coefficients in G* .

The groundwork is now in place to introduce the central concepts of homology theory.

Again, fix a cubical set $X \subseteq \mathbb{R}^d$. For $k \in \mathbb{Z}$, a k -cycle in X is an element of $Z_k(X, G) := \ker \partial_k$. A k -boundary in X is an element of $B_k(X, G) := \text{im } \partial_{k+1}$. Since the boundary map is a

group homomorphism, both $Z_k(X, G)$ and $B_k(X, G)$ are subgroups of $C_k(X, G)$. By fact 17, every boundary is a cycle. To see how the converse can fail, let X be the union of the four edges of a single plaquette (i.e., 2-cube.) Let $G = \mathbb{Z}/2\mathbb{Z}$ and assign $1 \in G$ to each edge. The 1-chain so obtained is a cycle, but not a boundary, because the plaquette itself is missing from X .

Speaking informally, the failure of a k -cycle to be a k -boundary indicates the presence of a hole in the complex. If $k = 2$ the hole is a 3-dimensional void (as in Swiss cheese); if $k = 0$ the hole is an additional connected component (this last part will be made clear later.) But actually, the situation is more subtle: the presence and absence of holes depends on the coefficient group. For an illustration of this phenomenon, see section 6.1.

The k th homology group of X with coefficients in G is the quotient group $H_k(X, G) := Z_k(X, G)/B_k(X, G)$. The homology of X with coefficients in G is the sequence $H_*(X, G) := (H_k(X, G))_{k \in \mathbb{Z}}$. An element of a homology group (that is, an equivalence class of cycles modulo boundaries) is called a *homology class*, and any two cycles belonging to the same homology class are said to be *homologous* to one another.

The cases $G = \mathbb{Z}$, $G = \mathbb{Q}$, and $G = \mathbb{Z}/q\mathbb{Z}$ (for $q \geq 2$) are called the integral, rational, and mod- q homology. In the literature on algebraic homology, the omission of the coefficient group—writing simply $H_k(X)$ —usually indicates the integral homology.⁸

Example 19 (Homology of a graph). Let X be a finite union of 1-cubes and 0-cubes in \mathbb{Z}^d . Then X may be viewed as an undirected graph $G_X = (V, E) = (\mathcal{K}_0(X), \mathcal{K}_1(X))$. (Every finite subgraph of the hypercubic lattice has this form.) Then $H_0(X, \mathbb{Z}) = \mathbb{Z}^n$ where n is the number of connected components of X . To see why, notice that every 0-chain is a 0-cycle, and the 0-boundaries are generated by the chains of the form $1_v - 1_w$ where v and w are vertices belonging to the same component; therefore, each element of $H_0(X, \mathbb{Z})$ is an equivalence class of chains that all have the same sum of coefficients within each component (more generally, see fact 25.) Also, $H_1(X, \mathbb{Z}) = \mathbb{Z}^m$ for some $m \geq 0$, because there are no 1-boundaries and every nontrivial 1-cycle has order ∞ in the group $C_1(X, \mathbb{Z})$. The rank m may be thought of as the maximal number of independent cycles in the graph. All other homology groups are trivial. It can be shown that $n - m = |\mathcal{K}_0(X)| - |\mathcal{K}_1(X)|$ (this is closely related to Euler’s polyhedron formula, $V - E + F = 2$.) We’ll see a generalization of this identity in fact 24. △

Example 20 (Homology of an elementary cube). Let $X = [0, 1] \times \{0\} \subseteq \mathbb{R}^2$. In the integral homology,

⁸By virtue of a result known as the universal coefficient theorem, the integral homology groups determine the homology groups with coefficients in any other abelian group [Mun84, p. 313].

the 1-chains have the form $n_{[0,1]}$ where $n \in \mathbb{Z}$ (X has only one 1-cube, $[0, 1]$, and a chain is an assignment of a coefficient n to that 1-cube). So the 0-boundaries are the 0-chains $n_{\{1\}} - n_{\{0\}}$ for $n \in \mathbb{Z}$. The 0-cycles are the 0-chains, which have the form $a_{\{0\}} + b_{\{1\}}$ for $a, b \in \mathbb{Z}$. Thus, the homology group $H_0(X, \mathbb{Z})$, defined as the quotient of the 0-cycles by the 0-boundaries, is isomorphic to \mathbb{Z} . The homology group $H_1(X, \mathbb{Z})$ is trivial because the only 1-cycle is $0_{[0,1]}$.

More generally, if X is a single elementary cube of any dimension in \mathbb{R}^d , then [KMM04, pp. 79-80]

$$H_k(X, \mathbb{Z}) = \begin{cases} \mathbb{Z}, & k = 0, \\ 0 \text{ (the trivial group)}, & k \neq 0. \end{cases} \quad \triangle$$

Example 21 (Homology of a sphere). For $0 \leq n < d$, Let X be the boundary (in the sense of general topology) of an $(n+1)$ -cube, so that X is homeomorphic to the n -sphere. It can be shown [KMM04, pp. 90, 303] that if $n > 0$ then $H_0(X, \mathbb{Z}) = H_n(X, \mathbb{Z}) = \mathbb{Z}$, and if $n = 0$ then $H_0(X, \mathbb{Z}) = \mathbb{Z}^2$, and for all n that $H_k(X, \mathbb{Z}) = 0$ for $k \neq 0, n$. (More concisely, $H_k(X, \mathbb{Z}) = \mathbb{Z}^{\llbracket k=0 \rrbracket + \llbracket k=n \rrbracket}$ for all k and n .)

The geometric intuition (for $n > 0$) is the following. The n -sphere for has one connected component; therefore, $H_0(X, \mathbb{Z}) = \mathbb{Z}$. It encloses one $(n+1)$ -dimensional void; therefore, $H_n(X, \mathbb{Z}) = \mathbb{Z}$ (because an n -cycle is determined completely by the coefficient on a single n -face of the void, and there are no nontrivial n -boundaries.) The fact that the other homology groups are trivial is more complicated to prove directly, but geometrically it has to do with the absence of voids of other dimensions. △

Example 22 (Homology of the 2-torus). Let $X \subseteq \mathbb{R}^3$ be a union of 32 plaquettes as shown in fig. 1, so that X is homeomorphic to the 2-dimensional torus.⁹

There are no 3-cubes in X , so there are no nontrivial 2-boundaries. To understand the 2-cycles, notice that since every edge belongs to exactly 2 plaquettes, the coefficient on a single plaquette uniquely determines the coefficients on all the remaining plaquettes. So $Z_2(X, G)$ is (isomorphic to) G . Thus, $H_2(X, G) = G$, too.

Every 1-cycle is homologous (that is, equal modulo a 1-boundary) to some 1-cycle whose support is a subset of the 8 edges shown in bold in fig. 1, which form two independent loops encircling the torus. This can be seen by algorithmically building up an appropriate 2-chain, one plaquette at a time, whose boundary cancels out the coefficients on all other edges. Moreover, it's not hard to see that every homology class contains only one such cycle. But a cycle supported on

⁹Alternately, though it's harder to visualize, we could build a torus using just 16 plaquettes in \mathbb{R}^4 . This discrete "Clifford torus" is a subset of the boundary of a 4-cube.

over \mathbb{Z} , we argue by contradiction: in any 2-cycle, the coefficient on a plaquette determines the coefficients on its adjacent plaquettes, but after wrapping all the way around the Klein bottle we end up with the wrong signs on the coefficients. This yields, however, a 1-boundary with even coefficients, which is the source of the $\mathbb{Z}/2\mathbb{Z}$ summand in $H_1(X, \mathbb{Z})$. To be more explicit, let c be the cycle that assigns 1 to each edge along some closed loop encircling the neck of the Klein bottle. Then c is not a boundary, but $2c$ is a boundary. This is an example of a *torsion* [Sti93, pp. 170-171].

We see that the mod-2 homology of the Klein bottle coincides with the mod-2 homology of the torus, but their integral homologies are distinct [Mun84, p. 52]. \triangle

The k th *Betti number* $\beta_k(X, \mathbb{Z})$ of X is the rank of the group $H_k(X, \mathbb{Z})$ (that is, the number of copies of \mathbb{Z} in its cyclic decomposition.) So for the Klein bottle $\beta_0(X, \mathbb{Z}) = \beta_1(X, \mathbb{Z}) = 1$ and $\beta_k(X, \mathbb{Z}) = 0$ for $k \neq 0, 1$. For $q \geq 2$ we have $H_k(X, \mathbb{Z}/q\mathbb{Z}) = (\mathbb{Z}/q\mathbb{Z})^{m_k} \oplus T$ for some m_k and some group T whose every element has order strictly less than q (according to the decomposition theorem for finite abelian groups), so we can define the k th *Betti number with coefficients in $\mathbb{Z}/q\mathbb{Z}$* to be $\beta_k(X, \mathbb{Z}/q\mathbb{Z}) := m_k$ (i.e., the rank of the k th homology group over $\mathbb{Z}/q\mathbb{Z}$.) In particular, for prime p we find $\beta_k(X, \mathbb{Z}/p\mathbb{Z}) = \log_p |H_k(X, \mathbb{Z}/p\mathbb{Z})|$. We can describe the homology groups, or for that matter any finitely-generated abelian groups, by their (integral) Betti numbers and *torsion coefficients*, which come from the cyclic decomposition. But arguably Betti numbers are passé, and the modern approach is to work with the homology groups directly. Historically, Betti numbers and torsion coefficients came first. The connection to group theory was made in 1926 by Emmy Noether [Sti93, p. 171].

We will, however, have use for the following result. See [Die08, pp. 308-310] for the details and the proof in the more general setting of cellular homology, and for historical references.

Fact 24 (Euler–Poincaré Theorem). *For every cubical set $X \subseteq \mathbb{R}^d$, and every integer $q \geq 2$,*

$$\sum_{k \geq 0} (-1)^k |\mathcal{K}_k(X)| = \sum_{k \geq 0} (-1)^k \beta_k(X, \mathbb{Z}) = \sum_{k \geq 0} (-1)^k \beta_k(X, \mathbb{Z}/q\mathbb{Z}).$$

The value of the sums in the Euler–Poincaré Theorem is called the *Euler characteristic* of X , and is denoted $\chi(X)$. The theorem states that the “combinatorial Euler characteristic” coincides with the “homological Euler characteristic”. For the plaquette 2-torus (example 22) the left-hand equality becomes $32 - 64 + 32 = 1 - 2 + 1$. For the Klein bottle (example 23), the right-hand equality becomes $1 - 1 = 1 + 2 - 1$.

Fact 25. For every cubical set $X \subseteq \mathbb{R}^d$, and every abelian group G ,

$$H_0(X, G) \cong G^{k(X)}$$

where $k(X)$ is the number of connected components in X . In particular, $\beta_0(X, \mathbb{Z}) = \beta_0(X, \mathbb{Z}/q\mathbb{Z}) = k(X)$ for all $q \geq 1$.

Proof. This was already shown in example 19 for the case where X is a graph and $G = \mathbb{Z}$. The proof of the general case is identical. \square

Often it's useful to add an "empty face" of dimension -1 to each cubical complex. That is, let \emptyset be the sole (-1) -cube, and consider it to be a face of every other elementary cube. This gives the so-called *reduced homology*, indicated by writing tildes above all symbols as below. The reduced homology eliminates many special cases in the statements of theorems and proofs in homology theory, though it's conventional (at least in introductory treatments) to work with the non-reduced homology. We'll use the reduced homology only to define cyclic boundary spin conditions in section 4, and in proposition 59, so these definitions can be omitted on the first reading.

The *augmented cubical chain complex* of X is the collection $(\tilde{C}_k(G), \tilde{\partial}_k)_{k \in \mathbb{Z}}$ where

$$\tilde{C}_k(X, G) = \begin{cases} G & \text{for } k = -1, \\ C_k(X, G) & \text{for } k \neq -1, \end{cases}$$

$$\tilde{\partial}_k = \begin{cases} c \mapsto \sum_{Q \in \mathcal{K}_0(X)} c(Q) & \text{for } k = 0, \quad (\text{i.e., sum of } c\text{'s coefficients on all vertices } Q) \\ 0 & \text{for } k = -1, \\ \partial_k & \text{for } k \neq 0, -1. \end{cases}$$

It's easy to augment the proof of the boundary relation (fact 17) to show $\tilde{\partial}_{k-1} \circ \tilde{\partial}_k = 0$. This leads to *reduced boundaries* $\tilde{B}_k(X, G) = \text{im } \tilde{\partial}_{k+1}$, *reduced cycles* $\tilde{C}_k(X, G) = \ker \tilde{\partial}_k$, *reduced homology groups* $\tilde{H}_k(X, G) = \tilde{C}_k(X, G) / \tilde{B}_k(X, G)$, and the *reduced homology* $\tilde{H}_*(X, G) := (\tilde{H}_k(X, G))_{k \in \mathbb{Z}}$. (We'll have no use for the reduced Betti numbers.) The change is not particularly significant, as the following result (stated as [KMM04, p. 90, Theorem 2.95] for $G = \mathbb{Z}$) indicates.

Fact 26. If X is a nonempty cubical set then

$$H_k(X, G) = \begin{cases} \tilde{H}_0(X, G) \oplus G, & k = 0, \\ \tilde{H}_k(X, G), & k \neq 0. \end{cases}$$

Proof. It's only necessary to examine $k = 0$ and -1 (the other cases are trivial.)

Let $k = 0$. A reduced 0-cycle is an assignment of coefficients to lattice points in X such that the sum of all coefficients is 0. The cosets modulo the 0-boundaries are characterized by the sums of coefficients on each component, and the sum on the last component is determined by the others. So $\tilde{H}_0(X, G) = G^{m-1}$ where m is the number of connected components.

Let $k = -1$. Every chain is a cycle, and every chain is a boundary because X is nonempty. So $\tilde{H}_{-1}(X, G)$ is trivial. \square

Next, we describe the cubical cohomology, which for us will be essential.

The *cubical cochain complex for X with coefficients in G* , $(C^k(X, G), \delta^k)_{k \in \mathbb{Z}}$, is defined by duality. Let $\text{hom}(A, B)$ be the group of all group homomorphisms from abelian group A to abelian group B (the group operation on $\text{hom}(A, B)$ is defined pointwise.) We define

$$C^k(X, G) := \text{hom}(C_k(X, G), G), \quad z \in \mathbb{Z}$$

and

$$\begin{aligned} \delta^k &: C^k(X, G) \rightarrow C^{k+1}(X, G), \\ c &\mapsto c \circ \partial_{k+1}, \quad k \in \mathbb{Z}. \end{aligned}$$

The elements of $C^k(X, G)$ are called *k-cochains*, and the map δ^k is called the *coboundary operator*. See fig. 2.

$$\begin{array}{ccccccc} \cdots & \xleftarrow{\partial_{k-1}} & C_{k-1} & \xleftarrow{\partial_k} & C_k & \xleftarrow{\partial_{k+1}} & C_{k+1} & \xleftarrow{\partial_{k+2}} & \cdots \\ & & \vdots & & \vdots & & \vdots & & \\ \cdots & \xrightarrow{\delta^{k-2}} & C^{k-1} & \xrightarrow{\delta^{k-1}} & C^k & \xrightarrow{\delta^k} & C^{k+1} & \xrightarrow{\delta^{k+1}} & \cdots \end{array}$$

Figure 2: Chain complex and cochain complex

It's common to use the angle-bracket notation for evaluation, $\langle a, b \rangle := a(b)$ for $a \in C^k(X, G)$

and $b \in C_k(X, G)$, and in this notation

$$\langle \delta^k c, d \rangle = \langle c, \partial_{k+1} d \rangle \quad \text{for } c \in C^k(X, G) \text{ and } d \in C_{k+1}(X, G).$$

This boundary–coboundary duality relation—for us, true by definition—is sometimes called the *discrete Stokes theorem*. It’s also possible to give a different definition of δ^k and derive this result as a genuine, albeit trivial, theorem (this is done in [FLV21, §2.3.2].)

Again, it’s most common to take $G = \mathbb{Z}$, as in [KM13]. But for us the most relevant coefficient group will be $G = \mathbb{Z}/q\mathbb{Z}$ with $q \geq 2$. Accordingly, recall the notation for group duality in section 2.1. When $G = \mathbb{Z}/q\mathbb{Z}$, we may identify $\text{hom}(C_k(X, G), G)$ with $C_k(X, G)^\wedge$ ($:= \text{hom}(C_k(X, G), \mathbb{T})$) via the embedding $G \hookrightarrow \mathbb{T}$, $[j] \mapsto e^{2\pi i j/q}$, because the order of every element of $C_k(X, G)$ divides q . So we may write

$$\begin{aligned} C^k(X, \mathbb{Z}/q\mathbb{Z}) &= C_k(X, \mathbb{Z}/q\mathbb{Z})^\wedge, \\ \delta^k &= \partial_{k+1}^*, \quad k \in \mathbb{Z}. \end{aligned}$$

The composition of adjacent coboundary operators is 0, just like with boundary operators.

Fact 27 (Coboundary relation). *For every $k \in \mathbb{Z}$,*

$$\delta^k \circ \delta^{k-1} = 0.$$

Proof. By fact 17,

$$\begin{aligned} \langle \delta^k \delta^{k-1} c, d \rangle &= \langle \delta^{k-1} c, \partial_{k+1} d \rangle = \langle c, \partial_k \partial_{k+1} d \rangle = \langle c, 0 \rangle = 0 \\ &\text{for } c \in C^{k-1}(X, G) \text{ and } d \in C_{k+1}(X, G). \quad \square \end{aligned}$$

A *k-cocycle* in X is an element of $Z^k(X, G) := \ker \delta^k$. A *k-coboundary* in X is an element of $B^k(X, G) := \text{im } \delta^{k-1}$. Every coboundary is a cocycle (fact 27) so we define the *kth cohomology group of X with coefficients in G* to be the quotient group $H^k(X, G) := Z^k(X, G)/B^k(X, G)$, and the *cohomology of X with coefficients in G* to be the sequence $H^*(X, G) := (H^k(X, G))_{k \in \mathbb{Z}}$.

Notice that if c is a k -boundary and d is a k -cocycle, then by the discrete Stokes theorem $\langle d, c \rangle = 0$.¹⁰

¹⁰In the (non-discrete) calculus of differential forms, the corresponding statement specializes to a classical result from

The following example is essential for the application to the Potts gauge theory.

Example 28. Let G be a (finite or infinite) cyclic group with generator 1 . Then a k -cochain, being a group homomorphism, is determined by the values it assumes on the chains 1_Q for all $Q \in \mathcal{X}_k(X)$. Thus, $C^k(X, G)$ may be identified with $C_k(X, G)$, and under this identification $\langle \cdot, \cdot \rangle$ is the dot product of coefficient vectors (with respect to the usual ring multiplication on G). For a chain $c \in C_k(X, G)$, write $\widehat{c} \in C^k(X, G)$ for the corresponding cochain under this identification.

If Q is an edge in X then the coboundary $\delta^1 \widehat{1}_Q$ is a 2-cochain that puts coefficients ± 1 on all plaquettes in X incident to Q , and 0 on all other plaquettes. \triangle

We can also define augmented cochains, reduced coboundaries, reduced cocycles, and reduced homology groups $\widetilde{H}^k(X, G)$: Let $\widetilde{C}^k(X, G)$ be the dual of $\widetilde{C}_k(X, G)$, let $\widetilde{\delta}^k$ be the dual of $\widetilde{\partial}_{k+1}$, and define the remaining objects analogously.

There is a result called the Poincaré Lemma that gives sufficient topological conditions on X so that all reduced cohomology groups are trivial, i.e., every reduced cocycle is a reduced coboundary. See, for example, [FLV21, Lemma 2.2] or [Des+05, §11].

Table 1 summarizes our notation. For a lengthier explanation of (integral) cubical homology theory, see [KMM04] and [KM13].

The machinery of cubical homology can be developed in several different ways. The approach we've followed, from [KMM04], is similar in flavor to simplicial homology. There is also a cubical variant of singular homology (see the remarks in [KMM04, §2.8] and [KM13, §1].) Another method (as explained in the introductory paragraphs to this section) is to use a notation reminiscent of the exterior calculus of differential forms, which may be more familiar to readers who have worked with de Rham cohomology. This discrete exterior calculus is used in many papers on lattice gauge theory, e.g., [Cha20; Cao20; FLV21]. For more comprehensive references see [Des+05] and [Cra23].

2.3 Probabilistic couplings and the stochastic ordering

This section reviews some basic probabilistic tools that will play a major role in what follows.

Let $(\Omega_1, \mathcal{A}_1, \mu_1)$ and $(\Omega_2, \mathcal{A}_2, \mu_2)$ be probability spaces. A *coupling* of μ_1 and μ_2 is a probability measure μ on the measurable space $(\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$ whose first and second marginals are μ_1 and μ_2 , respectively. In other words, the pushforwards of μ under the coordinate projections $\rho_1 : (\omega_1, \omega_2) \mapsto \omega_1$ and $\rho_2 : (\omega_1, \omega_2) \mapsto \omega_2$ are $\rho_1 \mu = \mu_1$ and $\rho_2 \mu = \mu_2$. This definition

vector calculus: A curl-free vector field in \mathbb{R}^3 has circulation 0 along any closed loop that bounds a surface.

Symbol	Meaning
$\mathcal{K}_k \subseteq \mathcal{P}(\mathbb{R}^d)$	Elementary k -cubes in \mathbb{R}^d for $0 \leq k \leq d$, empty otherwise
$\mathcal{K} \subseteq \mathcal{P}(\mathbb{R}^d)$	All elementary cubes in \mathbb{R}^d
$X \subseteq \mathbb{R}^d$	A cubical set: a finite union of elementary cubes from \mathcal{K}
$\mathcal{K}_k(X)$	The (finite) set of all k -cubes in X , i.e., $\{Q \in \mathcal{K}_k \mid Q \subseteq X\}$
$\mathcal{K}(X)$	All elementary cubes in X , i.e., $\{Q \in \mathcal{K} \mid Q \subseteq X\}$
$C_k(X, G)$	All k -chains: The group $G^{\mathcal{K}_k(X)}$; trivial for $k < 0$ and $k > d$
g_Q for $g \in G$ and $Q \in \mathcal{K}_k(X)$	The k -chain that takes value g on Q , and 0 on all other k -cubes
1_Q for $Q \in \mathcal{K}_k(X)$	The indicator of Q in $C_k(X, G)$, when G is the additive group of a ring with identity 1
$\overset{\circ}{Q}$ for $Q \in \mathcal{K}$	The associated elementary cell: The relative interior of Q
$\partial_k : C_k(X, G) \rightarrow C_{k-1}(X, G)$	Boundary map (a group homomorphism)
$B_k(X, G)$	All k -boundaries: The image of ∂_{k+1}
$Z_k(X, G)$	All k -cycles: The kernel of ∂_k
$H_k(X, G)$	The k th homology group: The quotient $Z_k(X, G)/B_k(X, G)$
$C^k(X, G)$	All k -cochains: The group $\text{hom}(C_k(X, G), G)$, equal to the dual $\widehat{C_k(X, G)}$ when $G = \mathbb{T}$ or when G is a finite subgroup of \mathbb{T}
$\delta^k : C^k(X, G) \rightarrow C^{k+1}(X, G)$	The coboundary map—the dual ∂_{k+1}^*
$B^k(X, G)$	All k -coboundaries: The image of δ^{k-1}
$Z^k(X, G)$	All k -cocycles: The kernel of δ^k
$H^k(X, G)$	The k th cohomology group: The quotient $Z^k(X, G)/B^k(X, G)$
$\beta_k(X, \mathbb{Z})$	The k th Betti number of X with integer coefficients: The rank of the group $H_k(X, \mathbb{Z})$ (i.e., the rank of its torsion-free part)
$\beta_k(X, \mathbb{Z}/q\mathbb{Z})$ for $q \geq 2$	The k th Betti number of X with coefficients in $\mathbb{Z}/q\mathbb{Z}$: The rank of $H_k(X, \mathbb{Z}/q\mathbb{Z})$ over $\mathbb{Z}/q\mathbb{Z}$
$\chi(X)$	The Euler characteristic $\sum_{k \geq 0} (-1)^k \mathcal{K}_k(X) $

Table 1: Notation for cubical homology and cohomology. Here $k \in \mathbb{Z}$ and G is a finite abelian group.

Discrete exterior calculus of differential forms	Cubical homology
oriented 0-cell x^+, x^- for $x \in \mathbb{Z}^d$	0-cube, or vertex, $Q \in \mathcal{K}_0$, not oriented
oriented edge dx_i , $1 \leq i \leq d$	1-cube, or edge, $Q \in \mathcal{K}_1$, not oriented
k-form $dx_{i_1} \wedge \cdots \wedge dx_{i_k}$	k-cube $Q \in \mathcal{K}_k$, not oriented
G-valued k-form f	k-cochain $c \in C^k(X, G)$
exterior derivative df of G-valued k-form f	coboundary $\delta^k c$ of k-cochain c
coderivative δf of G-valued k-form f	boundary $\partial_k c$ of k-chain c
closed k-form	k-cocycle
exact k-form	k-coboundary
closed surface	2-cycle
Hodge duality	dual cubical structure and Poincaré duality, not explained here

Table 2: Dictionary for translating between discrete exterior calculus and cubical homology

extends to arbitrary (finite and infinite) collections of spaces: a coupling of an indexed family of probability measures $((\Omega_i, \mathcal{A}_i, \mu_i))_{i \in \mathcal{J}}$ is a probability measure μ on the product measurable space $(\prod_{i \in \mathcal{J}} \Omega_i, \otimes_{i \in \mathcal{J}} \mathcal{A}_i)$ that satisfies $\rho_i \mu = \mu_i$ for every $i \in \mathcal{J}$.

Couplings can often be identified with (equivalence classes of) probability kernels. Recall that a *probability kernel* from $(\Omega_1, \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$ is a map $K : \Omega_1 \times \mathcal{A}_2 \rightarrow [0, 1]$ such that

- (i) the map $E \mapsto K(\omega_1, E)$ is a probability measure for each $\omega_1 \in \Omega_1$, and
- (ii) the map $\omega_1 \mapsto K(\omega_1, E)$ is measurable for each $E \in \mathcal{A}_2$.¹¹

Every probability kernel K from $(\Omega_1, \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$ and every probability measure μ_1 on $(\Omega_1, \mathcal{A}_1)$ together induce a probability measure $\mu_1 \otimes K$ on the product measurable space $(\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$, defined by

$$(\mu_1 \otimes K)(E) := \int_{\Omega_1} \mu_1(d\omega_1) \int_{\Omega_2} K(\omega_1, d\omega_2) \mathbf{1}_E(\omega_1, \omega_2), \quad E \in \mathcal{A}_1 \otimes \mathcal{A}_2.$$

The opposite procedure—starting with a probability measure μ on $(\Omega_1 \times \Omega_2, \mathcal{A}_1 \otimes \mathcal{A}_2)$ and obtaining from it a transition kernel—is called *disintegration of measure*. If $(\Omega_2, \mathcal{A}_2)$ is a standard Borel space (as tends to be the case in applications), then there exists a probability kernel K from $(\Omega_1, \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$ such that $\mu = (\rho_1 \mu) \otimes K$ (here, as before, $\rho_1 \mu$ is the marginal of μ on the

¹¹It suffices to check this condition over a generating π -system for \mathcal{A}_2 [Kle14, pp. 180–181].

first factor.) In the case that Ω_1 and Ω_2 are both finite, this result follows immediately from the definitions: let $K(x, \cdot)$ be the conditional measure $\mu(\cdot \mid \omega_1 = x)$ whenever $\mu(\omega_1 = x) \neq 0$ and pick $K(x, \cdot)$ arbitrarily for all x with $\mu(\omega_1 = x) = 0$. The general case is more subtle. The reader may consult [Çin11] or any other introductory text on measure-theoretic probability.

Once we have a probability kernel, we can use it to push forward probability measures from $(\Omega_1, \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$ or pull back functions from $(\Omega_2, \mathcal{A}_2)$ to $(\Omega_1, \mathcal{A}_1)$, and these two operations are compatible, as stated in the following well-known theorem [Çin11, Theorem 6.3].

Fact 29 (Measure–kernel–function theorem). *Let $(\Omega_1, \mathcal{A}_1)$ and $(\Omega_2, \mathcal{A}_2)$ be measurable spaces, and K a probability kernel from and $(\Omega_1, \mathcal{A}_1)$ to $(\Omega_2, \mathcal{A}_2)$. Define*

$$Kf_2(\omega_1) := \int_{\Omega_2} K(\omega_1, d\omega_2) f_2(\omega_2), \quad \omega_1 \in \Omega_1$$

for every measurable function f_2 on Ω_2 which is either nonnegative or bounded. For measurable $f_2 : \Omega_2 \rightarrow [0, \infty]$, this defines a measurable function $Kf_2 : \Omega_1 \rightarrow [0, \infty]$. For bounded measurable $f_2 : \Omega_2 \rightarrow \mathbb{C}$, this defines a bounded measurable function $Kf_2 : \Omega_1 \rightarrow \mathbb{C}$.

Define

$$\mu_1 K(E) := \int_{\Omega_1} \mu_1(d\omega_1) K(\omega_1, E), \quad E \in \mathcal{A}_2$$

for every probability measure μ_1 on Ω_1 . This defines a probability measure $\mu_1 K$ on $(\Omega_2, \mathcal{A}_2)$.

Furthermore,

$$(\mu_1 K)f_2 = \mu_1(Kf_2) = \int_{\Omega_1} \mu_1(d\omega_1) \int_{\Omega_2} K(\omega_1, d\omega_2) f_2(\omega_2)$$

for all μ_1 and f_2 considered above.

Proof. The case of $f_2 : \Omega_2 \rightarrow [0, \infty]$ is proved in [Çin11, Theorem 6.3] by a standard measure-theoretic argument. By subtracting an arbitrary constant, the result extends to every bounded $f_2 : \Omega_2 \rightarrow \mathbb{R}$. The case of bounded $f_2 : \Omega_2 \rightarrow \mathbb{C}$ follows by linearity. \square

It's immediate from the definitions that $\mu_1 K = \rho_2(\mu_1 \otimes K)$ for all μ_1 and K as above.

Readers familiar with dynamics might recognize fact 29 as the non-deterministic generalization of the identity $\int f_2 dT_*\mu_1 = \int f_2 \circ T d\mu_1$, where T is a measurable map from Ω_1 to Ω_2 . We can recover this deterministic identity by taking $K(x, \cdot)$ to be the point mass δ_{Tx} . In this case, the map

$\mu_1 \mapsto \mu_1 K$ is the pushforward operator with respect to T , and the map $f_2 \mapsto Kf_2$ is the Koopman operator $f_2 \mapsto f_2 \circ T$.

Furthermore, if both Ω_1 and Ω_2 are finite, then K may be identified with a right stochastic matrix. Writing μ_1 and f_2 as row and column vectors, respectively, the definitions given in fact 29 reduce to matrix multiplication, and the identity $(\mu_1 K)f_2 = \mu_1(Kf_2)$ is by associativity of matrix multiplication [Cin11, p. 46, Exercise 6.31]. This observation is relevant to section 3.4, where we discuss transferring observables between two models. (For the sake of completeness: The operation $\mu_1 \otimes K$ corresponds to taking the matrix product $\text{diag}(\mu_1)K$ and interpreting the resulting matrix as a joint probability mass function.)

Now we discuss the stochastic ordering. For a partially ordered set P , a function $f : P \rightarrow \mathbb{R}$ is *increasing* (or *monotone*¹²) if $x \leq y \implies f(x) \leq f(y)$. A set $S \subseteq P$ is *increasing* (or *monotone*) if its indicator $\mathbf{1}_S$ is an increasing function, that is, if $x \in S \implies y \in S$ whenever $x \leq y$. Given two probability measures μ and μ' on P , we say that μ is *stochastically smaller* than μ' , or that μ is *stochastically dominated* by μ' , and write $\mu \leq_{\text{st}} \mu'$, if the pair μ, μ' satisfies any of the equivalent conditions given in the following theorem.

Theorem 30 (Strassen's¹³ theorem, finite version). *For probability measures μ and μ' on a finite partially ordered set P , the following are equivalent.*

- (i) $\mu E \leq \mu' E$ for every increasing event $E \subseteq P$.
- (ii) $\mu f \leq \mu' f$ ¹⁴ for every increasing function $f : P \rightarrow \mathbb{R}$.
- (iii) There exists a coupling ν of μ and μ' such that $\nu\{(x, y) \in P \times P \mid x \leq y\} = 1$.

Proof. (iii) \implies (ii): For increasing $f : P \rightarrow \mathbb{R}$,

$$\begin{aligned} \mu f &= \int f(x) \, d\nu(x, y) = \int f(x) \mathbf{1}_{x \leq y} \, d\nu(x, y) \\ &\leq \int f(y) \mathbf{1}_{x \leq y} \, d\nu(x, y) = \int f(y) \, d\nu(x, y) = \mu' f. \end{aligned}$$

¹²The official terminology from order theory is *monotone* (or *order-preserving* or occasionally *isotone*) but in analysis and probability it's common to use the ambiguous term *increasing*. To say that $f : \mathbb{R} \rightarrow \mathbb{R}$ is "increasing" might mean, depending on the author and context, that f is monotone, or strictly monotone ($x < y \implies f(x) < f(y)$), or inflationary ($x \leq f(x)$ for all x).

¹³This theorem is traditionally attributed to Strassen [Str65, Theorems 6, 11], who proved it in the context of Polish spaces. But other variants had apparently [Str65, p. 432] already been published in 1961 by Kellerer [Kel61] and Dall'Aglio [Dal61]. Our theorem 30 and theorem 31 are considerably less general than the results in any of the aforementioned papers.

¹⁴We use the probabilist's de Finetti notation: For probability measure μ on Ω and random variable $X : \Omega \rightarrow \mathbb{C}$, we write $\mu X := E_\mu[X]$. A measurable set $E \subseteq \Omega$ is identified with its indicator $\mathbf{1}_E$.

(ii) \implies (i): Take $f = \mathbf{1}_E$.

(i) \implies (iii): See [LP16, Theorem 10.4] for a short proof of the equivalence (i) \iff (iii) via the max flow min cut theorem. \square

Besides the case of finite P , we'll also be interested in $P = \{0, 1\}^E$ for a countably infinite set E . Endow P with pointwise ordering, product topology, and Borel σ -algebra (which coincides with the cylinder σ -algebra.) We write $\mu \leq_{\text{st}} \mu'$ if the pair μ, μ' satisfies any of the equivalent conditions given in the following theorem.

Note that every increasing real-valued function on P is bounded because P has a least element and a greatest element, so the expectations in the theorem statement exist.

Theorem 31 (Strassen's theorem, countable version). *For probability measures μ and μ' on $P = \{0, 1\}^E$, the following are equivalent.*

(i) $\mu f \leq \mu' f$ for every measurable increasing function $f : P \rightarrow \mathbb{R}$.

(ii) $\mu f \leq \mu' f$ for every continuous increasing function $f : P \rightarrow \mathbb{R}$.

(iii) There exists a coupling ν of μ and μ' such that $\nu\{(x, y) \in P \times P \mid x \leq y\} = 1$.

Proof. (iii) \implies (i): Argue as in the proof of theorem 30.

(i) \implies (ii): Trivial.

(ii) \implies (iii): Take a sequence of finite sets $E_1 \subseteq E_2 \subseteq \dots \rightarrow E$, and for every $n \geq 1$ let μ_n and μ'_n be the respective marginals of μ and μ' . By (ii), every increasing function $f : E_n \rightarrow \mathbb{R}$ satisfies $\mu_n f \leq \mu'_n f$. By the finite version of this theorem (theorem 30), there exists a coupling ν_n of μ_n and μ'_n such that $\nu_n\{(x, y) \in \{0, 1\}^{E_n} \times \{0, 1\}^{E_n} \mid x \leq y\} = 1$. After passing to an appropriate subsequence we may assume that $\nu_1, \nu_2, \dots \rightarrow \nu$ weakly for some measure ν on $\{0, 1\}^E$. To be more precise, we may assume that $\rho_F \nu_n \rightarrow \rho_F \nu$ weakly for all finite $F \subseteq E$, where $\rho_F((x, y)) = (x|_F, y|_F)$. Writing $\rho_1(x, y) = x$, this implies $\rho_F \mu_n = \rho_1 \rho_F \nu_n \rightarrow \rho_1 \rho_F \nu$ weakly for all finite $F \subseteq E$. But $\rho_F \mu_n$ is eventually equal to $\rho_F \mu$, so $\rho_F \rho_1 \nu = \rho_F \mu$; thus, $\rho_1 \nu = \mu$ and likewise $\rho_2 \nu = \mu'$. Thus, ν is a coupling of μ and μ' . To prove $\nu\{(x, y) \in P \times P \mid x \leq y\} = 1$, express the event $\{(x, y) \in P \times P \mid x \leq y\}$ as the limit of the decreasing sequence of events $\{(x, y) \in P \times P \mid x|_{E_n} \leq y|_{E_n}\}$. (This compactness argument was suggested in [Lig05, p. 75].) \square

The set $\{0, 1\}$ in theorem 31 may be replaced with an arbitrary closed subset of \mathbb{R} , and the theorem statement still holds after replacing "function" with "bounded function" in (i) and (ii)

[GHM01, Theorem 4.6]. More generally still, a version of Strassen's theorem holds for Polish spaces [Lin92, p. 129].

For more on couplings and stochastic domination, refer to [GHM01, §4; Gri06, ch. 2; Lin92; Hol12].

3 The higher Potts and FK–Potts models and their coupling, in finite volume with free boundary condition

Sections 3 to 5 will use the following parameters. Take integers $1 \leq r \leq d$ (*cell dimension* and *ambient dimension*, respectively), real $p \in [0, 1]$, and integer $q \geq 1$ (which will never be assumed to be prime, except where stated explicitly.) Take $\beta \in [0, \infty]$ such that $p = 1 - e^{-\beta}$ (declaring $e^{-\infty} = 0$.) The gauge group G will be the additive group $\mathbb{Z}/q\mathbb{Z}$.

The usual practice in lattice gauge theory is instead to view G as the multiplicative group of complex q th roots of unity, but we will treat G as an additive group, which accords with the conventions of homology theory. We'll also consider the cochain groups (whose elements are characters) to be additive groups, denoting the trivial character by 0 , again, to keep in line with the conventions of homology theory over abelian groups. Our characters, however, are still maps into the circle $\mathbb{T} \subseteq \mathbb{C}$. So (confusingly) the character 0 is the constant 1 map, and (more confusingly) the sum of characters is their pointwise product (as already defined at the very start of section 2.1.) We will occasionally also need to add characters pointwise when discussing Fourier decompositions (of course, the group of cochains is not closed under pointwise sum.) The situation is far from ideal, but that is the price we pay for attempting to bridge the two theories.

Let Λ be a finite nonempty set of (elementary) r -cubes in \mathbb{R}^d , and let $X = \bigcup_{Q \in \Lambda} Q \subseteq \mathbb{R}^d$ (recall from section 2.2 that X includes, along with each r -cube $Q \in \Lambda$, all lower-dimension faces of Q .) A prototypical example is the N -box, $B_N = [-N, N]^d \subseteq \mathbb{R}^d$ for some $N \geq 1$, with $\Lambda = \mathcal{K}_r(B_N)$ (so that $X \subseteq B_N$, and $X \subsetneq B_N$ if $r < d$.)

3.1 The higher Potts model

The lattice spin system presented here encompasses both the classical Potts model (for parameter $r = 1$) and the Potts lattice gauge theory (for parameter $r = 2$) [AF84, §3], which assigns spins to the nearest-neighbor edges. The latter two models generalize the Ising model and Ising lattice gauge theory, respectively, which have $q = 2$. Spin systems for $r \geq 3$, that is, those which assign elements of the gauge group to cells of dimension 2 or greater, are called *higher lattice gauge theories*.

Let Σ_X^{15} be the additive group of $(r - 1)$ -cochains,

$$\Sigma_X := C^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) = (C_{r-1}(X, \mathbb{Z}/q\mathbb{Z}))^\wedge = ((\mathbb{Z}/q\mathbb{Z})^{\mathcal{K}_{r-1}(X)})^\wedge.$$

¹⁵In most cases our notation will explicitly specify X in order to avoid confusion once we begin discussing infinite-volume limits in section 5.

We identify the cochains with the chains as explained in eq. (3).¹⁶ Under this identification, a *configuration* $\sigma \in \Sigma_X$ is an assignment of a “spin” from $\mathbb{Z}/q\mathbb{Z}$ to each $(r-1)$ -cube in X . The $(r-1)$ -cube Potts model on X with parameters β and q has probability measure

$$\pi_{X,\beta,q}(\sigma) := \begin{cases} \frac{e^{-\beta H(\sigma)}}{Z_P(\beta,q)} & \text{where } H(\sigma) = -\sum_{Q \in \mathcal{K}_r(X)} \llbracket \sigma_Q = 1 \rrbracket & \text{for } 0 \leq \beta < \infty, \\ \frac{\llbracket \sigma \in Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \rrbracket}{|Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z})|} & & \text{for } \beta = \infty, \end{cases} \quad \sigma \in \Sigma_X, \quad (5)$$

where the *partition function* $Z_P(\beta, q)$ for $0 \leq \beta < \infty$ is the normalizing constant

$$Z_P(\beta, q) := \sum_{\sigma \in \Sigma_X} e^{-\beta H(\sigma)},$$

$\llbracket \cdot \rrbracket$ is the indicator function (page 3), and

$$\sigma_Q := \langle \sigma, \partial_r 1_Q \rangle, \quad Q \in \mathcal{K}_r(X).$$

Recall that this angle-bracket notation means $\sigma(\partial_r 1_Q)$; here σ is a cochain and $\partial_r 1_Q$ is a chain. So σ_Q is the sum of spins on the boundary of Q , taking orientation into account, considered as element of \mathbb{C} (that is, $\sigma_Q = e^{2\pi i k/q}$ for some integer k .) Thus, the Hamiltonian $H(\sigma)$ is the (negated) tally of the r -cubes with zero net boundary spin.

Note that to reconcile eq. (5) with the Ising model, where the summands in the Hamiltonian are ± 1 , the parameter β must be modified by a factor of 2 and the partition function must also be multiplied by a constant accordingly.

In the existing literature on lattice gauge theory, it’s more usual to describe the Hamiltonian $H(\sigma)$ in eq. (5) in terms of a unitary representation of the gauge group (see, for example, [Cao20, p. 1440].) Equation (5) may be expressed in this form by taking the unitary representation $\rho : \mathbb{Z}/q\mathbb{Z} \rightarrow \text{GL}(q, \mathbb{C})$ given in block matrix form as

$$\rho([k]) = \begin{bmatrix} \mathbf{0} & \mathbf{I}_{q-1} \\ 1 & \mathbf{0} \end{bmatrix}^k, \quad [k] \in \mathbb{Z}/q\mathbb{Z}$$

¹⁶Of course, it’s possible to simply define spin configurations to be chains. However, treating them as cochains results in a more aesthetically pleasing theory. If the motivation seems opaque at this point, consider that will always be taking the coboundaries of spin configurations and never their boundaries. See the beginning of section 3.4 for further motivation. Note that [HS16, p. 8] and [DS23, §5] also define spin configurations to be cochains.

where \mathbf{I}_{q-1} is the $(q-1) \times (q-1)$ identity matrix, so that $\rho([1])$ is a coordinate permutation of order q . Then $\text{tr}(\rho(g)) = q \llbracket g = [0] \rrbracket$ and (now considering σ_Q as an element of $\mathbb{Z}/q\mathbb{Z}$)

$$H(\sigma) = -\frac{1}{q} \sum_{Q \in \mathcal{K}_r(X)} \Re \text{tr}(\rho(\sigma_Q)).$$

The indicator in (5) for $\beta = \infty$ may be written as

$$\llbracket \sigma \in Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \rrbracket = \prod_{Q \in \mathcal{K}_r(X)} \llbracket \sigma_Q = 1 \rrbracket.$$

One way to understand this statement is to identify the cocycles with cycles via the standard dual basis, so that σ is (identified with) an assignment of spins (elements of $\mathbb{Z}/q\mathbb{Z}$) to the $(r-1)$ -cubes. Then, speaking loosely, the coboundary operator δ^{r-1} sends each $(r-1)$ cube to all its incident r -cubes, and if an r -cube is present then its incident $(r-1)$ -cubes must make zero net contribution to it.

This characterization of cocycles is useful enough to be stated as an explicit result for future reference.

Proposition 32. *For every cubical set $Y \subseteq \mathbb{R}^d$, integer k , and $\sigma \in C^{k-1}(Y, \mathbb{Z}/q\mathbb{Z})$,*

$$\sigma \in Z^{k-1}(Y, \mathbb{Z}/q\mathbb{Z}) \iff \prod_{Q \in \mathcal{K}_k(Y)} \llbracket \sigma_Q = 1 \rrbracket = 1.$$

Proof. $\sigma \in Z^{k-1}(Y, \mathbb{Z}/q\mathbb{Z}) \iff \delta^{k-1}\sigma = 0$
 $\iff \langle \delta^{k-1}\sigma, c \rangle = 1$ for every $c \in C_k(Y, \mathbb{Z}/q\mathbb{Z})$
 $\iff \langle \delta^{k-1}\sigma, 1_Q \rangle = 1$ for every $Q \in \mathcal{K}_k(Y)$
 $\iff \langle \sigma, \partial_k 1_Q \rangle = 1$ for every $Q \in \mathcal{K}_k(Y)$
 $\iff \prod_{Q \in \mathcal{K}_k(Y)} \llbracket \sigma_Q = 1 \rrbracket = 1. \quad \square$

3.2 The higher FK–Potts model

We now extend the FK–Potts, or *random-cluster*, model to arbitrary-dimension cells. The random-cluster model actually allows arbitrary real $q \in (0, \infty)$ [Gri06, §1.2], but we’ll be constrained to integer $q \geq 1$. The case $q = 1$ is independent Bernoulli(p) percolation on the r -cubes.

Let $\Omega_X := \{0, 1\}^{\mathcal{K}_r(X)}$. Elements of Ω_X will be called *configurations*. For a given configuration

$\omega \in \Omega_X$, we say an r -cube $Q \subseteq X$ is *open* if $\omega(Q) = 1$ and *closed* if $\omega(Q) = 0$. Each configuration ω gives rise to a cubical set consisting of all open r -cubes and all lower-dimension cubes,

$$\begin{aligned} X_\omega &:= \bigcup_{\substack{Q \in \mathcal{K}_r(X) \\ \omega(Q)=1}} Q \cup \bigcup_{Q \in \mathcal{K}_{r-1}(X)} Q \\ &= \bigcup_{\substack{Q \in \mathcal{K}_r(X) \\ \omega(Q)=1}} Q \cup \bigcup_{\substack{Q \in \mathcal{K}_k(X) \\ k < r}} Q \subseteq \mathbb{R}^d. \end{aligned}$$

To put it another way, X_ω is X with the relative interiors of all closed r -cubes removed,

$$X_\omega = X \setminus \bigcup_{\substack{Q \in \mathcal{K}_r(X) \\ \omega(Q)=0}} \overset{\circ}{Q}.$$

Note that $X = X_{\omega^1}$, where ω^1 is the configuration with all r -cubes open.

The r -cube FK–Potts model on X with parameters p and q has probability measure

$$\varphi_{X,p,q}(\omega) := \frac{1}{Z_{\text{FKP}}(p,q)} (1-p)^{c(\omega)} p^{o(\omega)} |Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|, \quad \omega \in \Omega_X, \quad (6)$$

where $Z_{\text{FKP}}(p,q)$ is the normalizing constant, the values $o(\omega)$ and $c(\omega)$ are the number of open and closed r -cubes in ω , respectively, and the last factor is the number of $(r-1)$ -cocycles in X_ω with coefficients in the group $\mathbb{Z}/q\mathbb{Z}$.

For $q = 1$ all (co)chain, (co)cycle, (co)boundary, and (co)homology groups are trivial, so each configuration ω occurs with probability $(1-p)^{c(\omega)} p^{o(\omega)}$.

Proposition 33 below is meant primarily as a reference to help reconcile various formulas in the existing literature, and secondarily to assist with proofs of some results that follow. In particular, for $r = 1$, after identifying the set $X_\omega \subseteq \mathbb{R}^d$ with a graph (example 19), eq. (8) reduces eq. (6) to the usual random-cluster measure $\varphi_{X,p,q}(\omega) = \frac{1}{Z_{\text{RC}}(p,q)} (1-p)^{c(\omega)} p^{o(\omega)} q^{k(\omega)}$. Equation (9) was presented as the “wrong” formula for $r = 2$ and $q \geq 2$ in [AF84, (3.7), (6.2)]. In general, when q is not prime, the order of the r th homology group $H_r(X_\omega, \mathbb{Z}/q\mathbb{Z})$ may fail to be a power of q , as demonstrated by the counterexamples in [AF84, §4].

Constant factors such as $q^{|\mathcal{K}_{r-1}(X)|}$ may, of course, be suppressed by absorbing into the partition function $Z_{\text{FKP}}(p,q)$, and likewise $q^{-o(\omega)}$ may appear as $q^{c(\omega)}$ (because $o(\omega) + c(\omega)$ is constant.)

Proposition 33 (Counting cocycles). *The dependence factor in eq. (6) satisfies*

$$|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| = \frac{|C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})|}{|B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|} = q^{|\mathcal{K}_{r-1}(X)| - o(\omega)} |H_r(X_\omega, \mathbb{Z}/q\mathbb{Z})|, \quad \omega \in \Omega_X. \quad (7)$$

The number of 0-cocycles is

$$|Z^0(X_\omega, \mathbb{Z}/q\mathbb{Z})| = q^{k(\omega)}, \quad \omega \in \Omega_X, \quad (8)$$

where $k(\omega)$ is the number of connected components in the graph $(V, E(\omega)) = (\mathcal{K}_0(X), \mathcal{K}_1(X_\omega))$.

If q is prime then

$$|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| = q^{|\mathcal{K}_{r-1}(X)| - o(\omega) + \beta_r(\omega)}, \quad \omega \in \Omega_X, \quad (9)$$

where $\beta_r(\omega) := \beta_r(X_\omega, \mathbb{Z}/q\mathbb{Z})$ (Betti numbers with coefficients were defined on page 29.)

If q is prime and $r = 2$ then

$$|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| = q^{|\mathcal{K}_0(X)| - k + \beta_1(\omega)} \quad (10)$$

where k is the number of connected components in X .

Proof. Let $\omega \in \Omega_X$. Since $\mathcal{K}_k(X_\omega) = \mathcal{K}_k(X)$ for all $k < r$,

$$C_k(X_\omega, \mathbb{Z}/q\mathbb{Z}) = C_k(X, \mathbb{Z}/q\mathbb{Z}) \quad \text{for all } k < r. \quad (11)$$

The first equality in eq. (7):

$$\begin{aligned} |Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| &= |\ker \delta^{r-1}| \\ &= |\ker \partial_r^*| \\ &= |\text{Ann}(\text{im } \partial_r)| && \text{(by fact 8)} \\ &= \frac{|C_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|}{|\text{im } \partial_r|} && \text{(by fact 4)} \\ &= \frac{|C_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|}{|B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|} \\ &= \frac{|C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})|}{|B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|} && \text{(by eq. (11).)} \end{aligned} \quad (12)$$

The second equality in eq. (7):

$$\begin{aligned}
|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| &= \frac{|C_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|}{|\operatorname{im} \partial_r|} && \text{(as in derivation (12))} \\
&= \frac{|C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})|}{|C_r(X_\omega, \mathbb{Z}/q\mathbb{Z})| / |\ker \partial_r|} && \text{(by the first isomorphism theorem)} \\
&= \frac{q^{|\mathcal{K}_{r-1}(X)|}}{q^{o(\omega)} / |Z_r(X_\omega, \mathbb{Z}/q\mathbb{Z})|} \\
&= q^{|\mathcal{K}_{r-1}(X)| - o(\omega)} |H_r(X_\omega, \mathbb{Z}/q\mathbb{Z})| && \text{(because } B_r(X_\omega, \mathbb{Z}/q\mathbb{Z}) \text{ is trivial.)}
\end{aligned}$$

As for eq. (8),

$$\begin{aligned}
|Z^0(X_\omega, \mathbb{Z}/q\mathbb{Z})| &= \frac{|C_0(X_\omega, \mathbb{Z}/q\mathbb{Z})|}{|B_0(X_\omega, \mathbb{Z}/q\mathbb{Z})|} && \text{(continuing from eq. (12) with } r = 1) \\
&= \frac{|Z_0(X_\omega, \mathbb{Z}/q\mathbb{Z})|}{|B_0(X_\omega, \mathbb{Z}/q\mathbb{Z})|} && \text{(because } C_0(X_\omega, \mathbb{Z}/q\mathbb{Z}) = Z_0(X_\omega, \mathbb{Z}/q\mathbb{Z})) \\
&= \left| \frac{Z_0(X_\omega, \mathbb{Z}/q\mathbb{Z})}{B_0(X_\omega, \mathbb{Z}/q\mathbb{Z})} \right| && \text{(by Lagrange's theorem for finite groups)} \\
&= |H_0(X_\omega, \mathbb{Z}/q\mathbb{Z})| \\
&= |(\mathbb{Z}/q\mathbb{Z})^{k(X_\omega)}| && \text{(by fact 25)} \\
&= q^{k(\omega)} && \text{(because } k(\omega) = k(X_\omega).)
\end{aligned}$$

If q is prime then $\mathbb{Z}/q\mathbb{Z}$ is a field, so all boundary maps are linear maps between vector spaces over $\mathbb{Z}/q\mathbb{Z}$. Hence the group $H_r(X_\omega, \mathbb{Z}/q\mathbb{Z})$ is also a vector space over $\mathbb{Z}/q\mathbb{Z}$, with

$$|H_r(X_\omega, \mathbb{Z}/q\mathbb{Z})| = q^{\beta_r(\omega)}.$$

This proves eq. (9).

If q is prime and $r = 2$ then the Euler characteristic formula (fact 24)

$$\chi(\omega) = \beta_2(\omega) - \beta_1(\omega) + \beta_0(\omega) = |\mathcal{K}_2(X_\omega)| - |\mathcal{K}_1(X_\omega)| + |\mathcal{K}_0(X_\omega)|$$

reduces by fact 25 to

$$\beta_2(\omega) - \beta_1(\omega) + k = o(\omega) - |\mathcal{K}_1(X)| + |\mathcal{K}_0(X)|,$$

where $k := k(X)$ is the number of connected components in X (we're using the fact that $k(X) =$

$k(X_\omega)$ whenever $r > 1$.) Rearranging,

$$|\mathcal{K}_1(X)| - o(\omega) + \beta_2(\omega) = |\mathcal{K}_0(X)| - k + \beta_1(\omega),$$

which proves eq. (10). □

FKG and its applications

The FKG machinery, introduced in the 1970s [Gri06, Appendix], consists of a correlation inequality together with a handful of related results that are important to statistical physics, percolation, and related areas. The FKG property was originally developed in the context of the random-cluster model, so it's reasonable to hope that it's shared by the higher FK–Potts model. Indeed, this turns out to be the case. After the proof, we'll explore a few consequences.

The notation and terminology will mostly follow [Gri06, §2.2].

For any finite set E define $\Omega_E := \{0, 1\}^E$ (with apologies for overloading the notation: the set Ω_X from before is Ω_E with $E = \mathcal{K}_r(X)$.) Endow Ω_E with pointwise ordering, $\omega_1 \leq \omega_2 \iff \omega_1(e) \leq \omega_2(e)$ for all $e \in E$. We may take joins and meets (i.e., respectively, least upper bounds and greatest lower bounds) of elements $\omega_1, \omega_2 \in \Omega_E$,

$$\begin{aligned} (\omega_1 \vee \omega_2)(e) &:= \max(\omega_1(e), \omega_2(e)), \\ (\omega_1 \wedge \omega_2)(e) &:= \min(\omega_1(e), \omega_2(e)) \quad \text{for } e \in E. \end{aligned}$$

A probability measure on a finite (or, more generally, discrete) measurable space is called (*strictly*) *positive* if every nonempty set has strictly positive probability.

Let μ be a positive probability measure on the measurable space $(\Omega_E, \mathcal{P}(\Omega_E))$. For $E' \subseteq E$ and $\xi \in \{0, 1\}^{E'}$, let μ_ξ be the probability measure on $\Omega_{E \setminus E'} = \{0, 1\}^{E \setminus E'}$ given by

$$\mu_\xi(\omega) := \mu(\omega \times \xi \mid \xi), \quad \omega \in \Omega_{E \setminus E'},$$

where $(\omega \times \xi)(e) = \begin{cases} \xi(e), & e \in E' \\ \omega(e), & e \notin E' \end{cases}$, and $\mu(\cdot \mid \xi)$ is shorthand for the conditional probability

$\mu(\cdot \mid \{\omega \times \xi \mid \omega \in \Omega_{E \setminus E'}\})$. Since μ is assumed to be positive, this conditioning event $\{\omega \times \xi \mid \omega \in \Omega_{E \setminus E'}\} \subseteq \Omega_E$ always has nonzero μ -measure, and μ_ξ is also positive.

We say that μ has the *weak FKG property*, or is *positively associated*, or has *positive correlations*, if

$$\mu(fg) \geq \mu(f)\mu(g) \quad \text{for all increasing functions } f, g : \Omega_E \rightarrow \mathbb{R}. \quad (13)$$

We say that μ has the *strong FKG property* if it satisfies any of the equivalent conditions in the following theorem [Gri06, Theorems 2.19, 2.24].

Theorem 34. *For a positive probability measure μ on $\Omega_E = \{0, 1\}^E$, where E is a finite set, the following are equivalent.*

(i) **Strong positive association:** *For every $E' \subseteq E$ and every $\xi \in \{0, 1\}^{E'}$, the measure μ_ξ is positively associated (inequality (13)).*

(ii) **FKG lattice condition (also known as log-supermodularity):** *For all pairs $\omega_1, \omega_2 \in \Omega_E$,*

$$\mu(\omega_1 \vee \omega_2) \mu(\omega_1 \wedge \omega_2) \geq \mu(\omega_1)\mu(\omega_2).$$

(iii) **2-Position FKG lattice condition:** *For all incomparable pairs $\omega_1, \omega_2 \in \Omega_E$ with Hamming distance 2 (that is, differing in precisely two positions $e, e' \in E$ with $\omega_1(e) + \omega_1(e') = \omega_2(e) + \omega_2(e') = 1$),*

$$\mu(\omega_1 \vee \omega_2) \mu(\omega_1 \wedge \omega_2) \geq \mu(\omega_1)\mu(\omega_2).$$

(iv) **Monotonicity:** *For every $E' \subseteq E$ and for all pairs $\xi, \zeta \in \{0, 1\}^{E'}$,*

$$\xi \leq \zeta \implies \mu_\xi \leq_{\text{st}} \mu_\zeta.$$

(v) **1-Monotonicity:** *For every $e \in E$ and for all pairs $\xi, \zeta \in \{0, 1\}^{E \setminus \{e\}}$,*

$$\xi \leq \zeta \implies \mu_\xi \leq_{\text{st}} \mu_\zeta$$

or, equivalently,

$$\xi \leq \zeta \implies \mu_\xi(e \mapsto 1) \leq \mu_\zeta(e \mapsto 1). \quad \square$$

The phrase *FKG Inequality* (which we won't use) may refer either to inequality (13) or to the implications (ii) \implies (13) [Gri06, Theorem 2.16] or (v) \implies (13) [Geo11, Theorem 4.11].

For the product Bernoulli(p) measure with $p \in (0, 1)$, inequality (13) holds and is named *Harris's Lemma* after Ted Harris, who published a proof in 1960 [BR06, pp. 39–42].

Inequality (13) may be contrasted with some analogous classical results for functions of a real variable. *Chebyshev's Association Inequality* states that nondecreasing functions $f, g : \mathbb{R} \rightarrow \mathbb{R}$ applied to a real-valued random variable X are positively correlated: $\mathbf{E}[f(X)g(X)] \geq \mathbf{E}[f(X)]\mathbf{E}[g(X)]$ (for two generalizations, see [BLM13, Theorems 2.14, 2.15].) A special case (by taking discrete uniform X) is *Chebyshev's Sum Inequality*: $\frac{1}{n} \sum a_i b_i \geq \left(\frac{1}{n} \sum a_i\right) \left(\frac{1}{n} \sum b_i\right)$ whenever $a_1 \leq \dots \leq a_n$ and $b_1 \leq \dots \leq b_n$.

Now, the result—a partial extension of [Gri06, Theorem 3.8], which covers the case $r = 1$ for arbitrary real $q \geq 1$. Recall that for us $1 \leq r \leq d$, and q is a positive integer. The result for general r and prime q was also given in [HS16, Theorem 5.1].¹⁷ Notice that $\varphi_{X,p,q}$ is positive for every $p \in (0, 1)$, so theorem 34 applies.

Theorem 35 (Strong FKG). *For every $p \in (0, 1)$, the higher FK–Potts measure (6) has the strong FKG property.*

Proof. We'll prove the FKG lattice condition

$$\varphi_{X,p,q}(\omega_1 \vee \omega_2) \varphi_{X,p,q}(\omega_1 \wedge \omega_2) \geq \varphi_{X,p,q}(\omega_1) \varphi_{X,p,q}(\omega_2), \quad \omega_1, \omega_2 \in \Omega_X.$$

By proposition 33, this reduces to

$$\frac{p^{o(\omega_1 \vee \omega_2) + o(\omega_1 \wedge \omega_2)} (1-p)^{c(\omega_1 \vee \omega_2) + c(\omega_1 \wedge \omega_2)}}{|B_{r-1}(X_{\omega_1 \vee \omega_2})| |B_{r-1}(X_{\omega_1 \wedge \omega_2})|} \geq \frac{p^{o(\omega_1) + o(\omega_2)} (1-p)^{c(\omega_1) + c(\omega_2)}}{|B_{r-1}(X_{\omega_1})| |B_{r-1}(X_{\omega_2})|}, \quad (14)$$

$\omega_1, \omega_2 \in \Omega_X.$

(For clarity, in this proof we'll forgo indicating the coefficient group $\mathbb{Z}/q\mathbb{Z}$.)

The numerators in (14) are equal¹⁸ because

$$\begin{aligned} o(\omega_1 \vee \omega_2) + o(\omega_1 \wedge \omega_2) &= o(\omega_1) + o(\omega_2), \\ c(\omega_1 \vee \omega_2) + c(\omega_1 \wedge \omega_2) &= c(\omega_1) + c(\omega_2), \quad \omega_1, \omega_2 \in \Omega_X. \end{aligned}$$

¹⁷Hiraoka and Shirai's proof is very short and uses a higher-level tool from algebraic topology. The proof presented here uses only elementary group theory. I do not know whether Hiraoka and Shirai's proof extends to arbitrary (non-prime) q .

¹⁸We are implicitly proving this general result: A positive probability measure \mathbf{P} has the strong FKG property if and only if any (or every) probability measure of the form $\mathbf{P}_p(\omega) = \frac{1}{Z_p} (1-p)^{c(\omega)} p^{o(\omega)} \mathbf{P}(\omega)$ for $p \in (0, 1)$ has the strong FKG property [Gri06, p. 33].

Our symbol for the boundary operator is ambiguous: We write both

$$\begin{aligned} \partial_r : C_r(X) &\rightarrow C_{r-1}(X) \quad \text{and} \\ \partial_r : C_r(X_\omega) &\rightarrow C_{r-1}(X_\omega) \quad (= C_{r-1}(X)) \quad \text{for every } \omega \in \Omega_X. \end{aligned}$$

In this proof, henceforth, the symbol ∂_r shall always refer to the former, “unconstrained” variant, while for each $\omega \in \Omega_X$ we’ll identify $C_r(X_\omega)$ with a subgroup of $C_r(X)$ (via the inclusion $\kappa_\omega : C_r(X_\omega) \rightarrow C_r(X)$, which sends an r -chain in X_ω to the same r -chain in X , assigning coefficient 0 to every r -cube missing from X_ω .) Clearly, under this identification the unconstrained map ∂_r satisfies $B_{r-1}(X_\omega) = \partial_r(C_r(X_\omega))$. Therefore,

$$\begin{aligned} |B_{r-1}(X_{\omega_1 \vee \omega_2})| |B_{r-1}(X_{\omega_1 \wedge \omega_2})| &= |\partial_r(C_r(X_{\omega_1 \vee \omega_2}))| |\partial_r(C_r(X_{\omega_1 \wedge \omega_2}))| \\ &= |\partial_r(C_r(X_{\omega_1}) + C_r(X_{\omega_2}))| |\partial_r(C_r(X_{\omega_1}) \cap C_r(X_{\omega_2}))| \\ &\leq |\partial_r(C_r(X_{\omega_1}))| |\partial_r(C_r(X_{\omega_2}))| \quad (\text{by fact 14}) \\ &= |B_{r-1}(X_{\omega_1})| |B_{r-1}(X_{\omega_2})|, \quad \omega_1, \omega_2 \in \Omega_X. \end{aligned}$$

This is the required relation between the denominators of inequality (14). □

The FKG properties have numerous applications. Here are a few of them.

Comparison inequality The proof of the next result is taken essentially unchanged from [Gri06, Theorem 3.21].¹⁹

Proposition 36 (Comparison inequality). *For $0 \leq p_1 \leq p_2 \leq 1$,*

$$\varphi_{X,p_1,q} \leq_{\text{st}} \varphi_{X,p_2,q}.$$

Proof. Let $f : \Omega_X \rightarrow \mathbb{R}$ be an increasing function. We must show that $\varphi_{p_1} f \leq \varphi_{p_2} f$ (for readability, in this proof we’ll suppress the parameters X and q in the notation $\varphi_{X,p,q}$ and $Z_{\text{FKP}}(p_2, q)$.)

Case 1: If $p_1 = 0$, then $\varphi_{p_1} f = f(\omega^0)$ where $\omega^0 \in \Omega_X$ is the all-closed configuration. Since $f(\omega^0) \leq f(\omega)$ for all $\omega \in \Omega_X$, it follows that $\varphi_{p_1} f \leq \varphi_{p_2} f$.

¹⁹For $r = 1$, the comparison inequalities also involve differing values $q_2 \leq q_1$, but the proof’s extension to $r \geq 1$ fails. There seems to be an analogous result assuming the stronger condition of divisibility $q_2 | q_1$, but we won’t pursue this further.

Case 2: If $p_2 = 1$, then $\varphi_{p_2} f = f(\omega^1)$ where $\omega^1 \in \Omega_X$ is the all-open configuration. Since $f(\omega) \leq f(\omega^1)$ for all $\omega \in \Omega_X$, it follows that $\varphi_{p_1} f \leq \varphi_{p_2} f$.

Case 3: If $0 < p_1 \leq p_2 < 1$, then

$$\begin{aligned} \varphi_{p_2} f &= \frac{1}{Z_{\text{FKP}}(p_2)} \sum_{\omega \in \Omega_X} f(\omega) (1-p_2)^{c(\omega)} p_2^{o(\omega)} |Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| \\ &= \frac{1}{Z_{\text{FKP}}(p_2)} \sum_{\omega \in \Omega_X} f(\omega) \left(\frac{1-p_2}{1-p_1} \right)^{c(\omega)} \left(\frac{p_2}{p_1} \right)^{o(\omega)} (1-p_1)^{c(\omega)} p_1^{o(\omega)} \\ &\quad \cdot |Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| \\ &= \frac{Z_{\text{FKP}}(p_1)}{Z_{\text{FKP}}(p_2)} \varphi_{p_1} f g \end{aligned}$$

where

$$g(\omega) = \left(\frac{1-p_1}{1-p_2} \right)^{-c(\omega)} \left(\frac{p_2}{p_1} \right)^{o(\omega)} > 0, \quad \omega \in \Omega_X.$$

The exponents $o(\cdot)$ and $-c(\cdot)$ are increasing, and $p_1 \leq p_2$ by assumption, so the function g is increasing. Setting $f = 1$ in the equality above, dividing, and applying the weak FKG property (inequality (13)),

$$\varphi_{p_2} f = \frac{\varphi_{p_2} f}{\varphi_{p_2} 1} = \frac{\frac{Z_{\text{FKP}}(p_1)}{Z_{\text{FKP}}(p_2)} \varphi_{p_1} f g}{\frac{Z_{\text{FKP}}(p_1)}{Z_{\text{FKP}}(p_2)} \varphi_{p_1} 1 g} = \frac{\varphi_{p_1} f g}{\varphi_{p_1} g} \geq \frac{\varphi_{p_1} f \cdot \varphi_{p_1} g}{\varphi_{p_1} g} = \varphi_{p_1} f. \quad \square$$

Thresholds for increasing events Recall the *Bernoulli model*: the family of product measures $\mu_p(\omega) = p^{o(\omega)}(1-p)^{c(\omega)}$, $p \in [0, 1]$, where $\omega \in \Omega_E = \{0, 1\}^E$ for some fixed finite set E . It's well-known that the function $p \mapsto \mu_p(A)$ is strictly increasing for every increasing event $A \neq \emptyset, \Omega_E$. Since $p \mapsto \mu_p(A)$ is continuous, with $\mu_0(A) = 0$ and $\mu_1(A) = 1$, there exists unique $p \in (0, 1)$ such that $\mu_p(A) = \frac{1}{2}$. This p is called the *threshold* for A . Thresholds play a key role in the theory of random graphs: E is the edge set, and A corresponds to the presence of some structure (for instance, a Hamiltonian cycle.) To be more accurate, random graph theory studies the asymptotic behavior of threshold sequences for increasing events $A_n \subseteq \{0, 1\}^{E_n}$ where $|E_0| < |E_1| < \dots$. Thresholds in random graphs were introduced in the late 1950s [ER60] and remain an active area of research. For instance, very recently a proof emerged for the Kahn-Kalai conjecture, a powerful tool for estimating thresholds in the Bernoulli model [KK06; PP22].

Just like the Bernoulli model, the higher FK-Potts model has unique thresholds:

The function $p \mapsto \varphi_{X,p,q}(\omega)$ is continuous for every $\omega \in \Omega_X = \{0, 1\}^{\mathcal{K}_r(X)}$, so the function $p \mapsto \varphi_{X,p,q}f$ is continuous for every function $f : \Omega_X \rightarrow \mathbb{R}$. In particular, if $A \subseteq \Omega_X$ is an increasing event with $A \neq \emptyset, \Omega_X$, then the function $\alpha : p \mapsto \varphi_{X,p,q}(A)$ is continuous, satisfies $\alpha(0) = 0$ and $\alpha(1) = 1$, and is weakly increasing by proposition 36. But α is a rational function (its denominator $p \mapsto Z_{\text{FKP}}(p, q)$ is a polynomial), so α must be strictly increasing on $[0, 1]$.²⁰ Thus, there exists unique $p \in (0, 1)$ with $\alpha(p) = \frac{1}{2}$.

More can be said. Theorem 35 allows us to put explicit bounds on the rate of growth of α , and in certain cases to quantify the sharpness of the threshold. See proposition 66.

The strong FKG property has many other useful consequences that we won't discuss [Gri06, Ch. 2]. Here is just one more result, quoted from [Gri06, §2.5, Theorem 2.53].

Proposition 37 (Exponential steepness). *For every $p \in (0, 1)$ and every nonempty event $A \subseteq \Omega_X$, the higher Fk–Potts measure (eq. (6)) satisfies*

$$\begin{aligned} \frac{d}{dp} \log \varphi_{X,p,q}(A) &\geq \frac{\varphi_{X,p,q}(H_A)}{p(1-p)} && \text{if } A \text{ is increasing and} \\ \frac{d}{dp} \log \varphi_{X,p,q}(A) &\leq -\frac{\varphi_{X,p,q}(H_A)}{p(1-p)} && \text{if } A \text{ is decreasing,} \end{aligned}$$

where H_A is the Hamming distance function

$$H_A(\omega) := \min_{\omega' \in A} |\{Q \in \mathcal{K}_r(X) \mid \omega(Q) \neq \omega'(Q)\}|, \quad \omega \in \Omega_X.$$

Proof. This is a special case of [Gri06, §2.5, Theorem 2.53], which applies to every probability measure of the form $\mathbf{P}_p(\omega) = \frac{1}{Z_p}(1-p)^{c(\omega)}p^{o(\omega)}\mathbf{P}(\omega)$ where \mathbf{P} is a probability measure with the strong FKG property. □

3.3 The Edwards–Sokal coupling

Here, we generalize the standard probabilistic coupling [ES88; Gri06, §1.4] of the Potts and FK–Potts (“random-cluster”) models.

The (finite-volume) *Edwards–Sokal coupling* of the $(r-1)$ -cube Potts model with the r -cube

²⁰The reasoning for this is as follows. Assume that $f : [0, 1] \rightarrow \mathbb{R}$ is a weakly increasing but not strictly increasing function, so that $f(x) = f(y) = c$ for some $0 \leq x < y \leq 1$ and some constant c . Then $f(t) = c$ for all $t \in [x, y]$. If also $f = \frac{p}{q}$ for nonzero polynomials p and q then $p(t) = cq(t)$ for all $t \in [x, y]$, so $p(t) - cq(t)$ is the zero polynomial. It follows that $f(t) = c$ for all $c \in [0, 1]$.

FK–Potts model is the probability measure

$$\begin{aligned} \mu_{X,p,q}(\sigma, \omega) &:= \frac{1}{Z_{ES}(p, q)} (1-p)^{c(\omega)} p^{o(\omega)} \mathbb{I}[\sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})], \\ (\sigma, \omega) &\in \Sigma_X \times \Omega_X = C^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \times \{0, 1\}^{\mathcal{K}_r(X)}, \end{aligned} \quad (15)$$

where again $o(\omega)$ and $c(\omega)$ are the number of open and closed r -cubes in ω , respectively, $Z_{ES}(p, q)$ is the normalizing constant, and $\mathbb{I}[\cdot]$ is the indicator.

Recall that we've fixed $p \in [0, 1]$ and $\beta \in [0, \infty]$ with $p = 1 - e^{-\beta}$ (first paragraphs of section 3).

Proposition 38 (Marginals). *The marginals of $\mu_{X,p,q}$ are*

$$\begin{aligned} \sum_{\omega \in \Omega_X} \mu_{X,p,q}(\sigma, \omega) &= \pi_{X,\beta,q}(\sigma), \quad \sigma \in \Sigma_X \quad \text{and} \\ \sum_{\sigma \in \Sigma_X} \mu_{X,p,q}(\sigma, \omega) &= \varphi_{X,p,q}(\omega), \quad \omega \in \Omega_X. \end{aligned}$$

Proof. If $p \in [0, 1)$ then for every $\sigma \in \Sigma_X$

$$\begin{aligned} \sum_{\omega \in \Omega_X} \mu_{X,p,q}(\sigma, \omega) &= \frac{1}{Z_{ES}(p, q)} \sum_{\omega \in \Omega_X} (1-p)^{c(\omega)} p^{o(\omega)} \prod_{Q \in \mathcal{K}_r(X_\omega)} \mathbb{I}[\sigma_Q = 1] \quad (\text{by proposition 32}) \\ &= \frac{1}{Z_{ES}(p, q)} \sum_{\omega \in \Omega_X} \left(\prod_{\substack{Q \in \mathcal{K}_r(X) \\ \omega(Q)=0}} (1-p) \right) \left(\prod_{\substack{Q \in \mathcal{K}_r(X) \\ \omega(Q)=1}} p \mathbb{I}[\sigma_Q = 1] \right) \\ &= \frac{1}{Z_{ES}(p, q)} \prod_{Q \in \mathcal{K}_r(X)} ((1-p) + p \mathbb{I}[\sigma_Q = 1]) \quad (\text{via expansion}) \\ &= \frac{1}{Z_{ES}(p, q)} (1-p)^{|\{Q \in \mathcal{K}_r(X) \mid \sigma_Q \neq 1\}|} \\ &= \frac{1}{Z_{ES}(p, q)} (e^{-\beta})^{|\mathcal{K}_r(X)| - \sum_{Q \in \mathcal{K}_r(X)} \mathbb{I}[\sigma_Q = 1]} \\ &\propto \frac{1}{Z_P(\beta, q)} e^{-\beta(-\sum_{Q \in \mathcal{K}_r(X)} \mathbb{I}[\sigma_Q = 1])} \\ &= \pi_{X,\beta,q}(\sigma). \end{aligned}$$

If $p = 1$ then for every $\sigma \in \Sigma_X$ (continuing the same derivation from the third line)

$$\begin{aligned}
\sum_{\omega \in \Omega_X} \mu_{X,p,q}(\sigma, \omega) &= \frac{1}{Z_{ES}(p, q)} \prod_{Q \in \mathcal{K}_r(X)} ((1-p) + p \llbracket \sigma_Q = 1 \rrbracket) \\
&= \frac{1}{Z_{ES}(p, q)} \prod_{Q \in \mathcal{K}_r(X)} \llbracket \sigma_Q = 1 \rrbracket \\
&= \frac{1}{Z_{ES}(p, q)} \llbracket \sigma \in Z^{k-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket \quad (\text{by proposition 32}) \\
&\propto \pi_{X,\beta,q}(\sigma).
\end{aligned}$$

For every $p \in [0, 1]$ and every $\omega \in \Omega_X$,

$$\begin{aligned}
\sum_{\sigma \in \Sigma_X} \mu_{X,p,q}(\sigma, \omega) &= \frac{1}{Z_{ES}(p, q)} (1-p)^{c(\omega)} p^{o(\omega)} \sum_{\sigma \in \Sigma_X} \llbracket \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket \\
&= \frac{1}{Z_{ES}(p, q)} (1-p)^{c(\omega)} p^{o(\omega)} |Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| \\
&\propto \frac{1}{Z_{FKP}(p, q)} (1-p)^{c(\omega)} p^{o(\omega)} |Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| \\
&= \varphi_{X,p,q}(\omega). \quad \square
\end{aligned}$$

Proposition 39 generalizes [Gri06, Theorem 1.10(c)].

Proposition 39 (Partition functions). *The partition functions satisfy*

$$\begin{aligned}
Z_{ES}(p, q) &= e^{-\beta |\mathcal{K}_r(X)|} Z_P(\beta, q), & 0 \leq p < 1; \\
Z_{ES}(p, q) &= Z_{FKP}(p, q), & 0 \leq p \leq 1.
\end{aligned}$$

Proof. In the derivation of the first marginal within the proof of proposition 38, the constant of proportionality must be 1 because on each side the sum over all $\sigma \in \Sigma_X$ is 1. So

$$\frac{1}{Z_{ES}(p, q)} (e^{-\beta})^{|\mathcal{K}_r(X)|} = \frac{1}{Z_P(\beta, q)}.$$

This proves the first equality. The second equality follows likewise from the derivation of the second marginal. \square

Next, we generalize [Gri06, Theorem 1.13]. Proposition 40 states that to sample from the $(r-1)$ -cube Potts model, we may first sample ω from the r -cube FK-Potts model, and then uniformly

select a cocycle σ compatible with ω ; whereas to sample from the r -cube FK–Potts model, we may first sample σ from the $(r - 1)$ -cube Potts model, leave closed every r -cube Q for which $\sigma_Q \neq 1$, and, for the remaining r -cubes, open each independently with probability p . Recall from section 1 that for $r = 1$ a uniform cocycle is an independent uniform choice of monochromatic spin for each connected component.

Proposition 40 (Conditionals).

The first conditional of $\mu_{X,p,q}$ is

$$\mu_{X,p,q}(\sigma \mid \omega) = \begin{cases} \frac{\llbracket \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket}{|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|} & \text{if } 0 \leq p < 1 \text{ or } \omega = \omega^1, \\ \text{undefined} & \text{if } p = 1 \text{ and } \omega \neq \omega^1, \end{cases} \quad \sigma \in \Sigma_X, \omega \in \Omega_X,$$

where $\omega^1 \in \Omega_X$ is the all-open configuration ($\omega^1(Q) = 1$ for every $Q \in \mathcal{K}_r(X)$.)

The second conditional of $\mu_{X,p,q}$ is

$$\mu_{X,p,q}(\omega \mid \sigma) = (1 - p)^{c(\omega) - v(\sigma)} p^{o(\omega)} \prod_{Q \in \mathcal{K}_r(X_\omega)} \llbracket \sigma_Q = 1 \rrbracket, \quad \sigma \in \Sigma_X, \omega \in \Omega_X,$$

where $v(\sigma)$ is the number of forbidden (v for “verboten”) r -cubes,

$$v(\sigma) := \left| \{Q \in \mathcal{K}_r(X) \mid \sigma_Q \neq 1\} \right|, \quad \sigma \in \Sigma_X.$$

Proof. If $p = 1$ and $\omega \neq \omega^1$, then $(1 - p)^{c(\omega)} = 0$ so $\mu_{X,p,q}(\sigma, \omega) = 0$ for every $\sigma \in \Sigma_X$, and we can’t condition on ω . Otherwise,

$$\begin{aligned} \mu_{X,p,q}(\sigma \mid \omega) &= \frac{\mu_{X,p,q}(\sigma, \omega)}{\sum_{\sigma' \in \Sigma_X} \mu_{X,p,q}(\sigma', \omega)} \\ &= \frac{(1 - p)^{c(\omega)} p^{o(\omega)} \llbracket \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket}{\sum_{\sigma' \in \Sigma_X} (1 - p)^{c(\omega)} p^{o(\omega)} \llbracket \sigma' \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket} \\ &= \frac{\llbracket \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket}{|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|}, \quad \sigma \in \Sigma_X, \omega \in \Omega_X. \end{aligned}$$

For $\sigma \in \Sigma_X$ let $\Omega_\sigma \subseteq \Omega_X$ be the set of configurations compatible with σ ,

$$\Omega_\sigma := \{\omega \in \Omega_X \mid \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})\} = \{\omega \in \Omega_X \mid \omega(Q) = 0 \text{ whenever } \sigma_Q \neq 1\}.$$

The second conditional is

$$\begin{aligned}
\mu_{X,p,q}(\omega \mid \sigma) &= \frac{\mu_{X,p,q}(\sigma, \omega)}{\sum_{\omega' \in \Omega_X} \mu_{X,p,q}(\sigma, \omega')} \\
&= \frac{(1-p)^{c(\omega)} p^{o(\omega)} \llbracket \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket}{\sum_{\omega' \in \Omega_X} (1-p)^{c(\omega')} p^{o(\omega')} \llbracket \sigma \in Z^{r-1}(X_{\omega'}, \mathbb{Z}/q\mathbb{Z}) \rrbracket} \\
&= \frac{(1-p)^{c(\omega)} p^{o(\omega)} \prod_{Q \in \mathcal{K}_r(X_\omega)} \llbracket \sigma_Q = 1 \rrbracket}{\sum_{\omega' \in \Omega_X} (1-p)^{c(\omega')} p^{o(\omega')} \prod_{Q \in \mathcal{K}_r(X_{\omega'})} \llbracket \sigma_Q = 1 \rrbracket} \quad (\text{by proposition 32}) \\
&= \frac{(1-p)^{c(\omega)} p^{o(\omega)} \prod_{Q \in \mathcal{K}_r(X_\omega)} \llbracket \sigma_Q = 1 \rrbracket}{\sum_{\omega' \in \Omega_\sigma} (1-p)^{c(\omega')} p^{o(\omega')}} \\
&= \frac{(1-p)^{c(\omega)} p^{o(\omega)} \prod_{Q \in \mathcal{K}_r(X_\omega)} \llbracket \sigma_Q = 1 \rrbracket}{(1-p)^{v(\sigma)} \sum_{\omega' \in \Omega_\sigma} (1-p)^{c(\omega')-v(\sigma)} p^{o(\omega')}} \\
&= (1-p)^{c(\omega)-v(\sigma)} p^{o(\omega)} \prod_{Q \in \mathcal{K}_r(X_\omega)} \llbracket \sigma_Q = 1 \rrbracket, \quad \sigma \in \Sigma_X, \omega \in \Omega_X,
\end{aligned}$$

where the last step is by expansion

$$\sum_{\omega' \in \Omega_\sigma} (1-p)^{c(\omega')-v(\sigma)} p^{o(\omega')} = \prod_{\substack{Q \in \mathcal{K}_r(X) \\ \sigma_Q = 1}} ((1-p) + p) = 1. \quad \square$$

3.4 Using the coupling

Recall that every spin configuration $\sigma \in \Sigma_X := C^{r-1}(X, \mathbb{Z}/q\mathbb{Z})$ is formally defined to be a character, that is, a homomorphism from the chain group $C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})$ into \mathbb{C} . For a chain $\gamma \in C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})$ we'll write $W_\gamma := \eta(\gamma)$, which is to say that W_γ is the evaluation map

$$W_\gamma : \Sigma_X \rightarrow \mathbb{C}, \quad \sigma \mapsto \sigma(\gamma).$$

Every point in the image of W_γ is a q th root of unity. We'll discuss the expectation of W_γ with respect to the higher Potts measure,

$$\langle W_\gamma \rangle_{X,\beta,q} := \pi_{X,\beta,q} W_\gamma.$$

For $r = 2$ and $d = 4$ (lattice gauge theory), a *Wilson loop* is a closed directed walk in the edge graph of X (for us a *closed directed walk* of length $n \geq 0$ is a sequence $(v_0, e_0, v_1, e_1, \dots, v_{n-1}, e_{n-1}, v_n)$ where v_i are vertices, $e_i = \{v_i, v_{i+1}\}$ are edges, and $v_0 = v_n$; in particular, repetitions of vertices and edges are permitted.) We associate with the walk a 1-cycle γ , in the sense that

γ picks up a ± 1 coefficient on an edge each time that edge is traversed.²¹ In this situation, W_γ is called a *Wilson loop variable*. The 1-cochain σ is meant to represent a random connection on the discretized “principal bundle” $\mathbb{Z}^4 \times \mathbb{Z}/q\mathbb{Z}$: To each edge, σ assigns an element of the gauge group $\mathbb{Z}/q\mathbb{Z}$, viewed as a multiplicative subgroup of \mathbb{C} (see the note on page 40 regarding the sum and product conventions.) The value $W_\gamma(\sigma)$ is the product of these elements along γ , that is, the holonomy of the connection σ along the Wilson loop. Its expectation $\langle W_\gamma \rangle_{X, \beta, q}$ is called a *Wilson loop expectation*. See [Cha19] for more details on Wilson loop variables in gauge theories. Theorem 41 states, in particular, that the Wilson loop expectation is equal to the probability that γ is the homological boundary of some surface in the FK–Potts model.

For $r = 1$ and $q \geq 2$ —the Ising and Potts models—the theorem implies that for any two vertices x and y the expected quotient of their spins (as complex q th roots of unity) coincides with the probability that x and y are connected by some path in the FK–Potts model. For these special cases, a number of equivalent formulations can be found in the literature. For instance, [Dum20, Corollary 1.2.1] instead takes another inner product of the vertex spins and a different function relating p to β . And [Gri06, Theorem 1.16] states that the two-point correlation function of the Potts measure is directly proportional to the two-point connectivity function of the FK–Potts measure,

$$\pi_{X, \beta, q}(\sigma(x) = \sigma(y)) - \frac{1}{q} = \left(1 - \frac{1}{q}\right) \varphi_{X, p, q}(x \text{ and } y \text{ are connected in } X_\omega), \quad x, y \in \mathcal{K}_0(X).$$

Theorem 41 for the special case of prime q appears in [DS23, Theorem 5]. See also the discussion in [DS23, §1.2, “Why does q need to be prime?”]. Theorem 41 generalizes the result to arbitrary integer $q \geq 1$.

We’re still in the finite-volume setting with free boundary condition. A matching result for more general boundary conditions will be given in theorem 64.

Theorem 41 (Expectation equals probability). *For every $(r - 1)$ -chain $\gamma \in C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})$,*

$$\langle W_\gamma \rangle_{X, \beta, q} = \varphi_{X, p, q}(\gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})).$$

Proof. Assume that either $p \in [0, 1)$ and $\omega \in \Omega_X$, or $p = 1$ and $\omega = \omega^1 \in \Omega_X$ (the all-open, or constant 1, configuration.) In either case, the conditional expectation $\mu_{X, p, q}(\sigma \mid \omega)$ is defined

²¹For definiteness, we may take $+1$ on the first edge of the walk, and subsequently pick signs so that for every two consecutive edges the boundaries cancel out on the connecting vertex. That is, arrange signs so that the resulting chain γ is a cycle.

(proposition 40). Let $w_\gamma(\omega) \in \mathbb{C}$ be the conditional expectation

$$\begin{aligned}
w_\gamma(\omega) &:= \mu_{X,p,q}(W_\gamma | \omega) = \sum_{\sigma \in \Sigma_X} W_\gamma(\sigma) \mu_{X,p,q}(\sigma | \omega) \\
&= \sum_{\sigma \in \Sigma_X} \sigma(\gamma) \frac{\llbracket \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket}{|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|} \\
&= \frac{1}{|\ker \delta^{r-1}|} \sum_{\sigma \in \ker \delta^{r-1}} \sigma(\gamma) \\
&= \frac{1}{|\text{Ann}(\text{im } \partial_r)|} \sum_{\sigma \in \text{Ann}(\text{im } \partial_r)} \sigma(\gamma) \quad (\text{by fact 8}) \\
&= \llbracket \gamma \in \text{im } \partial_r \rrbracket \quad (\text{by fact 12}) \\
&= \llbracket \gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket.
\end{aligned}$$

By the law of total expectation, $\langle W_\gamma \rangle_{X,\beta,q} = \varphi_{X,p,q} w_\gamma$ (noting that if $p = 1$ then $\varphi_{X,p,q}(\omega^1) = 1$, so this argument is valid for every $p \in [0, 1]$.) \square

It follows that Wilson loop expectation (in finite volume, with free boundary condition) is real, nonnegative, and increasing in β . We will eventually use this consequence to prove the analogous infinite-volume statement, corollary 72. See also corollary 65 for more general boundary conditions.

Corollary 42. *If $0 \leq \beta_1 \leq \beta_2 \leq \infty$, then*

$$0 \leq \langle W_\gamma \rangle_{X,\beta_1,q} \leq \langle W_\gamma \rangle_{X,\beta_2,q} \leq 1$$

for every $(r-1)$ -chain γ .

Proof. Immediate from theorem 41 and the comparison inequality, proposition 36: The event

$$\{\omega \in \Omega_X \mid \gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})\}$$

is increasing because if $\omega \leq \omega'$, and if γ is the boundary of an r -chain Γ in X_ω , then γ is also the boundary of the r -chain

$$\Gamma' : Q \mapsto \begin{cases} \Gamma(Q), & Q \in \omega, \\ 0, & Q \in \omega' \setminus \omega \end{cases}$$

in $X_{\omega'}$. \square

Notice that the variables W_γ for $\gamma \in C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})$ are precisely the characters of Σ_X (because of the natural isomorphism $\eta : C_{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \rightarrow \widehat{\Sigma_X}$ —see section 2.1.) So by the Fourier transform every complex-valued function of Σ_X is a linear combination of variables W_γ . This means that theorem 41 has more general uses (for instance, the proof of proposition 71.)

4 Boundary conditions and the spatial Markov property

To lay the groundwork for infinite-volume measures (section 5), the definitions and results in section 3 must be generalized from free boundary condition to arbitrary boundary conditions. But first, some general comments on factorization are called for.

The exponential-of-Hamiltonian formula of the finite-volume higher Potts model (eq. (5)) can be justified by the Hammersley–Clifford theorem²². Roughly speaking, this theorem states that a positive probability measure is a product of factors determined by its underlying graph’s cliques if and only if the measure satisfies one of several equivalent spatial Markov properties. In more concise language, the Gibbs random fields (or Gibbs ensembles) are precisely the Markov random fields. In our case, the underlying graph has the $(r - 1)$ -cubes in X as its vertices, and two vertices are adjacent if they’re both contained in the same r -cube in X . So a (maximal) clique is the set of all $(r - 1)$ -cubes bordering a single r -cube. Equation (5) for $\beta < \infty$ is easily seen to be a product of nonzero factors each of which is a function of the spins within a single clique. Consequently, a spatial Markov property holds (proposition 52.) For details on the Hammersley–Clifford theorem, see [Lau96, Chapter 3; Gri18, §7.2; Geo11, Theorem 2.30 and the bibliographical remarks on p. 454].

The higher FK–Potts model, on the other hand, has no Gibbs factorization (by cliques) and is therefore not a true Markov random field. However, a so-called “domain Markov property” is often described (for the bond percolation case, $r = 1$), which involves modifying the measure by identifying particular boundary vertices [Dum20, §1.2]. We generalize this property in proposition 53 and proposition 60. Loosely speaking, conditioning gives a boundary condition specified by a family of cycles supported on the boundary $(r - 1)$ -cubes.

4.1 Spin conditions

Take $r, d, p, q, \beta, G, X, \Sigma_X, \Omega_X$ as described in section 3. For readability, in this section we’ll often omit the coefficient group $G = \mathbb{Z}/q\mathbb{Z}$, writing $C^{r-1}(X) := C^{r-1}(X, \mathbb{Z}/q\mathbb{Z})$ and $Z^{r-1}(X) := Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z})$ and so on. Recall our configuration spaces, $\Omega_X = \{0, 1\}^{\mathcal{K}_r(X)}$ and $\Sigma_X = C^{r-1}(X)$.

²²Another justification is the variational principle, according to which these measures minimize the free energy [Geo11, pp. 308–309; Rue04, p. 4].

Definition 43. The *boundary* of X is the union of all boundary $(r - 1)$ -cubes,

$$\partial X := \bigcup_{Q \in \mathcal{S}} Q,$$

where

$$\begin{aligned} \mathcal{S} &= \{Q \in \mathcal{K}_{r-1}(X) \mid Q \subseteq R \text{ for some } R \in \mathcal{K}_r \setminus \mathcal{K}_r(X)\} \\ &= \{Q \in \mathcal{K}_{r-1} \mid Q \subseteq R \text{ for some } R \in \mathcal{K}_r \setminus \mathcal{K}_r(X) \text{ and } Q \subseteq R' \text{ for some } R' \in \mathcal{K}_r(X)\}. \end{aligned}$$

(Equality holds because X is a union of r -cubes.) Let $\Sigma_{\partial X} := C^{r-1}(\partial X) = ((\mathbb{Z}/q\mathbb{Z})^{\mathcal{S}})^{\wedge}$. Denote by $\rho_{X, \partial X} : \Sigma_X \rightarrow \Sigma_{\partial X}$ the coordinate projection (see definition 9 and fact 10.) \triangle

We'll extend all three families of measures $\pi_{X, \beta, q}$, $\varphi_{X, p, q}$, and $\mu_{X, p, q}$, by specifying the permissible spin configurations.

Definition 44 (Spin conditions).

- A *spin condition* (or SC) is a nonempty subset of Σ_X .
- A *subgroup spin condition* is a subgroup of Σ_X .
- A *boundary spin condition* (or BSC) is the preimage of a nonempty subset of $\Sigma_{\partial X}$ under $\rho_{X, \partial X}$.
- A *point boundary spin condition* is the preimage of a singleton $\{x\} \subseteq \Sigma_{\partial X}$ under $\rho_{X, \partial X}$. So a point BSC completely specifies the spins on ∂X , and leaves the remaining spins in X unspecified.
- A *cyclic boundary spin condition* is a set of the form $\rho_{X, \partial X}^{-1}(\text{Ann}_{C_{r-1}(\partial X)} \Xi)$ for some $\Xi \subseteq \tilde{Z}_{r-1}(\partial X)$ (recall from page 10 that $\text{Ann}_{C_{r-1}(\partial X)}$ sends subsets of $C_{r-1}(\partial X)$ to subgroups of $C^{r-1}(\partial X) = \Sigma_{\partial X}$.) By fact 7 and fact 10,

$$\rho_{X, \partial X}^{-1}(\text{Ann}_{C_{r-1}(\partial X)} \Xi) = \text{Ann}_{C_{r-1}(X)} \kappa_{X, \partial X}(\Xi)$$

where $\kappa_{X, \partial X} : C_{r-1}(\partial X) \rightarrow C_{r-1}(X)$ is the coordinate injection. So a cyclic BSC is the set of all spin configurations that kill a given family of reduced $(r - 1)$ -cycles supported on $\mathcal{S} = \mathcal{K}_{r-1}(\partial X)$. The use of *reduced* homology makes a difference only for $r = 1$; recall that in the non-reduced homology every 0-chain is a cycle, which is not what we want here.

- *Imprint boundary spin conditions* will be defined in definition 58. Roughly speaking, they are those that can be obtained by conditioning on external r -cubes: Let X^1 and X be cubical

sets with $X^1 \subseteq X$ and pick a configuration $\omega_2 \in \{0, 1\}^{\mathcal{K}_r(X) \setminus \mathcal{K}_r(X^1)}$. It will be shown in proposition 56 that the conditional FK–Potts measure $\varphi_{X,p,q}(\cdot \mid \omega_2)$ is an FK–Potts measure on X^1 with BSC, and a BSC of this form will be called an imprint BSC for X^1 . \triangle




Figure 3: Kinds of spin condition. All inclusions are strict.

The relationships between the various spin conditions are displayed in fig. 3. The only non-obvious relationship is the leftmost, which is demonstrated in proposition 59. The purpose of this rather elaborate classification scheme is to provide terminology that clarifies various statements in this section and the next. We illustrate with some simple examples.

The *free boundary condition* is $\xi = \Sigma_X$ (it’s *free* in the sense that all spin configurations are permitted.) For the free boundary condition the measures given below (eqs. (16), (17) and (18)) coincide with those given earlier (eqs. (5), (6) and (15), respectively.) The free boundary condition is a cyclic BSC (take $\Xi = \emptyset$ in the definition of cyclic BSC), and in fact an imprint BSC (by proposition 54), but not a point BSC (because $|\Sigma_{\partial X}| > 1$.)

In the Potts model ($r = 1$), the *wired boundary condition* is the set of all configurations that assign the same spin to every vertex in ∂X . This is clearly a subgroup SC and a BSC. Moreover, it is a cyclic BSC: take $\Xi = \{1_Q - 1_P \in C_0(\partial X) \mid Q, P \in \mathcal{K}_0(\partial X)\}$ (each element $1_Q - 1_P \in \Xi$ forces the spins on Q and P to be equal.) Actually, taking $\Xi = \tilde{Z}_0(\partial X)$ gives the same spin condition.²³

Therefore, we define the *wired boundary condition*²⁴ in the higher Potts model (any r) to be the cyclic BSC with $\Xi = \tilde{Z}_{r-1}(\partial X)$. Thus, in the gauge setting ($r = 2$), the wired boundary condition is the collection of all configurations where the product of spins around each (generalized) boundary loop is 0.

The *periodic boundary condition* in the case of a box, $X = \bigcup \mathcal{K}_r([-N, N]^d)$ for $N \geq 1$, is the set of all configurations that assign equal spins to each pair of opposing $(r - 1)$ -cubes in ∂X . This is not quite

²³For $r = 1$, if X is a box and $d \geq 2$ then the wired boundary condition is an imprint BSC, but this does not hold for general X . Suppose, for instance, that X is an annulus: start with a box in $d = 2$ and remove a single internal vertex and its 4 incident edges. Then there is no way to connect the internal boundary of the annulus to the external boundary (the complement of the annulus in the plane is not path-connected.) It may be possible to modify our definition of imprint BSCs to be more general so as to avoid such complications.

²⁴[Cha20, p. 17] calls this the *zero boundary condition*.

the same as working in a d -torus, because each boundary r -cube is duplicated, or quadruplicated (if it's a subset of a face of $[-N, N]^d$ of codimension 2), etc. To get the d -torus, modify X by removing the relative interior²⁵ of all but one r -cube from each such replicated collection, and then assign periodic boundary condition. In either case (X or modified X) the periodic boundary condition is a cyclic BSC: take $\Xi = \{1_Q - 1_P \in C_{r-1}(\partial X) \mid Q, P \in \mathcal{K}_{r-1}(\partial X) \text{ belong to opposing facets of } [-N, N]^d \text{ and are translates of each other along one coordinate}\}$. It may or may not be an imprint BSC: If $r = 1$ and $d = 2$ then it is not, but if $r = 1$ and $d = 3$ then it is (there is room to “wire up” each opposing pair of vertices by connecting them with an edge path outside X in \mathbb{R}^3 , but not in \mathbb{R}^2 , because some of the wires would have to intersect.)

Note that the more common definition of *periodic boundary condition* is for the modified, toroidal version of X (as in [FV17, p. 81].) Unfortunately, this toroidal version doesn't quite fit for us, in the sense that it's not a spin condition on X according to our definitions, so we won't investigate it further. However, it bears noting that the toroidal version is quite natural and enjoys many special properties such as invariance under translations, rotations, and reflections. A comprehensive study of such toroidal measures may be found in [Geo11, Part IV].

Proposition 45. *The intersection of two cyclic BSCs is a cyclic BSC.*

Proof. Let $\xi^1 = \rho_{X, \partial X}^{-1}(\text{Ann } \Xi^1)$ and $\xi^2 = \rho_{X, \partial X}^{-1}(\text{Ann } \Xi^2)$ where $\Xi^1, \Xi^2 \subseteq \tilde{Z}_{r-1}(\partial X)$. We may assume without loss of generality that Ξ^1 and Ξ^2 are subgroups of $\tilde{Z}_{r-1}(\partial X)$. The intersection is

$$\begin{aligned} \xi^1 \cap \xi^2 &= \rho_{X, \partial X}^{-1}(\text{Ann } \Xi^1) \cap \rho_{X, \partial X}^{-1}(\text{Ann } \Xi^2) \\ &= \rho_{X, \partial X}^{-1}(\text{Ann } \Xi^1 \cap \text{Ann } \Xi^2) \\ &= \rho_{X, \partial X}^{-1}(\text{Ann}(\Xi^1 + \Xi^2)) \quad \text{by fact 6.} \quad \square \end{aligned}$$

For every spin condition ξ (that is, every nonempty $\xi \subseteq \Sigma_X$) we define probability measures

$$\pi_{X, \beta, q}^\xi(\sigma) := \begin{cases} \frac{e^{-\beta H(\sigma)} \mathbb{I}[\sigma \in \xi]}{Z_P^\xi(\beta, q)} & \text{where } H(\sigma) = - \sum_{Q \in \mathcal{K}_r(X)} \mathbb{I}[\sigma_Q = 1] \text{ if } 0 \leq \beta < \infty, \\ \frac{\mathbb{I}[\sigma \in Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi]}{|Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi|} & \text{if } \beta = \infty \text{ and } Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi \neq \emptyset, \\ \text{undefined} & \text{if } \beta = \infty \text{ and } Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi = \emptyset, \end{cases} \quad (16)$$

²⁵So as to not remove any $(r-1)$ -cubes.

$$\varphi_{X,p,q}^\xi(\omega) := \begin{cases} \frac{1}{Z_{\text{FKP}}^\xi(p,q)} (1-p)^{c(\omega)} p^{o(\omega)} |Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi| & \text{if } p < 1 \text{ or } Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi \neq \emptyset, \\ \text{undefined} & \text{if } p = 1 \text{ and } Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi = \emptyset, \end{cases} \quad (17)$$

$$\mu_{X,p,q}^\xi(\sigma, \omega) := \begin{cases} \frac{1}{Z_{\text{ES}}^\xi(p,q)} (1-p)^{c(\omega)} p^{o(\omega)} \llbracket \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi \rrbracket & \text{if } p < 1 \text{ or } Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi \neq \emptyset, \\ \text{undefined} & \text{if } p = 1 \text{ and } Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi = \emptyset, \end{cases} \quad (18)$$

$$\sigma \in \Sigma_X, \quad \omega \in \Omega_X.$$

These definitions are motivated by proposition 46 and proposition 47. To understand the “undefined” lines, recall that $X = X_{\omega^1}$ where ω^1 is the all-open configuration. The condition $Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi \neq \emptyset$ ensures that at least one permissible spin is compatible with ω^1 , which is necessary because $\mu_{X,p,q}(\omega^1) = 1$ when $p = 1$ (eq. (15)). Here $Z_P^\xi(\beta, q)$, $Z_{\text{FKP}}^\xi(p, q)$, and $Z_{\text{ES}}^\xi(p, q)$ are normalizing constants, with $Z_P^\xi(\beta, q)$ defined only for $0 \leq \beta < \infty$ just like the free-boundary version $Z_P(\beta, q)$ earlier.

To keep the the notation reasonably clean, we’ll often write $\pi_{X,\beta,q}^\xi$, etc., even when $\xi \subseteq \Sigma_U$ for some other cubical set $U \subseteq X$ or $X \subseteq U$. This is to be understood in the following sense: If $X \subseteq U$ then take spin condition $\{\rho_{U,X}(x) \mid x \in \xi\} \subseteq \Sigma_X$, and if $U \subseteq X$ then take spin condition $\bigcup_{x \in \xi} \rho_{X,U}^{-1}\{x\} \subseteq \Sigma_X$. (Here $\rho_{U,X} : \Sigma_U \rightarrow \Sigma_X$ and $\rho_{X,U} : \Sigma_X \rightarrow \Sigma_U$ are the coordinate projections.) These conventions are easy to remember by keeping in mind that the role of ξ is to constrain the spins in X .

Our definitions of spin conditions, and the associated measures (16) to (18), are not standard. The usual approach is to specify spins on additional vertices (or $(r-1)$ -cubes) outside X , and to augment the Hamiltonian with boundary terms that describe interactions between X and these external spins (see, for example, [FV17, p. 81].) The reason we instead define boundary conditions as subsets of spin configurations on X is that it makes definitions (16) to (18) and various theorem statements in this section very clean. We do, however, pay a price: the definitions of Gibbs states in section 5 become somewhat more complicated.

Proposition 46. For every spin condition ξ ,

$$\begin{aligned}\pi_{X,\beta,q}^\xi(\sigma) &= \pi_{X,\beta,q}(\sigma \mid \sigma \in \xi) \quad \text{and} \\ \mu_{X,p,q}^\xi(\sigma, \omega) &= \mu_{X,p,q}(\sigma, \omega \mid \sigma \in \xi),\end{aligned}$$

unless $p = 1$ and $Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi = \emptyset$ (in which case all four quantities are undefined.)

Proof. Immediate from the definitions of these four measures (eqs. (5), (15), (16) and (18). □

Proposition 47 (Marginals). For every spin condition ξ , the marginals of $\mu_{X,p,q}^\xi$ are

$$\begin{aligned}\sum_{\omega \in \Omega_X} \mu_{X,p,q}^\xi(\sigma, \omega) &= \pi_{X,\beta,q}^\xi(\sigma), \quad \sigma \in \Sigma_X, \quad \text{and} \\ \sum_{\sigma \in \Sigma_X} \mu_{X,p,q}^\xi(\sigma, \omega) &= \varphi_{X,p,q}^\xi(\omega), \quad \omega \in \Omega_X,\end{aligned}$$

unless $p = 1$ and $Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi = \emptyset$ (in which case all these terms are undefined.)

Proof. Follow the proof for free boundary condition (proposition 38), carrying along a factor of $\llbracket \sigma \in \xi \rrbracket$. □

Alternate proof for first marginal. Start with proposition 38, condition on $\sigma \in \xi$, and apply proposition 46. □

Proposition 48 (Partition functions). For every spin condition ξ , the partition functions satisfy

$$Z_{\text{ES}}^\xi(p, q) = e^{-\beta|\mathcal{K}_r(X)|} Z_{\text{P}}^\xi(\beta, q) = Z_{\text{FKP}}^\xi(p, q) \quad \text{if } 0 \leq p < 1$$

and

$$Z_{\text{ES}}^\xi(p, q) = Z_{\text{FKP}}^\xi(p, q) \quad \text{if } p = 1 \text{ and } Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi \neq \emptyset.$$

(Recall that if $p = 1$ then $Z_{\text{P}}^\xi(\beta, q)$ is undefined, and if $p = 1$ and $Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) \cap \xi = \emptyset$ then $Z_{\text{ES}}^\xi(p, q)$ and $Z_{\text{FKP}}^\xi(p, q)$ are also undefined.)

Proof. Argue as in the proof of proposition 39. □

When a spin condition ξ is involved, the conditional measures work essentially the same way as before (see proposition 40 and the paragraph that precedes it), except now when picking a uniform cocycle σ compatible with ω we must pick among only those that belong to ξ .

Proposition 49 (Conditionals). *Let ξ be a spin condition.*

The first conditional of $\mu_{X,p,q}^\xi$ is

$$\mu_{X,p,q}^\xi(\sigma | \omega) = \begin{cases} \frac{|\{\sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi\}|}{|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi|} & \text{if } (0 \leq p < 1 \text{ or } \omega = \omega^1) \\ & \text{and } Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi \neq \emptyset, \\ \text{undefined} & \text{if } p = 1 \text{ and } \omega \neq \omega^1, \\ \text{undefined} & \text{if } Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi = \emptyset, \end{cases}$$

$\sigma \in \Sigma_X, \omega \in \Omega_X.$

where $\omega^1 \in \Omega_X$ is the all-open configuration ($\omega^1(Q) = 1$ for every $Q \in \mathcal{K}_r(X)$.)

The second conditional of $\mu_{X,p,q}^\xi$ is

$$\mu_{X,p,q}^\xi(\omega | \sigma) = \begin{cases} (1-p)^{c(\omega)-v(\sigma)} p^{o(\omega)} \prod_{Q \in \mathcal{K}_r(X_\omega)} \llbracket \sigma_Q = 1 \rrbracket & \text{if } \sigma \in \xi, \\ \text{undefined} & \text{if } \sigma \notin \xi, \end{cases} \quad \sigma \in \Sigma_X, \omega \in \Omega_X,$$

where

$$v(\sigma) := |\{Q \in \mathcal{K}_r(X) \mid \sigma_Q \neq 1\}|, \quad \sigma \in \Sigma_X.$$

Proof. For $\mu_{X,p,q}^\xi(\sigma | \omega)$: If $Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi = \emptyset$, then $\mu_{X,p,q}^\xi(\sigma, \omega) = 0$ for every $\sigma \in \Sigma_X$. Otherwise, follow the proof for free boundary condition (proposition 40).

For $\mu_{X,p,q}^\xi(\omega | \sigma)$: If $\sigma \notin \xi$, then $\mu_{X,p,q}^\xi(\sigma, \omega) = 0$ for every $\omega \in \Omega_X$. Otherwise, follow the proof for free boundary condition. \square

4.2 Spatial Markov properties

We'll examine several spatial Markov properties, starting with the higher Edwards–Sokal coupling itself (proposition 51.) Taking marginals will reveal the corresponding properties of the higher Potts and higher FK–Potts models (propositions 52 and 53.)

For the next few pages, we'll use the following notation.

Notation 50. Assume that X is the union of at least 2 r -cubes. Partition the family of r -cubes in X into $n \geq 2$ disjoint nonempty subfamilies, $\mathcal{K}_r(X) = \bigcup_{1 \leq i \leq n} \mathcal{A}_i$. Write $X^i = \bigcup_{Q \in \mathcal{A}_i} Q$ for $1 \leq i \leq n$, so that $X = \bigcup_i X^i$ (but this union is not necessarily disjoint, because distinct r -cubes

may have nonempty intersection.) The *interface* of the partition is

$$\mathcal{E} := \{B \in \mathcal{K}_{r-1}(X) \mid B \subseteq X^i \cap X^j \text{ for some distinct } i, j\}.$$

That is, the interface consists of all $(r-1)$ -cubes that are faces of r -cubes from at least two \mathcal{A}_i 's.

The collection of $(r-1)$ -cubes internal to \mathcal{A}_i is

$$\mathcal{B}_i := \{B \in \mathcal{K}_{r-1}(X) \mid B \subseteq X^i\} \setminus \mathcal{E} = \mathcal{K}_{r-1}(X^i) \setminus \mathcal{E} \quad \text{for } 1 \leq i \leq n.$$

Thus, the partition (\mathcal{A}_i) of $\mathcal{K}_r(X)$ induces a partition $\mathcal{K}_{r-1}(X) = \mathcal{E} \cup \bigcup_{1 \leq i \leq n} \mathcal{B}_i$ (but some of the sets $\mathcal{E}, \mathcal{B}_i$ may be empty.) Let

$$\begin{aligned} \Sigma_i &:= ((\mathbb{Z}/q\mathbb{Z})^{\mathcal{B}_i})^\wedge \quad \text{and} \quad \Omega_i := \{0, 1\}^{\mathcal{A}_i} \quad \text{for } 1 \leq i \leq n, \\ \Sigma_{\mathcal{E}} &:= ((\mathbb{Z}/q\mathbb{Z})^{\mathcal{E}})^\wedge. \end{aligned}$$

Identify Σ_X with $\Sigma_{\mathcal{E}} \times \prod_i \Sigma_i$ and Σ_{X^i} with $\Sigma_{\mathcal{E}} \times \Sigma_i$, and also Ω_X with $\prod_i \Omega_i$. The coordinate projection maps will be denoted

$$\begin{aligned} \rho_{\Omega, i} &: \Omega_X \rightarrow \Omega_i, \\ \rho_{X, X^i} &: \Sigma_X \rightarrow \Sigma_{X^i} (= \Sigma_{\mathcal{E}} \times \Sigma_i), \\ \rho_{\mathcal{E}} &: \Sigma_X \rightarrow \Sigma_{\mathcal{E}}. \end{aligned} \quad \triangle$$

Here's an example for $r = 1$: Let X be the union of all edges in the box $[-N, N]^d \subseteq \mathbb{R}^d$; let \mathcal{A}_1 be the family of all edges in a smaller box $[-M, M]^d$, including the boundary edges²⁶; and let \mathcal{A}_2 be the family of all edges that are in X but not in \mathcal{A}_1 . Then \mathcal{E} is the collection of vertices lying on the boundary of the smaller box, and \mathcal{B}_1 and \mathcal{B}_2 are the collection of vertices in the interior and exterior of the smaller box, respectively.

Proposition 51 (Spatial Markov property of Edwards–Sokal coupling).

If $\xi = \rho_{\mathcal{E}}^{-1}\{\sigma'_{\mathcal{E}}\} \subseteq \Sigma_X$ for some $\sigma'_{\mathcal{E}} \in \Sigma_{\mathcal{E}}$, then

$$\mu_{X, p, q}^{\xi}(\sigma, \omega) = \prod_{1 \leq i \leq n} \mu_{X^i, p, q}^{\xi}(\rho_{X, X^i}(\sigma), \rho_{\Omega, i}(\omega)).$$

²⁶Those edges that are subsets of $\text{bd}[-M, M]^d$, where bd is the boundary operator for the usual metric on \mathbb{R}^d .

In particular, with respect to $\mu_{X,p,q}$, the n σ -algebras generated by the factors $\Sigma_i \times \Omega_i$ are mutually independent conditional on the σ -algebra generated by the factor Σ_ε .

Proof. Identify each $(\sigma, \omega) \in \Sigma_X \times \Omega_X$ with $(\sigma_1, \omega_1, \dots, \sigma_n, \omega_n, \sigma_\varepsilon) \in \left(\prod_{1 \leq i \leq n} \Sigma_i \times \Omega_i\right) \times \Sigma_\varepsilon$. By proposition 32, for every $(\sigma, \omega) \in \Sigma_X \times \Omega_X$,

$$\begin{aligned} \mu_{X,p,q}(\sigma, \omega) &= \frac{1}{Z_{ES}(X, p, q)} (1-p)^{c(\omega)} p^{o(\omega)} \llbracket \sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \rrbracket \\ &= \frac{1}{Z_{ES}(X, p, q)} \left(\prod_{1 \leq i \leq n} (1-p)^{c(\omega_i)} p^{o(\omega_i)} \right) \prod_{Q \in \mathcal{K}_r(X_\omega)} \llbracket \sigma_Q = 1 \rrbracket \\ &= \frac{1}{Z_{ES}(X, p, q)} \prod_{1 \leq i \leq n} \left((1-p)^{c(\omega_i)} p^{o(\omega_i)} \prod_{\substack{Q \in \mathcal{A}_i \\ \omega_i(Q)=1}} \llbracket \sigma_Q = 1 \rrbracket \right). \end{aligned}$$

Each factor $\llbracket \sigma_Q = 1 \rrbracket$ is a function of σ_i and σ_ε because all $(r-1)$ -cubes bordering $Q \in \mathcal{A}_i$ belong to either \mathcal{B}_i or \mathcal{E} . In fact, we may write this as

$$\mu_{X,p,q}(\sigma, \omega) = \frac{\prod_{1 \leq i \leq n} Z_{ES}(X^i, p, q)}{Z_{ES}(X, p, q)} \prod_{1 \leq i \leq n} \mu_{X^i,p,q}(\sigma_i \times \sigma_\varepsilon, \omega_i), \quad (\sigma, \omega) \in \Sigma_X \times \Omega_X.$$

Now fix $\sigma'_\varepsilon \in \Sigma_\varepsilon$ and let $\xi = \{\sigma'_\varepsilon\} \times \prod_i \Sigma_i \times \Omega_i \subseteq \Sigma_X$ (that is, the spin condition ξ specifies the spins on the interface \mathcal{E} but doesn't restrict spins on the remaining $(r-1)$ -cubes.) By proposition 46,

$$\begin{aligned} \mu_{X,p,q}^\xi(\sigma, \omega) &= \llbracket \sigma_\varepsilon = \sigma'_\varepsilon \rrbracket \mu_{X,p,q}(\sigma_1, \dots, \sigma_n, \omega \mid \sigma_\varepsilon) \\ &\propto \llbracket \sigma_\varepsilon = \sigma'_\varepsilon \rrbracket \mu_{X,p,q}(\sigma, \omega) \\ &\propto \llbracket \sigma_\varepsilon = \sigma'_\varepsilon \rrbracket \prod_{1 \leq i \leq n} \mu_{X^i,p,q}(\sigma_i \times \sigma_\varepsilon, \omega_i) \\ &= \prod_{1 \leq i \leq n} \mu_{X^i,p,q}^\xi(\sigma_i \times \sigma_\varepsilon, \omega_i). \end{aligned}$$

Each factor in this product serves as a probability measure on $\Sigma_i \times \Omega_i$, so the net factor of proportionality is 1; that is,

$$\mu_{X,p,q}^\xi(\sigma, \omega) = \prod_{1 \leq i \leq n} \mu_{X^i,p,q}^\xi(\sigma_i \times \sigma_\varepsilon, \omega_i). \quad \square$$

A comment on proposition 51: The Edwards–Sokal coupling (eq. (15)) is a Gibbs random field

in the sense that its probability mass function is a product of locally-determined factors. So it must be a Markov random field (this is the easy direction of the Hammersley–Clifford theorem; see [Lau96, Proposition 3.8].) To wit, construct an undirected graph with vertex set $\mathcal{K}_r(X) \cup \mathcal{K}_{r-1}(X)$, where for each r -cube Q there’s an edge from Q to each of its boundary $(r - 1)$ -cubes, and edge between every pair of Q ’s boundary $(r - 1)$ -cubes. With respect to this graph, the measure $\mu_{\chi,p,q}$ has the global Markov property in the sense of [Lau96, §3.2]: If A, B, S are disjoint vertex sets such that every path from A to B passes through S , then A and B are independent conditional on S . Proposition 51 does not capture the full strength of this statement, as it conditions only on $(r - 1)$ -cubes. It seems futile to seek an analogous theorem for conditioning on a family of r -cubes because such a family cannot separate the graph. This also explains why the higher Potts model enjoys the spatial Markov property whereas the higher FK–Potts model does not, as we’ll see in propositions 52 and 53.

Proposition 52 (Spatial Markov property of higher Potts model). *With respect to $\pi_{\chi,\beta,q}$, the n σ -algebras generated by the factors Σ_i are mutually independent conditional on the σ -algebra generated by the factor Σ_ε .*

Proof. Follows directly from proposition 51 and proposition 47. □

The phrase *hidden Markov* below is meant to suggest at independence conditional on the spins, which aren’t explicitly present in the higher FK–Potts model but instead emerge as a new spin condition ξ on the interfacial $(r - 1)$ -cubes.

Proposition 53 (Spatial hidden Markov property of higher FK–Potts model). *For $n = 2$ in notation 50,*

$$\varphi_{\chi,p,q}(\omega_1, \omega_2) = \sum_{\xi=\{\sigma_\varepsilon\} \subseteq \Sigma_\varepsilon} \pi_{\chi,p,q}(\rho_\varepsilon^{-1} \xi) \varphi_{\chi^1,p,q}^\xi(\omega_1) \varphi_{\chi^2,p,q}^\xi(\omega_2), \quad (\omega_1, \omega_2) \in \Omega_1 \times \Omega_2,$$

where the sum is over all singletons containing an element of Σ_ε .

Proof. Omitting everywhere the subscript “ p, q ” for clarity,

$$\begin{aligned} \varphi_\chi(\omega_1, \omega_2) &= \sum_{\sigma \in \Sigma_\chi} \mu_\chi(\sigma, \omega_1, \omega_2) \\ &= \sum_{\sigma \in \Sigma_\chi} \sum_{\xi=\{\sigma_\varepsilon\} \subseteq \Sigma_\varepsilon} \pi_\chi(\rho_\varepsilon^{-1} \xi) \mu_\chi(\sigma, \omega_1, \omega_2 \mid \sigma \in \rho_\varepsilon^{-1} \xi) && \text{(by proposition 38)} \\ &= \sum_{\xi=\{\sigma_\varepsilon\} \subseteq \Sigma_\varepsilon} \pi_\chi(\rho_\varepsilon^{-1} \xi) \sum_{\sigma \in \Sigma_\chi} \mu_\chi^\xi(\sigma, \omega_1, \omega_2) && \text{(by proposition 46)} \end{aligned}$$

$$\begin{aligned}
&= \sum_{\xi=\{\sigma_\varepsilon\} \subseteq \Sigma_\varepsilon} \pi_X(\rho_\varepsilon^{-1} \xi) \sum_{\substack{\sigma'_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1 \\ \sigma_2 \in \Sigma_2}} \mu_{X^1}^\xi(\sigma_1, \sigma'_\varepsilon, \omega_1) \mu_{X^2}^\xi(\sigma_2, \sigma'_\varepsilon, \omega_2) \quad (\text{by proposition 51}) \\
&= \sum_{\xi=\{\sigma_\varepsilon\} \subseteq \Sigma_\varepsilon} \pi_X(\rho_\varepsilon^{-1} \xi) \sum_{\substack{\sigma_1 \in \Sigma_1 \\ \sigma_2 \in \Sigma_2}} \mu_{X^1}^\xi(\sigma_1, \sigma_\varepsilon, \omega_1) \mu_{X^2}^\xi(\sigma_2, \sigma_\varepsilon, \omega_2) \quad (\text{all other terms vanish}) \\
&= \sum_{\xi=\{\sigma_\varepsilon\} \subseteq \Sigma_\varepsilon} \pi_X(\rho_\varepsilon^{-1} \xi) \left(\sum_{\substack{\sigma'_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1}} \mu_{X^1}^\xi(\sigma_1, \sigma'_\varepsilon, \omega_1) \right) \left(\sum_{\substack{\sigma'_\varepsilon \in \Sigma_\varepsilon \\ \sigma_2 \in \Sigma_2}} \mu_{X^2}^\xi(\sigma_2, \sigma'_\varepsilon, \omega_2) \right) \\
&= \sum_{\xi=\{\sigma_\varepsilon\} \subseteq \Sigma_\varepsilon} \pi_X(\rho_\varepsilon^{-1} \xi) \varphi_{X^1}^\xi(\omega_1) \varphi_{X^2}^\xi(\omega_2) \quad (\text{by proposition 47}),
\end{aligned}$$

$$\omega_1 \in \Omega_1, \omega_2 \in \Omega_2. \quad \square$$

Notice that for $\varepsilon = \emptyset$, which is to say when X^1 and X^2 don't share any $(r-1)$ -cubes, proposition 53 reduces to independence (because Σ_ε is the trivial group, so the sum has only one term):

$$\varphi_{X,p,q}(\omega_1, \omega_2) = \varphi_{X^1,p,q}(\omega_1) \varphi_{X^2,p,q}(\omega_2), \quad (\omega_1, \omega_2) \in \Omega_1 \times \Omega_2.$$

4.3 Conditioning in the higher FK-Potts model

Conditioning in the random-cluster model ($r = 1$) is described in [Gri06, Lemma 4.13]. Our aim is to generalize this result. In the prototypical case, X is a box containing a smaller box X^1 , and we condition on the (open or closed status of the) r -cubes outside X^1 , that is, the elements of $\mathcal{K}_r(X) \setminus \mathcal{K}_r(X^1)$. Proposition 54 describes what happens when all external r -cubes are closed. It is a special case of proposition 56 and proposition 60, but a standalone proof is included to aid comprehension.

Proposition 54. *For $n = 2$ in notation 50, and $p \in [0, 1]$,*

$$\varphi_{X,p,q}(\omega_1 \mid \omega_2 = 0) = \varphi_{X^1,p,q}(\omega_1), \quad (\omega_1, \omega_2) \in \Omega_1 \times \Omega_2.$$

Proof. For every $\omega_1 \in \Omega_1$,

$$\begin{aligned}
\varphi_{X,p,q}(\omega_1 \mid \omega_2 = 0) &\propto \varphi_{X,p,q}(\omega_1, 0) \\
&\propto (1-p)^{c(\omega_1) + |\mathcal{A}_2|} p^{o(\omega_1) + 0} |Z^{r-1}(X_{(\omega_1, 0)})| \\
&\propto (1-p)^{c(\omega_1)} p^{o(\omega_1)} |Z^{r-1}(X_{(\omega_1, 0)})|
\end{aligned}$$

$$\begin{aligned}
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \sum_{\substack{(\sigma_\varepsilon, \sigma_1, \sigma_2) \\ \in \Sigma_\varepsilon \times \Sigma_1 \times \Sigma_2}} \mathbb{I}[(\sigma_1, \sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{(\omega_1, 0)})] \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \\
&\quad \cdot \sum_{\substack{(\sigma_\varepsilon, \sigma_1, \sigma_2) \\ \in \Sigma_\varepsilon \times \Sigma_1 \times \Sigma_2}} \mathbb{I}[(\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1)] \mathbb{I}[(\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_0^2)] \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \sum_{\substack{(\sigma_\varepsilon, \sigma_1, \sigma_2) \\ \in \Sigma_\varepsilon \times \Sigma_1 \times \Sigma_2}} \mathbb{I}[(\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1)] \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} |Z^{r-1}(X_{\omega_1}^1)| |\Sigma_2| \\
&\propto (1-p)^{c(\omega_1)} p^{o(\omega_1)} |Z^{r-1}(X_{\omega_1}^1)| \\
&\propto \varphi_{X^1, p, q}(\omega_1).
\end{aligned}$$

In the sixth line, we used the fact that every configuration $(\sigma_2, \sigma_\varepsilon) \in (\Sigma_2 \times \Sigma_\varepsilon) = \Sigma_{X^2}$ is a cocycle in the cubical set $X_{\omega_2}^2 = X_0^2$ which has no r -cubes. \square

Here is a simple consequence that we'll use in the proof of proposition 71.

Corollary 55. For $n = 2$ in notation 50, take two increasing events $E_1 \subseteq \Omega_1$ and $E \subseteq \Omega_1 \times \Omega_2$ that satisfy

$$\omega_1 \in E_1 \implies (\omega_1, 0) \in E.$$

Then their probabilities satisfy

$$\varphi_{X^1, p, q}(E_1) \leq \varphi_{X, p, q}(E).$$

Proof. Assume $p \in (0, 1)$ (the cases $p = 0, 1$ are trivial.) By conditioning,

$$\begin{aligned}
\varphi_{X, p, q}(E) &= \sum_{\omega'_2 \in \Omega_2} \varphi_{X, p, q}(\omega'_2) \varphi_{X, p, q}(E \mid \omega'_2) \\
&\geq \sum_{\omega'_2 \in \Omega_2} \varphi_{X, p, q}(\omega'_2) \varphi_{X, p, q}(E \mid \omega_2 = 0) \quad (\text{by monotonicity: theorem 35}) \\
&= \varphi_{X, p, q}(E \mid \omega_2 = 0) \\
&= \varphi_{X, p, q}((\omega_1, 0) \in E \mid \omega_2 = 0) \\
&\geq \varphi_{X, p, q}(\omega_1 \in E_1 \mid \omega_2 = 0) \quad (\text{by assumption on } E \text{ and } E_1) \\
&= \varphi_{X^1, p, q}(E_1) \quad (\text{by proposition 54.}) \quad \square
\end{aligned}$$

Now, a more general result. Again, the easiest scenario is a box $X = \mathcal{K}_r([-N, N]^d)$ containing a strictly smaller box $X^1 = \mathcal{K}_r([-M, M]^d)$.

Proposition 56 (Conditioning in the higher FK–Potts model with free boundary condition). *For $n = 2$ in notation 50, and $p \in (0, 1)$,*

$$\varphi_{X,p,q}(\omega_1 \mid \omega_2) = \varphi_{X^1,p,q}^{\xi'}(\omega_1), \quad (\omega_1, \omega_2) \in \Omega_1 \times \Omega_2$$

where $\xi' \subseteq \Sigma_{X^1}$ is a spin condition for X^1 given by

$$\begin{aligned} \xi' &= \{(\sigma_1, \sigma_\varepsilon) \in \Sigma_1 \times \Sigma_\varepsilon \mid (\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \text{ for some } \sigma_2 \in \Sigma_2\} \\ &= \rho_{X,X^1}(\rho_{X,X^2}^{-1}[Z^{r-1}(X_{\omega_2}^2)]). \end{aligned}$$

Proof. Observe that for every $\sigma = (\sigma_\varepsilon, \sigma_1, \sigma_2) \in \Sigma_\varepsilon \times \Sigma_1 \times \Sigma_2 = \Sigma_X$ and $\omega = (\omega_1, \omega_2) \in \Omega_1 \times \Omega_2 = \Omega_X$,

$$\begin{aligned} (\sigma_1, \sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_\omega) &\iff \prod_{Q \in \mathcal{K}_r(X_\omega)} \llbracket \sigma_Q = 1 \rrbracket \quad (\text{by proposition 32}) \\ &\iff \left(\prod_{Q \in \mathcal{K}_r(X_{\omega_1}^1)} \llbracket \sigma_Q = 1 \rrbracket \right) \left(\prod_{Q \in \mathcal{K}_r(X_{\omega_2}^2)} \llbracket \sigma_Q = 1 \rrbracket \right) \\ &\iff \left(\prod_{Q \in \mathcal{K}_r(X_{\omega_1}^1)} \llbracket (\sigma_1, \sigma_\varepsilon)_Q = 1 \rrbracket \right) \left(\prod_{Q \in \mathcal{K}_r(X_{\omega_2}^2)} \llbracket (\sigma_2, \sigma_\varepsilon)_Q = 1 \rrbracket \right) \\ &\iff (\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1) \quad \text{and} \quad (\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \end{aligned}$$

or, equivalently,

$$\llbracket (\sigma_1, \sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_\omega) \rrbracket = \llbracket (\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1) \rrbracket \llbracket (\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \rrbracket.$$

Take $\omega_2 \in \Omega_2$. Conditioning on ω_2 gives

$$\begin{aligned} \varphi_{X,p,q}(\omega_1 \mid \omega_2) &\propto \varphi_{X,p,q}(\omega_1, \omega_2) \\ &\propto (1-p)^{c(\omega_1)+c(\omega_2)} p^{o(\omega_1)+o(\omega_2)} |Z^{r-1}(X_\omega)| \quad \text{where } \omega = (\omega_1, \omega_2) \in \Omega_X \\ &\propto (1-p)^{c(\omega_1)} p^{o(\omega_1)} |Z^{r-1}(X_\omega)| \end{aligned}$$

$$\begin{aligned}
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1 \\ \sigma_2 \in \Sigma_2}} \mathbb{I}[(\sigma_1, \sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_\omega)] \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1 \\ \sigma_2 \in \Sigma_2}} \mathbb{I}[(\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1)] \mathbb{I}[(\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2)] \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \\
&\cdot \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1}} \left(\mathbb{I}[(\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1)] \sum_{\sigma_2 \in \Sigma_2} \mathbb{I}[(\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2)] \right), \quad \omega_1 \in \Omega_1.
\end{aligned}$$

To simplify this expression, we argue that the inner sum always evaluates to either 0 or a positive constant s that is independent of σ_ε and σ_1 . To see why this is so, let $\rho : Z^{r-1}(X_{\omega_2}^2) \rightarrow \Sigma_\varepsilon$ be the coordinate projection (to be more precise, ρ is the restriction to $Z^{r-1}(X_{\omega_2}^2) \subseteq \Sigma_2 \times \Sigma_\varepsilon$ of the projection $\Sigma_2 \times \Sigma_\varepsilon \rightarrow \Sigma_\varepsilon$.) The inner sum evaluates to $|\rho^{-1}\{\sigma_\varepsilon\}|$. But ρ is a group homomorphism, and by the first isomorphism theorem all nonempty preimages of singletons have equal number of elements. Let s be this common number of elements (s is a function of ω_2 but not of σ_ε .) Thus, for every $\sigma_\varepsilon \in \Sigma_\varepsilon$,

$$\sum_{\sigma_2 \in \Sigma_2} \mathbb{I}[(\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2)] = s \mathbb{I}[(\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \text{ for some } \sigma_2 \in \Sigma_2].$$

(The two sides are either both equal to s or both equal to 0, depending on σ_ε .) Pulling out the common factor s gives

$$\begin{aligned}
\varphi_{X,p,q}(\omega_1 | \omega_2) &\propto (1-p)^{c(\omega_1)} p^{o(\omega_1)} \\
&\cdot \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1}} \mathbb{I}[(\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1)] \mathbb{I}[(\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \text{ for some } \sigma_2 \in \Sigma_2] \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1}} \mathbb{I}[(\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1) \cap \xi'] \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} |Z^{r-1}(X_{\omega_1}^1) \cap \xi'| \\
&\propto \varphi_{X^1,p,q}^{\xi'}(\omega_1), \quad \omega_1 \in \Omega_1,
\end{aligned}$$

where ξ' is as described in the theorem statement above. □

Here's a simple illustration of proposition 56. It also serves as a prototypical example for

proposition 59.

Example 57. Take $d \geq 3$ and $r = 2$. Let X be the union of the 6 faces (plaquettes) of some 3-cube. Let \mathcal{A}_1 contain a single one of these plaquettes, and \mathcal{A}_2 the remaining 5 plaquettes. Then Σ_1 is trivial (there are no edges internal to the single plaquette in \mathcal{A}_1) and $\Sigma_\mathcal{E}$ is the configuration space of the spins on the 4 edges of the plaquette in \mathcal{A}_1 . Let ω_2 be the all-open configuration on Ω_2 (that is, condition on each of the 5 plaquettes in \mathcal{A}_2 being open.)

If $(\sigma_2, \sigma_\mathcal{E}) \in Z^{r-1}(X_{\omega_2}^2)$ then $\sigma_\mathcal{E} \in Z^{r-1}(X^1)$ (because the plaquettes form a closed surface, so that if the spins around the boundaries of 5 of them have sum 0 then the same is true of the 6th plaquette.) Conversely, it's not hard to see that for every $\sigma_\mathcal{E} \in Z^{r-1}(X^1)$ there exists $\sigma_2 \in \Sigma_2$ such that $(\sigma_2, \sigma_\mathcal{E}) \in Z^{r-1}(X_{\omega_2}^2)$. So ξ' is the set of all spin configurations on the edges incident to the single plaquette in \mathcal{A}_1 such that the sum of the 4 spins is 0. \triangle

We can now define imprint BSCs: those that can be obtained by starting with free boundary condition on some larger cubical set and then conditioning as in proposition 56.

Definition 58. Let X^1 be a nonempty union of r -cubes. An *imprint boundary spin condition* (or *imprint BSC*) for X^1 is a set $\xi' \subseteq \Sigma_{X^1}$ that satisfies the following condition. There exists some union of r -cubes²⁷ $X \supseteq X^1$, some partition of the r -cubes in X into 2 disjoint nonempty subfamilies \mathcal{A}_1 and \mathcal{A}_2 (that is, taking $n = 2$ in notation 50), with $X^1 = \bigcup_{Q \in \mathcal{A}_1} Q$, and some configuration $\omega_2 \in \Omega_2$, such that

$$\xi' = \rho_{X, X^1}(\rho_{X, X^2}^{-1}[Z^{r-1}(X_{\omega_2}^2)])$$

(where $X^2 = \bigcup_{Q \in \mathcal{A}_2} Q$ and other symbols as described in notation 50.) \triangle

For $r = 1$, there's a well-known description of imprint BSCs, often referred to as the *domain Markov property*²⁸ [Dum20, §1.2]: Every imprint BSC may be identified with a partition of \mathcal{E} . Within each block of the partition, all spins are required to be the same. There might be an analogous topological characterization of imprint BSCs for general r , but we will not investigate this.

Proposition 59. *Every imprint BSC is a cyclic BSC (that is, in proposition 56, ξ' is a cyclic BSC for X^1 .)*

²⁷It would not weaken this definition to take weak inclusion $X \supseteq X^1$ instead, because the free spin condition $\xi' = \Sigma_{X^1}$ can be obtained by conditioning on an all-closed external configuration as in proposition 54.

²⁸This terminology is questionable because there is no true conditional independence—see proposition 53.

Proof. Let ω_2 and ξ' be as in proposition 56. For this proof, we'll use the coordinate projections

$$\begin{aligned}\rho &:= \rho_{X^2, \mathcal{E}} : \Sigma_2 \times \Sigma_{\mathcal{E}} \rightarrow \Sigma_{\mathcal{E}}, \\ \rho' &:= \rho_{X^1, \mathcal{E}} : \Sigma_1 \times \Sigma_{\mathcal{E}} \rightarrow \Sigma_{\mathcal{E}}.\end{aligned}$$

In fact 7, take α to be the coordinate injection $\kappa : (\mathbb{Z}/q\mathbb{Z})^{\mathcal{E}} \rightarrow (\mathbb{Z}/q\mathbb{Z})^{\mathcal{E} \times \mathcal{B}_2}$. Its dual map is $\kappa^* = \rho$ (fact 10). The composition $\rho \circ \rho^{-1}$ is the identity, so the second part of fact 7 says

$$\text{Ann}_{(\mathbb{Z}/q\mathbb{Z})^{\mathcal{E}}} \circ \kappa^{-1} = \rho \circ \text{Ann}_{(\mathbb{Z}/q\mathbb{Z})^{\mathcal{E} \times \mathcal{B}_2}}. \quad (*)$$

By fact 8,

$$Z^{r-1}(X_{\omega_2}^2) = \ker \delta^{r-1} = \ker \partial_r^* = \text{Ann}_{(\mathbb{Z}/q\mathbb{Z})^{\mathcal{E} \times \mathcal{B}_2}} \text{im } \partial_r = \text{Ann}_{(\mathbb{Z}/q\mathbb{Z})^{\mathcal{E} \times \mathcal{B}_2}} B_{r-1}(X_{\omega_2}^2).$$

Applying ρ and combining with (*) gives

$$\rho(Z^{r-1}(X_{\omega_2}^2)) = \rho(\text{Ann}_{(\mathbb{Z}/q\mathbb{Z})^{\mathcal{E} \times \mathcal{B}_2}} B_{r-1}(X_{\omega_2}^2)) = \text{Ann}_{(\mathbb{Z}/q\mathbb{Z})^{\mathcal{E}}} \kappa^{-1}(B_{r-1}(X_{\omega_2}^2)).$$

Therefore,

$$\begin{aligned}\xi' &= \{(\sigma_1, \sigma_{\mathcal{E}}) \in \Sigma_1 \times \Sigma_{\mathcal{E}} \mid (\sigma_2, \sigma_{\mathcal{E}}) \in Z^{r-1}(X_{\omega_2}^2) \text{ for some } \sigma_2 \in \Sigma_2\} \\ &= (\rho')^{-1}(\rho(Z^{r-1}(X_{\omega_2}^2))) \\ &= (\rho')^{-1} \text{Ann}_{(\mathbb{Z}/q\mathbb{Z})^{\mathcal{E}}} \kappa^{-1}(B_{r-1}(X_{\omega_2}^2)).\end{aligned}$$

Recall that $B_{r-1}(X_{\omega_2}^2) = \text{im } \partial_r = \text{im } \tilde{\partial}_r = \tilde{B}_{r-1}(X_{\omega_2}^2)$ because $r \geq 1$. Take any $b \in (\mathbb{Z}/q\mathbb{Z})^{\mathcal{E}}$ such that

$$\kappa(b) \in B_{r-1}(X_{\omega_2}^2) = \tilde{B}_{r-1}(X_{\omega_2}^2) \subseteq \tilde{Z}_{r-1}(X_{\omega_2}^2)$$

(every reduced boundary is a reduced cycle.²⁹) Then $b \in \tilde{Z}_{r-1}(\bigcup_{E \in \mathcal{E}} E)$, because the extra 0 coefficients in $\kappa(b)$ don't contribute anything to the coefficients on $(r-2)$ -cubes after applying $\tilde{\partial}_{r-1}$. This proves that $\kappa^{-1}(B_{r-1}(X_{\omega_2}^2)) \subseteq \tilde{Z}_{r-1}(\bigcup_{E \in \mathcal{E}} E)$. It follows that ξ' is a cyclic BSC for X^1 . To be explicit: Take $\Xi = \bar{\kappa} \kappa^{-1} B_{r-1}(X_{\omega_2}^2)$ in the definition of cyclic BSC, where $\bar{\kappa} : (\mathbb{Z}/q\mathbb{Z})^{\mathcal{E}} \rightarrow$

²⁹But, if $r = 1$, not every $(r-1)$ -cycle is a reduced $(r-1)$ -cycle, which is the whole point of using the reduced homology here.

$(\mathbb{Z}/q\mathbb{Z})^{\mathcal{X}_{r-1}(\partial X^1)} = C_{r-1}(\partial X^1)$ is the coordinate injection (this follows from the first part of fact 7 with $\alpha = \bar{\kappa}$.)

Actually, we've proved slightly more: ξ' is induced by a family of cycles in $\bigcup_{E \in \mathcal{E}} E$, which may be a strict subset of ∂X . \square

The converse to proposition 59 does not hold in general: It's not true that every cyclic BSC is an imprint BSC. As a counterexample, take the Potts model with $r = 1$ and $q = 4$, and let X be a single 1-cube; that is, an edge joining two vertices v and w . In the definition of cyclic BSC let $\Xi = \{c\}$ where $c = 2_v + 2_w$. Then ξ is the set of all vertex spin configurations where the sum of spins has even parity (there are 8 such configurations.) But any imprint BSC will either leave both spins free or will require both spins to be equal (so an imprint BSC will have either 4 or 16 configurations.)

Proposition 56 described conditioning starting with free boundary condition, but we can be somewhat more general.

Proposition 60 (Conditioning in the higher FK–Potts model). *Take $n = 2$ in notation 50 and $p \in (0, 1)$. Let ξ be a subgroup spin condition on X that has the form $(\rho_{X, X^2})^{-1}(\xi^2)$ for some subgroup $\xi^2 \subseteq \Sigma_{X^2}$ (meaning that ξ is permitted to restrict spins in \mathcal{B}_2 and \mathcal{E} but not in \mathcal{B}_1 .) Then*

$$\varphi_{X, p, q}^{\xi}(\omega_1 \mid \omega_2) = \varphi_{X^1, p, q}^{\xi'}(\omega_1), \quad (\omega_1, \omega_2) \in \Omega_1 \times \Omega_2$$

where $\xi' \subseteq \Sigma_{X^1}$ is a subgroup boundary spin condition for X^1 given by

$$\begin{aligned} \xi' &= \{(\sigma_1, \sigma_{\mathcal{E}}) \in \Sigma_1 \times \Sigma_{\mathcal{E}} \mid (\sigma_2, \sigma_{\mathcal{E}}) \in Z^{r-1}(X_{\omega_2}^2) \cap \xi^2 \text{ for some } \sigma_2 \in \Sigma_2\} \\ &= \rho_{X, X^1}(\rho_{X, X^2}^{-1}[Z^{r-1}(X_{\omega_2}^2) \cap \xi^2]) \\ &= \rho_{X, X^1}(\rho_{X, X^2}^{-1}[Z^{r-1}(X_{\omega_2}^2)] \cap \xi). \end{aligned}$$

Proof. Take $\omega_2 \in \Omega_2$. Similarly to the proof of proposition 56, conditioning on ω_2 gives

$$\begin{aligned} \varphi_{X, p, q}^{\xi}(\omega_1 \mid \omega_2) &\propto \varphi_{X, p, q}^{\xi}(\omega_1, \omega_2) \\ &\propto (1-p)^{c(\omega_1)+c(\omega_2)} p^{o(\omega_1)+o(\omega_2)} |Z^{r-1}(X_{(\omega_1, \omega_2)}) \cap \xi| \\ &\propto (1-p)^{c(\omega_1)} p^{o(\omega_1)} |Z^{r-1}(X_{(\omega_1, \omega_2)}) \cap \xi| \end{aligned}$$

$$\begin{aligned}
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \\
&\quad \cdot \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1 \\ \sigma_2 \in \Sigma_2}} \llbracket (\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1) \rrbracket \llbracket (\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \rrbracket \llbracket (\sigma_1, \sigma_2, \sigma_\varepsilon) \in \xi \rrbracket \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \\
&\quad \cdot \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1 \\ \sigma_2 \in \Sigma_2}} \llbracket (\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1) \rrbracket \llbracket (\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \cap \xi^2 \rrbracket \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \\
&\quad \cdot \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1}} \left(\llbracket (\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1) \rrbracket \sum_{\sigma_2 \in \Sigma_2} \llbracket (\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \cap \xi^2 \rrbracket \right), \quad \omega_1 \in \Omega_1.
\end{aligned}$$

Next, argue as in the proof of proposition 56, but with the group $Z^{r-1}(X_{\omega_2}^2) \cap \xi^2$ in place of $Z^{r-1}(X_{\omega_2}^2)$. We arrive at

$$\begin{aligned}
\varphi_{X,p,q}^{\xi}(\omega_1 | \omega_2) &\propto (1-p)^{c(\omega_1)} p^{o(\omega_1)} \\
&\quad \cdot \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1}} \llbracket (\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1) \rrbracket \llbracket (\sigma_2, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_2}^2) \cap \xi^2 \text{ for some } \sigma_2 \in \Sigma_2 \rrbracket \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} \sum_{\substack{\sigma_\varepsilon \in \Sigma_\varepsilon \\ \sigma_1 \in \Sigma_1}} \llbracket (\sigma_1, \sigma_\varepsilon) \in Z^{r-1}(X_{\omega_1}^1) \cap \xi' \rrbracket \\
&= (1-p)^{c(\omega_1)} p^{o(\omega_1)} |Z^{r-1}(X_{\omega_1}^1) \cap \xi'| \\
&\propto \varphi_{X^1,p,q}^{\xi'}(\omega_1), \quad \omega_1 \in \Omega_1,
\end{aligned}$$

where ξ' is as described in the theorem statement above.

The projections are group homomorphisms, so ξ' is a subgroup of Σ_{X^1} . Moreover, ξ' restricts only the Σ_ε component, not the Σ_1 component, and every $(r-1)$ -cube $E \in \mathcal{E}$ satisfies $E \subseteq \partial X^1$, so ξ' is a boundary spin condition for Σ_{X^2} . \square

4.4 Further results

Here, we generalize several statements from sections 3.2 and 3.4. Proposition 61 (which generalizes the first part of proposition 33) is needed for the proof of theorem 62.

Proposition 61 (Counting cocycles, with spin condition). *Let ξ be a subgroup spin condition. The*

dependence factor in eq. (17) satisfies

$$|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi| = \frac{|C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})|}{|B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) + \text{Ann}^{-1} \xi|}, \quad \omega \in \Omega_X$$

where $\text{Ann}^{-1} \xi = \eta^{-1} \text{Ann} \xi \subseteq C_{r-1}(X, \mathbb{Z}/q\mathbb{Z}) = C_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})$ (see facts 5 and 6, taking $H = \text{Ann}^{-1} \xi$ in fact 5.)

Proof. In this proof, the boundary (∂_r) and coboundary (δ^{r-1}) maps will be those of the cubical set X_ω , and Ann will be the induced bijection from the collection of subgroups of $C_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})$ to the collection of subgroups of $\Sigma_X = C^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})$ (see fact 6, noting in particular that $\xi = \text{Ann} \text{Ann}^{-1} \xi$.) For every $\omega \in \Omega_X$,

$$\begin{aligned} |Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi| &= |\ker \delta^{r-1} \cap \xi| \\ &= |\ker \partial_r^* \cap \xi| \\ &= |\text{Ann}(\text{im } \partial_r) \cap \text{Ann} \text{Ann}^{-1} \xi| && \text{(by fact 8)} \\ &= |\text{Ann}(\text{im } \partial_r + \text{Ann}^{-1} \xi)| && \text{(by fact 6)} \\ &= \frac{|C_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|}{|\text{im } \partial_r + \text{Ann}^{-1} \xi|} && \text{(by fact 4)} \\ &= \frac{|C_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|}{|B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) + \text{Ann}^{-1} \xi|} \\ &= \frac{|C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})|}{|B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) + \text{Ann}^{-1} \xi|} && \text{(by eq. (11).)} \quad \square \end{aligned}$$

Theorem 62 generalizes theorem 35. Observe that if ξ is a subgroup of Σ_X then $\varphi_{X,p,q}^\xi$ is (strictly) positive when $p \in (0, 1)$, because $Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi \neq \emptyset$ for every $\omega \in \Omega_X$. So theorem 34 applies to $\varphi_{X,p,q}^\xi$ just as it does to $\varphi_{X,p,q}$.

Theorem 62 (Strong FKG, with spin condition). *For every $p \in (0, 1)$ and every subgroup spin condition ξ , the higher FK–Potts measure (17) has the strong FKG property.*

Proof. Follow the proof of theorem 35, except now use the cocycle count formula from proposition 61, and at the end apply the more general result from fact 14 taking $D = \text{Ann}^{-1} \xi$. \square

Proposition 63 generalizes proposition 36.

Proposition 63 (Comparison inequality, with spin condition). *For every subgroup spin condition ξ , if $0 \leq p_1 \leq p_2 \leq 1$ then*

$$\varphi_{X,p_1,q}^\xi \leq_{\text{st}} \varphi_{X,p_2,q}^\xi.$$

Proof. Follow the proof of proposition 36, replacing $|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|$ with $|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi|$. \square

Recall the notation W_γ from section 3.4. Let $\langle W_\gamma \rangle_{X,\beta,q}^\xi := \pi_{X,\beta,q}^\xi W_\gamma$. Theorem 64 generalizes theorem 41.

Theorem 64 (Expectation equals probability, with spin condition). *For every subgroup spin condition ξ and every $(r-1)$ -chain $\gamma \in C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})$,*

$$\langle W_\gamma \rangle_{X,\beta,q}^\xi = \varphi_{X,p,q}^\xi (\gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) + \text{Ann}^{-1} \xi).$$

Proof. We first consider the case $0 \leq p < 1$. Let $\xi' = \text{Ann}^{-1} \xi$. The set ξ' is a subgroup of $C_{r-1}(X, \mathbb{Z}/q\mathbb{Z}) (= C_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}))$ (see fact 6.) Define w_γ to be the conditional expectation

$$\begin{aligned} w_\gamma(\omega) &:= \mu_{X,p,q}^\xi(W_\gamma | \omega) = \sum_{\sigma \in \Sigma_X} W_\gamma(\sigma) \mu_{X,p,q}^\xi(\sigma | \omega) \\ &= \sum_{\sigma \in \Sigma_X} \sigma(\gamma) \frac{|\{\sigma \in Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi\}|}{|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi|} \\ &= \frac{1}{|\xi \cap \ker \delta^{r-1}|} \sum_{\sigma \in \xi \cap \ker \delta^{r-1}} \sigma(\gamma) \\ &= \frac{1}{|\xi \cap \text{Ann}(\text{im } \partial_r)|} \sum_{\sigma \in \xi \cap \text{Ann}(\text{im } \partial_r)} \sigma(\gamma) && \text{(by fact 8)} \\ &= \frac{1}{|\text{Ann}(\text{im } \partial_r + \xi')|} \sum_{\sigma \in \text{Ann}(\text{im } \partial_r + \xi')} \sigma(\gamma) && \text{(by fact 6)} \\ &= \llbracket \gamma \in \text{im } \partial_r + \xi' \rrbracket && \text{(by fact 12)} \\ &= \llbracket \gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) + \text{Ann}^{-1} \xi \rrbracket, && \omega \in \Omega_X. \end{aligned}$$

(This is valid for all $\omega \in \Omega_X$, because ξ is a subgroup and therefore $Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) \cap \xi \neq \emptyset$.)

Thus, by the law of total expectation,

$$\langle W_\gamma \rangle_{X,\beta,q}^\xi = \varphi_{X,p,q}^\xi w_\gamma = \varphi_{X,p,q}^\xi (\gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) + \text{Ann}^{-1} \xi).$$

That takes care of $0 \leq p < 1$. If $p = 1$ then $\varphi_{X,p,q}^\xi(\omega^1) = 1$ (where ω^1 is the all-open configuration), and the above derivation is valid for $\omega = \omega^1$. Thus, the equality of functions

$$w_\gamma(\omega) = \llbracket \gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) + \text{Ann}^{-1} \xi \rrbracket.$$

holds with probability 1, and we may take the expectation on both sides just as before. \square

A simple example for theorem 64 is the Ising model (spins 0 and 1) with wired boundary condition, that is, ξ is the set of all spin configurations that assign equal spin to each boundary vertex. The group $\text{Ann}^{-1} \xi$ consists of all 1-chains in which an even number of boundary vertices have spin 1. Let $\gamma = 1_v + 1_w$ for vertices v, w . The expectation of W_γ is equal to the probability that γ differs from a boundary by some element of $\text{Ann}^{-1} \xi$, and is always equal to 1, even if v and w lie in distinct connected components of X . But for free boundary condition the expectation of W_γ is 0 if v and w lie in distinct components. In general, to see why it's necessary to assume that ξ be a subgroup, consider the following (trivial) counterexample: Let $q = 3$, take $\gamma = 1_Q$ for some $Q \in \mathcal{K}_{r-1}(X)$, and let $\xi = \{c\}$ for some $c \in \Sigma_X$ such that $c(Q) = e^{2\pi i/3}$. Then $\langle W_\gamma \rangle_{X,\beta,q}^\xi \notin \mathbb{R}$.

Corollary 65 generalizes corollary 42.

Corollary 65. *For every subgroup spin condition ξ , if $0 \leq \beta_1 \leq \beta_2 \leq \infty$ then*

$$0 \leq \langle W_\gamma \rangle_{X,\beta_1,q}^\xi \leq \langle W_\gamma \rangle_{X,\beta_2,q}^\xi \leq 1$$

for every $(r-1)$ -chain γ .

Proof. As in the proof of corollary 42, this follows immediately from theorem 64 and proposition 63, now taking instead the increasing event

$$\{\omega \in \Omega_X \mid \gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) + \text{Ann}^{-1} \xi\}. \quad \square$$

4.5 Sharp threshold

Lastly, as promised on page 51, here is a result about the sharpness of thresholds. Its proof is a straightforward extension of the proof of [Gri06, Theorem 3.16],³⁰ which applies to the random-

³⁰It appears there may be a small error in the statement and proof of [Gri06, Theorem 3.16]: in place of our factor $\frac{1}{p+q(1-p)}$ it contains $\min \left\{ 1, \frac{q}{(p+q(1-p))^2} \right\}$.

cluster model on arbitrary graphs. A few consequences regarding thresholds may be found in [Gri06, p. 42].

As a simple example, take $d \geq 2$ and $r = 2$ and $q \geq 2$, take $\mathcal{A} = \mathcal{G}$ (the group of all lattice translations, rotations, reflections, and compositions thereof), let γ be a nontrivial 1-cycle supported on some 5×5 square, and let A be the event that some translate (under \mathcal{G}) of γ belongs to $B_1(X_\omega, \mathbb{Z}/q\mathbb{Z})$. Since $\frac{c}{p+q(1-p)} \geq \frac{c}{q}$, the derivative $\frac{d}{dp} \varphi_{X,p,q}^\xi(A)$ at the threshold (p such that $\varphi_{X,p,q}^\xi(A) = \frac{1}{2}$) is bounded below by $\frac{c}{2q} \log |\mathcal{K}_r(X)|$.

Proposition 66 (Sharp threshold). *There exists a constant $c \in (0, \infty)$, independent of all parameters, such that the following holds. Let X be the modified box with periodic boundary spin condition ξ , as described on page 61, and identify X with a subset of the d -torus $\mathbb{T}^d = \mathbb{R}^d / 2\pi\mathbb{Z}^d$. Let \mathcal{G} be the group of all isometries on \mathbb{T}^d that send integer lattice points to integer lattice points, and let \mathcal{A} be a subgroup of \mathcal{G} such that $\mathcal{K}_r(X)$ is \mathcal{A} -transitive. Let $A \subseteq \Omega_X$ be an increasing event that is invariant under \mathcal{A} . Then, for every $p \in (0, 1)$,*

$$\frac{d}{dp} \varphi_{X,p,q}^\xi(A) \geq \frac{c}{p+q(1-p)} \min \{ \varphi_{X,p,q}^\xi(A), 1 - \varphi_{X,p,q}^\xi(A) \} \log |\mathcal{K}_r(X)|.$$

Proof. The periodic boundary spin condition and the modification of the box ensure that $\varphi_{X,p,q}^\xi$ is \mathcal{G} -invariant. According to [Gri06, Theorem 2.48] there exists c such that, for all X, A , and p as described above,

$$\frac{d}{dp} \varphi_{X,p,q}^\xi(A) \geq c \frac{\varphi_{X,p,q}^\xi(J_Q)(1 - \varphi_{X,p,q}^\xi(J_Q))}{p(1-p)} \min \{ \varphi_{X,p,q}^\xi(A), 1 - \varphi_{X,p,q}^\xi(A) \} \log |\mathcal{K}_r(X)|$$

where J_Q is the event that Q is open for some given $Q \in \mathcal{K}_r(X)$ (by invariance, the particular choice of Q doesn't matter.)

We produce bounds on the probability $\varphi_{X,p,q}^\xi(J_Q)$ by conditioning on all r -cubes in $\mathcal{K}_r(X) \setminus \{Q\}$. For $\omega \in \Omega_X$ write ω^+ and ω^- for the configuration ω modified on Q to make Q open and closed, respectively. Let $M_\omega = \{\omega^+, \omega^-\}$. Then

$$\begin{aligned} \varphi_{X,p,q}^\xi(J_Q | M_\omega) &= \varphi_{X,p,q}^\xi(\omega^+ | M_\omega) \\ &= \frac{\varphi_{X,p,q}^\xi(\omega^+)}{\varphi_{X,p,q}^\xi(\omega^+) + \varphi_{X,p,q}^\xi(\omega^-)} \\ &= \frac{p|Z^{r-1}(X_{\omega^+})|}{p|Z^{r-1}(X_{\omega^+})| + (1-p)|Z^{r-1}(X_{\omega^-})|}. \end{aligned}$$

For every $\omega \in \Omega_X$ the ratio $|Z^{r-1}(X_{\omega^-})|/|Z^{r-1}(X_{\omega^+})|$ is a natural number that divides q , because

$$Z^{r-1}(X_{\omega^+}) = \{c \in Z^{r-1}(X_{\omega^-}) \mid \sigma_Q(c) = 1\}$$

and σ_Q is a group homomorphism into the group of complex roots of q . So

$$\begin{aligned} \varphi_{X,p,q}^\xi(J_Q \mid M_\omega) &\in \left[\frac{p}{p+q(1-p)}, p \right] \quad \text{and} \\ 1 - \varphi_{X,p,q}^\xi(J_Q \mid M_\omega) &\in \left[1-p, \frac{q(1-p)}{p+q(1-p)} \right]. \end{aligned}$$

By the law of total probability, the same bounds hold for $\varphi_{X,p,q}^\xi(J_Q)$ and $1 - \varphi_{X,p,q}^\xi(J_Q)$, respectively. Therefore,

$$\frac{\varphi_{X,p,q}^\xi(J_Q)(1 - \varphi_{X,p,q}^\xi(J_Q))}{p(1-p)} \geq \frac{1}{p+q(1-p)},$$

which combined with the inequality above completes the proof. □

5 Infinite volume

Sections 3 and 4 describe the higher Potts and FK–Potts models in a finite region, but of course statistical physics is concerned with phenomena that emerge as the number of interacting elements tends to infinity. In lattice gauge theories, a major open problem is to understand the decay of a Wilson loop expectation with the size of the loop. It has been argued that if the Wilson loop expectation decays exponentially in the area enclosed by the loop, then the gauge theory has *quark confinement*, meaning that quarks do not appear in isolation [Cha21]. See also [Aiz+83], which studies the sharpness of the transition from exponential-in-area decay to exponential-in-perimeter decay in the case of (independent) Bernoulli plaquette percolation.

We won't discuss the decay of Wilson loop expectations, and we won't even carry out a complete investigation of infinite-volume limits—indeed, even for the Ising model on \mathbb{Z}^3 the infinite-volume limits are not fully understood [Bov06, p. 72]. We can, however, give some definitions and a few preliminary results.

We will focus on the infinite-volume higher Potts model, using the (finite-volume) coupling with the higher FK–Potts model as a proof device.

5.1 Gibbs states of the higher Potts model

For the basic setup, we'll take the DLR approach (due to Dobrushin, Lanford, and Ruelle [Bov06, p. 51].) It would take us too far off-track to explain the DLR machinery in general, so the definitions given below are specific to our model. Unfortunately, despite efforts to keep this section reasonably self-contained, a full understanding may be difficult without knowledge of the general case. The interested reader can find introductions in [FV17, ch. 6; Bov06, ch. 4; EFS93, §2] and more comprehensive treatments in [Rue04; Geo11].

Roughly speaking, a (DLR) Gibbs state is a measure on the space of all spin configurations on an infinite lattice, whose conditionals on all finite sublattices are described by *Gibbs ensembles*: probability measures of the form $\frac{1}{Z} e^{-\mathcal{H}}$ with appropriate Hamiltonians \mathcal{H} . (This kind of roundabout definition is needed because there's no way to define a Hamiltonian on the entire infinite lattice at once.) Thus, a Gibbs state can be thought of as describing a macroscopic physical system that is everywhere at microscopic equilibrium. The Gibbs states include all infinite-volume weak limits of the (finite-volume) Gibbs ensembles [Rue04, §1.9]. Moreover, the *variational principle* dictates that the Gibbs states that are translation-invariant are precisely the measures that maximize the

topological pressure [Rue04, §4.2], a quantity that can be interpreted as the negation of free energy density. So, in essence, the DLR framework captures the idea that a spin system is globally at equilibrium if and only if it is locally at equilibrium. See [EFS93, pp. 933–934] for more on the physical interpretation of Gibbs states.

Gibbs states are by no means the last word on lattice spin systems. A more general framework is that of *specifications*, described in [Geo11]. In loose terms, non-Gibbsian specifications are those where the local conditional measures cannot be described by a Hamiltonian based on a spin interaction that decays sufficiently rapidly with distance. Non-Gibbsian specifications are not uncommon. They often crop up when taking scaling limits [EFS93]. Also, even though the higher Potts model can be described by a Gibbsian specification, the same cannot be said of the random-cluster model or our higher FK–Potts model. Infinite-volume FK–Potts measures can be defined via non-Gibbsian specifications (as is done in [Gri06, §4.4]), but we will not investigate this approach here. Instead, we’ll focus on the Gibbs states of the higher Potts model because, after all, the ultimate objective is to understand gauge theories.

Our first task is to define the configuration space Σ . A configuration $\sigma \in \Sigma$ will be a simultaneous assignment of a spin to each $(r - 1)$ -cube in \mathbb{R}^d (the parameters r, d, p, q, β are as described at the beginning of section 3.) Formally, recall that in finite volume (section 3.1) we took configuration space $\Sigma_X = C^{r-1}(X, \mathbb{Z}/q\mathbb{Z}) = (C_{r-1}(X, \mathbb{Z}/q\mathbb{Z}))^\wedge = ((\mathbb{Z}/q\mathbb{Z})^{\mathcal{K}_{r-1}(X)})^\wedge$. Our chains, cochains, homology, etc. were defined only for finite volume, so we now define

$$\Sigma := \left((\mathbb{Z}/q\mathbb{Z})^{\mathcal{K}_{r-1}(\mathbb{R}^d)} \right)^\wedge,$$

where $(\mathbb{Z}/q\mathbb{Z})^{\mathcal{K}_{r-1}(\mathbb{R}^d)}$ is the direct sum of countably infinitely many copies of the group $\mathbb{Z}/q\mathbb{Z}$, one for each $(r - 1)$ -cube (recall from definition 15 that the parentheses in the exponent indicate direct sum, as opposed to direct product.) That is, $(\mathbb{Z}/q\mathbb{Z})^{\mathcal{K}_{r-1}(\mathbb{R}^d)}$ is the group of all finitely-supported assignments of coefficients in $\mathbb{Z}/q\mathbb{Z}$ to the $(r - 1)$ -cubes. Thus, by fact 16 and the identification $\mathbb{Z}/q\mathbb{Z} \cong \widehat{\mathbb{Z}/q\mathbb{Z}}$, the configuration space Σ is identified with the direct product

$$\Sigma \cong (\mathbb{Z}/q\mathbb{Z})^{\mathcal{K}_{r-1}(\mathbb{R}^d)}$$

(all assignments of coefficients to $(r - 1)$ -cubes, now *without* the requirement of finite support.)

More generally, let

$$\Sigma_\Lambda := ((\mathbb{Z}/q\mathbb{Z})^{(\wedge)})^\wedge \cong (\mathbb{Z}/q\mathbb{Z})^\wedge \quad \text{for all } \Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d).$$

Endow Σ_Λ with the product topology; that is, take the discrete topology $\mathbb{Z}/q\mathbb{Z}$ and the product topology on $(\mathbb{Z}/q\mathbb{Z})^\wedge$. The space Σ_Λ is compact for every $\Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d)$. Let $\rho_\Lambda : \Sigma \rightarrow \Sigma_\Lambda$ be the coordinate projection.

The following terminology is standard [EFS93, pp. 895–896; GHM01, §2.3].

Definition 67. An *observable* is a Borel-measurable function $f : \Sigma \rightarrow \mathbb{C}$. A *local observable* is a function $f : \Sigma \rightarrow \mathbb{C}$ that may be written as $f = f_\Lambda \circ \rho_\Lambda$ for some finite $\Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d)$ and some function $f_\Lambda : \Sigma_\Lambda \rightarrow \mathbb{C}$. That is, a local observable is one that depends on the spins of only finitely many $(r-1)$ -cubes. \triangle

Any element $\gamma \in (\mathbb{Z}/q\mathbb{Z})^{(\mathcal{K}_{r-1}(\mathbb{R}^d))}$ (such as a Wilson loop: section 3.4) thus gives a local observable $W_\gamma : \Sigma \rightarrow \mathbb{C}, \sigma \mapsto \sigma(\gamma)$. Once we have a measure on Σ , we may define the *observed value* of γ to be the expectation of W_γ .

For the exposition below, we'll introduce some special notation.

Notation 68. The reader may refer to fig. 4 as a visual guide. Take a finite nonempty set $\Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d)$. Let

$$X_\Lambda = \bigcup \{Q \in \mathcal{K}_r(\mathbb{R}^d) \mid \exists P \in \Lambda : P \subseteq Q\}.$$

That is, X_Λ is the cubical set consisting of all r -cubes incident to some $(r-1)$ cube in Λ . Write

$$\begin{aligned} \Sigma_\Lambda &= ((\mathbb{Z}/q\mathbb{Z})^\wedge)^\wedge, \\ \bar{\Lambda} &= \mathcal{K}_{r-1}(X_\Lambda), & \Sigma_{\bar{\Lambda}} &= ((\mathbb{Z}/q\mathbb{Z})^{\bar{\Lambda}})^\wedge, \\ \Lambda' &= \mathcal{K}_{r-1}(X_\Lambda) \setminus \Lambda, \\ \Lambda^c &= \mathcal{K}_{r-1}(\mathbb{R}^d) \setminus \Lambda, & \Sigma_{\Lambda^c} &= ((\mathbb{Z}/q\mathbb{Z})^{(\Lambda^c)})^\wedge. \end{aligned}$$

The set Λ' includes, but may be strictly larger than, the set $\mathcal{K}_{r-1}(\partial X_\Lambda)$ of all boundary $(r-1)$ -cubes of X_Λ . Note that $\Sigma = \Sigma_{\Lambda^c} \oplus \Sigma_\Lambda$, and $\Sigma_{X_\Lambda} = \Sigma_{\bar{\Lambda}}$. Here Σ_{X_Λ} is defined as in section 3.1 taking X_Λ in place of X ; notice the distinction between Σ_Λ and Σ_{X_Λ} . Let

$$\xi_\eta = \{s \in \Sigma_{\bar{\Lambda}} \mid s|_{\Lambda'} = \eta|_{\Lambda'}\} \quad \text{for every } \eta \in \Sigma_{\Lambda^c}.$$

The set ξ_η is a spin condition for X_Λ . In the terminology of section 4, if $\Lambda' = \mathcal{K}_{r-1}(\partial X_\Lambda)$ (which is not the case in fig. 4) then ξ_η is a point BSC for X_Λ . \triangle




Figure 4: Example set $X_\Lambda \subseteq \mathbb{R}^2$ for $r = 1$. Here Λ consists of 8 vertices, and X_Λ is the union of all displayed edges together with their endpoints.

Definition 69. A (DLR) Gibbs state in the higher Potts model is a Borel³¹ probability measure $\pi_{\beta, q}$ on Σ that satisfies either, hence both, of the following conditions (which are equivalent by [Rue04, Theorem 1.8].)

(a) (Conditionals)

For every nonempty finite $\Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d)$,

$$\pi_{\beta, q} = (\rho_{\Lambda^c} \pi_{\beta, q}) \otimes K_\Lambda \quad \text{where } K_\Lambda(\eta, \cdot) := (\rho_\Lambda \pi_{X_\Lambda, \beta, q}^{\xi_\eta})(\cdot), \eta \in \Sigma_{\Lambda^c}.$$

That is, the operation of (i) taking the Σ_{Λ^c} -marginal of $\pi_{\beta, q}$, followed by (ii) taking this marginal's product with the probability kernel K_Λ defined as the Σ_Λ -marginal of the conditioned Potts measure (eq. (16)), gives back the original measure $\pi_{\beta, q}$.

(b) (Marginals)

For every nonempty finite $\Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d)$, there exists a probability measure λ^c on Σ_{Λ^c} such that

$$\rho_\Lambda \pi_{\beta, q} = \int_{\Sigma_{\Lambda^c}} \lambda^c(d\eta) \rho_\Lambda \pi_{X_\Lambda, \beta, q}^{\xi_\eta}(\cdot).$$

That is, the marginals are convex combinations of finite-volume higher Potts measures with boundary conditions. (Actually, these convex combinations are always finite, because there are only $q^{|\Lambda^c|}$ distinct sets ξ_η .) \triangle

³¹The Borel σ -algebra on Σ is the σ -algebra generated by the cylinders (the finitely-supported events).

Although we have no motive to do so, it's also possible to define Gibbs states $\mu_{p,q}$ for the Edwards–Sokal coupling on the joint configuration space

$$\left\{ (\sigma, \omega) \in \left((\mathbb{Z}/q\mathbb{Z})^{\mathcal{K}_{r-1}(\mathbb{R}^d)} \right) \wedge \times \{0, 1\}^{\mathcal{K}_r(\mathbb{R}^d)} \mid (\sigma, \omega) \in F \right\}$$

where $F = \{(\sigma, \omega) \in \Sigma_X \times \Omega_X \mid \forall Q \in \mathcal{K}_r(\mathbb{R}^d) : \omega(Q) = 1 \implies \sigma_Q = 1\}$. This configuration space is defined as a cartesian product restricted by a finite set of forbidden patterns, analogously to shifts of finite type. See [Rue04, §1.1] for the general theory of Gibbs states on such “restricted” configuration spaces.

The reason that one cannot define Gibbs states on the higher FK–Potts model is that (in finite volume) the change in probability by opening or closing a single r -cube can depend on individual r -cubes arbitrarily far away. Consider the case $r = 1$ and $q > 1$, where closing a single edge can increase the probability of a configuration by a factor of either $\frac{1-p}{p}q$ or $\frac{1-p}{p}$, depending on if doing so increases the number of components. But whether the number of components increases can depend on whether the two components are linked by a single edge arbitrarily far away. These long-distance effects were already hinted at in section 4 where we discussed the Hammersley–Clifford theorem: The spatial Markov property does not hold for FK–Potts, no matter how fat we make the wall (on which we’re conditioning) between two regions, so there is no Hamiltonian with finite-range interactions.³²

We could try to work around this difficulty by defining higher FK–Potts “hidden Gibbs states”: take a Gibbs state in the higher Potts model, open permissible r -cubes independently with probability p , and then forget the spins.³³ This would be in line with our goal of understanding the higher Potts Gibbs states. Unfortunately, it isn’t obvious what the connection is between (i) these hidden Gibbs states, (ii) the infinite-volume random fields given by the non-Gibbsian specification of our finite-volume higher FK–Potts model, and (iii) the thermodynamic limits of the higher FK–Potts model. The connection between the latter two is poorly understood even for $r = 1$ [Gri06, p. 79]. For these reasons, we won’t discuss higher FK–Potts in infinite volume. There is, however, some work done in this area in [DS23, §4.2, §5.2], which instead first defines the infinite-volume higher FK–Potts model and then uses it to define the infinite-volume higher Potts model (this requires

³²Actually, for a Gibbs specification it’s not necessary to have finite-range interactions, but it is necessary that the interactions satisfy a certain summability condition (which, in the FK–Potts model, they do not.) See [Rue04, §1.2] for details.

³³Formally, we’re pushing forward a Gibbs state through a probability kernel from Σ to Ω . To prove that it’s a legitimate kernel, use the result quoted in footnote 11.

uniformly picking a cocycle from an infinite collection of compatible cocycles).

5.2 Thermodynamic limits of the higher Potts model

Now, a few words on the thermodynamic limits of the higher Potts model. We'll follow precisely the definitions in [Rue04, ch. 1]. Let (Λ_n) be a sequence of finite subsets of $\mathcal{K}_{r-1}(\mathbb{R}^d)$ with $\Lambda_n \rightarrow \mathcal{K}_{r-1}(\mathbb{R}^d)$, in the sense that every $(r-1)$ -cube belongs to Λ_n for all but finitely many n . For every n let μ_n be a probability measure on Σ_{Λ_n} . For all finite $A \subseteq B \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d)$ let $\rho_{AB} : \Sigma_B \rightarrow \Sigma_A$ denote the projection onto the A marginal, and likewise $\rho_A : \Sigma \rightarrow \Sigma_A$. It can be shown by a diagonalization argument [Rue04, Proposition 1.4] that there exists a subsequence (Λ_{n_i}) such that the limit

$$\lim_{i \rightarrow \infty} \rho_{\Lambda_{n_i}} \mu_{n_i} = \psi_{\Lambda}$$

exists for every finite $\Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d)$ (here, as before, $\rho_{\Lambda_{n_i}} \mu_{n_i}$ is the Σ_{Λ} -marginal of μ_{n_i} , and convergence is in the usual sense of weak limits of probability measures on Σ_{Λ} .) For every such subsequence (Λ_{n_i}) , there exists a unique probability measure ψ on Σ such that

$$\psi_{\Lambda} = \rho_{\Lambda} \psi$$

for every finite $\Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d)$. Such a measure ψ is called a *thermodynamic limit* of the sequence (μ_n) . Of course, a sequence of measures can have many distinct thermodynamic limits (by passing to distinct subsequences.)

Definition 70.

- Take a sequence of cubical sets X_n and let $\Lambda_n = \mathcal{K}_{r-1}(X_n)$. Assume that
 - (i) each X_n is a finite union of r -cubes (as described at the beginning of section 3);
 - (ii) for each n , if Q is an r -cube whose every $((r-1)$ -dimensional) facet belongs to Λ_n , then $Q \subseteq X_n$ ³⁴ and
 - (iii) $\Lambda_n \rightarrow \mathcal{K}_{r-1}(\mathbb{R}^d)$ (therefore, $X_n \rightarrow \bigcup \mathcal{K}_r(\mathbb{R}^d)$) as $n \rightarrow \infty$.

A thermodynamic limit of the corresponding higher Potts measures with free boundary condition, $\pi_{X_n, \beta, q}$ on Σ_{Λ_n} (defined in eq. (5)), will be called a *higher Potts thermodynamic limit with free boundary condition*.

³⁴This condition ensures that the Hamiltonian takes into account every possible interaction between elements of Λ_n , and not merely the ones associated with those r -cubes that happened to be included in X_n .

- Take a sequence of finite sets $\Lambda_n \rightarrow \mathcal{K}_{r-1}(\mathbb{R}^d)$ and a sequence (η_n) where $\eta_n \in \Sigma_{\Lambda_n^c}$. Let X_{Λ_n} ³⁵ and ξ_{η} be as defined in notation 68. A thermodynamic limit of the higher Potts measures $\pi_{X_n, \beta, q}^{\xi_{\eta_n}}$ on Σ_{Λ_n} (eq. (16)) will be called a *higher Potts thermodynamic limit*. \triangle

Definition 70 is slightly awkward, and could probably be replaced by a equivalent definition that is simpler. It's presented in this form so as to be obviously compatible with [Rue04, ch. 1]. There is only one slight difference: Our sets Λ_n are not just any finite sets of $(r - 1)$ -cubes, but rather they always arise from the sets X_n which were declared to be unions of r -cubes. However, this doesn't cause any loss of generality. For if we defined a more general higher Potts measure $\pi_{\Lambda_n, \beta, q}$ on Σ_{Λ_n} according to the general treatment in [Rue04, ch. 1], then this measure would put uniform independent spin on any isolated $(r - 1)$ -cube (that is, one that is not part of an r -cube all of whose facets are in Λ_n). For that reason, once Λ_n is large enough, the marginals ρ_{Λ} would be no more general.

Importantly, it's a general result [Rue04, Theorem 1.9] that every thermodynamic limit with free boundary condition is a Gibbs state, and that the set of all Gibbs states is the closed convex hull (in the usual weak topology) of the set of all thermodynamic limits (i.e., for all possible sequences of boundary conditions (η_n) .)

What happens when we take an infinite-volume limit with a boundary condition that isn't a point BSC as in definition 70, but instead is a more general BSC (for example, wired and periodic boundary conditions in the Ising and Potts models?) Then still every thermodynamic limit is a Gibbs state. Such situations are discussed in great generality in [Geo11, ch. 4], which calls them "random boundary conditions".

We'll now prove that the thermodynamic limit with free boundary condition does not require passing to a subsequence (Λ_{n_i}) when taking the weak limit. In particular, the higher Potts thermodynamic limit with free boundary condition is unique. The proof uses the coupling to the higher FK-Potts model along with its strong FKG property, which is invoked through a proxy, corollary 55.

Proposition 71. *There exists a unique probability measure $\pi_{\beta, q}$ on Σ such that for every sequence (X_n) (and $\Lambda_n := \mathcal{K}_{r-1}(X_n)$) satisfying the three conditions given in the first paragraph of definition 70,*

$$\lim_{n \rightarrow \infty} \rho_{\Lambda_n} \pi_{X_n, \beta, q} = \rho_{\Lambda} \pi_{\beta, q} \quad \text{for each finite } \Lambda \subseteq \mathcal{K}_{r-1}(\mathbb{R}^d). \quad (19)$$

³⁵These X_{Λ_n} , on the other hand, will not necessarily satisfy the special condition discussed in the previous footnote, although enlarging them to do so would not modify the measures $\pi_{X_n, \beta, q}^{\xi_{\eta_n}}$.

Furthermore, this unique $\pi_{\beta,q}$ is invariant under all symmetries of \mathbb{Z}^d (viz., all translations, rotations, and reflections, and compositions thereof.)

Proof. Recall X_Λ , $\bar{\Lambda}$, and $\Sigma_{\bar{\Lambda}} = \Sigma_{X_\Lambda}$ from notation 68. Pick a sequence (X_n) satisfying the three conditions.

Let $\pi_{\beta,q}$ be some thermodynamic limit of $(\pi_{X_n,\beta,q})$ (i.e., after passing to some subsequence, eq. (19) holds.) To show that $\pi_{\beta,q}$ satisfies eq. (19) *without* passing to a subsequence, it suffices to prove the following statement. For each finite nonempty $\Lambda \subseteq \mathcal{X}_{r-1}(\mathbb{R}^d)$, every function $f : \Sigma_{\bar{\Lambda}} \rightarrow \mathbb{C}$ has converging expectations

$$\lim_{n \rightarrow \infty} (\rho_{\bar{\Lambda} \wedge \Lambda_n} \pi_{X_n,\beta,q})f = (\rho_{\bar{\Lambda}} \pi_{\beta,q})f.$$

The reason it suffices to use $\bar{\Lambda}$ is that any $g : \Sigma_\Lambda \rightarrow \mathbb{C}$ may be composed with the projection $\rho_{\Lambda \bar{\Lambda}}$ to give a function $f = g \circ \rho_{\Lambda \bar{\Lambda}} : \Sigma_{\bar{\Lambda}} \rightarrow \mathbb{C}$.

Fix some finite nonempty $\Lambda \subseteq \mathcal{X}_{r-1}(\mathbb{R}^d)$. The characters of Σ_{X_Λ} are the evaluation maps $W_\gamma : \Sigma_{X_\Lambda} \rightarrow \mathbb{C}$ where $\gamma \in C_{r-1}(X_\Lambda, \mathbb{Z}/q\mathbb{Z})$, as was observed on page 58. By Fourier decomposition, f may be written as a linear combination of these functions W_γ . Therefore, by linearity of expectation, it suffices to prove

$$\lim_{n \rightarrow \infty} (\rho_{\bar{\Lambda} \wedge \Lambda_n} \pi_{X_n,\beta,q})W_\gamma = (\rho_{\bar{\Lambda}} \pi_{\beta,q})W_\gamma, \quad \gamma \in C_{r-1}(X_\Lambda, \mathbb{Z}/q\mathbb{Z}).$$

We know already that there is some subsequence along which convergence holds for every γ . We'll show that passing to a subsequence is unnecessary, by proving that the expectation $(\rho_{\bar{\Lambda} \wedge \Lambda_n} \pi_{X_n,\beta,q})W_\gamma$ is monotonically increasing in X_n for every γ : that is, the condition $X_n \subseteq X_m$ (equivalently, $\Lambda_n \subseteq \Lambda_m$) implies $(\rho_{\bar{\Lambda} \wedge \Lambda_n} \pi_{X_n,\beta,q})W_\gamma \leq (\rho_{\bar{\Lambda} \wedge \Lambda_m} \pi_{X_m,\beta,q})W_\gamma$.³⁶

Take n large enough that $\bar{\Lambda} \subseteq \Lambda_n$. Then $X_\Lambda \subseteq X_n$ (because of the second assumption on (X_n)), so we can identify each $\gamma \in C_{r-1}(X_\Lambda, \mathbb{Z}/q\mathbb{Z})$ with $\kappa\gamma \in C_{r-1}(X_n, \mathbb{Z}/q\mathbb{Z})$ where κ is the injection (definition 9.) Thus, we may consider each W_γ to also be a function on a larger domain, $W_\gamma : \Sigma_{X_n} \rightarrow \mathbb{C}$, and (by the expectation-of-expectation "tower law") $(\rho_{\bar{\Lambda} \wedge \Lambda_n} \pi_{X_n,\beta,q})W_\gamma = \pi_{X_n,\beta,q}W_\gamma$.

³⁶A minor technical clarification: Although $\lim_n \Lambda_n = \mathcal{X}_{r-1}(\mathbb{R}^d)$ and thus $\lim_n X_n = \bigcup \mathcal{X}_r(\mathbb{R}^d)$, we never assumed $n \leq m \implies X_n \subseteq X_m$. Still, for every n there does exist M such that $M \leq m \implies X_n \subseteq X_m$, and that is enough for the existence of the limit $\lim_{n \rightarrow \infty} (\rho_{\bar{\Lambda} \wedge \Lambda_n} \pi_{X_n,\beta,q})W_\gamma$.

According to theorem 41, this expectation coincides with the probability

$$\varphi_{X_n, p, q}(\gamma \in B_{r-1}(X_n, \omega, \mathbb{Z}/q\mathbb{Z})).$$

These probabilities are indeed monotonically increasing in X_n by corollary 55.

This proves that $\pi_{\beta, q}$ satisfies eq. (19) for our chosen sequence (X_n) . As $\pi_{\beta, q}$ is uniquely determined by its marginals $\rho_{\Lambda} \pi_{\beta, q}$, it follows that there's exactly one thermodynamic limit associated with each sequence (X_n) . To see that $\pi_{\beta, q}$ does not depend on (X_n) , observe that for any *other* sequence of cubical sets Y_n we may interleave with X_n to get the sequence $(X_1, Y_1, X_2, Y_2, \dots)$, along which there is still only one limit.

To show symmetry-invariance, let $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be an isometry satisfying $T(\mathbb{Z}^d) = \mathbb{Z}^d$. Let T act on Σ , and more generally for all $\Lambda \subseteq \mathcal{X}_{r-1}$ let T map Σ_{Λ} into $\Sigma_{T\Lambda}$, as $T(\sigma)(Q) = \sigma(T^{-1}(Q))$ for $\sigma \in \Sigma_{\Lambda}$ and $Q \in \Lambda$. We must prove $T\pi_{\beta, q} = \pi_{\beta, q}$. From the symmetries in the definition of $\pi_{X_n, \beta, q}$ it follows that for sufficiently large n every function $f : \Sigma_{X_{\Lambda}} \rightarrow \mathbb{C}$ satisfies

$$(\rho_{(T\bar{\Lambda})(T\Lambda_n)} \pi_{TX_n, \beta, q})(f \circ T^{-1}) = (\rho_{\bar{\Lambda}\Lambda_n} \pi_{X_n, \beta, q})f.$$

Sending $n \rightarrow \infty$ gives

$$(\rho_{T\bar{\Lambda}} \pi_{\beta, q})(f \circ T^{-1}) = (\rho_{\bar{\Lambda}} \pi_{\beta, q})f$$

or, equivalently,

$$\pi_{\beta, q}(f \circ T^{-1} \circ \rho_{T\bar{\Lambda}}) = \pi_{\beta, q}(f \circ \rho_{\bar{\Lambda}}).$$

But $T^{-1} \circ \rho_{T\bar{\Lambda}} = \rho_{\bar{\Lambda}} \circ T^{-1}$, so this implies

$$T^{-1} \pi_{\beta, q}(f \circ \rho_{\bar{\Lambda}}) = \pi_{\beta, q}(f \circ \rho_{\bar{\Lambda}}).$$

Therefore, $T\pi_{\beta, q} = \pi_{\beta, q}$, because every function depending on only finitely many spins may be expressed as $f \circ \rho_{\bar{\Lambda}}$ for some finite Λ and some f . \square

In particular, proposition 71 proves the existence and translation-invariance of infinite-volume limits of Wilson loop expectations in the higher Potts model (with free boundary condition).

Corollary 72 is connected to two recent results:

- [Cha20, Theorem 5.4], which shows existence and translation-invariance of free-boundary

infinite-volume limits of local observables in the Ising lattice gauge theory ($r = 2$ and $d = 4$) at weak coupling (that is, sufficiently large β), and whose proof involves a specific estimate on the decay of correlations.

- [FLV21, Theorem 4.1], which shows existence and translation invariance of free-boundary infinite-volume limits of local observables in the clock (planar Potts) lattice gauge theory ($r = 2$ and $d = 4$, and general q) for all $\beta \geq 0$, and whose proof uses an estimate known as Ginibre's inequality but otherwise is very similar to the proof here. It's also mentioned [FLV21, p. 3] that the proof can be extended to a general finite abelian group and a general unitary faithful irreducible representation.

Similar results on the existence of infinite-volume limits for free boundary condition have been known since at least 1982 (see the discussion and further references in [Cao20, p. 1441].) Corollary 72 extends these results to arbitrary cell dimension r in the special case of the higher Potts lattice gauge theory.

Corollary 72. *Take (X_n) as in proposition 71, and let $\gamma \in C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})$ for any cubical set $X \subseteq \mathbb{R}^d$. Identify γ with the corresponding element of $C_{r-1}(X_n, \mathbb{Z}/q\mathbb{Z})$ for all $X_n \supseteq X$ (by assigning 0 to all new r -cubes as before.) Writing $\langle W_\gamma \rangle_{X_n, \beta, q} := \pi_{X_n, \beta, q} W_\gamma$ as in section 3.4, the limit of expectations*

$$\lim_{n \rightarrow \infty} \langle W_\gamma \rangle_{X_n, \beta, q}.$$

exists, is real, is independent of the particular choice of the sequence (X_n) , and is invariant under lattice isometries (in the sense that the limit is unchanged when γ is translated, rotated, or reflected.) Moreover, if $0 \leq \beta_1 \leq \beta_2 \leq \infty$, then

$$0 \leq \lim_{n \rightarrow \infty} \langle W_\gamma \rangle_{X_n, \beta_1, q} \leq \lim_{n \rightarrow \infty} \langle W_\gamma \rangle_{X_n, \beta_2, q} \leq 1.$$

Proof. Immediate from proposition 71 and corollary 42. □

Next, a correlation inequality. For the Ising model in finite volume, it overlaps with two of Griffith's inequalities [Gri67, Theorems 2, 3] and is also referred to as one of the GKS inequalities [Geo11, p. 456]. Note that $W_{\gamma_1 + \gamma_2} = W_{\gamma_1} W_{\gamma_2}$ (pointwise product on the right-hand side.)

Corollary 73. *Take (X_n) as in proposition 71, and let $\gamma_1, \gamma_2 \in C_{r-1}(X, \mathbb{Z}/q\mathbb{Z})$ for any cubical set*

$X \subseteq \mathbb{R}^d$. Then, for every $p \in (0, 1)$,

$$\lim_{n \rightarrow \infty} \langle W_{\gamma_1 + \gamma_2} \rangle_{X_n, \beta, q} \geq \left(\lim_{n \rightarrow \infty} \langle W_{\gamma_1} \rangle_{X_n, \beta, q} \right) \left(\lim_{n \rightarrow \infty} \langle W_{\gamma_2} \rangle_{X_n, \beta, q} \right).$$

Proof. Take n large enough that $X_n \supseteq X$. For every $\omega \in \Omega_{X_n}$, if $\gamma_1, \gamma_2 \in B_{r-1}(X_{n, \omega}, \mathbb{Z}/q\mathbb{Z})$ then $\gamma_1 + \gamma_2 \in B_{r-1}(X_{n, \omega}, \mathbb{Z}/q\mathbb{Z})$ because the boundaries form a group. The two events $\{\omega : \gamma_1 \in B_{r-1}(X_{n, \omega}, \mathbb{Z}/q\mathbb{Z})\}$ and $\{\omega : \gamma_2 \in B_{r-1}(X_{n, \omega}, \mathbb{Z}/q\mathbb{Z})\}$ are increasing. Thus, by positive association (see theorem 35, which entails inequality (13)),

$$\begin{aligned} \varphi_{X_n, p, q}(\gamma_1 \in B_{r-1}(X_{n, \omega}, \mathbb{Z}/q\mathbb{Z})) \\ \cdot \varphi_{X_n, p, q}(\gamma_2 \in B_{r-1}(X_{n, \omega}, \mathbb{Z}/q\mathbb{Z})) &\leq \varphi_{X_n, p, q}(\gamma_1, \gamma_2 \in B_{r-1}(X_{n, \omega}, \mathbb{Z}/q\mathbb{Z})) \\ &\leq \varphi_{X_n, p, q}(\gamma_1 + \gamma_2 \in B_{r-1}(X_{n, \omega}, \mathbb{Z}/q\mathbb{Z})). \end{aligned}$$

By theorem 41 it follows that

$$\langle W_{\gamma_1 + \gamma_2} \rangle_{X_n, \beta, q} - \langle W_{\gamma_1} \rangle_{X_n, \beta, q} \langle W_{\gamma_2} \rangle_{X_n, \beta, q} \geq 0.$$

Now take the limit as $n \rightarrow \infty$. □

Incidentally, there exists another, very general way to prove invariance under symmetries. It's a general property of Gibbs states that for *every* Gibbsian specification that is invariant under a symmetry group of the lattice, every free-boundary thermodynamic limit is invariant under the same symmetry group. And analogous results hold for non-free boundary conditions, provided the boundary conditions themselves have symmetries. See [Geo11, pp. 91–92, Examples (5.20)(1)–(2)].³⁷ The same reference also shows invariance under symmetries of the spin group: in our case of the higher Potts model, the free-boundary thermodynamic limit $\pi_{\beta, q}$ is invariant under the operation of spin reversal (negating the coefficient of each $(r-1)$ -cube simultaneously.)

For general (non-free) boundary conditions, thermodynamic limits of the Potts model are not necessarily translation-invariant. The best-known example is the non-translation-invariant *Dobrushin states* in the Ising model in dimension $d \geq 3$, which are obtained by taking the infinite-volume limit of boxes with spin $+1$ on boundary vertices whose first coordinate is nonnegative, and spin -1 on boundary vertices whose first coordinate is negative [FV17, §3.10.7; Bov06, Remark

³⁷The cited examples apply directly only to $r = 1$ but the theorems they reference apply to all r .

4.3.19].

It's well-known, however, that for strong coupling (that is, when β is sufficiently close to 0) there exists a unique Gibbs state and hence only one thermodynamic limit—which is therefore invariant under the symmetries of \mathbb{Z}^d by the arguments above. This is known as Dobrushin's uniqueness criterion [Geo11, ch. 8]. From [Geo11, Proposition 8.8] we directly calculate that the higher Potts model has precisely one Gibbs state whenever

$$0 \leq \beta < \frac{1}{(d-r+1)(2^r-1)}.$$

For weak coupling (large β), the situation is more complicated. In the Potts model ($r = 1$ and $d \geq 2$) it is known that there exists β_c such that when $0 \leq \beta < \beta_c$ there is a unique Gibbs state and when $\beta_c < \beta$ there exist q mutually singular Gibbs states [GHM01, Theorem 3.2].

6 Odds and ends

6.1 Examples of pathological surfaces

This informal section is meant to illustrate the difference between boundaries in the sense of homotopy and boundaries in the sense of homology, to help clarify the concepts of section 2.2 for readers unfamiliar with homology. The examples here are an elaboration on [AF84, §4].

As the first example, take a large flat rectangular sheet of plaquettes embedded in \mathbb{Z}^4 , remove two plaquettes from it, and join the resulting holes by a tube, to get a plaquette surface homeomorphic to that pictured in fig. 5. Such an orientation-flipping tube is called a *cross-handle* in the theory of classification of surfaces [Wee20, ch. 5]. (The middle of the cross-handle does not actually intersect the rectangular sheet: move it out into the fourth dimension to avoid this.) The drawing shows a smooth cross-handle, but of course it will actually be a rectangular tube with sharp corners. This surface is homeomorphic to a Klein bottle with a point removed. It is a non-orientable 2-manifold with boundary, whose boundary is the loop γ bounding the large rectangle.




Figure 5: Plaquette surface with a cross-handle, overhead schematic (left) and oblique view (right).

The outer loop γ is not contractible within the surface. To see why, embed the surface in \mathbb{R}^3 by letting the middle of the cross-handle pass outside the rectangular sheet. Run a wire through the cross-handle and extend the wire's ends upward and downward (perpendicular to the rectangular sheet) to infinity. The wire together with γ form a nontrivial link, and the wire does not intersect the surface, so there is no way to shrink γ to a point within the surface.

However, now consider a cycle over $\mathbb{Z}/2\mathbb{Z}$ supported on the edges of γ , assigning coefficient $1 \in \mathbb{Z}/2\mathbb{Z}$ to each edge in γ . The chain thus defined is a boundary in the homological sense: It is the boundary of the 2-chain that assigns 1 to each plaquette in the surface.

On the other hand, if the coefficient group is either $G = \mathbb{Z}$ or $G = \mathbb{Z}/q\mathbb{Z}$ with $q \geq 3$, then the cycle that assigns coefficient $1 \in G$ (or -1 as orientation demands, because as explained in section 2.2 edges are not oriented) to each edge in γ is no longer a homological boundary. To see why, we argue by contradiction. Suppose that this cycle is the boundary of some 2-chain c . The coefficient in c of each plaquette in the surface is already uniquely determined, because each edge

in γ determines the coefficient of its incident plaquette, and, since the cycle is supported on γ , each plaquette has equal (or opposite, again as orientation demands) coefficient to all its neighboring plaquettes. But consider the two plaquettes that were removed from the sheet when adding the cross-handle. Their boundary edges all have coefficient equal to 1, so the cross-handle forces the identity $1 = -1$. This can hold only for coefficient group $\mathbb{Z}/2\mathbb{Z}$. However, for $G = \mathbb{Z}/q\mathbb{Z}$ and q even, putting a coefficient of $q/2$ on each edge in γ does give a boundary.

As a side note, for coefficients in \mathbb{Z} it's still not the case that every loop that is the support of a boundary is contractible. One counterexample: In the double torus (i.e., the connected sum of two tori: the genus-2 surface pictured in [Wee20, p. 253]), assign coefficient 1 to the plaquettes of one torus and 0 to those of the other. The support of the boundary of this 2-chain is a loop running around the waist of the double torus, which can be shown to not be contractible.

Again, take a rectangular sheet of plaquettes in \mathbb{Z}^4 , and now remove *three* plaquettes, joining the first and second hole with a cross-handle and joining the second and third hole with another cross-handle, as pictured in fig. 6. Such a “surface” is not a 2-manifold, because the four edges bounding the center hole are each incident to three plaquettes (one plaquette in the sheet and one in each cross-handle.)




Figure 6: Plaquette surface with double cross-handle, overhead schematic (left) and oblique view (right). The centres of the two bights do not intersect the sheet, but instead pass beside it in the fourth dimension.

Again, the outer loop γ is not contractible within the surface, by an argument analogous to before.

But take a cycle of $\mathbb{Z}/q\mathbb{Z}$ supported on the edges of γ , assigning coefficient $j \in \mathbb{Z}/q\mathbb{Z}$ to each edge. In order for this cycle to be a boundary, it's necessary and sufficient for the plaquettes in the flat rectangular sheet to have coefficient j ; those in the first cross-handle, j (to satisfy the zero-boundary constraint around the left hole); those in the second cross-handle, $2j$ (due to middle hole); and (again) those in the sheet, $-2j$ (due to right hole). So $j = -2j$, or $3j = 0$. Thus, there exists a nontrivial boundary supported on γ if and only if $q \equiv 3 \pmod{3}$.

Now glue together the two previous examples along γ , to obtain a set of plaquettes pictured

schematically in fig. 7.

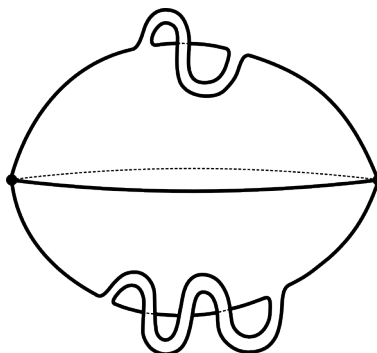


Figure 7: The previous two examples glued together along the outer loop.

Take coefficient group $\mathbb{Z}/6\mathbb{Z}$. Again we will consider boundaries supported on γ . The contribution to γ due to the top sheet (with the single cross-handle) is either 0 or 3; the contribution due to the bottom sheet is 0, 2, or 4. So every cycle supported on γ is the boundary of exactly one 2-chain. For other coefficient groups $\mathbb{Z}/q\mathbb{Z}$, γ supports a nontrivial boundary if and only if q is divisible by either 2 or 3.

The upshot is that we shouldn't expect it to be geometrically obvious whether a particular loop is a (homological) boundary in a given cubical set. But we can always answer this question computationally using the methods described in section 6.2.

Such pathological situations do not occur in the classical case, $r = 1$ (the spins-on-vertices Potts model.) If a pair of distinct points $\{x, y\}$ is a homological boundary, i.e., $1_x - 1_y \in B_0(X_\omega, \mathbb{Z}/q\mathbb{Z})$, then x and y belong to the same component—that is, the associated 0-sphere is contractible in the graph induced by the edges open in ω . To see why, take a chain $c \in C_1(X_\omega, \mathbb{Z}/q\mathbb{Z})$ with $\partial_1 c = 1_x - 1_y$. The support of c is a subset of the set of edges open in ω ; these edges induce a subgraph G_c of X_ω . We may assume that G_c is acyclic because if d is a 1-cycle then $\partial_1(c - d) = \partial_1 c = 1_x - 1_y$. The support of $\partial_1 c$ contains all leaf vertices (meaning: those incident to exactly one edge) of G_c . Thus, G_c is an acyclic graph whose only leaves are x and y . It follows that the edges of G_c form a path connecting x and y .

Therefore, there's no essential distinction between 0-boundaries in the sense of homology and in the sense of homotopy. This, combined with theorem 41, might be one reason that the decay of correlations is easier to analyze in the classical Potts model than in the Potts gauge (and higher Potts gauge) theories.

6.2 Algorithms

The models described above are amenable to computation via matrix arithmetic.

Counting cocycles

To compute the probability mass function in eq. (6) we must find $|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| = |\ker \delta^{r-1}|$. The chain groups are modules over the principal ideal domain $\mathbb{Z}/q\mathbb{Z}$, and δ^{r-1} is a module homomorphism between chain groups. So δ^{r-1} can be described by a matrix over a principal ideal domain, $M \in M_{m \times n}(\mathbb{Z}/q\mathbb{Z})$. Perform Gaussian elimination on M to obtain a Smith normal form, $M' = \begin{bmatrix} D & 0 \\ 0 & 0 \end{bmatrix}$, where D is the diagonal matrix $\text{diag}(d_1, d_2, \dots, d_r)$ for some $d_1, \dots, d_r \in \mathbb{Z}/q\mathbb{Z}$ with $d_1 \mid d_2 \mid \dots \mid d_r$ and $r \geq 0$. Writing $d_j = [c_j]$ with representatives $c_j \in \mathbb{Z}$, the number of cocycles is

$$|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})| = |\ker M'| = q^{n-r} |\ker \text{diag}(d_1, d_2, \dots, d_r)| = q^{n-r} \prod_{1 \leq j \leq r} \gcd(c_j, q).$$

If we have a more general boundary condition ξ (as per eq. (17)), where ξ is a subgroup of $\Sigma_X = C^{r-1}(X, \mathbb{Z}/q\mathbb{Z})$, then we are tasked with finding the size of the intersection of two $\mathbb{Z}/q\mathbb{Z}$ -submodules of Σ_X . This procedure is hardly any more difficult: First, express ξ as the kernel of some $\mathbb{Z}/q\mathbb{Z}$ -module homomorphism $\alpha : \Sigma_X \rightarrow A$ (as can always be done: the projection onto the group quotient Σ/ξ is such a homomorphism.) Then, compute the size of the kernel of the $\mathbb{Z}/q\mathbb{Z}$ -module homomorphism $(\delta^{r-1}, \alpha) : \Sigma_X \rightarrow C^r(X, \mathbb{Z}/q\mathbb{Z}) \oplus A$ via the method detailed above.

For information about Gaussian elimination and the Smith normal form of a matrix over a principal ideal domain, see [Gor16, §14.2].

Conditional sampling

The first conditional, $\mu_{X,p,q}(\sigma \mid \omega)$ (proposition 40), is uniform and therefore can be computed immediately once we find the size of its support, $|Z^{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})|$, using the method just explained in section 6.2. And again, the variant with boundary conditions (proposition 49) presents no additional challenges when ξ is a subgroup of Σ_X .

To compute the second conditional, $\mu_{X,p,q}(\omega \mid \sigma)$, no special techniques are needed.

To sample σ conditional on ω , perhaps the easiest way is to first reduce to Smith normal form as described in section 6.2, then sample each component independently and uniformly, and then

transform back to the original basis. This works equally well when there are boundary conditions (again, ξ should be a subgroup of Σ_X .) Contrast this to the the classical random-cluster model (spins on vertices, $r = 1$), where each connected component in X_ω is independently and uniformly assigned a spin.

Sampling ω conditional on σ is even easier: open each allowable r -cube Q (i.e., each for which $\sigma_Q = 1$) independently with probability p , and leave all other r -cubes closed.

Coupling from the past

The strong FKG property proved in section 3.2 for the higher FK–Potts measure $\varphi_{X,p,q}$ —in particular, 1-monotonicity—is precisely what’s needed for monotonicity [Gri06, Inequality (8.9)] of the Gibbs sampler (i.e., Glauber dynamics, or the single-site heat bath algorithm.) Monotonicity allows us to use the “coupling from the past” technique for perfect sampling [Thö00; Gri06, §8.4]. Without monotonicity, coupling from the past would be computationally infeasible because it’s necessary to simulate a separate Markov chain for each possible initial configuration; with monotonicity, two chains are enough (one starting from the all- r -cubes-open configuration, another starting from the all-closed configuration.)

This method lets us estimate Wilson loop expectations for small finite-volume lattices. If the expectation of a Wilson loop W_γ is w , then $\varphi_{X,p,q}(\gamma \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z})) = w$ (theorem 41.) So the variance of an estimate \hat{w} for w , computed as the mean of an independent sample of size n , is $\text{Var}(\hat{w}) = \text{Var}(\text{Bin}(n, w)) = \frac{w(1-w)}{n} \leq \frac{1}{4n}$ (where $\text{Bin}(n, w)$ is the binomial distribution on n elements.) A single sampling unit ω can be obtained by simulating a pair of Markov chains on Ω_X until they coalesce, and then the predicate $\omega \in B_{r-1}(X_\omega, \mathbb{Z}/q\mathbb{Z}) = \text{im } \partial_r$ can be computed via Gaussian elimination, analogously to section 6.2.

Simulations of spin models have long been studied in the physics literature. See, for instance, [ES88; KG95; KO12].

6.3 Ground states of the random-cluster model

In statistical physics, the *ground states* of a model (that is given as a discrete probability measure on a configuration space) are the configurations with maximal probability mass.

Determining the ground states is often a useful early step toward understanding a model’s behaviour. As a brief diversion, we’ll demonstrate by reviewing the simplest interesting example, the iid Bernoulli model. Take configuration space $\Omega = \{0, 1\}^E$ for some finite set E , and let

$\mathbf{P}_p(\omega) = p^{o(\omega)}(1-p)^{c(\omega)}$ where $o(\omega) = |\{e \in E \mid \omega(e) = 1\}|$ and $c(\omega) = |\{e \in E \mid \omega(e) = 0\}|$ for $\omega \in \Omega$. The ground states are

$$\begin{array}{ll} \omega^0 \text{ (the all-0 configuration)} & \text{for } 0 \leq p < \frac{1}{2}, \\ \text{all configurations} & \text{for } p = \frac{1}{2}, \\ \omega^1 \text{ (the all-1 configuration)} & \text{for } \frac{1}{2} < p \leq 1. \end{array}$$

On the other hand, if $0 < p < \frac{1}{2}$ and E is large then, despite \mathbf{P}_p being maximized at ω^0 , in absolute terms $\mathbf{P}_p(\omega^0)$ can be quite small. Consider the density observable $d_p(\omega) = \frac{o(\omega)}{|\Omega|}$, whose distribution is normalized binomial, $\frac{1}{|\Omega|} \text{Bin}(2^{|E|}, p)$. If d_p is the only available observable then the macrostates are the sets $M_k = \{\omega \in \Omega \mid d_p(\omega) = \frac{k}{|\Omega|}\}$ for $k = 0, 1, \dots, |E|$. Thus, the ground state, considered as the macrostate $M_0 = \{\omega^0\}$ among all macrostates, is actually very unlikely to occur when $|E|$ is large. Instead, there is another highest-probability state M_K where $K := \text{argmax}_k \mathbf{P}_p(M_k)$. We can consider the members of M_K to be the *typical configurations*.

Borrowing language from statistical physics, define the *energy* $U_p(\omega) := -o(\omega) \log p - c(\omega) \log(1-p)$ and *entropy* $S(\omega) := -\frac{o(\omega)}{|\Omega|} \log \frac{o(\omega)}{|\Omega|} - \frac{c(\omega)}{|\Omega|} \log \frac{c(\omega)}{|\Omega|}$. The ground state M_0 has minimum energy, the state(s) $M_{\lfloor |E|/2 \rfloor}, M_{\lceil |E|/2 \rceil}$ have maximum entropy, and the equilibrium state M_K exhibits a kind of energy–entropy trade-off (in fact, $K/|E| \rightarrow p$ as $|E| \rightarrow \infty$ ³⁸; note however that the variational principle concerning the free energy $F_p = U_p - S$ [Rue04, p. 4] doesn't quite apply because in that context the equilibrium macrostate is the distribution \mathbf{P}_p itself.)

We go on to investigate the ground states of the random-cluster model (not the general higher FK–Potts model, but only the case $r = 1$.) Although the results here use well-known tools, they have never before (to my knowledge) been published.

Recall from section 1 the random-cluster model on a finite graph $G = (V, E)$,

$$\varphi_{G,p,q}(\omega) \propto (1-p)^{c(\omega)} p^{o(\omega)} q^{k(\omega)}, \quad p \in (0, 1), \quad q \in (0, \infty), \quad \omega \in \Omega := \{0, 1\}^E$$

where $o(\cdot)$ and $c(\cdot)$ are the number of open and closed edges, respectively, and $k(\cdot)$ is the number of open clusters (that is, connected components in the subgraph induced by the open edges, including isolated vertices.)

³⁸To prove this, let $f(k) = \binom{n}{k} p^k (1-p)^{n-k}$ for $0 \leq k \leq n$ and $p \in (0, 1)$, and calculate $\frac{f(k+1)}{f(k)} > 1 \iff \frac{k+1}{n+1} < p$ for $0 \leq k < n$.

The next result is a partial characterization of the ground states of $\varphi_{G,p,q}$ in the case of planar G . It also assumes another property of G : the face density of a configuration never exceeds its edge density. More explicitly, let F and f be the number of faces in the graphs G and $G_\omega := (V, \{x \in E \mid \omega(x) = 1\})$, respectively, not counting the unbounded exterior face, and let E and e be the number of edges in the graphs G and G_ω , respectively. The property is

$$\frac{f}{F} \leq \frac{e}{E} \quad \text{for every } \omega \in \Omega. \quad (\dagger)$$

Note that for ω^0 and ω^1 the inequality in (\dagger) reduces to $0 = 0$ and $1 = 1$, respectively.

We restrict ourselves to $q > 1$. If $q = 1$ then $\varphi_{G,p,q}$ reduces to iid Bernoulli, whose ground states we've already examined. The method of proof in proposition 74 does in fact also handle all $0 < q \leq 1$, and the case $0 < q < 1$ is interesting from the perspective of graph theory: for certain values of p the ground states may be the spanning trees, the forests, or the connected subgraphs. Incidentally, this also provides a geometric explanation for some of the weak limits as $q \rightarrow 0$ described in [Gri06, §1.5]. But further discussion for $0 < q \leq 1$ is omitted due to time constraints.

Proposition 74.

Let $G = (V, E)$ be a finite connected planar graph with at least one edge. Let $p \in (0, 1)$ and $q \in (1, \infty)$.

(i) If $p \leq \frac{1}{2}$ then the all-closed configuration ω^0 is the unique ground state of $\varphi_{G,p,q}$.

(ii) If $p > \frac{1}{2}$ and (\dagger) holds, then:

(a) If $p < \left(1 + \exp\left[-\frac{V-1}{E} \log q\right]\right)^{-1}$ then the all-closed configuration ω^0 is the unique ground state.

(b) If $p > \left(1 + \exp\left[-\frac{V-1}{E} \log q\right]\right)^{-1}$ then the all-open configuration ω^1 is the unique ground state.

(c) If $p = \left(1 + \exp\left[-\frac{V-1}{E} \log q\right]\right)^{-1}$ then ω^0 and ω^1 are both ground states. If moreover (\dagger) holds with strict inequality for all $\omega \in \Omega \setminus \{\omega^0, \omega^1\}$, then there are no other ground states.

Proof. Maximizing $\varphi_{G,p,q}$ amounts to maximizing its logarithm, which (since $c(\omega) + o(\omega) = |E|$) satisfies

$$\log \varphi_{G,p,q}(\omega) \propto o(\omega) \log \frac{p}{1-p} + k(\omega) \log q = (o(\omega), k(\omega)) \cdot \left(\log \frac{p}{1-p}, \log q \right)$$

where \cdot is the Euclidean inner product on \mathbb{R}^2 . Therefore, the ground states are the configurations

$\omega \in \Omega$ for which the point $\mathbf{x}_\omega := (o(\omega), k(\omega)) \in \mathbb{R}^2$ has greatest orthogonal projection onto the vector $\mathbf{v}_{p,q} := \left(\log \frac{p}{1-p}, \log q \right) \in \mathbb{R}^2$.

We will examine the geometry of the set $S := \{\mathbf{x}_\omega \mid \omega \in \Omega\}$, which depends only on the graph G . This set S is pictured for several graphs in fig. 9 and fig. 10. Most points in S correspond to many different configurations; the number of configurations associated with each point is indicated by a numeral beside it.

Evidently, $\mathbf{x}_{\omega^0} = (0, V)$ and $\mathbf{x}_{\omega^1} = (E, 1)$, and these are the only configurations associated with these two points. Every other configuration ω must satisfy $1 \leq o(\omega) \leq E-1$ and $1 \leq k(\omega) \leq V-1$, and also $o(\omega) + k(\omega) \geq V$ (by Euler's formula, eq. (21).) If ω is a spanning tree then $\mathbf{x}_\omega = (V-1, 1)$; removing edges one by one from this spanning tree (in any order) yields points $(V-k, k)$ for all $1 \leq k \leq V$. This explains the lower horizontal boundary and the lower-left diagonal boundary in the displayed figures.

By assumption, $q > 1$, so $\log q > 0$.

For (i), the condition $p \leq \frac{1}{2}$ implies $\log \frac{p}{1-p} \leq 0$, so the vector $\mathbf{v}_{p,q}$ lies in the third quadrant. In this case, it's clear geometrically that ω^0 is the unique ground state.

For (ii), the condition $p > \frac{1}{2}$ implies $\log \frac{p}{1-p} > 0$, so the vector $\mathbf{v}_{p,q}$ lies in the interior of the first quadrant.

Let ℓ be the line passing through \mathbf{x}_{ω^0} and \mathbf{x}_{ω^1} . We will show that if (†) holds then all points \mathbf{x}_ω lie below ℓ , and if (†) holds strictly as specified in (c) then all points \mathbf{x}_ω for $\omega \notin \{\omega^0, \omega^1\}$ lies strictly below ℓ .

Recall Euler's formula (the early form of fact 24),

$$V - E + F = 1 \tag{20}$$

where (overloading the symbols as usual) V is the number of vertices, E is the number of edges, and F is the number of faces in G , not counting the unbounded exterior face. For the induced subgraph $G_\omega := (V, \{x \in E \mid \omega(x) = 1\})$, to keep the notation compact we'll write e ($= o(\omega)$) for the number of edges, f for the number of faces, and also k ($= k(\omega)$) for the number of components. Euler's formula for G_ω is

$$V - e + f = k. \tag{21}$$

The line ℓ is given in slope–intercept form by the equation $y = \frac{1-V}{E}x + V$. Thus, for every $\omega \in \Omega$,

$$\begin{aligned}
x_\omega \text{ lies (weakly) below } \ell &\iff k \leq \frac{1-V}{E}e + V \\
&\iff f - e \leq \frac{1-V}{E}e && \text{(by eq. (21))} \\
&\iff f \leq \frac{1-V+E}{E}e \\
&\iff f \leq \frac{F}{E}e && \text{(by eq. (20))} \\
&\iff \frac{f}{F} \leq \frac{e}{E}. && (22)
\end{aligned}$$

Likewise, x_ω lies strictly below ℓ if and only if $\frac{f}{F} < \frac{e}{E}$.

Assume (+), so that by the arguments above S lies within the convex hull of $\{x_{\omega^0}, x_{\omega^1}, x_{\omega^s}\}$ where ω^s is a spanning tree (to be explicit: we’ve shown that every $x_\omega \in S$ lies on the correct side of each face of this simplex.) The slope of the normal to ℓ (in the outward, top-right, direction) is $\frac{E}{V-1}$ (this is valid because we’ve assumed $E \geq 1$ and thus $V > 1$.) Recalling that $\mathbf{v}_{p,q}$ lies in the interior of the first quadrant, we see geometrically that if the slope of $\mathbf{v}_{p,q}$ is strictly greater than $\frac{E}{V-1}$ then ω^0 is the unique ground state; if the slope is strictly less than $\frac{E}{V-1}$ then ω^1 is the unique ground state, and if the slope is equal to $\frac{E}{V-1}$ then ω^0 and ω^1 are ground states. Moreover, in the case of equality, under the strict inequality condition given in (c) there are no other ground states.

The slope of $\mathbf{v}_{p,q}$ is $\frac{\log q}{\log \frac{p}{1-p}}$. Rearranging the comparisons,

$$\begin{aligned}
\frac{\log q}{\log \frac{p}{1-p}} \leq \frac{E}{V-1} &\iff \frac{V-1}{E} \log q \leq \log \frac{p}{1-p} && \text{(since } p > \frac{1}{2}\text{)} \\
&\iff \frac{p}{1-p} \geq \exp \left[\frac{V-1}{E} \log q \right] \\
&\iff \frac{1-p}{p} \leq \exp \left[-\frac{V-1}{E} \log q \right] \\
&\iff \frac{1}{p} \leq 1 + \exp \left[-\frac{V-1}{E} \log q \right] \\
&\iff p \geq \left(1 + \exp \left[-\frac{V-1}{E} \log q \right] \right)^{-1}. && \square
\end{aligned}$$

It bears mentioning that this method of proof has other applications, such as identifying the lowest-probability configurations, or efficiently generating a graph that displays the entropy of $\varphi_{G,p,q}$ as a function of (p, q) (pre-compute the set S together with the labels as in figs. 9 and 10 and

project the data orthogonally onto $\text{span}\{\mathbf{v}_{p,q}\}$ for each (p, q) ; time may be saved by parametrizing $\mathbf{v}_{p,q}$ in polar coordinates.)

The code used to generate figs. 9 and 10 (for the proof of proposition 74) is listed in section 6.3. Generating fig. 10 for S_4 took 9 minutes on an AMD Ryzen 3 3200G desktop CPU, with 12 GB memory allocated. A more conscientious implementation³⁹ would finish in seconds and allocate less than one kilobyte. But because of combinatorial explosion it would not be useful to optimize the code. The graph S_n has $2^E = 2^{2n(n-1)}$ configurations, so if we could compute $k(\omega)$ using a mere 226 clock cycles per configuration then (on a 3.8 GHz processor, single-threaded) analyzing S_4 would take one second, S_5 would take 18 hours, and S_6 would take two millennia.

The complete graph K_n is not planar for $n \geq 5$, so fig. 9 displays S only for K_3 and K_4 . Observe from these images that for some graphs all points of S lie within the simplex, and for other graphs this is not the case. Specifically, by the equivalence (22), the condition (+) holds for K_3 and K_4 but not for the graphs $K_3 + K_3$ and $K_4 + K_4$ obtained by joining two copies of the complete graph with a single edge.

We will show, however, that in the case of the finite planar square lattice, (+) always holds (proposition 75).




Figure 8: The square lattice S_4

³⁹Iterate over the configurations in a Gray code ordering, storing only the final tally.


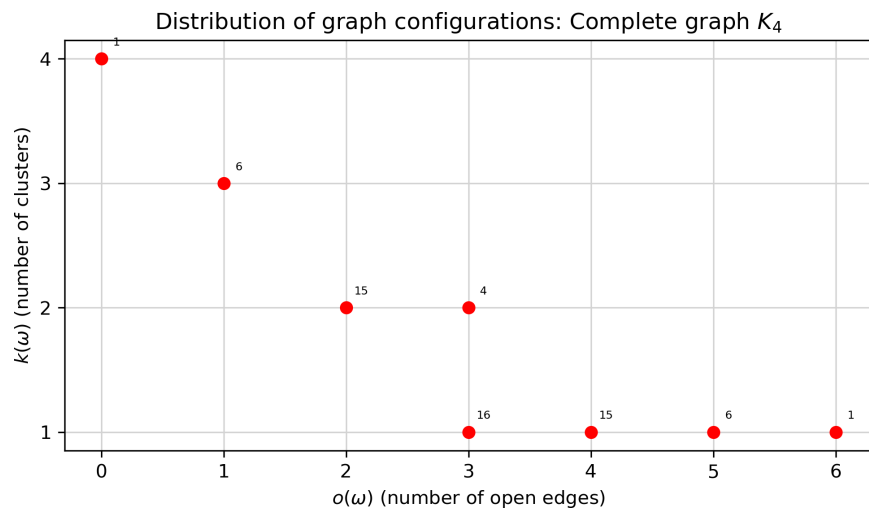




Figure 9: The set S for various graphs. The “Siamese graph” $K_n + K_n$ consists of two copies of K_n joined by a single edge.




Figure 10: The set S for the 4×4 square lattice S_4 , pictured in fig. 8.

Proposition 75. Let $n \geq 2$ and let G be the finite planar square lattice S_n (i.e., on $n \times n$ vertices.) The condition (\dagger) holds with strict inequality for every $\omega \in \Omega \setminus \{\omega^0, \omega^1\}$.

Proof. Fix $n \geq 2$. The vertex, edge, and face counts for S_n are respectively

$$V = n^2, \quad E = 2n(n - 1), \quad F = (n - 1)^2.$$

Substituting these values into (\dagger) tells us that we must prove

$$f < \frac{n-1}{2n}e \tag{23}$$

for every $\omega \in \Omega \setminus \{\omega^0, \omega^1\}$. Our assumption $n \geq 2$ implies $\frac{1}{4} \leq \frac{n-1}{2n} < \frac{1}{2}$. Thus, if for a configuration $\omega \in \Omega$ the graph G_ω satisfies $f \leq \frac{n-1}{2n}e$, and if ω' is obtained from ω by the removal of one face and at most two edges, then $G_{\omega'}$ satisfies (23). Thus, we may assume (by induction, filling in faces one plaquette at a time) that each face of G_ω consists of a single plaquette. Moreover, the removal of an edge that isn't incident to a face can only make the inequality (23) tighter, so we may assume that every edge is incident to a face. In summary, we may assume without loss of generality that the set of open edges is a union of boundaries of plaquettes, as for example pictured in fig. 11. (Note a subtlety in the above argument: the all-open ω^1 and all-closed ω^0 configurations have equality $f = \frac{n-1}{2n}e$.) Take an arbitrary such plaquette boundary union




Figure 11: A union-of-plaquettes configuration in S_6 . The displayed edges are open.




Figure 12: Left: All southeast edges in S_5 . Right: Some plaquettes in S_5 together with their southeast edges.

$\omega \notin \{\omega^0, \omega^1\}$. Forget (23): we will prove directly the original form, $\frac{f}{F} < \frac{e}{E}$. Let $E_{se} = 2(n-1)^2$ be the number of edges that are directly southeast of some plaquette, all of which are shown in the first image in fig. 12. Let e_{se} be the number of edges directly southeast of the faces in G_ω ; an example is shown in the second image in fig. 12. We have $e_{se} = 2f$ and $E_{se} = 2F$, and therefore

$$\frac{f}{F} = \frac{e_{se}}{E_{se}}. \quad (24)$$

Let e_w be the number of vertical edges open in ω that are *not* directly east of any face of G_ω , and e_n the number of horizontal edges not directly south of any face. Thus, $e = e_{se} + e_w + e_n$. The second image in fig. 12 has $e_{se} = 12$, $e_w = 4$, and $e_n = 5$. Likewise, let E_w and E_n be the number of vertical and horizontal (respectively) edges in S_n not already counted by E_{se} , so that $E_w = E_n = n-1$ and $E = E_{se} + E_w + E_n$.

Each of the nonempty horizontal rows of plaquettes in G_ω contributes at least 1 to e_w and at most $n-1$ to f , whereas each horizontal row of plaquettes in $G = S_n$ contributes exactly 1 to E_w and exactly $n-1$ to F . Thus, $\frac{e_w}{f} \geq \frac{1}{n-1} = \frac{E_w}{F}$, and therefore

$$\frac{f}{F} \leq \frac{e_w}{E_w}. \quad (25)$$

Likewise,

$$\frac{f}{F} \leq \frac{e_n}{E_n}. \quad (26)$$

Since we're assuming $\omega \notin \{\omega^0, \omega^1\}$, there is either some nonempty row or some nonempty column of faces in G_ω that has strictly fewer than $n-1$ faces, so at least one of the inequalities (25) and (26) holds strictly. Combining (24) to (26) via the median inequality⁴⁰ gives the required

$$\frac{f}{F} < \frac{e_{se} + e_w + e_n}{E_{se} + E_w + E_n} = \frac{e}{E}. \quad \square$$

Incidentally, the ground states of the random-cluster model are unrelated to the ground states of the Potts model—the Edwards–Sokal coupling doesn't give any useful connection. But the ground states of the Potts model are easy to identify (for positive interactions $\beta > 0$): Assign equal spins to the vertices within each connected component of G . More generally, in the higher Potts model with free boundary condition, the ground states are the cocycles (elements of $Z^{r-1}(X, \mathbb{Z}/q\mathbb{Z})$.)

⁴⁰The median inequality is: $\frac{a}{b} \leq \frac{c}{d} \implies \frac{a}{b} \leq \frac{a+c}{b+d} \leq \frac{c}{d}$ and $\frac{a}{b} < \frac{c}{d} \implies \frac{a}{b} < \frac{a+c}{b+d} < \frac{c}{d}$, where $a, b, c, d \in \mathbb{R}$ and $b, d > 0$.

References

- [AF84] M. Aizenman and J. Fröhlich. “Topological anomalies in the n dependence of the n -states Potts lattice gauge theory”. In: *Nuclear Phys. B* 235.1, FS11 (1984), pp. 1–18.
- [Aiz+83] M. Aizenman et al. “On a sharp transition from area law to perimeter law in a system of random surfaces”. In: *Comm. Math. Phys.* 92 (1983), pp. 19–69.
- [BH11] John C. Baez and John Huerta. “An invitation to higher gauge theory”. English. In: *Gen. Relativ. Gravitation* 43.9 (2011), pp. 2335–2392.
- [BR06] Béla Bollobás and Oliver Riordan. *Percolation*. Cambridge University Press, 2006.
- [BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- [Bov06] Anton Bovier. *Statistical Mechanics of Disordered Systems: A Mathematical Perspective*. Cambridge University Press, 2006.
- [Cao20] Sky Cao. “Wilson loop expectations in lattice gauge theories with finite gauge groups”. In: *Comm. Math. Phys.* 380.3 (2020), pp. 1439–1505.
- [CST18] Tullio Ceccherini-Silberstein, Fabio Scarabotti, and Filippo Tolli. *Discrete Harmonic Analysis*. Cambridge University Press, 2018.
- [Cha19] Sourav Chatterjee. “Yang–Mills for probabilists”. In: *Probability and Analysis in Interacting Physical Systems*. Springer Nature, 2019, pp. 1–16.
- [Cha20] Sourav Chatterjee. “Wilson loops in Ising lattice gauge theory”. In: *Comm. Math. Phys.* 377.1 (2020), pp. 307–340.
- [Cha21] Sourav Chatterjee. “A probabilistic mechanism for quark confinement”. In: *Comm. Math. Phys.* 385.1 (2021), pp. 1007–1039.
- [Çin11] Erhan Çinlar. *Probability and Stochastics*. Springer-Verlag, 2011.
- [Cra23] Keenan Crane. *Discrete Differential Geometry: An Applied Introduction*. [<https://www.cs.cmu.edu/~kmc Crane/Projects/DDG/>; accessed 2023-05-29]. 2023.
- [Dal61] Giorgio Dall’Aglio. “Sulle distribuzioni doppie con margini assegnati soggette a delle limitazioni”. In: *Giorn. Ist. Ital. Attuari* 94 (1961), pp. 94–108.
- [Des+05] Mathieu Desbrun et al. *Discrete exterior calculus*. 2005. arXiv: [math/0508341](https://arxiv.org/abs/math/0508341) [[math.DG](https://arxiv.org/abs/math/0508341)].

- [Die08] Tammo tom Dieck. *Algebraic Topology*. European Mathematical Society, 2008.
- [DW82] K. Drühl and H. Wagner. “Algebraic formulation of duality transformations for abelian lattice models”. In: *Ann. Physics* 141.2 (1982), pp. 225–253.
- [Dum16] Hugo Duminil-Copin. *Graphical Representations of Lattice Spin Models: Cours Peccot, Collège de France*. Spartacus supérieur. Spartacus IDH, 2016.
- [Dum20] Hugo Duminil-Copin. “Lectures on the Ising and Potts models on the hypercubic lattice”. In: *Random graphs, phase transitions, and the Gaussian free field*. Vol. 304. Springer Proc. Math. Stat. Switzerland: Springer Nature, 2020, pp. 35–161.
- [DS23] Paul Duncan and Benjamin Schweinhart. *Topological phases in the plaquette random-cluster model and Potts lattice gauge theory*. 2023. arXiv: [2207.08339v2](https://arxiv.org/abs/2207.08339v2).
- [ES88] Robert G. Edwards and Alan D. Sokal. “Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm”. In: *Phys. Rev. D* (3) 38.6 (1988), pp. 2009–2012.
- [EFS93] Aernout C. D. van Enter, Roberto Fernández, and Alan D. Sokal. “Regularity properties and pathologies of position-space renormalization-group transformations: scope and limitations of Gibbsian theory”. In: *Journal of Statistical Physics* 72.5/6 (1993), pp. 879–1167.
- [ER60] Paul Erdős and Alfréd Rényi. “On the evolution of random graphs”. In: *Magyar Tud. Akad. Mat. Kutató Int. Közl.* 5 (1960), pp. 17–61.
- [FLV21] Malin P. Forsström, Jonatan Lenells, and Fredrik Viklund. *Wilson loops in finite Abelian lattice gauge theories*. 2021. arXiv: [2001.07453](https://arxiv.org/abs/2001.07453) [math.PR].
- [FV17] Sacha Friedli and Yvan Velenik. *Statistical Mechanics of Lattice Systems*. Cambridge University Press, 2017.
- [Geo11] Hans-Otto Georgii. *Gibbs Measures and Phase Transitions*. Second Edition. Vol. 9. De Gruyter Studies in Mathematics. Berlin: Walter de Gruyter & Co., 2011.
- [GHM01] Hans-Otto Georgii, Olle Häggström, and Christian Maes. “The random geometry of equilibrium phases”. In: *Phase transitions and critical phenomena, Vol. 18*. Vol. 18. Phase Transit. Crit. Phenom. San Diego, CA: Academic Press, 2001, pp. 1–142.
- [Gor16] Alexey L. Gorodentsev. *Algebra I*. Springer International Publishing, 2016.

- [Gri67] Robert B. Griffiths. “Correlations in Ising ferromagnets. I”. In: *Journal of Mathematical Physics* 8 (1967), pp. 478–483.
- [Gri06] Geoffrey R. Grimmett. *The Random-Cluster Model*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [Gri18] Geoffrey R. Grimmett. *Probability on Graphs*. Vol. 8. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2018.
- [HS16] Yasuaki Hiraoka and Tomoyuki Shirai. “Tutte polynomials and random-cluster models in Bernoulli cell complexes”. In: *RIMS Kôkyûroku Bessatsu B59* (2016), pp. 289–304.
- [Hol12] Frank den Hollander. *Probability Theory: The Coupling Method*. [https://mathematicaster.org/teaching/lcs22/hollander_coupling.pdf; accessed 2023-05-30]. 2012.
- [KMM04] Tomasz Kaczynski, Konstantin Mischaikow, and Marian Mrozek. *Computational Homology*. New York: Springer-Verlag, 2004.
- [KM13] Tomasz Kaczynski and Marian Mrozek. “The cubical cohomology ring: an algorithmic approach”. In: *Found. Comput. Math.* 13.5 (2013), pp. 789–818.
- [KK06] Jeff Kahn and Gil Kalai. *Thresholds and expectation thresholds*. 2006. arXiv: [math/0603218v2](https://arxiv.org/abs/math/0603218v2).
- [KG95] N. Kawashima and J. E. Gubernatis. “Generalization of the Fortuin—Kasteleyn transformation and its application to quantum spin simulations”. In: *Journal of Statistical Physics* 80.1/2 (1995), pp. 169–221.
- [Kel61] Hans G. Kellerer. “Funktionen auf Produkträumen mit vorgegebenen Marginal-Funktionen”. In: *Math. Ann.* 144 (1961), pp. 323–344.
- [Kle14] Achim Klenke. *Probability Theory: A Comprehensive Course*. Second Edition. Springer-Verlag, 2014.
- [KO12] Yukihiro Komura and Yutaka Okabe. “GPU-based Swendsen–Wang multi-cluster algorithm for the simulation of two-dimensional classical spin systems”. In: *Computer Physics Communications* 183 (2012), pp. 1155–1161.
- [Lau96] Steffen L. Lauritzen. *Graphical Models*. Vol. 17. Oxford Statistical Science Series. Oxford Science Publications. New York: The Clarendon Press, 1996.
- [Lig05] Thomas M. Liggett. *Interacting Particle Systems*. Classics in Mathematics. Reprint of the 1985 original. Springer, 2005.

- [Lin92] Torngny Lindvall. *Lectures on the Coupling Method*. New York: John Wiley & Sons, 1992.
- [LP16] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*. Vol. 42. Cambridge Series in Statistical and Probabilistic Mathematics. New York NY: Cambridge University Press, 2016.
- [Mun84] James R. Munkres. *Elements of Algebraic Topology*. Westview Press, 1984.
- [PP22] Jinyoung Park and Huy Tuan Pham. *A proof of the Kahn-Kalai conjecture*. 2022. arXiv: [2203.17207](https://arxiv.org/abs/2203.17207).
- [Pey20] Gabriel Peyré. *The Discrete Algebra of the Fourier Transform*. [<https://mathematical-tours.github.io/daft/>; accessed 2021-12-13]. 2020.
- [Pfe03] Hendryk Pfeiffer. “Higher gauge theory and a non-Abelian generalization of 2-form electrodynamics”. In: *Ann. Physics* 308.2 (2003), pp. 447–477.
- [PR15] Paweł Pilarczyk and Pedro Real. “Computation of cubical homology, cohomology, and (co)homological operations via chain contraction”. In: *Adv. Comput. Math.* 41.1 (Feb. 2015), pp. 253–275.
- [Pot52] R. B. Potts. “Some generalized order-disorder transformations”. In: *Proceedings of the Cambridge Philosophical Society* 48 (1952), pp. 106–109.
- [Rom12] Steven Roman. *Fundamentals of Group Theory: An Advanced Approach*. Boston: Birkhäuser, 2012.
- [Rue04] David Ruelle. *Thermodynamic Formalism: The Mathematical Structures of Equilibrium Statistical Mechanics*. Second Edition. Cambridge University Press, 2004.
- [Sti93] John Stillwell. *Classical Topology and Combinatorial Group Theory*. Second Edition. Springer-Verlag, 1993.
- [Str65] Volker Strassen. “The existence of probability measures with given marginals”. In: *Ann. Math. Statist.* 36 (1965), pp. 423–439.
- [Ter99] Audrey Terras. *Fourier analysis on finite groups and applications*. Vol. 43. London Mathematical Society Student Texts. Cambridge University Press, 1999.
- [Thö00] Erik Thönnnes. “A primer on perfect simulation”. In: *Statistical Physics and Spatial Statistics*. Springer-Verlag, 2000, pp. 349–379.
- [Wee20] Jeffrey R. Weeks. *The Shape of Space*. Third Edition. Taylor & Francis, 2020.

- [Wu82] F. Y. Wu. "The Potts model". In: *Reviews of Modern Physics* 54.1 (1982), pp. 235–268.
- [Zie94] Günter M. Ziegler. *Lectures on Polytopes*. Springer-Verlag, 1994.

Appendix: Code listings

This Python 3 code was used to generate figs. 9 and 10 in section 6.3.

```
1 from collections import defaultdict
import matplotlib
import matplotlib.pyplot as plt
import numpy as np

6 # Graph data format:
# (vertices, edges) where:
# vertices is a set
# edges is a set of unordered pairs (as frozensets) of vertices
# Only simple connected finite graphs are supported.

11
# Return graph: Square lattice with side length n (i.e., n*n vertices.)
def make_square_graph(n):
    def z2dist(edge1, edge2):
        return abs(edge1[0] - edge2[0]) + abs(edge1[1] - edge2[1])
16 vertices = set([(x, y) for x in range(n) for y in range(n)])
edges = set([frozenset({v, w})
              for v in vertices for w in vertices
              if z2dist(v,w) == 1])
    return (vertices, edges)

21
# Return graph: Complete graph on n vertices.
def make_complete_graph(n):
    assert 1 <= n and n <= 4 # K_n is not planar for n > 4.
    vertices = set(range(n))
26 edges = set([frozenset({v,w}) for v in vertices for w in vertices if v != w])
    return (vertices, edges)

# Return graph: Two copies of complete graph K_n, joined by a single edge.
def make_siamese_graph(n):
31 assert n >= 1 # Otherwise, can't join them.
vertices_a = set(range(0, n))
edges_a = set([frozenset({v,w}) for v in vertices_a for w in vertices_a if v != w])
vertices_b = set(range(n, 2*n))
edges_b = set([frozenset({v,w}) for v in vertices_b for w in vertices_b if v != w])
36 joiner = frozenset((0, n))
```

```

    return (set.union(vertices_a, vertices_b),
            set.union(edges_a, edges_b, {joiner}))

# Return the number of clusters.
41 def num_clusters(graph):
    remaining_vertices = set(graph[0])
    remaining_edges = set(graph[1])
    count = 0
    while(len(remaining_vertices) > 0):
46     incident_vertices = set((remaining_vertices.pop(),))
        current_cluster = set(incident_vertices)
        while(incident_vertices != {}):
            incident_edges = set()
            for v in incident_vertices:
51                 new_incident_edges = set(filter(lambda e: v in e, remaining_edges))
                    incident_edges = incident_edges.union(new_incident_edges)
                    remaining_edges = remaining_edges.difference(new_incident_edges)
            current_cluster = current_cluster.union(incident_vertices)
            if len(incident_edges) == 0:
56                 incident_vertices = {}
            else:
                incident_vertices = set.difference(
                    set(frozenset.union(*incident_edges)),
                    incident_vertices
61                )
            count += 1
            remaining_vertices = remaining_vertices.difference(current_cluster)
    return count

66 # Return a list of all subsets of s.
def powerset(s):
    def _powerset(s):
        if len(s) == 0:
            return [set()]
71     x = s.pop()
        tail = powerset(s)
        return tail + [t.union((x,)) for t in tail]
    return _powerset(set(s))

```

```

76 def graph_configurations(graph):
    vertices, edges = graph
    for edge_subset in powerset(edges):
        yield (vertices, edge_subset)

81 # Return a (default)dict with keys (e, c) and values n, such that n is the number
    # of configurations that have e edges and c clusters. All entries will have n >= 1.
    def ledger(graph):
        l = defaultdict(lambda: 0)
        for g in graph_configurations(graph):
86         l[(len(g[1]), num_clusters(g))] += 1;
        return l

    if __name__ == "__main__":
91     # Generate the data

        data = ledger(make_square_graph(4))
        graph_name = "Square graph $S_4$"

96     #data = ledger(make_complete_graph(4))
        #graph_name = "Complete graph $K_4$"

        #data = ledger(make_siamese_graph(4))
        #graph_name = "Siamese graph $K_4 + K_4$"

101     # Plot the data
        plt.figure(dpi=300) # High resolution
        plt.rc('axes', axisbelow=True) # Grid behind everything else
        plt.grid(True, which="both", color="lightgrey")

106     data_flat = [(k[0], k[1], v) for (k, v) in data.items()]
        e_values = np.array([i[0] for i in data_flat])
        c_values = np.array([i[1] for i in data_flat])
        plt.scatter(e_values, c_values, color='red')

111     for e,c,n in data_flat:
        label = "{:g}".format(n)
        plt.annotate(n,

```

```

    (e,c),
116     textcoords="offset points",
        xytext=(8,7), # label offset
        ha='center',
        arrowprops=None,
        fontsize=6
121     )

plt.gca().xaxis.set_major_locator(matplotlib.ticker.MultipleLocator(1))
plt.gca().yaxis.set_major_locator(matplotlib.ticker.MultipleLocator(1))
plt.title("Distribution of graph configurations: " + graph_name)
plt.xlabel("$o(\omega)$ (number of open edges)")
126 plt.ylabel("$k(\omega)$ (number of clusters)")
plt.show()

```