The Berezinskii–Kosterlitz–Thouless transition at the critical point

Piet Lammers

Abstract. The theoretical physicists Berezinskii [Ber72] and independently Kosterlitz and Thouless [KT73] described a new type of phase transition in the 1970s. They argued that the motor behind this phase transition, which occurs in the classical XY model (or classical plane rotor model) in two dimensions, is the changing influence of topological defects known as *vortices* and *antivortices*. Roughly speaking, such vortices and antivortices may be interpreted as particles in their own right, which are bound into pairs of opposing charge at low temperature, but not at high temperature. Fröhlich and Spencer rigorously established the *Berezinskii–Kosterlitz–Thouless (BKT) transition* in the 1980s [FS81].

Ideas which appeared in the new millennium enabled the development of an original perspective on the situation. The current lecture notes, which serve as the basis for the *cours Peccot* (to be) taught at *Collège de France* in 2024, aim to rigorously present some of these developments in a self-contained fashion. The following topics are covered:

- A phase transition for height functions is established through *symmetry breaking*,
- This transition is shown to be sharp by using *renormalisation inequalities*,
- The *Brydges–Fröhlich–Spencer random walk* is used to show that this phase transition coincides with the BKT transition.

The existence of the BKT transition is derived as a corollary of those three results. We highlight that these ideas enable the derivation of new results at and around the critical point, where the work of Fröhlich and Spencer [FS81] is perturbative and therefore gives less information at the critical point.

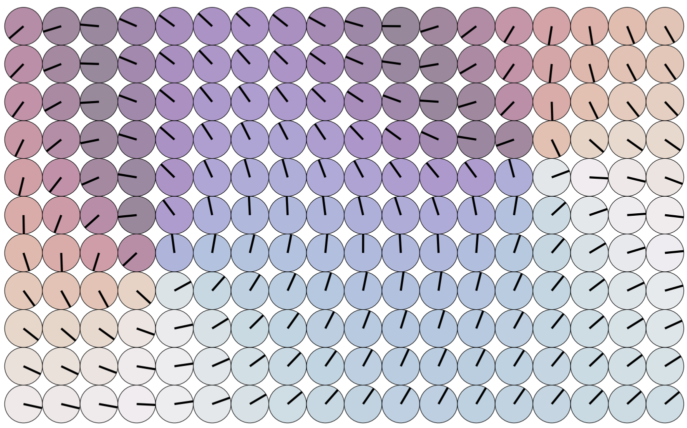

Date: 23 January 2024.

Key words and phrases. XY model, Berezinskii–Kosterlitz–Thouless transition, vortex-antivortex binding, symmetry breaking, percolation theory, Brydges–Fröhlich–Spencer random walk.

About these notes

These notes are work in progress, and because of time constraints I focussed on the pedagogy of the mathematical arguments. I envision to update these notes during the course on a rolling basis, then take some time to finetune the details and add further references and remarks. Please do not hesitate to get in touch if you have any questions or remarks regarding these notes.

Acknowledgement

I would like to warmly thank Hugo Duminil-Copin for being a mentor and collaborator. The generosity with which he has supervised me during my postdoc has had a decisive influence on the work exhibited here.

Contents

About these notes 2
Acknowledgement 2
1. The XY model and the BKT transition 3
1.1. Definition of the XY model 3
1.2. Spin lattice models and phase transitions 4
1.3. Physics overview 5
1.4. The Fröhlich–Spencer approach 7
1.5. Expansion of the XY model 10
1.6. The height function 11
1.7. The new approach 14
2. Review of the FKG inequality 14
2.1. The Harris inequality and the FKG inequality 15
2.2. The FKG lattice condition 17
2.3. The Ising model 20
2.4. Boundary conditions and the infinite-volume limit 23
2.5. Percolation: uniqueness of the infinite cluster 27
2.6. Percolation: planar graphs 28
2.7. Percolation: the noncoexistence theorem in 2D (Zhang's argument) 29
2.8. Percolation: the noncoexistence theorem in 2D (general argument) 31
References 35

1. The XY model and the BKT transition

We are interested in the *XY model*, which is also known in the literature as the *plane rotor model*, the *classical XY model*, or the *classical plane rotor model*.

1.1. Definition of the XY model.

Let $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ denote a finite simple graph, and let $\mathbb{S}^1 \subset \mathbb{C}$ denote the unit circle embedded in the complex plane.

Definition 1.1 (The XY model)**.** Let $J \in [0,\infty)^{\mathbf{E}}$ denote a family of *coupling constants*. The *Hamiltonian* $H^{\mathrm{XY}}_{\mathbf{G},J}(\sigma) \in \mathbb{R}$ associated to a *spin configuration* $\sigma \in (\mathbb{S}^1)^{\mathbf{V}}$ is defined by

$$H^{\mathrm{XY}}_{\mathbf{G},J}(\sigma) := -\sum_{xy\in\mathbf{E}} J_{xy}(\sigma_x,\sigma_y); \qquad (\sigma_x,\sigma_y) := \frac{1}{2}(\sigma_x\bar{\sigma}_y + \bar{\sigma}_x\sigma_y).$$

The *XY model* on the graph $\mathbf{G}$ is the probability measure $\langle\,\cdot\,\rangle^{\mathrm{XY}}_{\mathbf{G},J}$ on $(\mathbb{S}^1)^{\mathbf{V}}$ defined through its expectation functional

$$\langle A\rangle^{\mathrm{XY}}_{\mathbf{G},J} := \frac{1}{Z^{\mathrm{XY}}_{\mathbf{G},J}}\int A(\sigma)e^{-2H^{\mathrm{XY}}_{\mathbf{G},J}(\sigma)}\mathrm{d}\sigma; \qquad Z^{\mathrm{XY}}_{\mathbf{G},J} := \int e^{-2H^{\mathrm{XY}}_{\mathbf{G},J}(\sigma)}\mathrm{d}\sigma,$$

where $\mathrm{d}\sigma$ denotes the Haar measure on $(\mathbb{S}^1)^{\mathbf{V}}$. Subscripts are omitted when they are clear from the context. Typically we take $J \equiv \beta \in [0,\infty)$, in which case β is called the *inverse temperature*.

Recall that the Haar measure is a probability measure in which each spin is oriented independently and uniformly random in the unit circle. Notice that each term in the definition of the Hamiltonian equals the cosine of the angle difference between σ_x and σ_y, in the sense that

$$H^{\mathrm{XY}}(\sigma) = -\sum_{xy\in\mathbf{E}} J_{xy}\cos(\theta_y-\theta_x); \qquad \theta := -i\log\sigma \in (\mathbb{R}/2\pi\mathbb{Z})^{\mathbf{V}}. \tag{1}$$

The Hamiltonian therefore favours configurations with small angle differences along the edges. The inverse temperature parameter β regulates how strongly small angle differences are favoured over large angle differences, see Figure 1. Let us also write $T := 1/\beta$ for the *temperature*; it is linguistically convenient to define this parameter although it is usually the inverse temperature β that appears in mathematical equations.

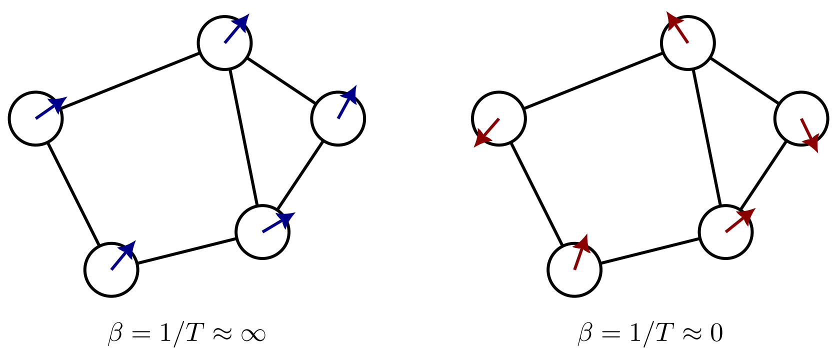

Figure 1. The XY model on a small graph. On the left, the temperature is low, which means that alignment of the spins is strongly favoured. On the right, the temperature is high, which means that the interaction is almost negligible and as a consequence the spins behave almost independently.

The most important observable is the *two-point function*, defined, for a fixed pair of vertices $x, y \in \mathbf{V}$, as

$$\langle\sigma_x\bar{\sigma}_y\rangle^{\mathrm{XY}} = \langle\bar{\sigma}_x\sigma_y\rangle^{\mathrm{XY}} = \langle(\sigma_x,\sigma_y)\rangle^{\mathrm{XY}} = \langle\cos(\theta_y-\theta_x)\rangle^{\mathrm{XY}}.$$

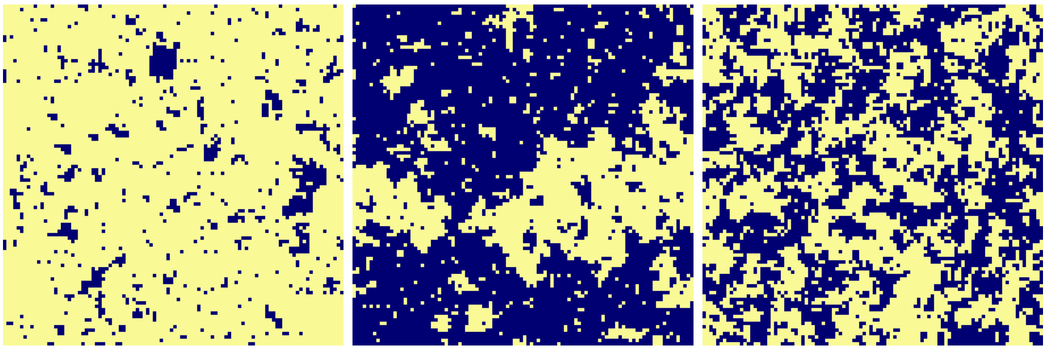

$\beta > \beta_c$ $\beta = \beta_c$ $\beta < \beta_c$

FIGURE 2. The Ising model on a 100×100 torus. On the left, the interaction is so strong that one of the two spin states in $\{\pm 1\}$ dominates the entire picture, in the sense that there is a unique monochromatic infinite component (in this case the large yellow cluster). The middle picture exhibits a sample from the Ising model at the critical point. One can prove rigorously that there is no infinite monochromatic cluster; instead, large, finite clusters of either colour appear. On the right the interaction is too weak to induce long-range correlations, and it can be proved that in this case all monochromatic clusters are small.

Notice that the vertices x and y do not only interact through the edge xy (if this edge exists), but also indirectly through all other edges in the graph.

1.2. **Spin lattice models and phase transitions.** For $\mathbf{G}$ we often take large subgraphs of the square lattice graph $\mathbb{Z}^d$, and in those instances the XY model is called a *lattice model*. Lattice models serve as discrete models for physical experiments. They also serve as discretisations of Euclidean- and quantum field theories. There exists a large family of lattice models, which facilitates the mathematical study of a vast range of physical phenomena. The XY model is a *lattice spin model*, which means that the model consists of random spins which are assigned to the vertices of the graph. The oldest lattice spin model is undoubtedly the *Ising model*. Its definition is identical to that of the XY model, except that the Haar measure $\mathrm{d}\sigma$ is replaced by the Haar measure on the set $\{\pm 1\}^{\mathbf{V}}$, and $2H^{\mathrm{XY}}$ is replaced by $H^{\mathrm{Ising}}_{\mathbf{G},J}(\sigma) := -\sum_{xy \in \mathbf{E}} J_{xy}\sigma_x\sigma_y$.

We say that a lattice model undergoes a *phase transition* when its qualitative properties change when a continuous parameter of the model, often the inverse temperature β, passes a specific value, which is then called the *critical point* or *critical temperature*. For example, the Ising model in dimension $d = 2$ is known to undergo a phase transition at some critical point $\beta_c \in (0, \infty)$, see Figure 2. Below the critical temperature, the Ising model in dimension $d = 2$ exhibits an infinite monochromatic cluster (the yellow cluster on the left in Figure 2) that consists of vertices that are all in the same state. This phenomenon is called *long-range order* or *spontaneous magnetisation* in this context.

A common objective for physicists and mathematicians is the identification and classification of phase transitions. Thus, interesting questions include:

(1) Does a model undergo a phase transition?
(2) What is the qualitative behaviour of the model at and around the critical point?
(3) What are the fundamental features of statistical mechanics models determining the qualitative behaviour of the phase transition?

These questions depend starkly on the model of interest, but also on the dimension of the graph on which they are studied. Peierls was amongst the first to prove the existence of a

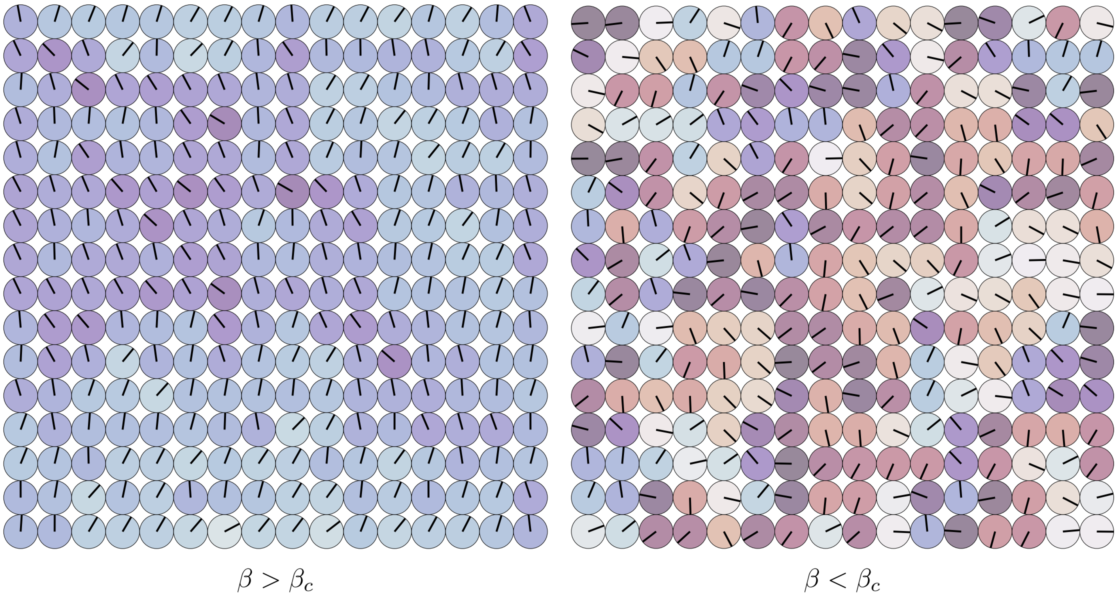

FIGURE 3. The XY model on a 32×32 torus at low and high temperature

phase transition in a lattice spin model, by deriving the magnetisation transition in the Ising model in dimension $d = 2$ in 1936 [Pei36]. The Ising model is now perhaps the best understood of all lattice models: we refer to two works of Duminil-Copin for a pedagogical introduction [Dum17] and an overview [Dum22].

1.3. **Physics overview.** Let us start with a physical introduction to the Berezinskii–Kosterlitz–Thouless (BKT) transition. The BKT transition concerns the XY model in dimension $d = 2$ (that is, on the two-dimensional square lattice graph). The first question is to ask if the model undergoes a phase transition, that is, if the model behaves differently at low and high temperature. See Figure 3 for two samples from the XY model at different temperatures. Let us make some preliminary observations.

- At high temperature (weak interaction), the spins seem to quickly decorrelate over large distances. The two-point function $\langle \sigma_x \bar{\sigma}_y \rangle^{\mathrm{XY}}_{\mathbb{Z}^2,\beta}$ decays exponentially fast in $\|y - x\|_2$ when $\beta \in (0, \infty)$ is sufficiently small. This is in fact easy to prove mathematically in several different ways, such as for example by analysis of the high-temperature expansion, or by comparing the two-point function of the XY model to an Ising model at high temperature or a percolation model with a small percolation parameter.
- Our focus will therefore be on understanding the complementary regime, that is, the regime where the interaction strength $\beta \in (0, \infty)$ is so large that it significantly affects the system. Consider for a moment the situation where β is very large. Neighbouring spins then tend to point in the same approximate direction. However, the spins are sampled from a continuous measure, and therefore the alignment is never perfect. Moreover, by expanding the cosine in (1) to second order, we observe that small angles are penalised quadratically, and in particular angles of order $o(1/\sqrt{\beta})$ are not penalised so much. This gives the model a certain flexibility and sets it apart from the Ising model; in the low-temperature Ising model, a huge energy gap must be bridged for neighbouring spins to be different.

Mermin and Wagner proved in 1966 that there is no continuous symmetry breaking in two dimensions [MW66]. This essentially means that at any fixed inverse temperature $\beta \in (0, \infty)$, the flexibility described above is sufficient to avoid a situation in which most

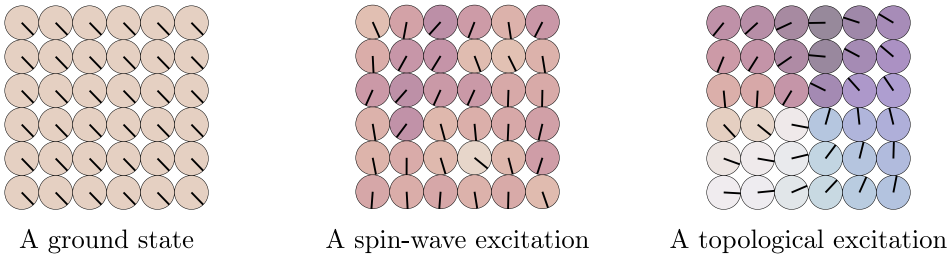

FIGURE 4. The ground state and the two types of excitations

spins point in the same direction. This seems to contradict Figure 3, LEFT, where it looks like most spins point towards the north, but the latter is only an artefact of the finite size of the torus (32×32) which is small compared to the value of β. More precisely, for each value of β, there is a minimum size n_β for the torus above which the spins do not exhibit a preferential direction. This result is a negative result, because it tells us that the two-dimensional XY model cannot have a phase transition similar to the phase transition of the two-dimensional Ising model.

Thus, the same question remains: does the model undergo any kind of phase transition? There is a completely generic method for detecting phase transitions: namely, phase transitions correspond to the points $\beta_c \in (0, \infty)$ where some appropriately chosen quantities $Q(\beta)$ fail to be analytic. To illustrate this point, let us consider the partition function. Of course, it is not possible to study the partition function of the full square lattice $\mathbb{Z}^2$ directly. Write $\Lambda_n := \{-n, \ldots, n\}^2 \subset \mathbb{Z}^2$, and abuse notation to also write Λ_n for the subgraph of the square lattice induced by Λ_n. Write E_n for the corresponding edge set. We consider the *pressure*, that is, the renormalised quantity

$$P(\beta) := \lim_{n\to\infty} \frac{1}{|\Lambda_n|} \log Z^{\mathrm{XY}}_{\Lambda_n,\beta}.$$

The Maclaurin series of $P(\,\cdot\,)$ is the formal series

$$\sum_{k=0}^{\infty} \frac{P^{(k)}(0)}{k!} \beta^k.$$

Each derivative of $P(\beta)$ at $\beta = 0$ can be calculated explicitly through a combinatorial algorithm that runs in finite time. Thus, in order to gain intuition, one may calculate a finite number of coefficients in the above series, and then heuristically extrapolate this sequence of coefficients in order to guess an approximate value for the radius of convergence. This radius of convergence is then conjectured to coincide with the first point where $P(\beta)$ fails to be analytic in β.

This procedure has drawbacks: the extrapolation of coefficients is entirely nonrigorous, and there is no way to know the number of coefficients that need to be calculated in order to see the correct pattern. Moreover, the radius of convergence may in fact indicate the presence of *complex* singularities without physical meaning. Nevertheless, Stanley [Sta68] and Moore [Moo69] used it to predict that the classical XY model undergoes a phase transition in two dimensions. It must be said that the same method was used to predict a phase transition for the classical Heisenberg model in two dimensions (see [Moo69] and the references therein) which is now generally expected to be false (there do not exist rigorous proofs in either direction).

In the 1970s, Berezinskii and independently Kosterlitz and Thouless proposed a mechanism for a phase transition in the classical XY model in two dimensions. In order to understand this mechanism, let us concentrate for a moment on the Hamiltonian $H^{\mathrm{XY}} : (\mathbb{S}^1)^{\Lambda_n} \to \mathbb{R}$. A *ground state* is a global minimum of H^{XY}. The ground states of H^{XY} coincide precisely with the spin configurations in which all spins point in the exact same direction, see Figure 4, LEFT. Deviations from the grounds state may be decomposed into deviations of two types: first, it is possible that there are other local minima than the ground state, second, a configuration may be a perturbation from a local minimum. Perturbations away from a local minimum are called *spin-wave excitations*, see Figure 4, MIDDLE. Local minima other than the ground states are called *topological excitations*, see Figure 4, RIGHT. Topological excitations exist because the spin space $\mathbb{S}^1$ is not simply connected, and correspond to faces of the square lattice graph where the spin configuration makes a full turn.

Let us assume that each spin configuration σ can be written as an independent product $\sigma = \sigma^{\mathrm{sw}}\sigma^{\mathrm{t}}$, where σ^{sw} and σ^{t} capture the spin-wave deviations from the ground state and the topological excitations respectively. This is not exactly true for the classical XY model, but this assumption may be convincingly justified in several ways (we do not do this here). The spin-wave σ^{sw} should be thought of as a continuous object as it captures the (continuous) perturbations of the ground state, while σ^{t} is in its very nature a discrete object counting the turns that the configuration makes around each face of the square lattice graph. The previously mentioned Mermin–Wagner argument essentially asserts that the spin-wave σ^{sw} does not exhibit magnetisation, which immediately implies that σ does not exhibit magnetisation. The spin-wave may in fact be identified with another object: the *discrete Gaussian free field*. Let us not go into details, but rather mention that the latter is a well-understood object that is known not to undergo a phase transition. Thus, if the XY model undergoes a phase transition, then this phase transition must be driven by the topological content σ^{t} of the model.

Consider the regime of very low temperatures. Since topological excitations are costly, they cannot occur with a large density. Nevertheless, they must occur with some positive density due to entropy considerations. We interpret the faces where the configuration makes a full turn (as in Figure 4, RIGHT) as particles in their own right, and call them *vortices* and *antivortices*. Indeed, the particles come in two flavours, depending on whether the configuration turns *in the same direction* as the face (as in Figure 4, RIGHT or the bottom-left quarter of Figure 5), or *in the opposite direction* (as in the top-right quarter of Figure 5). The particles are thought of as carrying some electric charge. At very low temperature a single vortex is very energetically costly, but its cost is greatly diminished if an antivortex appears nearby. Therefore it is natural that vortices and antivortices appear in pairs of opposite sign (we say that they are *bound* into pairs), see Figure 5.

Berezinskii [Ber72] and independently Kosterlitz and Thouless [KT73] propose the binding of vortices and antivortices as the motor behind the phase transition in the XY model. They argue that while binding is natural at low temperature, there is some critical temperature $\beta_c \in (0, \infty)$ below which the interaction strength is so small that the vortices and antivortices are not bound together and rather appear more or less independently. While Berezinskii was the first to describe this mechanism, Kosterlitz and Thouless (in addition to describing it) proposed a precise picture for the behaviour of the model around the critical point β_c using renormalisation group methods. This includes a prediction for the behaviour of the two-point function. Kosterlitz and Thouless received the 2016 Nobel prize in physics for their work (Berezinskii had passed away at the moment of the award), and the phase transition is named the Berezinskii–Kosterlitz–Thouless transition in their honour.

1.4. **The Fröhlich–Spencer approach.** Mathematically, we frame the BKT transition as a qualitative change in the behaviour of the two-point function. As a first step, we prove that the two-point function satisfies some regularity properties. These regularity properties follow from so-called *correlation inequalities*, which are inequalities involving correlation functions. They form the basic building blocks in the analysis of statistical mechanics

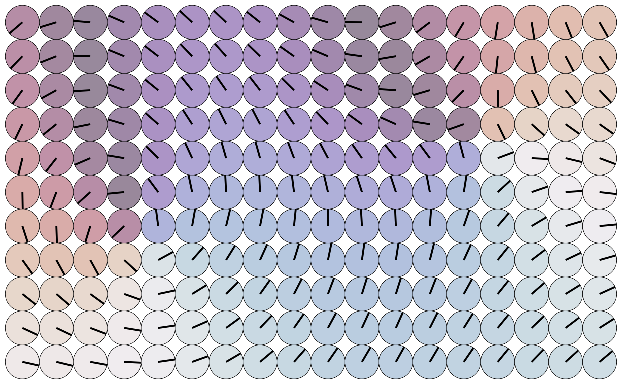

FIGURE 5. A vortex-antivortex pair at very low temperature. The fact that they appear close to each other greatly reduces the total energy cost, which explains why they are bound in pairs.

models such as the XY model, and they were already known at the time that Fröhlich and Spencer published their breakthrough work on the BKT transition.

Proposition (Ginibre inequality [Gin70])**.** For any finite graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$, vertices $x, y, u, v \in \mathbf{V}$, and coupling constants $J \in [0, \infty)^{\mathbf{E}}$, the random variables (σ_x, σ_y) and (σ_u, σ_v) have a nonnegative covariance in $\langle \cdot \rangle^{\mathrm{XY}}$. In other words, the following inequality holds true:

$$\langle (\sigma_x, \sigma_y)(\sigma_u, \sigma_v) \rangle^{\mathrm{XY}} \geq \langle (\sigma_x, \sigma_y) \rangle^{\mathrm{XY}} \langle (\sigma_u, \sigma_v) \rangle^{\mathrm{XY}}.$$

The proof is somewhat orthogonal to the rest of the theory; we leave it until later. We now state two simple corollaries.

TODO Add the proof in a later section.

Lemma (Monotonicity in the coupling strengths)**.** *For any finite graph* $\mathbf{G} = (\mathbf{V}, \mathbf{E})$*, vertices* $x, y \in \mathbf{V}$*, and families of coupling constants* $J, J' \in [0, \infty)^{\mathbf{E}}$*, we have*

$$J \leq J' \qquad \Longrightarrow \qquad \langle (\sigma_x, \sigma_y) \rangle_J^{\mathrm{XY}} \leq \langle (\sigma_x, \sigma_y) \rangle_{J'}^{\mathrm{XY}}.$$

Proof sketch. Define the interpolation

$$f(\alpha) := \langle (\sigma_x, \sigma_y) \rangle_{J + \alpha(J' - J)}^{\mathrm{XY}} = \frac{\int (\sigma_x, \sigma_y) e^{2 \sum_{uv \in \mathbf{E}} (J_{uv} + \alpha(J'_{uv} - J_{uv}))(\sigma_u, \sigma_v)} \mathrm{d}\sigma}{\int e^{2 \sum_{uv \in \mathbf{E}} (J_{uv} + \alpha(J'_{uv} - J_{uv}))(\sigma_u, \sigma_v)} \mathrm{d}\sigma},$$

differentiate in α, then apply Ginibre to see that the derivative is nonnegative. □

Lemma (Monotonicity in the graph)**.** *Recall that* (Λ_n, E_n) *is the subgraph of the square lattice induced by* $\{-n, \ldots, n\}^2$*. For any* $x, y \in \Lambda_n$ *and* $\beta \in [0, \infty)$*, we have*

$$\langle (\sigma_x, \sigma_y) \rangle_{\Lambda_n, \beta}^{\mathrm{XY}} \leq \langle (\sigma_x, \sigma_y) \rangle_{\Lambda_{n+1}, \beta}^{\mathrm{XY}}.$$

Proof. Apply the previous lemma, noting that $\langle (\sigma_x, \sigma_y) \rangle_{\Lambda_n, \beta}^{\mathrm{XY}} = \langle (\sigma_x, \sigma_y) \rangle_{\Lambda_{n+1}, \beta \cdot \mathbb{1}_{E_n}}^{\mathrm{XY}}$. □

Definition 1.2. For any $x, y \in \mathbb{Z}^2$ and $\beta \in [0, \infty)$, define

$$\langle (\sigma_x, \sigma_y) \rangle_{\mathbb{Z}^2, \beta}^{\mathrm{XY}} := \lim_{n \to \infty} \langle (\sigma_x, \sigma_y) \rangle_{\Lambda_n, \beta}^{\mathrm{XY}}.$$

This limit exists due to monotonicity in the graph, and this quantity is increasing (that is, nondecreasing) in β due to monotonicity in the coupling strengths. The monotonicity in the graph implies that the limit is invariant under replacing Λ_n by $\Lambda_n - v$ for some fixed $v \in \mathbb{Z}^2$, and therefore the quantity is invariant under shifting x and y by the same vector v.

Definition 1.3 (Definition of the critical point)**.** There exists a unique value $\beta_c \in [0, \infty]$ such that:

- For $\beta < \beta_c$, $\langle (\sigma_x, \sigma_y) \rangle^{\mathrm{XY}}_{\mathbb{Z}^2,\beta}$ decays exponentially fast in $\|y - x\|_2$,
- For $\beta > \beta_c$, $\langle (\sigma_x, \sigma_y) \rangle^{\mathrm{XY}}_{\mathbb{Z}^2,\beta}$ does not decay exponentially fast in $\|y - x\|_2$.

The constant β_c is called the *critical inverse temperature* or the *BKT point*.

This is a practical definition for a critical point for the following reasons.

- It clearly encodes a qualitative change in the behaviour of the two-point function.
- It is easy to prove that $\beta_c > 0$ (we also do this later).
- Derivatives of many interesting quantities, such as the pressure $P(\beta)$ defined in Subsection 1.3, can be expressed in terms of two-point functions. If the two-point function exhibits exponential decay, then one expects to be able to bound those derivatives, implying that those quantities are analytic. This suggests that there cannot be any phase transitions as long as the two-point function decays exponentially fast. Thus, β_c should in fact coincide with the *smallest* critical point. In particular, if it so happens to be that $\beta_c = \infty$, then one would expect that the model does not undergo any kind of phase transition.

TODO Add reference to the location.

The highlight of these notes is that we provide a new, rigorous proof of the following fact.

Theorem 1 (Existence of the BKT transition [FS81]). *On $\mathbb{Z}^2$, we have $\beta_c < \infty$.*

In fact, Fröhlich and Spencer prove much more precise results on the behaviour of the two-point function for large values of β, and the above theorem is a mere corollary of that.

Theorem ([FS81]). *There exists a constant $\kappa \in (0, \infty)$ such that*

$$\langle (\sigma_x, \sigma_y) \rangle^{\mathrm{XY}}_{\mathbb{Z}^2,\beta} \geq \frac{1}{\kappa(1 + \|y - x\|_2)^{\kappa/\beta}} \qquad \forall \beta \in [\kappa, \infty),\ \forall x, y \in \mathbb{Z}^2.$$

We now give a very rough sketch of how Fröhlich and Spencer proved this result. At the first step of the proof, they transform the model into a model of *height functions*. These are random integer- or real-valued functions on the square lattice $\mathbb{Z}^2$. The XY model itself is dual to an integer-valued height function; see Subsection 1.6 below. One would like to understand how much this integer-valued height function deviates from a real-valued Gaussian height function called the Gaussian free field. As was already mentioned in the physics introduction, this Gaussian free field is understood well from the mathematical perspective. In order to compare the integer-valued height function to the Gaussian free field, one introduces a family of interpolations between the two, which essentially consists of Gaussian free fields but with an extra interaction term which makes the heights prefer to be close to an integer. This interaction term may be interpreted as a magnetic field. The interaction strengths of this magnetic field at each vertex are interpreted as particles or charges in their own right. At very low temperature, one can show that those particles or charges are sufficiently well-behaved, which allows one to derive the desired bounds on the two-point function. We refer to [KP17] for a pedagogical introduction to the proof. While the proof requires the temperature of the XY model to be very low (very high β), it is quite robust. For example, the result of Fröhlich and Spencer is stable under introducing arbitrary magnetic fields of a certain kind, as is demonstrated in the recent work of Garban and Sepúlveda [GS20]. This is remarkable because this particular magnetic field destroys any kind of symmetry, thus demonstrating that the proof does not depend on it.

The relation with the height function is still central in this course, but our analysis of the height function is entirely distinct from the one of Fröhlich and Spencer. In particular, we focus on probabilistic methods which also allow us to derive results on the height function at and around the critical point β_c defined above.

1.5. **Expansion of the XY model.** For any functions a and b defined on some finite set X, we use the following shorthand notations:

$$a^b := \prod_{x\in X} a_x^{b_x}; \qquad a! := \prod_{x\in X} a_x!.$$

We first state a trivial result concerning the infinite-temperature XY model.

Proposition. Consider the XY model on $\mathbf{G}$ at inverse temperature $\beta = 0$. Then

$$\forall a \in \mathbb{Z}^{\mathbf{V}}, \qquad \langle \sigma^a \rangle^{\mathrm{XY}}_{\mathbf{G},0} = \int \sigma^a \mathrm{d}\sigma = \mathbb{1}_{\{a\equiv 0\}}.$$

Proof. Recall that $\mathrm{d}\sigma$ denotes the Haar measure (a probability measure) in which all spins are independent and uniformly random. If $a \equiv 0$ then we are integrating the integrand 1 with respect to a probability measure. By independence, the integral decomposes as a product over the vertices. If $a_x \neq 0$ then we are averaging $\sigma_x^{a_x}$ with respect to the uniform distribution $\mathrm{d}\sigma_x$ on the unit circle $\mathbb{S}^1 \subset \mathbb{C}$, which clearly produces zero. □

Recall from Definition 1.1 for the XY model that the expectation $\langle A \rangle^{\mathrm{XY}}_{\mathbf{G},J}$ of any random variable A is defined through

$$Z^{\mathrm{XY}} \langle A \rangle^{\mathrm{XY}} := \int A(\sigma) e^{-2H^{\mathrm{XY}}(\sigma)} \mathrm{d}\sigma.$$

We are principally interested in the partition function ($A \equiv 1$) and in multipoint correlation functions ($A(\sigma) = \sigma^a$ for some $a \in \mathbb{Z}^{\mathbf{V}}$). Both cases are treated simultaneously; the partition function just corresponds to the choice $a \equiv 0$. We start our analysis by first writing $e^{-2H^{\mathrm{XY}}(\sigma)}$ as a product over the terms in the definition of the Hamiltonian, then expanding the exponential in each factor so appearing. This yields

$$\begin{aligned} Z^{\mathrm{XY}} \langle \sigma^a \rangle^{\mathrm{XY}} &= \int \sigma^a e^{-2H^{\mathrm{XY}}(\sigma)} \mathrm{d}\sigma \\ &= \int \sigma^a \prod_{xy \in \vec{\mathbf{E}}} e^{J_{xy}\sigma_x \bar{\sigma}_y} \mathrm{d}\sigma \\ &= \sum_{\mathbf{n} \in (\mathbb{Z}_{\geq 0})^{\vec{\mathbf{E}}}} \frac{J^{\mathbf{n}}}{\mathbf{n}!} \int \sigma^a \prod_{xy\in\vec{\mathbf{E}}} (\sigma_x \bar{\sigma}_y)^{\mathbf{n}_{xy}} \mathrm{d}\sigma. \end{aligned} \tag{2}$$

For the last equation, we expanded each exponential, but we also interchanged the sum and the integral. This is easily justified: the combinatorial terms in the denominator grow so fast that we may apply Fubini's theorem without blinking an eye. The final expression (2) may be simplified in two steps: first, we rewrite the sum on the left in terms of a measure (in some sense this step is purely cosmetic); second, we evaluate the integral over $\mathrm{d}\sigma$ using the previous proposition.

For the first step, we first introduce the product measure on $\mathbf{n} \in (\mathbb{Z}_{\geq 0})^{\vec{\mathbf{E}}}$ defined through

$$\mathbb{M}_{\mathbf{G},J}[\{\mathbf{n} = \mathbf{m}\}] := \frac{J^{\mathbf{m}}}{\mathbf{m}!} = \prod_{xy\in\vec{\mathbf{E}}} \frac{J_{xy}^{\mathbf{m}_{xy}}}{\mathbf{m}_{xy}!}.$$

Notice that $\mathbb{M}_{\mathbf{G},J}$ is the nonnormalised version of the probability measure in which $\mathbf{n}$ is a family of independent Poisson random variables whose parameters are given by J. Again, we drop subscripts when they are obvious from the context. The random function $\mathbf{n}$ is called a *directed random current*. We may now rewrite (2) into

$$Z^{\mathrm{XY}} \langle \sigma^a \rangle^{\mathrm{XY}} = \mathbb{M}\left[\int \sigma^a \prod_{xy\in\vec{\mathbf{E}}} (\sigma_x \bar{\sigma}_y)^{\mathbf{n}_{xy}} \mathrm{d}\sigma \right].$$

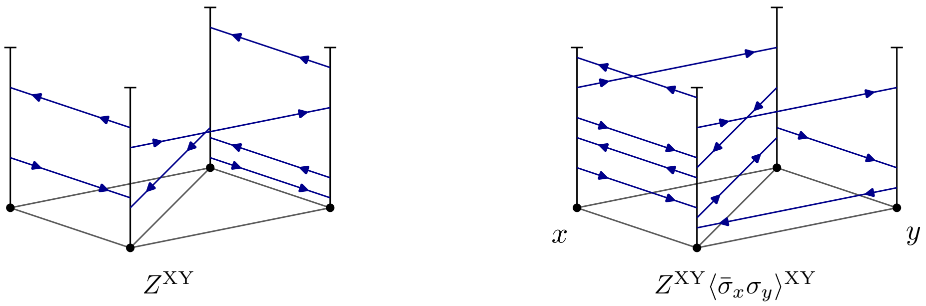

FIGURE 6. Two multigraphs $(\mathbf{V},\mathbf{n})$ contributing to the expansions of the two quantities. The graph $\mathbf{G}$ consists of four vertices, five edges, and is drawn in grey. Each multigraph $(\mathbf{V},\mathbf{n})$ is visualised as follows: each multi-edge corresponding to an edge xy is assigned a uniformly random *time* in the interval $[0,J_{xy}]$. The time axis points in the vertical direction. For example, on the left, there are $\sum_{xy}\mathbf{n}_{xy}=7$ multi-edges, and therefore there are 7 time-decorated directed edges, drawn in blue.

We interpret $(\mathbf{V},\mathbf{n})$ as a random directed multigraph, where $\mathbf{n}_{xy}$ counts the number of multi-edges pointing from x to y.

We now focus on the second step, namely the integral over $\mathrm{d}\sigma$ appearing in the previous display. Observe that

$$\prod_{xy\in\vec{\mathbf{E}}}(\sigma_x\bar\sigma_y)^{\mathbf{n}_{xy}}=\sigma^{\partial\mathbf{n}};\qquad \partial\mathbf{n}:\mathbf{V}\to\mathbb{Z},\ x\mapsto\sum_{y\sim x}\mathbf{n}_{xy}-\mathbf{n}_{yx}.$$

Thus, $(\partial\mathbf{n})_x$ equals the out-degree minus the in-degree of the vertex x in the random graph $(\mathbf{V},\mathbf{n})$. The function $\partial\mathbf{n}$ is called the *source function*, and if $\partial\mathbf{n}\equiv 0$ then the directed random current $\mathbf{n}$ is called *sourceless*. Using the proposition for the infinite-temperature XY model, we get

$$\int\sigma^a\prod_{xy\in\vec{\mathbf{E}}}(\sigma_x\bar\sigma_y)^{\mathbf{n}_{xy}}\mathrm{d}\sigma=\int\sigma^{a+\partial\mathbf{n}}\mathrm{d}\sigma=\mathbb{1}_{\{a+\partial\mathbf{n}=0\}}.$$

We have now proved the following theorem, illustrated by Figure 6.

Theorem 1.4 (The Poisson expansion of the XY model)**.** *Consider the XY model on a finite graph* $\mathbf{G}$ *with coupling constants* $J\in[0,\infty)^{\mathbf{E}}$*. Then*

$$\forall a\in\mathbb{Z}^{\mathbf{V}},\qquad Z^{\mathrm{XY}}\langle\sigma^a\rangle^{\mathrm{XY}}=\mathbb{M}[\{\partial\mathbf{n}=-a\}].$$

In particular,

$$Z^{\mathrm{XY}}=\mathbb{M}[\{\partial\mathbf{n}=0\}].$$

Thus, we have rewritten the partition and correlation functions of the XY model in terms of independent Poisson random variables, with a condition on the net-degree of the multigraph $(\mathbf{V},\mathbf{n})$ at each vertex.

1.6. **The height function.** Now let $\mathbf{G}=(\mathbf{V},\mathbf{E},\mathbf{F})$ denote a *finite planar graph* (Figure 7); this is just a finite graph $(\mathbf{V},\mathbf{E})$ embedded in the plane $\mathbb{R}^2\cong\mathbb{C}$ such that no two edges cross; $\mathbf{F}$ denotes the set of *faces* including a special *outer face* $f_\infty\in\mathbf{F}$. The *dual edge* xy^* of an edge $xy\in\mathbf{E}$ is the set containing the two faces adjacent to the edge xy. Let $\mathbf{E}^*$ denote the set of dual edges, and let $\mathbf{G}^*:=(\mathbf{F},\mathbf{E}^*,\mathbf{V})$ denote the dual planar graph.

We now define the height function associated to sourceless directed currents on finite planar graphs. The definition, contained in the following theorem, is illustrated by Figure 8.

Theorem 1.5 (The height function of a sourceless current)**.** *Let $\mathbf{n} \in (\mathbb{Z}_{\geq 0})^{\vec{\mathbf{E}}}$ denote a sourceless directed current on the directed edges of the finite planar graph $\mathbf{G} = (\mathbf{V}, \mathbf{E}, \mathbf{F})$. Then there exists a unique function $h_{\mathbf{n}} \in \mathbb{Z}^{\mathbf{F}}$ with the following two properties:*

- *For any directed edge $xy \in \vec{\mathbf{E}}$ having the face $f_r(xy) \in \mathbf{F}$ on its right and the face $f_\ell(xy) \in \mathbf{F}$ on its left, we have*
$$h_{\mathbf{n}}(f_r(xy)) - h_{\mathbf{n}}(f_\ell(xy)) = \mathbf{n}_{xy} - \mathbf{n}_{yx},$$
- *We have $h_{\mathbf{n}}(f_\infty) = 0$.*

This function is called the height function *of the sourceless current* $\mathbf{n}$.

Proof. We first focus on uniqueness of the height function $h_{\mathbf{n}}$. The first property guarantees uniqueness of the *gradient* of $h_{\mathbf{n}}$; the second property fixes the global constant. It suffices to prove existence of the height function $h_{\mathbf{n}}$. We would like to define $h_{\mathbf{n}}(f)$ as the integral of the gradient (given by the first property) along a path from f_∞ to f. For existence, it suffices to prove that this definition is independent of the chosen path. In fact, it suffices to prove that for each vertex $x \in \mathbf{V}$, the path integral of the gradient over the closed path visiting the faces around x yields zero. This requirement is precisely equivalent to the requirement that $(\partial \mathbf{n})_x = 0$, which is equivalent to sourcelessness of $\mathbf{n}$. $\square$

Definition 1.6 (The law on height functions)**.** For fixed $\beta \in [0, \infty)$, let $V_\beta : \mathbb{Z} \to [0, \infty]$ denote the *Poisson potential* defined by
$$V_\beta(k) := -\log \mathbb{P}_\beta[\{A - B = k\}] = 2\beta - \log \sum_{a,b \in \mathbb{Z}_{\geq 0},\, a+b=k} \frac{\beta^{a+b}}{a! \cdot b!},$$
where A and B are independent Poisson random variables of parameter β in the probability measure $\mathbb{P}_\beta$. For any finite planar graph $\mathbf{G} = (\mathbf{V}, \mathbf{E}, \mathbf{F})$ and any family of coupling constants $J \in [0, \infty)^{\mathbf{E}}$, let $\mu^*_{\mathbf{G},J}$ denote the probability measure on height functions $h \in \mathbb{Z}^{\mathbf{F}}$ defined by
$$\mu^*_{\mathbf{G},J}[\{h = \zeta\}] := \frac{\mathbb{1}_{\{\zeta(f_\infty)=0\}}}{Z^*_{\mathbf{G},J}} e^{-H^*_{\mathbf{G},J}(\zeta)},$$
where
$$H^*_{\mathbf{G},J}(\zeta) := \sum_{fg \in \mathbf{E}^*} V_{J_{fg}}(\zeta(g) - \zeta(f)); \qquad Z^*_{\mathbf{G},\beta} := \sum_{\zeta \in \mathbb{Z}^{\mathbf{F}}} \mathbb{1}_{\{\zeta(f_\infty)=0\}} e^{-H^*_{\mathbf{G},J}(\zeta)}.$$
Here we slightly abuse notation by writing $J_{fg} := J_{(fg)^*}$.

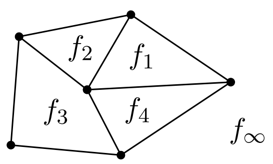

FIGURE 7. A small planar graph with faces $\mathbf{F} = \{f_1, f_2, f_3, f_4, f_\infty\}$.

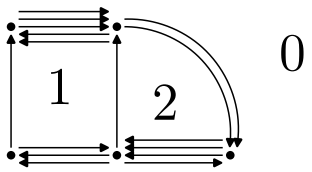

FIGURE 8. The height function $h_{\mathbf{n}} \in \mathbb{Z}^{\mathbf{F}}$ of the sourceless current $\mathbf{n}$.

The function e^{-V_β} is the probability mass function of the difference of two independent Poisson random variables of parameter β. This function e^{-V_β} is log-concave, being the convolution of two log-concave distributions. It is also symmetric around $0 \in \mathbb{Z}$, and therefore attains its minimum at zero. This means that the height function $h \equiv 0$ is the ground state.

Each function V_β is convex and symmetric. We think of $H^*_\beta(h)$ as penalising large gradients for h; it "pulls" heights at neighbouring faces together. When $\beta = 0$ the probability mass function e^{-V_β} equals $\mathbb{1}_{\{0\}}$, which means that this pull is "infinitely strong"; only the ground state $h \equiv 0$ has finite energy. As β increases, the distribution e^{-V_β} becomes more spread out, which means that the favouring of small gradients is softer.

Theorem 1.7 (The XY–heights correspondence)**.** *Let $\mathbf{G} = (\mathbf{V}, \mathbf{E}, \mathbf{F})$ denote a finite planar graph and let $J \in [0,\infty)^{\mathbf{E}}$. Then*

$$\forall \zeta \in \mathbb{Z}^{\mathbf{F}}, \qquad Z^* \cdot \mu^*[\{h = \zeta\}] = \mathbb{M}[\{h_{\mathbf{n}} = \zeta\} \cap \{\partial \mathbf{n} \equiv 0\}].$$

In particular:

- *The two models share the same partition function, that is, $Z^* = Z^{\mathrm{XY}}$,*
- *The law of $h_{\mathbf{n}}$ in $\mathbb{M}[\,\cdot\,|\{\partial\mathbf{n} \equiv 0\}]$ is precisely μ^*.*

Here we use the notation $M[\,\cdot\,|A] := M[(\,\cdot\,)\mathbb{1}_A]/M[A]$ for the measure M conditioned on the event A. This measure is a probability measure, even if M is not.

Proof. Fix ζ. We have an explicit expression for $Z^* \cdot \mu^*[\{h = \zeta\}]$. To find an expression for $\mathbb{M}[\{h_{\mathbf{n}} = \zeta\} \cap \{\partial \mathbf{n} \equiv 0\}]$, one must sum the mass of all currents contributing to the event. At each edge $xy \in \mathbf{E}$, the value of $\mathbf{n}_{xy} - \mathbf{n}_{yx}$ of any contributing current can be expressed directly in terms of ζ. By summing the possible values for the current at each edge, one finds exactly the same expression as for the height functions measure. $\square$

Remark 1.8 (Phase transitions and partition functions)**.** As discussed in the physics introduction, phase transitions are sometimes defined as the points where the pressure $P(\beta)$ (or some derived quantity) fails to be analytic in the parameter β. From this perspective, Theorem 1.7 suggests that in the planar setting, the phase transition of the XY model is identical to that of the dual height function.

One can also define the phase transition as the critical value for β where a qualitative change in the behaviour of the model occurs. Of course, the two definitions of phase transition are believed to coincide almost always, but this is not always easy to make rigorous mathematically.

In fact, in these lectures, we shall prove that the transition in the qualitative behaviour occurs at the same point for the two models, but we *do not* prove that this point coincides with the "partition function transition point". This conjecture remains, to the best knowledge of the author, open.

Remark 1.9 (The XY-height correspondence as a Fourier transform)**.** We view the height function as a Fourier transform of the XY model. There are several reasons for this.

- The spins of the XY model take values in the unit circle. The group $\mathbb{Z}$ of integers is the Fourier dual of the circle. Thus, it is no surprise that the dual model is a model of height functions.
- We observe a form of *temperature inversion*. This means that when the interactions of the XY model are strong (high β), the interactions in the height functions model are weak (the Hamiltonian does not penalise gradients so much). On the other hand, when $\beta = 0$, the spins in the XY model are independent (no interaction), while the gradients along all edges are deterministically forced to be zero. Temperature inversion may be viewed as a form of the Heisenberg uncertainty principle.

Remark 1.10 (The planar Ising model and the self-dual point)**.** One can execute a similar analysis for the Ising model. In this case, the dual model is also an Ising model, since the Fourier dual to the group $\mathbb{Z}/2\mathbb{Z}$ is $\mathbb{Z}/2\mathbb{Z}$. Suppose that the underlying graph $\mathbf{G} = (\mathbf{V}, \mathbf{E}, \mathbf{F})$ is the square lattice graph, which is self-dual. Temperature inversion means that the dual model is an Ising model at some "dual" temperature $T^* := T^*(T)$, which is strictly decreasing in the value of the temperature T of the original Ising model. If one believes that the model undergoes *exactly one* phase transition which can furthermore be read off from the partition function, then this phase transition must necessarily occur at the self-dual point, that is, the unique value for T such that $T^* = T$. Since the relation $T^*(T)$ between the primal and dual temperature is completely transparent, the above yields a very simple way to calculate the exact value of the critical temperature.

There exist several methods to justify the assumptions made above and to demonstrate that the self-dual point coincides with the transition point (Onsager was the first to calculate the precise value of the transition temperature in 1944 [Ons44]). In the case of the XY model we are not so lucky: the dual model is not an XY model (but a height functions model), and therefore there is obviously no special temperature exhibiting self-duality. In fact, although we are going to prove the *existence* of a critical point, its precise value remains unknown. On the other hand, this also makes the model more interesting: for the Ising model, the duality relation implies that the subcritical and the supercritical phase exhibit similar behaviour. This argument does not apply to the XY model, and in fact the subcritical phase $\beta < \beta_c$ and the supercritical phase $\beta > \beta_c$ are fundamentally different.

1.7. **The new approach.** The new approach to Theorem 1 consists of three steps.

(1) First, we prove that the height function undergoes a phase transition. This was already known in the work of Fröhlich and Spencer, but our method is entirely original. We use basic ideas from percolation theory, in place of a perturbative expansion with charged particles.
(2) Second, we prove that the phase transition for height functions is *sharp*. This means that two transition points coincide, namely the transition point where correlations start decaying exponentially fast, and the transition point where the fluctuations of the height function become bounded. This is the first time that the behaviour of the height function is described at and around the transition point.
(3) Third, we use the Brydges–Fröhlich–Spencer walk in order to derive a direct relation between the decay rate (also known as the *mass*) in the XY model, and the decay rate in the dual height function. In particular, the mass of the XY model vanishes if and only if the mass of the height function vanishes. In other words, the phase transitions of the two models coincide. This means that the existence of the BKT transition is now a corollary of the existence of the phase transition of the height function, which was established in the first step.

Several tangential results are obtained *en passant*.

Shortly after the first step appeared online, two teams proved (independently of one another) that the first step implies the BKT transition [AHPS21, EL23]. The second step outlined above is *a priori* unrelated to the BKT transition and to [AHPS21, EL23], as it is exclusively concerned with the phase transition of the height function. The third step outlined above (which is about the relation between the two models) relies on a somewhat different analysis than the ones in [AHPS21, EL23], and may thus be considered an alternative to [AHPS21, EL23]. The advantage of this third step is that it proves the equivalence (beyond the mere existence) of the two phase transitions.

2. Review of the FKG inequality

The purpose of this section is to develop some ideas around the FKG inequality. To keep the discussion light, we sometimes omit mentioning the σ-algebra, in which case its presence is implicit. In those instances, it is implicitly understood that we only consider observables which are measurable. For example, when we work with probability measures supported on countably many elements, then we can just consider the complete σ-algebra on the underlying sample space, and there is not really a need to mention it specifically.

2.1. **The Harris inequality and the FKG inequality.**

Definition 2.1 (FKG inequality)**.** Let $(\Omega, \preceq)$ denote a partially ordered set and μ a probability measure on Ω. We say that μ satisfies the *FKG inequality* if

$$\mu[fg] \geq \mu[f]\mu[g] \tag{3}$$

for any bounded $\preceq$-nondecreasing functions $f, g : \Omega \to \mathbb{R}$.

For simplicity, we often call $\preceq$-nondecreasing functions *increasing* and $\preceq$-nonincreasing functions *decreasing*. Notice that f is increasing if and only if $-f$ is decreasing. The FKG inequality may therefore be formulated in terms of decreasing functions, or in terms of a mixture of increasing and decreasing functions.

Remark 2.2. The FKG inequality is named after Fortuin, Kasteleyn, and Ginibre, because they proved Theorem 2.13 and Corollary 2.14 stated below in 1971 [FKG71]. Thus, it would be most correct to only call Equation (3) the *FKG inequality* if one of those results is used to establish it. However, we shall simply use the term *FKG inequality* for any situation where Equation (3) holds true in order to avoid confusion.

The inequality in Equation (3) had already been obtained by Harris in the context of independent percolation models in 1960 [Har60, Lemma 4.1]. Thus, in the context of independent percolation, Equation (3) is called the *Harris inequality*. We state and prove the Harris inequality below in Theorem 2.6.

The FKG inequality is closely related to *stochastic domination*.

Definition 2.3 (Stochastic domination)**.** Let μ and ν denote two probability measures on some partially ordered set $(\Omega, \preceq)$. We say that μ is *stochastically dominated* by ν, and write $\mu \preceq_{\text{stoch}} \nu$, if for any bounded increasing function $f : \Omega \to \mathbb{R}$, we have $\mu[f] \leq \nu[f]$.

Recall that for any event $E \subset \Omega$ of positive probability, the *conditional expectation* $\mu[\,\cdot\,|E]$ is defined through

$$\mu[\,\cdot\,|E] := \mu[(\,\cdot\,)\mathbb{1}_E]/\mu[E].$$

We call an event E *increasing* or *decreasing* whenever its characteristic function $\mathbb{1}_E$ is increasing or decreasing respectively.

Lemma 2.4. *Let $(\Omega, \preceq)$ denote a partially ordered set and μ a probability measure on Ω satisfying the FKG inequality. Suppose that the events E_+ and E_- are increasing and decreasing respectively, and that both occur with positive probability. Then*

$$\mu[\,\cdot\,|E_-] \preceq_{\text{stoch}} \mu \preceq_{\text{stoch}} \mu[\,\cdot\,|E_+].$$

Proof. Using the definitions of the FKG inequality and of conditional measure, we get that for any bounded increasing function f, we have

$$\mu[f] \leq \mu[f\mathbb{1}_{E_+}]/\mu[\mathbb{1}_{E_+}] = \mu[f|E_+].$$

We conclude that $\mu \preceq_{\text{stoch}} \mu[\,\cdot\,|E_+]$. The other inequality is proved similarly. □

Exercise 2.5 (The FKG inequality on totally ordered sets)**.** Prove that if $(\Omega, \preceq)$ is a totally ordered set (such as $\mathbb{R}$), then any probability measure on it satisfies the FKG inequality.

As partially ordered set $(\Omega, \preceq)$ we often take the set of real-valued functions on some fixed countable index set I, endowed with the natural partial ordering on this set of functions (that is, $\omega \preceq \eta$ whenever $\omega_i \leq \eta_i$ for all $i \in I$).

Theorem 2.6 (Harris inequality [Har60, Lemma 4.1]). *Let I denote a finite index set, and let μ_i denote a measure on $\mathbb{R}$ for each $i \in I$. Then the product measure $\prod_i \mu_i$ on $\mathbb{R}^I$ satisfies the FKG inequality.*

In order to prove this result, we first state a useful property.

Lemma 2.7 (Tower property). *Let $(\Omega, \preceq)$ and $(\tilde{\Omega}, \tilde{\preceq})$ denote two partially ordered sets. Suppose given a probability measure μ on Ω and, for each $\omega \in \Omega$, a probability measure κ_ω on $\tilde{\Omega}$. Assume the following:*

- *μ satisfies the FKG inequality on $(\Omega, \preceq)$,*
- *κ_ω satisfies the FKG inequality on $(\tilde{\Omega}, \tilde{\preceq})$ for any $\omega \in \Omega$,*
- *The map $\omega \mapsto \kappa_\omega$ is increasing, that is,*

$$(\omega \preceq \eta) \implies (\kappa_\omega \preceq_{\mathrm{stoch}} \kappa_\eta) \qquad \forall \omega, \eta \in \Omega.$$

Then the measure $\nu := \mu\kappa := \int d\mu(\omega)\kappa_\omega[\,\cdot\,]$ satisfies the FKG inequality on $(\tilde{\Omega}, \tilde{\preceq})$.

Proof. Consider two bounded increasing functions $f, g : \tilde{\Omega} \to \mathbb{R}$, and write $F, G : \Omega \to \mathbb{R}$ for the functions $\omega \mapsto \kappa_\omega[f]$ and $\omega \mapsto \kappa_\omega[g]$ respectively. Notice that F and G are increasing because the map $\omega \mapsto \kappa_\omega$ is increasing. Now

$$\nu[fg] = \int d\mu(\omega)\kappa_\omega[fg] \geq \int d\mu(\omega)\kappa_\omega[f]\kappa_\omega[g] = \mu[FG] \geq \mu[F]\mu[G] = \nu[f]\nu[g].$$

The first and second inequality are due to the FKG inequality for κ_ω and μ respectively. $\square$

The Harris inequality is an immediate corollary of Exercise 2.5 and the product rule.

Lemma 2.8 (Product rule). *Let μ and μ' denote probability measures satisfying the FKG inequality on the partially ordered sets $(\Omega, \preceq)$ and $(\Omega', \preceq')$ respectively. Then $\mu \otimes \mu'$ satisfies the FKG inequality on the partially ordered set $(\Omega \times \Omega', \preceq \times \preceq')$.*

Proof. For $\omega \in \Omega$, we define $\kappa_\omega := \delta_\omega \otimes \mu'$. The map $\omega \mapsto \kappa_\omega$ is clearly increasing, and the measure κ_ω satisfies the FKG inequality because μ' satisfies the FKG inequality. The tower property implies that $\nu = \mu\kappa = \mu \times \tilde{\mu}$ satisfies the FKG inequality. $\square$

The Harris inequality is often applied in the context of *independent percolation*. In the remainder of this subsection, let $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ denote a simple graph.

Definition 2.9 (Edge percolation measure). An *edge percolation measure* is a probability measure on $\Omega := \{0,1\}^{\mathbf{E}}$ endowed with the product σ-algebra. Given a percolation configuration $\omega \in \Omega$, we call an edge $xy \in \mathbf{E}$ *open* when $\omega_{xy} = 1$ and *closed* when $\omega_{xy} = 0$. We typically identify ω with the set of open edges, and study the random graph $(\mathbf{V}, \omega)$. We say that ω *percolates* if this graph contains an infinite connected component. Write $\omega^{\mathsf{c}} := 1 - \omega$ for the complementary percolation; we can also study the random graph $(\mathbf{V}, \omega^{\mathsf{c}}) = (\mathbf{V}, \mathbf{E} \setminus \omega)$ of closed edges.

The FKG inequality is very useful in the context of percolation measures, because connectivity events such as

$$\{x \leftrightarrow y\} := \{\omega : x \text{ and } y \text{ belong to the same connected component of } (\mathbf{V}, \omega)\}$$

are often increasing. Other interesting events include:

$$\{x \overset{S}{\leftrightarrow} y\} := \{\omega : \text{there exists an } \omega\text{-open path from } x \text{ to } y \text{ through } S\} \qquad \forall S \subset \mathbf{V};$$
$$\{x \leftrightarrow \infty\} := \{\omega : x \text{ belongs to an infinite connected component of } (\mathbf{V}, \omega)\}.$$

Definition (Independent edge percolation model). Consider some function $p \in [0,1]^{\mathbf{E}}$. Then the *independent edge percolation measure* on $\mathbf{G}$ with opening probabilities p is defined as the measure

$$\mu^{\mathrm{edge}}_{\mathbf{G},p} := \prod_{xy \in \mathbf{E}} (p\delta_0 + (1-p)\delta_1).$$

This means that each edge xy is open with a probability p_{xy} and closed with a probability $1-p_{xy}$, independently of the states of all other edges.

Being a product measure, the Harris inequality applies to $\mu^{\mathrm{edge}}_{\mathbf{G},p}$ when $\mathbf{G}$ is finite.

Exercise 2.10. Suppose that $\mathbf{G}$ is a finite connected simple graph, and consider some function $p \in (0,1)^{\mathbf{E}}$. Prove that the map

$$d : \mathbf{V}\times\mathbf{V} \to \mathbb{R},\; xy \mapsto -\log \mu^{\mathrm{edge}}_{\mathbf{G},p}[\{x \leftrightarrow y\}]$$

defines a metric on $\mathbf{V}$. Notice in particular that the Harris inequality implies the triangular inequality for d.

Definition (Site percolation)**.** A *site percolation measure* is a probability measure on $\Omega := \{0,1\}^{\mathbf{V}}$ endowed with the product σ-algebra. Given a percolation configuration $\omega \in \Omega$, we call a site $x \in \mathbf{V}$ *open* when $\omega_x = 1$ and *closed* when $\omega_x = 0$. We typically identify ω with the set of open sites, and study the random graph $(\omega, \omega_{\mathbf{E}})$, where $\omega_{\mathbf{E}}$ denotes the set of edges which are entirely contained in ω. We say that ω *percolates* if this graph contains an infinite connected component. Write $\omega^{\mathsf{c}} := 1-\omega$; we can also study the random graph $(\omega^{\mathsf{c}}, (\omega^{\mathsf{c}})_{\mathbf{E}}) = (\mathbf{V}\setminus\omega, (\mathbf{V}\setminus\omega)_{\mathbf{E}})$ of closed sites. For $p \in [0,1]^{\mathbf{V}}$, define the *independent site percolation measure*

$$\mu^{\mathrm{site}}_{\mathbf{G},p} := \prod_{x\in\mathbf{V}} (p\delta_0 + (1-p)\delta_1).$$

The Harris inequality also applies to independent site percolation on finite graphs. The definitions of

$$\{x \leftrightarrow y\}; \qquad \{x \overset{S}{\leftrightarrow} y\}; \qquad \{x \leftrightarrow \infty\}$$

adapt to site percolation in the obvious way, and these events are obviously increasing in this context as well.

2.2. **The FKG lattice condition.** Fortuin, Kasteleyn, and Ginibre proposed a more general strategy for deriving Equation (3). This strategy is explained in this subsection. We closely follow [FKG71].

Definition 2.11 (Distributive lattices)**.** A *distributive lattice* is a tuple $(\Omega, \preceq, \vee, \wedge)$ where $(\Omega, \preceq)$ is a partially ordered set and where $\vee, \wedge : \Omega\times\Omega \to \Omega$ are binary operators satisfying the following properties for any $x, y, z \in \Omega$:

(1) $x \vee y$ equals the least upper bound of x and y with respect to $\preceq$,
(2) $x \wedge y$ equals the greatest lower bound of x and y with respect to $\preceq$,
(3) The following two *distribution equations*:
  - $x \wedge (y \vee z) = (x\wedge y) \vee (x \wedge z)$,
  - $x \vee (y \wedge z) = (x\vee y) \wedge (x \vee z)$.

The tuple is called *finite* or *countable* whenever Ω has these respective properties.

Definition 2.12 (FKG lattice condition)**.** Let $X : \Omega \to [0,\infty)$ denote a function defined on some distributive lattice $(\Omega, \preceq, \vee, \wedge)$. We say that X satisfies the *FKG lattice condition* if

$$X(\omega\vee\eta)\cdot X(\omega\wedge\eta) \geq X(\omega)\cdot X(\eta) \qquad \forall \omega,\eta \in \Omega.$$

Theorem 2.13. *Let $(\Omega, \preceq, \vee, \wedge)$ denote a finite distributive lattice, and let $X : \Omega \to [0,\infty)$ denote a strictly positive function satisfying the FKG lattice condition. Then the probability measure μ defined by its expectation functional*

$$\mu[f] := \frac{1}{Z}\sum_{\omega\in\Omega} X(\omega) f(\omega); \qquad Z := \sum_{\omega\in\Omega} X(\omega)$$

satisfies the FKG inequality on $(\Omega, \preceq)$.

Proof. Induct on $|\Omega|$: the induction basis is trivial, and we assume that $|\Omega| \geq 2$. The proof of the induction step consists of two parts: first, we decompose the lattice Ω into a number of strictly smaller sublattices, then we use the induction hypothesis (the FKG inequality on these sublattices) and the FKG lattice condition to derive the FKG inequality on Ω.

Let $o \in \Omega$ denote the unique $\preceq$-minimal element, and let a denote some $\preceq$-minimal element in $\Omega \setminus \{o\}$. Define

$$\begin{aligned}
\Omega_+ &:= \{\omega \in \Omega : a \preceq \omega\}; & \mathbb{1}_+ &:= \mathbb{1}_{\Omega_+}; & \mu_+[\,\cdot\,] &:= \mu[\,\cdot\,|\Omega_+];\\
\Omega_- &:= \{\omega \in \Omega : a \not\preceq \omega\}; & \mathbb{1}_- &:= \mathbb{1}_{\Omega_-}; & \mu_-[\,\cdot\,] &:= \mu[\,\cdot\,|\Omega_-];\\
\Omega_-' &:= \{\omega \vee a : \omega \in \Omega_-\}; & \mathbb{1}_-' &:= \mathbb{1}_{\Omega_-'}; & \mu_-'[\,\cdot\,] &:= \mu[\,\cdot\,|\Omega_-'].
\end{aligned}$$

Claim. The following hold true:

(1) The three subsets of Ω are in fact sublattices,
(2) The map $\omega \mapsto \omega \vee a$ is an isomorphism from Ω_- to Ω_-',
(3) The set Ω_-' is a decreasing subset of $(\Omega_+, \preceq)$.

Notice that the induction hypothesis implies the FKG inequality for the expectation functionals on the three smaller lattices.

Proof of the claim. The proof is an elementary exercise involving distributive lattices.

(1) It is straightforward to work out that Ω_+ is a sublattice; it is also straightforward to see that Ω_-' is a sublattice whenever Ω_- is a sublattice. We therefore focus on Ω_-. The fact that o and a are $\preceq$-minimal in Ω and $\Omega \setminus \{o\}$ respectively imply that
$$(a \not\preceq \omega) \iff (\omega \wedge a = o) \qquad \forall \omega \in \Omega.$$
In other words,
$$\Omega_- = \{\omega \in \Omega : \omega \wedge a = o\},$$
which is clearly a sublattice.

(2) The map $\omega \mapsto \omega \vee a$ is a homomorphism of lattices; it suffices to show that it is injective. Suppose that $\omega, \omega' \in \Omega_-$ have the same image. Then
$$\omega = \omega \wedge (\omega \vee a) = \omega \wedge (\omega' \vee a) = (\omega \wedge \omega') \vee (\omega \wedge a) = (\omega \wedge \omega') \vee o = \omega \wedge \omega'.$$
By symmetry of ω and ω', this means that $\omega = \omega'$, thus proving injectivity.

(3) Let $x \in \Omega_-$ and $x' := x \vee a \in \Omega_-'$. Choose $y' \in \Omega_+$ with $y' \preceq x'$; it suffices to prove that $y' \in \Omega_-'$. Notice that
$$y' = y' \wedge x' = y' \wedge (x \vee a) = (y' \wedge x) \vee (y' \wedge a) = (y' \wedge x) \vee a.$$
But $y' \wedge x \in \Omega_-$ which indeed proves that $y' \in \Omega_-'$.

Assertion. For any increasing function f on $(\Omega, \preceq)$ we have $\mu_-[f] \leq \mu_+[f]$.

Proof of the assertion. We shall in fact prove the stronger statement that
$$\mu_+[f] \geq \mu_-'[f] \geq \mu_-[f]. \tag{4}$$

The inequality on the left follows from the FKG inequality for $\mu_+[\,\cdot\,]$. Indeed,
$$\mu_+[f] \geq \mu_+[f|\Omega_-'] = \mu_-'[f].$$
The inequality is the FKG inequality applied to the increasing function f and the decreasing subset Ω_-' of $(\Omega_+, \preceq)$ (see the claim).

Now focus on the right inequality in (4). Since $\Omega_- \to \Omega_-'$, $\omega \mapsto \omega \vee a$ is a bijection, we have
$$\mu_-'[f] = \frac{\sum_{\omega \in \Omega_-} X(\omega \vee a) f(\omega \vee a)}{\sum_{\omega \in \Omega_-} X(\omega \vee a)}.$$

Writing $X(\omega \vee a) = X(\omega)X'(\omega)$ where $X'(\omega) := \frac{X(\omega\vee a)}{X(\omega)}$, we get

$$\mu'_-[f] = \frac{\mu_-[f(\omega \vee a)X']}{\mu_-[X']} \geq \frac{\mu_-[fX']}{\mu_-[X']}.$$

For the inequality in this display we just used that $f(\omega \vee a) \geq f(\omega)$. To conclude that the right hand side equals at least $\mu_-[f]$ we apply the FKG inequality to $\mu_-[\,\cdot\,]$, observing that f is increasing by assumption and that X' is increasing due to the FKG lattice condition. This establishes the assertion.

Let f and g denote increasing functions on $(\Omega, \preceq)$. Our goal is to prove that

$$\Delta := \mu[fg\mathbb{1}_\Omega]\mu[\mathbb{1}_\Omega] - \mu[f\mathbb{1}_\Omega]\mu[g\mathbb{1}_\Omega] \geq 0.$$

Writing $\mathbb{1}_\Omega = \mathbb{1}_+ + \mathbb{1}_-$, this is equivalent to demonstrating nonnegativity of

$$\begin{aligned}\Delta = {} & \mu[fg\mathbb{1}_+]\mu[\mathbb{1}_+] - \mu[f\mathbb{1}_+]\mu[g\mathbb{1}_+] \\ & + \mu[fg\mathbb{1}_-]\mu[\mathbb{1}_-] - \mu[f\mathbb{1}_-]\mu[g\mathbb{1}_-] \\ & + \mu[fg\mathbb{1}_+]\mu[\mathbb{1}_-] - \mu[f\mathbb{1}_+]\mu[g\mathbb{1}_-] \\ & + \mu[fg\mathbb{1}_-]\mu[\mathbb{1}_+] - \mu[f\mathbb{1}_-]\mu[g\mathbb{1}_+].\end{aligned}$$

The FKG inequalities for $\mu_\pm[\,\cdot\,]$ imply that:

- The first and second line are nonnegative,
- For the third and fourth line, we have $\mu[fg\mathbb{1}_\pm] \geq \mu[f\mathbb{1}_\pm]\mu[g\mathbb{1}_\pm]/\mu[\mathbb{1}_\pm]$.

Applying these inequalities and multiplying both sides by $\mu[\mathbb{1}_+]\mu[\mathbb{1}_-]$ yields

$$\Delta\mu[\mathbb{1}_+]\mu[\mathbb{1}_-] \geq (\mu[f\mathbb{1}_+]\mu[\mathbb{1}_-] - \mu[f\mathbb{1}_-]\mu[\mathbb{1}_+])(\mu[g\mathbb{1}_+]\mu[\mathbb{1}_-] - \mu[g\mathbb{1}_-]\mu[\mathbb{1}_+]).$$

Both factors on the right are nonnegative due to the assertion. $\square$

It was quite useful in the previous proof that the lattice was finite and that each of its elements had a positive probability of occurring. It is trivial to remove these requirements *a posteriori*.

Corollary 2.14. *Let $(\Omega, \preceq, \vee, \wedge)$ denote a distributive lattice, and let $X : \Omega \to [0, \infty)$ denote a function satisfying the FKG lattice condition and such that* $\mathrm{Support}(X) := \{X > 0\} \subset \Omega$ *is nonempty and countable, and such that $Z := \sum_{\omega \in \mathrm{Support}(X)} X(\omega)$ is finite. Then the probability measure μ defined by its expectation functional*

$$\mu[f] := \frac{1}{Z} \sum_{\omega \in \mathrm{Support}(X)} X(\omega)f(\omega)$$

satisfies the FKG inequality on $(\Omega, \preceq)$.

Proof. The set $\mathrm{Support}(X)$ is closed under $\vee$ and $\wedge$ because of the FKG lattice condition, and therefore it forms a countable sublattice. Let $(\omega_k)_{k\geq 0}$ denote an enumeration of its elements, let Ω_n denote the finite sublattice generated by $\omega_0, \dots, \omega_n$, and let μ_n denote the probability measure defined by

$$\mu_n[f] := \frac{1}{Z_n} \sum_{\omega \in \Omega_n} X(\omega)f(\omega); \qquad Z_n := \sum_{\omega \in \Omega_n} X(\omega).$$

Now each measure μ_n satisfies the FKG inequality, and since $\mu_n \to \mu$ in the total variation metric, this also implies the FKG inequality for μ. $\square$

We proved the FKG inequality in the context of general distributive lattices. In practice, however, all our distributive lattices $(\Omega, \preceq, \vee, \wedge)$ have the following structure:

- Ω is a subset of $\mathbb{R}^I$ where I is some countable index set,
- $\preceq$ is the restriction to Ω of the standard partial ordering on functions in $\mathbb{R}^I$,
- $\vee$ and $\wedge$ denote the pointwise maximum and minimum operations respectively.

Lemma 2.15 (Basic properties)**.** *Let Ω denote a distributive lattice.*

- *If Ω is totally ordered, then any $X:\Omega\to[0,\infty)$ satisfies the FKG lattice condition.*
- *Let Ω_i denote a distributive lattice and $X_i:\Omega_i\to[0,\infty)$ a function satisfying the FKG lattice condition for $i\in\{1,2\}$. Then for each $i\in\{1,2\}$, the function*

$$\tilde{X}_i:\Omega_1\times\Omega_2,\ (\omega_1,\omega_2)\mapsto X_i(\omega_i)$$

  satisfies the FKG lattice condition on $\Omega_1\times\Omega_2$.
- *If $X,Y:\Omega\to[0,\infty)$ satisfy the FKG lattice condition, then so does XY.*
- *If Ω is a sublattice of $\mathbb{R}^I$, then for any convex function V and for any $i,j\in I$, the function $\omega\mapsto e^{-V(\omega_j-\omega_i)}$, $\Omega\mapsto[0,\infty)$ satisfies the FKG lattice condition.*

The proofs of these basic properties are left to the reader as an exercise.

Lemma 2.16 (Monotonicity in boundary conditions)**.** *Let Ω and Ω' denote two distributive lattices and $X:\Omega\times\Omega'\to[0,\infty)$ a function satisfying the FKG lattice condition. For each $\eta\in\Omega'$, let μ_η be the probability measure on Ω defined by*

$$\mu_\eta[f]:=\frac{1}{Z_\eta}\sum_{\omega\in\Omega,\,X(\omega,\eta)>0}f(\omega)X(\omega,\eta);\qquad Z_\eta:=\sum_{\omega\in\Omega,\,X(\omega,\eta)>0}X(\omega,\eta)$$

whenever $\{\omega\in\Omega:X(\omega,\eta)>0\}$ is countable and $Z_\eta\in(0,\infty)$. Then, for any $\eta,\eta'\in\Omega'$ making μ_η and $\mu_{\eta'}$ well-defined, we have

$$(\eta\preceq\eta')\qquad\Longrightarrow\qquad(\mu_\eta\preceq_{\text{stoch}}\mu_{\eta'}).$$

Proof. Choose η and η' making μ_η and $\mu_{\eta'}$ well-defined and with $\eta\preceq\eta'$. Let $\tilde\Omega:=\Omega\times\{\eta,\eta'\}$. This is a distributive lattice on which X satisfies the FKG lattice condition. Thus, the probability measure μ defined by

$$\mu[f]:=\frac{1}{Z_\eta+Z_{\eta'}}\sum_{(\omega,\eta'')\in\tilde\Omega,\,X(\omega,\eta'')>0}f(\omega,\eta'')X(\omega,\eta'')$$

satisfies the FKG inequality. In particular, Lemma 2.4 implies that for any bounded increasing function $f:\Omega\to\mathbb{R}$, we have

$$\mu_\eta[f]=\mu[f(\omega)|\Omega\times\{\eta\}]\le\mu[f(\omega)]\le\mu[f(\omega)|\Omega\times\{\eta'\}]=\mu_{\eta'}[f].$$

This proves the desired stochastic domination. □

2.3. **The Ising model.** Throughout this subsection, $\mathbf{G}=(\mathbf{V},\mathbf{E})$ denotes a fixed finite simple graph. We also fix a family of *coupling constants* $J\in[0,\infty)^{\mathbf{E}}$. Rather than going for a soft landing, we immediately introduce the Fortuin–Kasteleyn coupling, then derive the Ising model as a marginal of this measure that is initially constructed on the product space.

Definition 2.17 (Fortuin–Kasteleyn coupling)**.** Let $\mu^{\mathrm{FK}}_{\mathbf{G},J}$ denote the probability measure on $(\sigma,\omega)\in\{\pm1\}^{\mathbf{V}}\times\{0,1\}^{\mathbf{E}}$ defined by

$$\mu^{\mathrm{FK}}_{\mathbf{G},J}[\{(\sigma,\omega)=(\bar\sigma,\bar\omega)\}]=\frac{1}{Z^{\mathrm{Ising}}_{\mathbf{G},J}}\mathbb{1}_{\{\bar\sigma\perp\bar\omega\}}(e^J-e^{-J})^{\bar\omega}(e^{-J})^{(1-\bar\omega)}.\tag{5}$$

Here $\bar\sigma\perp\bar\omega$ means that $\bar\sigma$ is constant on each connected component of $(\mathbf{V},\bar\omega)$, and we recall that $a^b:=\prod_{x\in X}a_x^{b_x}$ when a and b are functions defined on some finite set X. The constant $Z^{\mathrm{Ising}}_{\mathbf{G},J}$ is the partition function.

Having defined this coupling, $2+2=4$ questions immediately arise: we can study the two marginals of this product measure, as well as the two conditional measures after conditioning on the other element.

We start with the first marginal. Fix some spin configuration $\sigma \in \{\pm 1\}^{\mathbf{V}}$. In order to calculate the weight of σ, we must sum the weight of the pair (σ, ω) given by Equation (5) over all percolation configurations $\omega \in \{0,1\}^{\mathbf{E}}$ which satisfy $\sigma \perp \omega$. We view this sum as a product over edges in $xy \in \mathbf{E}$. The factor corresponding to each edge $xy \in \mathbf{E}$ is:

$$\mathbb{1}_{\{\sigma_x\sigma_y=+1\}}(e^{J_{xy}} - e^{-J_{xy}}) + e^{-J_{xy}} = e^{J_{xy}\sigma_x\sigma_y}.$$

Indeed, if $\sigma_x\sigma_y = -1$ then the edge is closed so that the only contribution is $e^{-J_{xy}}$, while if $\sigma_x\sigma_y = +1$ the edge can be both open and closed so that both terms contribute. We have now proved the following lemma.

Lemma 2.18 (Definition of the Ising model)**.** *Let $\mu^{\mathrm{Ising}}_{\mathbf{G},J}$ denote the first marginal of $\mu^{\mathrm{FK}}_{\mathbf{G},J}$, that is, the probability measure on $\sigma \in \{\pm 1\}^{\mathbf{V}}$ defined by*

$$\mu^{\mathrm{Ising}}_{\mathbf{G},J}[\{\sigma = \bar\sigma\}] := \mu^{\mathrm{FK}}_{\mathbf{G},J}[\{\sigma = \bar\sigma\}] = \frac{1}{Z^{\mathrm{Ising}}_{\mathbf{G},J}} e^{-H^{\mathrm{Ising}}_{\mathbf{G},J}(\bar\sigma)}; \qquad Z^{\mathrm{Ising}}_{\mathbf{G},J} := \sum_{\sigma\in\{\pm1\}^{\mathbf{V}}} e^{-H^{\mathrm{Ising}}_{\mathbf{G},J}(\sigma)},$$

where $H^{\mathrm{Ising}}_{\mathbf{G},J} : \{\pm 1\}^{\mathbf{V}} \to \mathbb{R}$ is the Hamiltonian *of the Ising model defined by*

$$H^{\mathrm{Ising}}_{\mathbf{G},J}(\sigma) := -\sum_{xy\in\mathbf{E}} J_{xy}\sigma_x\sigma_y.$$

We also write $\langle \cdot \rangle^{\mathrm{Ising}}_{\mathbf{G},J} := \mu^{\mathrm{Ising}}_{\mathbf{G},J}[\,\cdot\,]$ for the expectation functional.

For the second marginal of $\mu^{\mathrm{FK}}_{\mathbf{G},J}$, we observe that σ only appears in the weight of each configuration in Equation (5) through the indicator $\mathbb{1}_{\{\bar\sigma\perp\bar\omega\}}$. Thus, in order to calculate the weight of each configuration ω, we must simply add an extra factor $f(\omega)$ which counts the number of configurations σ that are consistent with ω. It is easy to see that this number is equal to 2^k where k is the number of connected components of the graph $(\mathbf{V}, \omega)$. Indeed, by definition of $\sigma \perp \omega$, we observe that σ can assign precisely two spins to each connected component of this graph, independently of the spin value at other connected components. We have now proved the following lemma.

Lemma 2.19 (Definition of the random-cluster model)**.** *For any $\omega \in \{0,1\}^{\mathbf{E}}$, we let $k(\omega)$ denote the number of connected components of the graph $(\mathbf{V}, \omega)$. For $q \in (0, \infty)$, let $\mu^{\mathrm{RCM}}_{\mathbf{G},J,q}$ denote the probability measure on $\omega \in \{0,1\}^{\mathbf{E}}$ defined by*

$$\mu^{\mathrm{RCM}}_{\mathbf{G},J,q}[\{\omega = \bar\omega\}] := \frac{1}{Z^{\mathrm{RCM}}_{\mathbf{G},J,q}} q^{k(\bar\omega)}(e^J - e^{-J})^{\bar\omega}(e^{-J})^{(1-\bar\omega)}. \tag{6}$$

Then for $q = 2$, the measure $\mu^{\mathrm{RCM}}_{\mathbf{G},J,2}$ is the second marginal of $\mu^{\mathrm{FK}}_{\mathbf{G},J}$, and $Z^{\mathrm{RCM}}_{\mathbf{G},J,2} = Z^{\mathrm{Ising}}_{\mathbf{G},J}$.

Before making some interesting observations about these two measures, we finish answering our four questions by describing the conditional distributions.

Lemma 2.20 (Conditional laws in the Fortuin–Kasteleyn coupling)**.** *The following are true.*

- ***The law of σ conditional on ω.*** *The law of σ in $\mu^{\mathrm{FK}}_{\mathbf{G},J}[\,\cdot\,|\{\omega = \bar\omega\}]$ consists in flipping an independent fair coin for each connected component of $(\mathbf{V}, \bar\omega)$ in order to determine the spin at each vertex.*
- ***The law of ω conditional on σ.*** *For fixed σ, define $p(\sigma) \in [0,1]^{\mathbf{E}}$ by*

$$p(\sigma)_{xy} = \mathbb{1}_{\{\sigma_x\sigma_y=+1\}}\frac{e^{J_{xy}} - e^{-J_{xy}}}{e^{J_{xy}}}. \tag{7}$$

Then the law of ω in $\mu^{\mathrm{FK}}_{\mathbf{G},J}[\,\cdot\,|\{\sigma = \bar\sigma\}]$ is precisely equal to $\mu^{\mathrm{edge}}_{\mathbf{G},p(\sigma)}$.

Proof. The proofs of both statements are already subtly contained in the above reasoning. For the first statement, observe that the conditional law of σ is uniformly random in the set of configurations of $\bar{\sigma} \in \{\pm 1\}^{\mathbf{V}}$ that satisfy $\bar{\sigma} \perp \bar{\omega}$. For the second statement, we already observed above that to compute the weight of a particular configuration for σ in $\mu^{\mathrm{FK}}_{\mathbf{G},J}$, we could decompose over the edges $xy \in \mathbf{E}$. If $\sigma_x\sigma_y = -1$ then that edge is always closed; if $\sigma_x\sigma_y = +1$ then the relative weights of being open and closed are given by $e^{J_{xy}} - e^{-J_{xy}}$ and $e^{-J_{xy}}$ respectively, leading to the opening probability given by Equation (7). $\square$

While the Fortuin–Kasteleyn coupling does not satisfy the FKG inequality, its two marginals and its two conditional measures do satisfy this inequality. This is obvious for the two conditional measures, and therefore we only state precise results for the marginals.

Lemma 2.21 (FKG inequalities)**.** *The following hold true.*

- *The FKG lattice condition is satisfied by* $\{0,1\}^{\mathbf{E}} \to [0,\infty)$, $\omega \mapsto q^{k(\omega)}$ *for* $q \in [1,\infty)$.
- *The measure* $\mu^{\mathrm{Ising}}_{\mathbf{G},J}$ *satisfies the FKG inequality.*
- *The measure* $\mu^{\mathrm{RCM}}_{\mathbf{G},J,q}$ *satisfies the FKG inequality for* $q \in [1,\infty)$.

Proof.

- It is easy to see that it suffices to check the *Holley criterion*, that is, the FKG lattice condition but only for configurations $\omega, \eta \in \{0,1\}^{\mathbf{E}}$ which differ in only two places. The verification of the Holley criterion is a classical exercise that we leave to the reader.
- By the basic properties of the FKG lattice condition, it suffices to prove the FKG lattice condition for $\sigma \mapsto e^{J_{xy}\sigma_x\sigma_y}$. But this map may just be written $\sigma \mapsto e^{-V_{xy}(\sigma_y-\sigma_x)}$ where V_{xy} is the convex function $V_{xy}(a) = J_{xy}(a^2-2)/2$, so that the basic properties yield the result.
- Again, by the basic properties, our strategy is again to decompose the weight on the right in Equation (6) into factors, and to then prove the FKG lattice condition for each factor. For the factor $q^{k(\omega)}$ this follows from the first statement in this lemma; the other factors depend on one edge at a time so that the basic properties yield the result. $\square$

Until now we have treated the two marginals of the Fortuin–Kasteleyn coupling separately. The following lemma uses the coupling to derive an essential relation between those two marginals. For any $A \subset \mathbf{V}$, we write $\sigma_A := \prod_{x\in A}\sigma_x$. Let $A\Delta B$ denote the symmetric difference of the sets A and B. This means that $\sigma_A\sigma_B = \sigma_{A\Delta B}$. Let $\mathcal{E}_A \subset \{0,1\}^{\mathbf{E}}$ denote the event that each connected component of $(\mathbf{V},\omega)$ intersects an even number of vertices in A. Notice that $\mathcal{E}_A$ is an increasing event, and that $\mathcal{E}_A \cap \mathcal{E}_B \subset \mathcal{E}_{A\Delta B}$.

Lemma 2.22. *We have* $\langle\sigma_A\rangle^{\mathrm{Ising}} = \mu^{\mathrm{RCM}}_2[\mathcal{E}_A]$ *for any* $A \subset \mathbf{V}$.

Proof. We calculate the expectation of σ_A in the measure $\mu^{\mathrm{FK}}_{\mathbf{G},J}$. Conditioning on ω, we observe that:

- If $\omega \in \mathcal{E}_A$, then $\sigma_A = 1$ almost surely,
- If $\omega \notin \mathcal{E}_A$, then we may find a connected component of $(\mathbf{V},\omega)$ intersecting an odd number of vertices in A, and in this case the flip symmetry of the coins tells us that the conditional expectation of σ_A is zero.

This implies the lemma. $\square$

This representation of $\langle\sigma_A\rangle^{\mathrm{Ising}}$ allows us to prove some important inequalities.

Theorem 2.23 (Griffiths inequalities)**.** *(1) The following hold true for any* $A, B \subset \mathbf{V}$:

- ***First Griffiths inequality:*** *we have* $\langle\sigma_A\rangle^{\mathrm{Ising}} \geq 0$;
- ***Second Griffiths inequality:*** *we have* $\langle\sigma_A\sigma_B\rangle^{\mathrm{Ising}} \geq \langle\sigma_A\rangle^{\mathrm{Ising}}\langle\sigma_B\rangle^{\mathrm{Ising}}$.

(2) The following hold true for any $K, K' \in [0,\infty)^{\mathbf{E}}$:

- $\langle e^{\sum_{xy\in\mathbf{E}} K_{xy}\sigma_x\sigma_y}\rangle^{\mathrm{Ising}} \geq 1$;

- $\langle e^{\sum_{xy\in\mathbf{E}}(K_{xy}+K'_{xy})\sigma_x\sigma_y}\rangle^{\mathrm{Ising}} \geq \langle e^{\sum_{xy\in\mathbf{E}}K_{xy}\sigma_x\sigma_y}\rangle^{\mathrm{Ising}}\langle e^{\sum_{xy\in\mathbf{E}}K'_{xy}\sigma_x\sigma_y}\rangle^{\mathrm{Ising}}$.

(3) Consider the lattice $[0,\infty)^{\mathbf{E}}$. The map

$$[0,\infty)^{\mathbf{E}} \to [0,\infty),\ J \mapsto Z_J^{\mathrm{Ising}}$$

satisfies, for any $J, K, K' \in [0,\infty)^{\mathbf{E}}$, the following two inequalities:

- $Z_{J+K}^{\mathrm{Ising}} \geq Z_J^{\mathrm{Ising}}$,
- $Z_{J+K+K'}^{\mathrm{Ising}} \cdot Z_J^{\mathrm{Ising}} \geq Z_{J+K}^{\mathrm{Ising}} \cdot Z_{J+K'}^{\mathrm{Ising}}$.

In particular, this map is monotone in J and satisfies the FKG lattice condition.

Proof. (1) For the first Griffiths inequality, simply observe that the expectation may be expressed as a probability through the previous lemma. Now focus on the second Griffiths inequality. We have

$$\begin{aligned}\langle\sigma_A\sigma_B\rangle^{\mathrm{Ising}} = \mu_2^{\mathrm{RCM}}[\mathcal{E}_{A\Delta B}] \geq \mu_2^{\mathrm{RCM}}[\mathcal{E}_A\cap\mathcal{E}_B]& \\ \geq \mu_2^{\mathrm{RCM}}[\mathcal{E}_A]\mu_2^{\mathrm{RCM}}[\mathcal{E}_B] = \langle\sigma_A\rangle^{\mathrm{Ising}}\langle\sigma_B\rangle^{\mathrm{Ising}}&.\end{aligned}$$

The first inequality is due to inclusion of events; the second inequality is the FKG inequality for the random-cluster model with $q = 2$.

(2) Develop the exponentials and compare terms using the Griffiths inequalities.

(3) Apply the previous corollary, observing that $Z_{J+K}^{\mathrm{Ising}} = Z_J^{\mathrm{Ising}}\langle e^{\sum_{xy\in\mathbf{E}}K_{xy}\sigma_x\sigma_y}\rangle_J^{\mathrm{Ising}}$. □

Lemma 2.24. *If $q \geq 1$, then the random-cluster model $\mu_{\mathbf{G},J,q}^{\mathrm{RCM}}$ is monotone in J.*

Proof. Apply Lemma 2.16, observing that the map

$$(\omega, J) \mapsto q^{k(\omega)}(e^J - e^{-J})^{\omega}(e^{-J})^{(1-\omega)}$$

satisfies the FKG lattice condition. □

COMMENT May want to state more refined relations here; see the comments in the TeX file.

2.4. **Boundary conditions and the infinite-volume limit.** Phase transitions usually occur on infinite graphs such as the hypercubic lattice $\mathbb{Z}^d$, but we have only considered finite graphs thus far. Throughout this subsection, $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ denotes a locally finite countable simple graph. We use the symbol $\subset\subset$ to say that a set is a finite subset of another set. A finite subset $\Lambda \subset\subset \mathbf{V}$ is called *domain*. Although the ideas in this subsection apply in a greater generality, we shall only consider models whose spins are associated to the *vertices* $\mathbf{V}$. Thus, we generally consider the lattice $\Omega = \mathbb{R}^{\mathbf{V}}$ endowed with the product σ-algebra $\mathcal{F}$. Write $\mathcal{P}(\Omega, \mathcal{F})$ for the set of probability measures on this measurable space. All measures in this section are viewed as measures in $\mathcal{P}(\Omega, \mathcal{F})$. For any $\Lambda \subset \mathbf{V}$, we write $\mathcal{F}_\Lambda := \sigma(\sigma_x : x \in \Lambda)$. An object is called *locally measurable* if it is measurable with respect to $\mathcal{F}_\Lambda$ for some domain Λ.

2.4.1. *Monotonicity.* The FKG inequality often implies two notions of monotonicity called *monotonicity in boundary conditions* and *monotonicity in domains*. Rather than stating the results in the largest generality possible, we simply state them for the Ising model in a transparent way, making it straightforward to apply similar ideas to other models.

For a fixed domain Λ and for some spin configuration $\eta \in \{\pm 1\}^{\mathbf{V}}$, we let

$$\Omega_{\mathbf{G},\Lambda,\eta} := \{\sigma \in \{\pm 1\}^{\mathbf{V}} : \sigma|_{\mathbf{V}\setminus\Lambda} = \eta|_{\mathbf{V}\setminus\Lambda}\}.$$

Notice that $\Omega_{\mathbf{G},\Lambda,\eta}$ is a countable distributive lattice (in fact, it is even finite), so that the theory developed above applies to it. For $J \in [0,\infty)^{\mathbf{E}}$, we also define

$$H_{\mathbf{G},J,\Lambda}^{\mathrm{Ising}}(\sigma) := -\sum_{xy\in\mathbf{E}(\Lambda)} J_{xy}\sigma_x\sigma_y,$$

where $\mathbf{E}(\Lambda) \subset \mathbf{E}$ denotes the set of edges having at least one endpoint in Λ. We define

$$\mu^{\text{Ising}}_{\mathbf{G},J,\Lambda,\eta}[f] := \frac{1}{Z^{\text{Ising}}_{\mathbf{G},J,\Lambda,\eta}} \sum_{\sigma \in \Omega_{\mathbf{G},\Lambda,\eta}} e^{-H^{\text{Ising}}_{\mathbf{G},J,\Lambda}(\sigma)} f(\sigma) \qquad Z^{\text{Ising}}_{\mathbf{G},J,\Lambda,\eta} := \sum_{\sigma \in \Omega_{\mathbf{G},\Lambda,\eta}} e^{-H^{\text{Ising}}_{\mathbf{G},J,\Lambda}(\sigma)}.$$

We often drop $\mathbf{G}$ and J from the notations. Notice that $H^{\text{Ising}}_{\Lambda}$ satisfies the FKG lattice condition and that $\mu^{\text{Ising}}_{\Lambda,\eta} \in \mathcal{P}(\Omega, \mathcal{F})$ satisfies the FKG inequality.

Remark 2.25. Notice that the definition of $\mu^{\text{Ising}}_{\Lambda,\eta}$ would be the same if we replaced $\mathbf{E}(\Lambda)$ in the definition of the Hamiltonian by a larger, finite set of edges, since those extra terms would only depend on η. Those extra terms would thus be constant, and not contribute to the relative probability of each configuration in $\Omega_{\Lambda,\eta}$.

Lemma 2.26 (Monotonicity in boundary conditions)**.** *For any domain $\Lambda \subset\subset V$ and boundary conditions $\eta, \eta' \in \{\pm 1\}^{\mathbf{V}}$, we have*

$$(\eta \leq \eta') \qquad \implies \qquad (\mu^{\text{Ising}}_{\Lambda,\eta} \preceq_{\text{stoch}} \mu^{\text{Ising}}_{\Lambda,\eta'}).$$

Proof. The proof is essentially identical to that of Lemma 2.16. □

Sometimes there exist special boundary conditions, which are in a certain sense *minimal* or *maximal*. In that case, the FKG inequality may be leveraged to obtain monotonicity in domains.

Lemma 2.27 (Monotonicity in domains)**.** *Let σ^- denote the unique spin configuration with $\sigma^- \equiv -1$. Then for any $\Lambda, \Lambda' \subset\subset \mathbf{V}$, we have*

$$(\Lambda \subset \Lambda') \qquad \implies \qquad (\mu^{\text{Ising}}_{\Lambda,\sigma^-} \preceq_{\text{stoch}} \mu^{\text{Ising}}_{\Lambda',\sigma^-}).$$

Proof. Notice that $\mu^{\text{Ising}}_{\Lambda,\sigma^-} = \mu^{\text{Ising}}_{\Lambda',\sigma^-}[\,\cdot\,|\Omega_{\Lambda,\sigma^-}]$. Since the conditioning event is a decreasing subset of $\Omega_{\Lambda',\sigma^-}$, the FKG inequality for $\mu^{\text{Ising}}_{\Lambda',\sigma^-}$ immediately implies the result. □

2.4.2. *Abstract properties of finite-domain measures.* The objective of the remainder of this section is to construct and describe the infinite-volume limit of families of finite-domain measures.

Definition 2.28 (Properties of families of finite-domain measures)**.** Consider a family $(\mu_\Lambda)_{\Lambda \subset\subset \mathbf{V}} \subset \mathcal{P}(\Omega, \mathcal{F})$ of finite-domain measures.

- **Monotonicity in the domain.** We call the family *monotone in the domain* if
$$\forall \Lambda, \Lambda' \subset\subset \mathbf{V}, \qquad (\Lambda \subset \Lambda') \quad \implies \quad (\mu_\Lambda \preceq_{\text{stoch}} \mu_{\Lambda'}).$$
- **Tight.** We call the family *tight* if
$$\forall x \in \mathbf{V}, \qquad \lim_{K \to \infty} \sup_{\Lambda \subset\subset \mathbf{V}} \mu_\Lambda[\{|\sigma_x| \geq K\}] = 0.$$
- **Boundary Markov property.** Consider a domain $\Lambda \subset\subset \mathbf{V}$ and its partition $(\Lambda_i)_i$ into connected components. We say that the family satisfies the *boundary Markov property* if for any Λ and for any family of $\mathcal{F}_{\Lambda_i}$-measurable functions f_i, we have
$$\mu_\Lambda[\textstyle\prod_i f_i] = \prod_i \mu_{\Lambda_i}[f_i].$$

The fourth property is only defined when the underlying graph $\mathbf{G}$ embedded in $\mathbb{R}^d$ in a way that is symmetric under the action of some full-rank lattice $\mathcal{L} \subset \mathbb{R}^d$.

- **Shift-invariance.** Let $\mathcal{L} \subset \mathbb{R}^d$ denote some full-rank lattice. For $\theta \in \mathcal{L}$ and for any observable f, let θf denote the observable defined by $(\theta f)(\sigma) = f(\sigma \circ \theta)$. The family $(\mu_\Lambda)_{\Lambda \subset\subset \mathbb{Z}^d}$ is called *$\mathcal{L}$-invariant* or *shift-invariant* if $\mu_\Lambda[f] = \mu_{\theta\Lambda}[\theta f]$ for any $\Lambda \subset\subset \mathbb{Z}^d$, for any $\theta \in \mathcal{L}$, and for any observable f.

2.4.3. *The infinite-volume limit: construction.* If a family of finite-domain measures is monotone in the domain and tight, then we can make sense of an infinite-volume limit. However, for this limit to make formal sense, we must first introduce the appropriate topologies on $\mathcal{P}(\Omega, \mathcal{F})$.

Definition 2.29 (The strong topology and the weak topology)**.** An observable f is called *local* if it is $\mathcal{F}_\Lambda$-measurable for some $\Lambda \subset\subset \mathbf{V}$. It is additionally called *continuous* if it may be written $f(\sigma) = f'(\sigma|_\Lambda)$ for some continuous function $f' : \mathbb{R}^\Lambda \to \mathbb{C}$.

- **The strong topology.** The *strong topology* is the coarsest topology on $\mathcal{P}(\Omega, \mathcal{F})$ making the map $\mu \mapsto \mu[f]$ continuous for any bounded local observable f.
- **The weak topology.** The *weak topology* is the coarsest topology on $\mathcal{P}(\Omega, \mathcal{F})$ making $\mu \mapsto \mu[f]$ continuous for any bounded continuous local observable f.

In practice, when a sequence $(\mu_n)_n$ tends to some limit μ in the weak topology, it can often also be proved that the limit tends to the same limit in the strong topology. The distinction between the two topologies is not really of great interest to us.

Theorem 2.30 (Existence of the infinite-volume limit)**.** *If a family of finite-domain measures $(\mu_\Lambda)_\Lambda$ is monotone in the domain and tight, then there is a unique weak limit point $\mu \in \mathcal{P}(\Omega, \mathcal{F})$ such that $\mu_\Lambda \to \mu$ as $\Lambda \nearrow \mathbf{V}$, that is, such that $\mu_{\Lambda_n} \to \mu$ in the weak topology for any sequence $(\Lambda_n)_n$ such that any vertex $x \in \mathbf{V}$ belongs to Λ_n for sufficiently large n.*

Proof. Let $\Delta \subset\subset \mathbf{V}$ denote a fixed domain. Claim that the restriction of μ_Λ to $\mathcal{F}_\Delta$ converges weakly to some probability measure ν_Δ on $\mathbb{R}^\Delta$ as $\Lambda \nearrow \mathbf{V}$. The claim implies the theorem as follows: if the claim is true, then the family $(\nu_\Delta)_\Delta$ extends to a unique measure $\nu \in \mathcal{P}(\Omega, \mathcal{F})$ by the Kolmogorov extension theorem, and we may choose $\mu = \nu$.

If f is any bounded increasing $\mathcal{F}_\Delta$-measurable observable, then $\mu_\Lambda[f]$ is increasing in Λ, and therefore tends to some limit $\nu_\Delta[f]$. It is a straightforward exercise to see that the family $(\nu_\Delta[f])_f$ extends to a probability measure ν_Δ on $\mathbb{R}^\Delta$, and that $\mu_\Lambda[g] \to \nu_\Delta[g]$ as $\Lambda \nearrow \mathbf{V}$ for any bounded continuous $\mathcal{F}_\Delta$-measurable observable g. $\square$

2.4.4. *The infinite-volume limit: ergodicity.*

Definition 2.31 (Shift-invariance and ergodicity)**.** Suppose that the graph $\mathbf{G}$ is embedded in $\mathbb{R}^d$ in a way that is symmetric under the action of some full-rank lattice $\mathcal{L} \subset \mathbb{R}^d$. Let $\mu \in \mathcal{P}(\Omega, \mathcal{F})$ denote a probability measure.

- **Shift-invariance.** The measure μ is called *shift-invariant* or *$\mathcal{L}$-invariant* if $\mu[f] = \mu[\theta f]$ for any observable f and for any $\theta \in \mathcal{L}$.
- **Ergodicity.** The measure μ is called *ergodic* or *$\mathcal{L}$-ergodic* if it satisfies a zero-one law on any event A which is $\mathcal{L}$-invariant in the sense that $\mathbb{1}_A(\sigma) = \mathbb{1}_A(\sigma \circ \theta)$ for any $\sigma \in \Omega$ and $\theta \in \mathcal{L}$.

Theorem 2.32 (Ergodicity of the infinite-volume limit)**.** *Suppose that a family of finite-domain measures $(\mu_\Lambda)_\Lambda$ has the four properties of Definition 2.28 (that is, it is monotone in the domain, tight, has the boundary Markov property, and is $\mathcal{L}$-invariant). Then there is a unique weak limit point $\mu \in \mathcal{P}(\Omega, \mathcal{F})$ such that $\mu_\Lambda \to \mu$ as $\Lambda \nearrow \mathbf{V}$, and this limit point μ is $\mathcal{L}$-ergodic.*

TODO Add remark about adding i.i.d. randomness.

Proof. Existence of the limit point μ was proved in theorem 2.30, and this limit is $\mathcal{L}$-invariant by symmetry. It suffices to prove the zero-one law for shift-invariant events. Without loss of generality, $\mathbf{G}$ is the square lattice graph $\mathbb{Z}^d$.

Fix a domain $\Delta_0 \subset\subset \mathbf{V}$, an increasing $\mathcal{F}_{\Delta_0}$-measurable event E_0, and a nontrivial shift $\theta \in \mathcal{L} \setminus \{0\}$. Define $\Delta_n := \theta^n \Delta_0$, and let E_n denote the increasing $\mathcal{F}_{\Delta_n}$-measurable event such that $\mathbb{1}_{E_n}(\sigma) = \mathbb{1}_{E_0}(\sigma \circ \theta^n)$. Let $p := \mu[E_0]$, and notice that $\mu[E_n] = p$ for any $n \in \mathbb{Z}$ due to shift-invariance.

We first claim that the distribution of $(\mathbb{1}_{E_0}, \mathbb{1}_{E_n})$ in μ tends to the distribution of two *independent* Bernoulli random variables (with parameter p) as $n \to \infty$. To see that this is true, let $B_m := [-m, m]^d \cap \mathbf{V}$, and let $p_m := \mu_{B_m}[E_0]$. Fix $m_0 \in \mathbb{Z}_{\geq 1}$ so large that $\Delta_0 \subset B_{m_0}$. Notice that:

- By monotonicity in domains, we have $p_m \to p$ as $m \to \infty$,
- By monotonicity in domains and the boundary Markov property, we have, for $m_0 \leq m < n/8$,
$$\mu[E_0 \cap E_n] \geq \mu_{B_m \cup \theta^n B_m}[E_0 \cap E_n] = p_m^2,$$
- Similarly, for $m_0 \leq m < n/8$,
$$\mu[E_0 \cup E_n] \geq \mu_{B_m \cup \theta^n B_m}[E_0 \cup E_n] = 1 - (1 - p_m)^2.$$

Thus, we get

$$\liminf_{n\to\infty} \mu[E_0 \cap E_n] \geq p^2; \qquad \liminf_{n\to\infty} \mu[E_0 \cup E_n] \geq 1 - (1 - p)^2.$$

But since $\mu[E_0] = \mu[E_n] = p$, it is immediate that the limits in this display exist and that the inequalities are in fact equalities. This proves the claim.

In fact, by similar reasoning, it is easy to see that for fixed $n \in \mathbb{Z}_{\geq 1}$, the distribution of $(\mathbb{1}_{E_0}, \mathbb{1}_{E_k}, \ldots, \mathbb{1}_{E_{(n-1)k}})$ tends to the distribution of n independent Bernoulli trials with parameter p as $k \to \infty$.

In order to derive a contradiction, we suppose that A is an $\mathcal{L}$-invariant event such that $0 < \mu[A] < 1$. Then $\mu[\,\cdot\,|A]$ is well-defined and not equal to μ. In particular, one may choose Δ_0 and E_0 such that $p' := \mu[E_0|A] \neq \mu[E_0] =: p$. Since A is $\mathcal{L}$-invariant, this means that $\mu[E_k|A] = p'$ for all k. This is clearly inconsistent with our knowledge on the asymptotic distribution of $(\mathbb{1}_{E_0}, \mathbb{1}_{E_k}, \ldots, \mathbb{1}_{E_{(n-1)k}})$. More precisely, if we have an arbitrary number of independent Bernoulli trials with fixed parameter p, then we cannot change the conditional expectation of all those trials to some fixed $p' \neq p$ by conditioning on a single event of fixed probability $\mu[A]$. □

2.4.5. *The infinite-volume limit: FKG inequality.*

Theorem 2.33 (FKG inequality of the infinite-volume limit)**.** *Suppose given a family of finite-domain measures $(\mu_\Lambda)_\Lambda$ which all satisfy the FKG inequality. Suppose moreover that $\mu \in \mathcal{P}(\Omega, \mathcal{F})$ is some weak subsequential limit of this family as $\Lambda \nearrow \mathbf{V}$. Then μ satisfies the FKG inequality.*

Proof. Notice first that $\mu[fg] \geq \mu[f]\mu[g]$ if f and g are bounded increasing continuous local functions by definition of the weak topology and the FKG inequality for each finite-domain measure μ_Λ. The dominated convergence theorem implies that this inequality extends to bounded increasing local functions.

Using the dominated convergence theorem again, it is easy to see that this inequality extends to bounded increasing functions which may be written as pointwise limits of bounded increasing local functions. This already includes all bounded increasing functions that we typically encounter: for example, the observable

$$\mathbb{1}_{\{\text{the subgraph induced by } \{\sigma \geq 1\} \text{ contains an infinite cluster}\}}$$

is of this type. To prove that the inequality may be extended to all bounded increasing functions one should apply Strassen's theorem [Str65], we also refer to [GHM01, Theorem 4.6] for details. □

2.4.6. *Application to the Ising model.* Consider $\mathbf{G} = \mathbb{Z}^d$, and let $J \equiv \beta \in [0,\infty)$. Consider

$$\mu_\Lambda := \mu^{\text{Ising}}_{\mathbf{G},\beta,\Lambda,\sigma^-}.$$

It is easy to see that the family of finite-domain measures $(\mu_\Lambda)_\Lambda$ satisfies the four criteria of Definition 2.28; in particular, tightness is obvious because each spin belongs to $\{\pm 1\}$, and the boundary Markov property is easy to establish once observing that each interaction term in the Hamiltonian involves the value of σ at the two endpoints of an edge. In particular, $\mu_\Lambda \to \mu^-$ in the weak topology for some $\mathbb{Z}^d$-ergodic measure $\mu^- \in \mathcal{P}(\Omega,\mathcal{F})$ as $\Lambda \nearrow \mathbf{V}$. Moreover, since each spin takes values in the discrete set $\{\pm 1\}$, it is easy to see that the convergence in fact occurs in the strong topology. Each measure μ_Λ satisfies the FKG inequality, and therefore μ^- satisfies the FKG inequality as well.

2.5. **Percolation: uniqueness of the infinite cluster.** We now explain the *Burton–Keane argument.* The Burton–Keane argument is a robust argument in statistical mechanics. It does not rely on the FKG inequality, but rather on another notion called *insertion tolerance.* We shall phrase the argument in terms of site percolation.

Definition 2.34 (Insertion tolerance)**.** Consider a site percolation measure μ. We say that μ has *insertion tolerance* if, for any domain $\Lambda \subset\subset \mathbf{V}$ and for any $A \in \mathcal{F}_{\mathbf{V}\setminus\Lambda}$ of positive probability,

$$\mu[A \cap \{\text{all sites in } \Lambda \text{ are open}\}] > 0.$$

Theorem 2.35 (Burton–Keane)**.** *Suppose that* $\mathbf{G}$ *is a connected locally finite simple graph embedded in* $\mathbb{R}^d$ *in a way that is invariant under the action of some lattice* $\mathcal{L} \subset \mathbb{R}^d$ *of rank* d*. If* μ *is an* $\mathcal{L}$*-ergodic site percolation measure on* $\mathbf{G}$ *with insertion tolerance, then the number of infinite connected components is almost surely equal to either zero or one.*

Proof. Without loss of generality, $\mathbf{G}$ is the square lattice graph $\mathbb{Z}^d$. By ergodicity, we may find some $N \in \{\infty, 0, 1, \dots\}$ such that the number of infinite connected components is almost surely equal to N. Our goal is to prove that $N \leq 1$.

Suppose first that $1 < N < \infty$. Let $B_n := [-n,n]^d \cap \mathbf{V}$, and let ∂B_n denote the set of vertices in $\mathbf{V} \setminus B_n$ which are adjacent to B_n. Then for $n \in \mathbb{Z}_{\geq 1}$ sufficiently large, the event A that all infinite connected components intersect ∂B_n has a positive probability. Moreover, this event is $\mathcal{F}_{\mathbf{V}\setminus B_n}$ measurable. By insertion tolerance,

$$\mu[A \cap \{\text{all sites in } B_n \text{ are open}\}] > 0.$$

But $A \cap \{\text{all sites in } B_n \text{ are open}\}$ is included in the event that there exists at most one infinite connected component, contradicting this number is almost surely equal to $N > 1$.

Suppose now that $N = \infty$. Then for $n \in \mathbb{Z}_{\geq 1}$ sufficiently large, the event that ∂B_n intersects at least *three* infinite connected components, is strictly positive. For any shift $\theta \in \mathcal{L}$, write $T(\theta)$ for the event

$$T(\theta) := \{\text{at least three infinite } (\omega \setminus \theta B_n)\text{-components intersect } \theta\partial B_n\} \\ \cap \{\text{all sites in } \theta B_n \text{ are open}\}.$$

If the event $T(\theta)$ occurs then θB_n is called a *trifurcation box*, because it may be thought of a box where an infinite cluster splits into three branches. By insertion tolerance, $p := \mu[T(0)] > 0$, and by shift-invariance, $\mu[T(\theta)] = p > 0$ for any $\theta \in \mathcal{L}$.

In the final part of the proof, we derive a geometric contradiction from the fact that $\inf_{\theta\in\mathcal{L}} \mu[T(\theta)] > 0$. First, let $\mathcal{L}' \subset \mathcal{L}$ be a full-rank sublattice such that for any $\theta \in \mathcal{L}' \setminus \{0\}$, the boxes B_n and θB_n do not have a point in common. For $m \in \mathbb{Z}_{\geq 0}$ very large, let

$$\#B_m := |\{\theta \in \mathcal{L}' : \theta B_n \subset B_m\}|; \qquad \#T_m := |\{\theta \in \mathcal{L}' : \theta B_n \subset B_m \text{ and } T(\theta) \text{ occurs}\}|;$$

Since $\mu[\#T_m] = p \cdot \#B_m$, we know that $\mu[\{\#T_m \geq p \cdot \#B_m\}] > 0$.

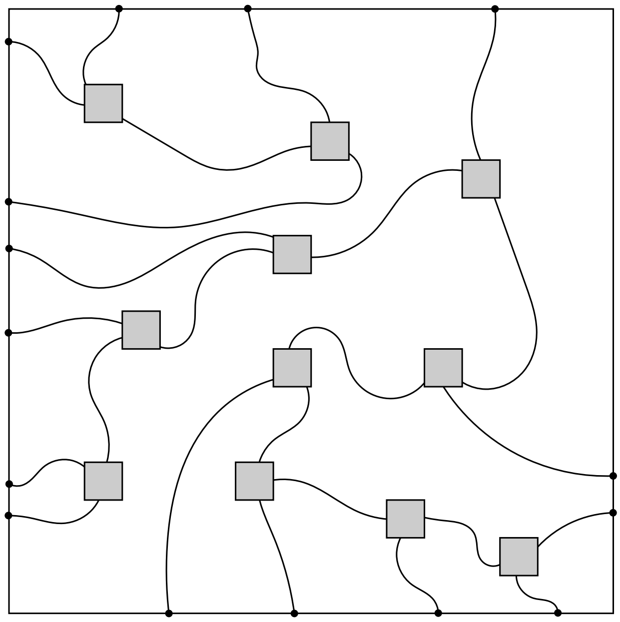

FIGURE 9. The geometric contradiction in the Burton–Keane argument

For the contradiction, we prove that $\{\#T_m \geq p \cdot \#B_m\} = \varnothing$ for m sufficiently large. The argument is illustrated by Figure 9. We may view the boxes $\theta B_n \subset B_m$ such that $\theta \in \mathcal{L}'$ and such that $T(\theta)$ occurs as the vertices of a finite forest (in the graph sense) with a degree of at least three. Finite forests have the property that the number of leafs equals at least the number of vertices having degree three or higher. The leafs, in this case, correspond to the branches of the infinite connected components pointing through the boundary of B_m. Writing ω' for the set ω with all the trifurcation boxes $\theta B_n \subset B_m$ removed, it can be proved rigorously that ω' has at least $\#T_m$ connected components intersecting ∂B_m. In particular, $\#T_m \leq |\partial B_m|$, and therefore $\{\#T_m \geq p \cdot \#B_m\} = \varnothing$ as soon as $|\partial B_m| < p \cdot \#B_m$. But the left hand side grows as m^{d-1} as $m \to \infty$, while the right hand side grows as m^d. This makes the inequality true for m sufficiently large. $\square$

2.6. **Percolation: planar graphs.** The Burton–Keane argument is extremely robust: it works in any dimension and does not even rely on the FKG inequality. The BKT transition is a two-dimensional phenomenon, and one key ingredient for deriving it is a two-dimensional percolation result that we call the *noncoexistence theorem*. The two subsections following this one are dedicated to deriving this theorem under slightly different assumptions.

By a *planar graph*, we mean triple $\mathbf{G} = (\mathbf{V}, \mathbf{E}, \mathbf{F})$ where $(\mathbf{V}, \mathbf{E})$ is a simple graph embedded in the plane $\mathbb{R}^2 \cong \mathbb{C}$ in such a way that no two edges cross. All planar graphs in this text are tacitly assumed to be *locally finite*, meaning that all vertices have a finite degree and that compact subsets of the plane contain finitely many vertices. The letter $\mathbf{F}$ denotes the set of *faces*. If $\mathbf{G}$ is infinite (such as the square lattice graph), then it is possible that $\mathbf{F}$ does not contain a distinguished outer face. The dual graph $\mathbf{G}^* = (\mathbf{F}, \mathbf{E}^*, \mathbf{V})$ is defined as before.

Consider a deterministic site percolation configuration $\omega \in \{0,1\}^{\mathbf{V}}$. Planarity imposes certain constraints on the geometry of the pair $(\omega, \omega^{\mathsf{c}})$. For example, it is impossible that a rectangle is crossed horizontally by ω and vertically by ω^{c}; see Figure 10. We say that this event cannot occur *by planarity* or *due to planarity constraints*. One may think informally of the noncoexistence theorem as an extension of the example in Figure 10 to the infinite volume. Of course, there is no such thing as "crossing an infinite rectangle", and therefore we must slightly change the events under consideration.

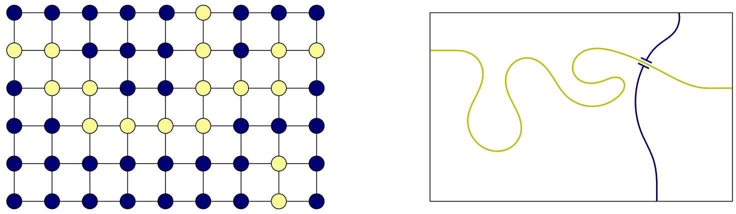

FIGURE 10. LEFT: It is impossible that the rectangle is crossed from left to right by an open (yellow) path, and simultaneously from top to bottom by a closed (blue) path. In this case the horizontal yellow crossing is present, thus barring the existence of a vertical blue crossing. RIGHT: This event is sometimes impressionistically drawn like this.

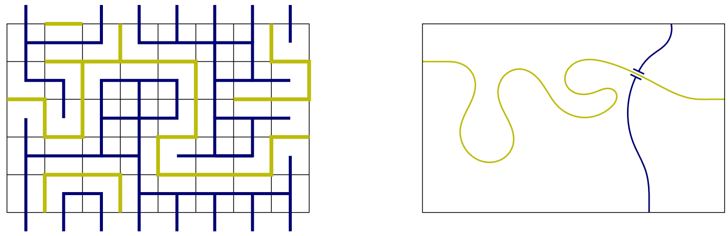

FIGURE 11. Adaptation to edge percolation of Figure 10

Definition 2.36 (Coexistence)**.** We say that two percolation configurations ω and η *coexist* when both percolate (that is, when both contain an infinite connected component). This definitions makes sense for both site- and edge percolation configurations.

While the results are first stated in terms of site percolation, they also apply to edge percolation. The natural dual object to a site ω is the complement ω^{c} of the set of open sites. The natural dual object to an edge percolation ω, is the *dual percolation* $\omega^\dagger$.

Definition 2.37 (Dual edge percolation)**.** For any edge percolation $\omega \in \{0,1\}^{\mathbf{E}}$, let $\omega^\dagger \in \{0,1\}^{\mathbf{E}^*}$ denote the *dual percolation*, defined by $\omega^\dagger = \{xy^* \in \mathbf{E} : xy \notin \omega\}$.

2.7. **Percolation: the noncoexistence theorem in 2D (Zhang's argument).**

Definition 2.38 (Planar graphs with a rotation symmetry)**.** For $\alpha \in \mathbb{R}$, let R_α denote the rotation of the plane $\mathbb{R}^d \cong \mathbb{C}$ by an angle of α. An *R_α-invariant planar graph* is a connected locally finite planar graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ embedded in the plane, and which is invariant under the action of R_α. Just like for shift-invariance, a measure μ is called *R_α-invariant* when $\mu[f(\sigma \circ R_\alpha)] = \mu[f(\sigma)]$ for any observable f.

Theorem 2.39 (Noncoexistence theorem via Zhang's argument)**.** *Consider an R_α-invariant planar graph $\mathbf{G}$ for some $\alpha \in \{\pi/2, 2\pi/3\}$. There does not exist an R_α-invariant site percolation measure μ such that:*

- *μ-almost surely, there is exactly one open infinite cluster,*
- *μ-almost surely, there is exactly one closed infinite cluster,*
- *μ satisfies the FKG inequality.*

Example 2.40 (Independent percolation)**.** The theorem is illustrated by the following application. Let $\mathbf{G}$ denote the triangular lattice, and let μ denote the percolation measure which flips an independent fair coin for each vertex to decide if it is open or closed. See Figure **??** for a sample from this measure. In this figure, the hexagons having an open vertex at its centre are coloured blue; the hexagons corresponding to closed vertices are yellow. The measure μ is ergodic and satisfies the FKG inequality (see Theorems 2.32

TODO Add figure.

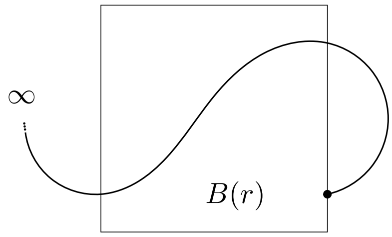

FIGURE 12. The event $C(\text{open}, \text{Left})$ in the proof of Theorem 2.39

and 2.33). The Burton–Keane argument implies that there is at most one infinite open component, and at most one infinite closed component. Since open and closed vertices play the same role, there are two possibilities:

- There is a one infinite open component and one infinite closed component,
- There are no infinite components of either type.

The noncoexistence theorem (Theorem 2.39) rules out the first possibility. In Figure **??**, this means that all clusters are finite. Thus, every blue cluster is surrounded by yellow hexagons, and vice versa.

TODO Add figure.

Proof of Theorem 2.39. For notational simplicity, we suppose that $\mathbf{G}$ is the square lattice graph $\mathbb{Z}^2$ and that $\alpha = \pi/2$. We assume the existence of μ and aim for a contradiction. The main ingredients of the proof are:

- The planarity says that open and closed clusters do not cross (see below for details),
- The FKG inequality,
- The *square root trick*, which says that for any family of increasing events $(A_k)_{1 \leq k \leq n}$,

$$\max_k \mu[A_k] \geq 1 - \sqrt[n]{1 - \mu[\cup_k A_k]}.$$

The square root trick is an immediate corollary of the FKG inequality.

Fix $\varepsilon := 1/7$. For fixed $r \in \mathbb{Z}_{\geq 0}$, we define:

$$B(r) := [-r, r]^2 \cap \mathbb{Z}^2 \subset \mathbf{V}; \qquad \partial B(r) := \{x \in \mathbf{V} \setminus B(r) : x \text{ is adjacent to } B(r)\}.$$

The set $\partial B(r)$ may naturally be written as a disjoint union of its four *sides*, which we denote Right, Top, Left, and Bottom in the obvious way. Write Sides for the set containing those four sets.

We now proceed as follows.

- By existence of the infinite clusters, we may choose r so large that the infinite open and the infinite closed cluster intersect $B(r)$ with a probability of at least $1 - \varepsilon^4$. Define

$$I(\text{open}) := \{\partial B(r) \leftrightarrow \infty\};$$
$$I(\text{closed}) := \{\omega : 1 - \omega \in I(\text{open})\}.$$

- For any $S \in \text{Sides}$, write

$$C(\text{open}, S) := \{S \leftrightarrow^{\mathbf{V} \setminus B(r)} \infty\};$$
$$C(\text{closed}, S) := \{\omega : 1 - \omega \in C(\text{open}, S)\},$$

  see Figure 12. Notice that $I(\text{open}) := \cup_S C(\text{open}, S)$. By R_α-invariance of μ, the four events in this union have the same probability, and therefore the square root trick implies that, for any side $S \in \text{Sides}$, we have

$$\mu[C(\text{open}, S)] \geq 1 - \varepsilon = 1 - \tfrac{1}{7}.$$

  Identical reasoning applies to $C(\text{closed}, S)$.

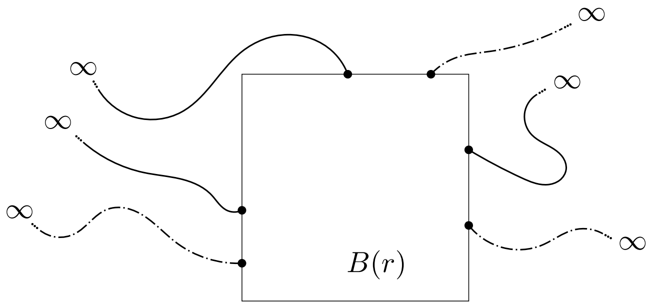

FIGURE 13. The event E in the proof of Theorem 2.39. Open paths are solid, closed paths are dashed. All six paths are infinitely long and avoid $B(r)$. It is impossible that the three open paths belong to a single open cluster, and that simultaneously all closed paths belong to a single closed cluster; this would contradict the planarity of the plane.

- By a union bound, the event
$$E := \cap_{S\in\{\text{Right},\text{Top},\text{Left}\}}(C(\text{open}, S) \cap C(\text{closed}, S))$$
occurs with a probability of at least $1 - \frac{6}{7} > 0$.

However, it is easy to see that if the event E occurs, then there must be at least two infinite open clusters, or at least two infinite closed clusters (see Figure 13). Almost sure uniqueness of the infinite clusters thus implies that $\mu[E] = 0$, contradicting that $\mu[E] > \frac{1}{7}$.

The proof is the same in the case of $\alpha = 2\pi/3$ (for example, when $\mathbf{G}$ is the triangular lattice); the only difference is that one defines $B(r)$ to be a large triangle rather than a large square. ☐

2.8. **Percolation: the noncoexistence theorem in 2D (general argument).** This subsection is included for general culture only. Theorem 2.42 was originally used in the original proof of height function delocalisation in [Lam22], but the simplified proof in these notes relies on the simpler Zhang argument (Theorem 2.39) instead.

Definition 2.41 (Doubly periodic planar graphs)**.** A *doubly periodic planar graph* is a connected locally finite planar graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ embedded in the plane, and which is invariant under the action of some lattice isomorphic to $\mathbb{Z} \times \mathbb{Z}$, see Figure **??**.

TODO Add figure.

Theorem 2.42 (Noncoexistence theorem)**.** *Consider a doubly periodic planar graph* $\mathbf{G}$ *invariant under the action of some lattice* $\mathcal{L} \cong \mathbb{Z} \times \mathbb{Z}$*. There does not exist an* $\mathcal{L}$*-invariant site percolation measure* μ *such that:*

- μ*-almost surely, there is exactly one open infinite cluster,*
- μ*-almost surely, there is exactly one closed infinite cluster,*
- μ *satisfies the FKG inequality.*

Proof. Let μ denote the measure whose existence we aim to contradict. The proof relies on four ingredients:

- Qualitative properties of infinite clusters coming from shift-invariant measures,
- *Duality relations* for box crossings (described below),
- The *intermediate value theorem* (and discrete versions of it),
- The *square root trick*, which says that for any family of increasing events $(A_k)_{1\leq k\leq n}$,
$$\max_k \mu[A_k] \geq 1 - \sqrt[n]{1 - \mu[\cup_k A_k]}.$$

The square root trick is an immediate corollary of the FKG inequality.

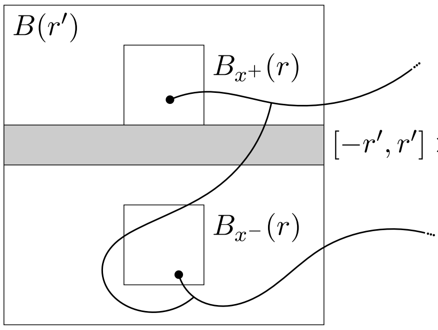

FIGURE 14. Proof of the preliminary observations

For notational simplicity, we suppose that $\mathbf{G}$ is the square lattice graph $\mathbb{Z}^2$ and that $\mathcal{L} = \mathbb{Z}^2$. For any sets $A, B, S \subset \mathbb{R}^2$, let

$$\{A \leftrightarrow B\}; \qquad \{A \leftrightarrow^S B\}; \qquad \{A \leftrightarrow \infty\}$$

denote respectively: the event that there is an open path from a vertex in A to a vertex in B, and the event that there is an open path from A to B which is entirely contained in S, and the event that there is an open path starting from A that visits an infinite number of sites. Let $R = [0, i] \times [0, j]$ denote a rectangle whose four sides are denoted Right, Top, Left, and Bottom. The following *duality relations* are important:

$$\{\text{Top} \leftrightarrow^R \text{Bottom}\} \cap \{\omega : 1 - \omega \in \{\text{Left} \leftrightarrow^R \text{Right}\}\} = \varnothing;$$
$$\{\text{Left} \leftrightarrow^R \text{Right}\} \cap \{\omega : 1 - \omega \in \{\text{Top} \leftrightarrow^R \text{Bottom}\}\} = \varnothing.$$

Define

$$v(i, j) := \mu[\{\text{Top} \leftrightarrow^R \text{Bottom}\}]; \qquad h(i, j) := \mu[\{\text{Left} \leftrightarrow^R \text{Right}\}].$$

Notice that these probabilities would be invariant under shifting the rectangle R. Throughout this proof we write $B_x(r) := [x_1 - r, x_1 + r] \times [x_2 - r, x_2 + r]$, and think of it as a ball of radius $r \in \mathbb{Z}_{\geq 0}$ centred at $x \in \mathbb{Z}^2$. Also set $B(r) := B_{(0,0)}(r)$.

Claim (Preliminary observations)**.** All of the following are true:

- $v(i, j)$ is increasing in i and decreasing in j,
- $h(i, j)$ is decreasing in i and increasing in j,
- Uniqueness of the infinite open cluster implies that
$$\lim_{i\to\infty} v(i, j) = 1; \qquad \lim_{j\to\infty} h(i, j) = 1.$$
- Uniqueness of the infinite closed cluster and the duality relations imply that
$$\lim_{j\to\infty} v(i, j) = 0; \qquad \lim_{i\to\infty} h(i, j) = 0.$$

Proof of the preliminary observations. The first two items are obvious by inclusion of events. We elaborate on the other two by proving that $\lim_{i\to\infty} v(i, j) = 1$. Fix j and $\varepsilon > 0$, we aim to prove that $v(i, j) \geq 1 - 3\varepsilon$ for sufficiently large i. See Figure 14 for an illustration.

Fix r so large that $\mu[\{B(r) \leftrightarrow \infty\}] \geq 1 - \varepsilon$. Set $x^\pm := (0, \pm(j + r))$, and note that

$$\mu[\{B_{x^+}(r) \leftrightarrow B_{x^-}(r)\}] \geq \mu[\{B_{x^+}(r) \leftrightarrow \infty\} \cap \{B_{x^-}(r) \leftrightarrow \infty\}] \geq 1 - 2\varepsilon.$$

On the left, we use that there is one infinite cluster almost surely; if the two balls are connected to infinity, then they must be connected to each other. We may now pick r' so large that

$$\mu[\{B_{x^+}(r) \leftrightarrow^{B(r')} B_{x^-}(r)\}] \geq 1 - 3\varepsilon.$$

But if this event occurs, then clearly the rectangle $[-r', r'] \times [0, j]$ is crossed vertically, proving that $v(2r', j) \geq 1 - 3\varepsilon$. This proves the preliminary observations.

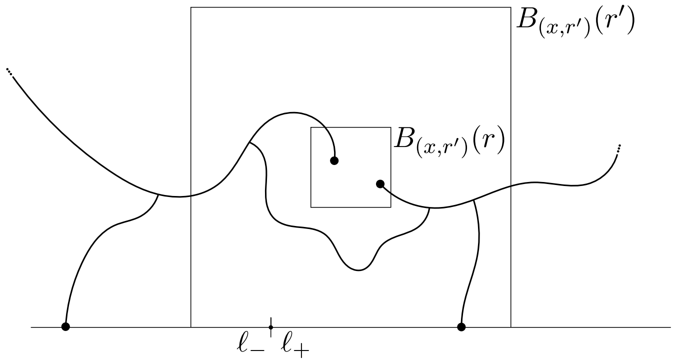

FIGURE 15. Proof of the claim

If R is very tall, then we expect that it is easier to cross R horizontally than vertically. If R is very wide, then we expect the vertical crossing to be easier. Thus, by continuously varying the aspect ratio of R and using the intermediate value theorem, there should be some special point where crossings in the two directions have an equal probability. To formalise this idea, we let j_i denote (for each i) the unique integer such that

$$v(i, j_i) \geq h(i, j_i); \qquad v(i, j_i + 1) < h(i, j_i + 1).$$

The intermediate value theorem and the preliminary observations imply that j_i is well-defined and that $j_i \to \infty$ as $i \to \infty$. The definition of j_i implies immediately that μ satisfies the following statement:

$$\left(\limsup_{i\to\infty} \big(h(i, j_i) \vee v(i, j_i + 1)\big) = 1\right)$$
$$\implies \left(\limsup_{i\to\infty} \Big(\big(v(i, j_i) \wedge h(i, j_i)\big) \vee \big(v(i, j_i + 1) \wedge h(i, j_i + 1)\big)\Big) = 1\right).$$

To conclude the proof of the theorem, we derive the following contradictory assertion.

Assertion. As $i, j \to \infty$ (in the sense that $i \wedge j \to \infty$), we have:
- $h(i, j) \vee v(i, j + 1) \to 1$,
- $h(i, j) \wedge v(i, j) \to 0$.

We only prove the first item; the second item follows by the same proof and the duality relations. In fact, the first statement is slightly stronger than the second, because the height of the vertical crossing is $j + 1$, but this really plays no role in the proof. Thus, we essentially aim to prove that if i and j are large, then the maximum of the horizontal and vertical crossing probabilities of R is close to one. We first prove another claim before proving the assertion.

Claim. Let $\mathbb{H} := \mathbb{R} \times [0, \infty)$ denote the upper half plane. Then

$$\mu[\{\mathbb{H} \leftrightarrow^{\mathbb{H}} \infty\}] = \mu[\{\omega : 1 - \omega \in \{\mathbb{H} \leftrightarrow^{\mathbb{H}} \infty\}\}] = 0.$$

Similar statements hold true if we rotate $\mathbb{H}$ by an angle in $\pi\mathbb{Z}/2$.

Proof of the claim. We prove that the complement $1 - \omega$ of ω does not percolate in the upper half plane; the proof for the primal percolation is the same. Define

$$\ell_- := (-\infty, 0) \times \{0\}; \qquad \ell_+ := [0, \infty) \times \{0\}; \qquad \ell := \mathbb{R} \times \{0\}.$$

If $(1 - \omega) \cap \mathbb{H}$ contains an infinite cluster with positive probability, then each infinite cluster intersects ℓ almost surely. In particular, one such cluster intersects the point $(0, 0) \in \mathbb{Z}^2$ with positive probability. To obtain the desired contradiction, it suffices to prove that $\mu[\{\ell_- \leftrightarrow^{\mathbb{H}} \ell_+\}] = 1$; see Figure 15.

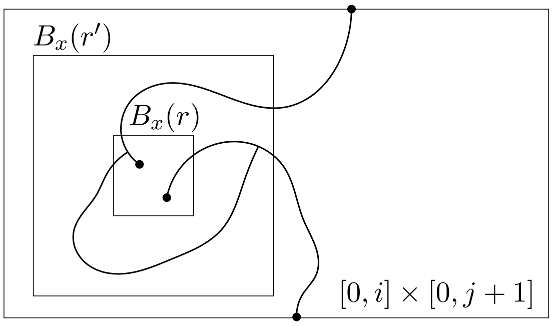

FIGURE 16. Proof of the assertion

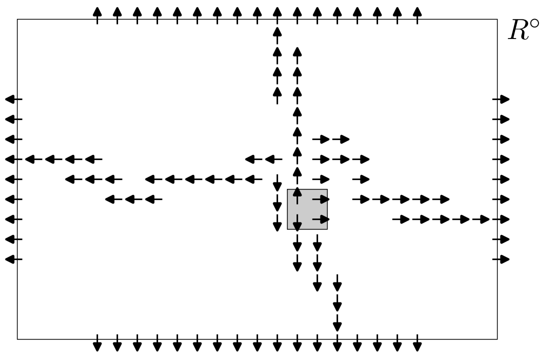

FIGURE 17. The map m and the existence of the square F

Fix $\varepsilon > 0$; we aim to prove that $\mu[\{\ell_- \leftrightarrow^{\mathbb{H}} \ell_+\}] \geq 1 - 3\varepsilon$. Fix $r \geq 2$ so large that $\mu[\{B(r-2) \leftrightarrow \infty\}] \geq 1 - \varepsilon^4$, and fix $r' > r$ so large that

$$\mu[\{\omega \cap B(r') \text{ has two distinct connected components connecting } B(r) \text{ to } \partial B(r'-1)\}] \leq \varepsilon.$$

Such values for r and r' exist because ω contains a unique infinite connected component almost surely. Notice now that for any $x \in \mathbb{Z}$, we have:

- $\mu[\{B_{(x,r')}(r-2) \leftrightarrow^{\mathbb{H}} \ell\}] \geq 1 - \varepsilon^4$,
- $\mu[\{B_{(x,r')}(r-2) \leftrightarrow^{\mathbb{H}} \ell_-\}] \vee \mu[\{B_{(x,r')}(r-2) \leftrightarrow^{\mathbb{H}} \ell_+\}] \geq 1 - \varepsilon$.

The second statement follows from the first using the square root trick. By considering the limits $x \to -\infty$ and $x \to +\infty$ as well as the intermediate value theorem, we deduce that there exists some x such that

$$\mu[\{B_{(x,r')}(r-2) \leftrightarrow^{\mathbb{H}} \ell_-\}],\ \mu[\{B_{(x+1,r')}(r-2) \leftrightarrow^{\mathbb{H}} \ell_+\}] \geq 1 - \varepsilon.$$

By inclusion of events and a union bound, we obtain

$$\mu[\{B_{(x,r')}(r) \leftrightarrow^{\mathbb{H}} \ell_-\} \cap \{B_{(x,r')}(r) \leftrightarrow^{\mathbb{H}} \ell_+\}] \geq 1 - 2\varepsilon.$$

By definition of r', the probability that both events occur but that the two events are not realised by the same connected component is at most ε, which implies that

$$\mu[\{\ell_- \leftrightarrow^{\mathbb{H}} \ell_+\}] \geq 1 - 3\varepsilon.$$

This is the claim.

We now focus on the first item in the assertion. Fix $\varepsilon > 0$; we aim to prove that

$$h(i,j) \vee v(i,j+1) \geq 1 - 3\varepsilon$$

for sufficiently large i and j. The proof is illustrated by Figure 16.

Fix r and r' as in the proof of the previous claim. Define $R^\circ := [r', i-r'] \times [r', j-r']$. By definition of r and the square root trick, there is, for each $x \in R^\circ$, some side I of R such that

$$\mu[\{B_x(r-2) \leftrightarrow^R I\}] \geq 1-\varepsilon.$$

Write $d(\cdot,\cdot)$ for the Euclidean metric on $\mathbb{R}^2$. Define the map

$$m : R^\circ \to \{\text{Right}, \text{Top}, \text{Left}, \text{Bottom}\}$$

such that:

- $m(x)$ satisfies $\mu[\{B_x(r-2) \leftrightarrow^R m(x)\}] \geq 1-\varepsilon$,
- $m(x)$ breaks ties by preferring the side which is the d-closest to x,
- $m(x)$ breaks the remaining ties in an arbitrary way.

If x is close to an edge of R but not to a corner of R, then the claim implies that $m(x)$ chooses the closest edge, see Figure 17. Formally, there exists a constant $n \geq r'$ such that, for sufficiently large i, j, if $d(x, \partial R) \leq r'$ but there is only one side I with $d(x, I) \leq n$, then $m(x) = I$. Now suppose that $i, j \geq 8n$. The intermediate value theorem implies the existence of a face F of the square lattice graph which is contained in R° and such that $m(F)$ contains two opposite sides of R. We suppose that Top and Bottom belong to $m(F)$. Let $x \in F$ denote the bottom-left corner of F. Then by inclusion of events and a union bound we get

$$\mu[\{B_x(r-1) \leftrightarrow^R \text{Bottom}\} \cap \{B_x(r-1) \leftrightarrow^R \text{Top}\}] \geq 1-2\varepsilon,$$

which also implies

$$\mu[\{B_x(r) \leftrightarrow^R \text{Bottom}\} \cap \{B_x(r) \leftrightarrow^{R+(0,1)} \text{Top}+(0,1)\}] \geq 1-2\varepsilon.$$

The probability that the two paths realising the above events without being connected within the box $[0,i] \times [0,j+1]$ is at most ε by definition of r', which leads to $v(i,j+1) \geq 1-3\varepsilon$. If instead Left and Right belonged to $m(F)$, then it can be proved that $h(i,j) \geq 1-3\varepsilon$ by a similar argumentation. $\square$

References

[AHPS21] Michael Aizenman, Matan Harel, Ron Peled, and Jacob Shapiro, *Depinning in integer-restricted Gaussian fields and BKT phases of two-component spin models*, arXiv preprint arXiv:2110.09498 (2021).

[Ber72] V.L. Berezinskii, *Destruction of long-range order in one-dimensional and two-dimensional systems possessing a continuous symmetry group. II. Quantum systems*, Sov. Phys. JETP **34** (1972), no. 3, 610–616.

[Dum17] Hugo Duminil-Copin, *Lectures on the Ising and Potts models on the hypercubic lattice*, PIMS-CRM Summer School in Probability, Springer, 2017, pp. 35–161.

[Dum22] ———, *100 years of the (critical) Ising model on the hypercubic lattice*, Proc. Int. Cong. Math, vol. 1, 2022.

[EL23] Diederik van Engelenburg and Marcin Lis, *An elementary proof of phase transition in the planar XY model*, Communications in Mathematical Physics **399** (2023), no. 1, 85–104.

[FKG71] C.M. Fortuin, P.W. Kasteleyn, and J. Ginibre, *Correlation inequalities on some partially ordered sets*, Communications in Mathematical Physics **22** (1971), 89–103.

[FS81] Jürg Fröhlich and Thomas Spencer, *The Kosterlitz-Thouless transition in two-dimensional abelian spin systems and the Coulomb gas*, Communications in Mathematical Physics **81** (1981), no. 4, 527–602.

[GHM01] Hans-Otto Georgii, Olle Häggström, and Christian Maes, *The random geometry of equilibrium phases*, Phase transitions and critical phenomena, vol. 18, Elsevier, 2001, pp. 1–142.

[Gin70] Jean Ginibre, *General formulation of Griffiths' inequalities*, Communications in mathematical physics **16** (1970), 310–328.

[GS20] Christophe Garban and Avelio Sepúlveda, *Statistical reconstruction of the Gaussian free field and KT transition*, arXiv preprint arXiv:2002.12284 (2020).

[Har60] T.E. Harris, *A lower bound for the critical probability in a certain percolation process*, Mathematical Proceedings of the Cambridge Philosophical Society, vol. 56, Cambridge University Press, 1960, pp. 13–20.

[KP17] Vital Kharash and Ron Peled, *The Fröhlich–Spencer proof of the Berezinskii-Kosterlitz-Thouless transition*, arXiv preprint arXiv:1711.04720 (2017).

[KT73] John Michael Kosterlitz and David James Thouless, *Ordering, metastability and phase transitions in two-dimensional systems*, Journal of Physics C: Solid State Physics **6** (1973), no. 7, 1181–1203.

[Lam22] Piet Lammers, *Height function delocalisation on cubic planar graphs*, Probability Theory and Related Fields **182** (2022), no. 1-2, 531–550.

[Moo69] MA Moore, *Additional evidence for a phase transition in the plane-rotator and classical Heisenberg models for two-dimensional lattices*, Physical Review Letters **23** (1969), no. 15, 861.

[MW66] N David Mermin and Herbert Wagner, *Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models*, Physical Review Letters **17** (1966), no. 22, 1133.

[Ons44] Lars Onsager, *Crystal statistics. I. A two-dimensional model with an order-disorder transition*, Physical Review **65** (1944), no. 3-4, 117.

[Pei36] Rudolf Peierls, *On Ising's model of ferromagnetism*, Mathematical Proceedings of the Cambridge Philosophical Society, vol. 32, Cambridge University Press, 1936, pp. 477–481.

[Sta68] H Eugene Stanley, *Dependence of critical properties on dimensionality of spins*, Physical Review Letters **20** (1968), no. 12, 589.

[Str65] Volker Strassen, *The existence of probability measures with given marginals*, The Annals of Mathematical Statistics **36** (1965), no. 2, 423–439.

Laboratoire de Probabilités, Statistique et Modélisation (LPSM)

Sorbonne Université

Centre national de la recherche scientifique (CNRS)
Email address: plammers@lpsm.paris