BULLETIN (New Series) OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 61, Number 1, January 2024, Pages 55–71
https://doi.org/10.1090/bull/1802
Article electronically published on August 2, 2023

A STROLL AROUND THE CRITICAL POTTS MODEL

MARTIN HAIRER

ABSTRACT. Over the past decade or so, a broad research programme spearheaded by H. Duminil-Copin and his collaborators has vastly increased our understanding of a number of critical or near-critical statistical mechanics models. Most prominently, these include the q-state Potts models and, essentially equivalently, the FK cluster models. In this short review, we present a small selection of recent results from this research area.

CONTENTS

1. INTRODUCTION

One of the simplest, yet extremely rich, models of statistical mechanics is the Ising model, which has historically been introduced as a toy model for the behaviour of ferromagnets. (This model was actually first invented by Wilhelm Lenz in 1920, who then gave it to his student Ernst Ising to study.) The definition of the model goes as follows. Given a finite connected graph G, identified here with its set of vertices, we consider the configuration space $\Omega = \{-1, 1\}^G$ and define on Ω an energy functional $E(\sigma) = -\frac{1}{2}\sum_{x\sim y} \sigma_x\sigma_y$, where $x \sim y$ if and only if the vertices x and y are connected by an edge in G. One should think here of the vertices of G as indexing spatial locations (for example, of individual atoms in a metallic solid) of the graph structure as indicating which locations are neighbours in space, and of σ_x as denoting a spin variable associated to such a location. The energy is then defined in such a way that states of low energy are those where many pairs of neighbouring spins are aligned.

Given an inverse temperature β, one then defines a probability measure μ_β on Ω by setting $\mu_\beta(\{\sigma\}) = Z^{-1}\exp(-\beta E(\sigma))$, where Z is such that $\mu_\beta(\Omega) = 1$. For

Received by the editors April 5, 2023.
2020 *Mathematics Subject Classification.* Primary 82B26, 82B43, 82B20.
Key words and phrases. Ising model, Potts model, percolation.
This work was supported by the Royal Society through a research professorship, grant number RP/R1/191065. Sections 3 and 4 of this review were previously published in the 2022 ICM laudatio for Hugo Duminil-Copin's Fields Medal.

definiteness, when we talk about "the Ising model on G at inverse temperature β", we mean the measure μ_β as just described. The interpretation of the model in terms of spins and atoms suggests that an interesting special case is that where G is a large piece of a lattice, for example $G = \Lambda_N = \{-N, \ldots, N\}^d$ or $G = \mathbf{Z}^d \cap N\mathcal{O}$ for some open set $\mathcal{O} \subset \mathbf{R}^d$ with smooth boundary, with edges between nearest neighbours. Writing μ_β^N for the Ising model on G_N, it turns out that the limit $\mu_\beta = \lim_{N\to\infty} \mu_\beta^N$ exists and can therefore be interpreted as the Ising model on $\mathbf{Z}^d$.

One very interesting qualitative feature of this model is that it exhibits a *phase transition* in every dimension $d \geq 2$: there exists a critical (dimension-dependent) value β_c which delineates two regimes in which the measure μ_β behaves very differently. At *high temperature*, namely for $\beta < \beta_c$, the *spontaneous magnetisation*, namely the random quantity $M = N^{-d} \sum_{i\in\Lambda_N} \sigma_i$, converges to 0 in probability as $N \to \infty$. For $\beta > \beta_c$ on the other hand, it converges in probability to a limiting random variable that can take exactly two possible values $\pm h_\beta \neq 0$ with equal probabilities. The actual value of β_c is only known in dimension 2 where it equals $\beta_c = \log\big(1 + \sqrt{2}\big)$ [Ons44]. (There is no phase transition at all in dimension 1 and the spontaneous magnetisation M always vanishes, so in some sense $\beta_c = +\infty$ there.)

The expression and just mentioned result for the spontaneous magnetisation M has the flavour of a *law of large numbers*, so it is natural to ask whether there is an associated *central limit theorem* describing the fluctuations of the magnetisation. In other words, does the law of the quantity $N^{-d/2} \sum_{i\in\Lambda_N} (\sigma_i - M)$ converge to that of a normal distribution? This is indeed the case when $\beta \neq \beta_c$, but the corresponding variance diverges as $\beta \to \beta_c$. The behaviour *at* the critical temperature is highly nontrivial and it is not even clear at first sight how such an expression should be normalised. In other words, does there exist a value α such that the law of

$$N^{-\alpha} \sum_{i\in\Lambda_N} (\sigma_i - M)$$

admits a nondegenerate limit distribution as $N \to \infty$ when $\beta = \beta_c$? It was shown in a recent series of works [CGN15, CGN16] that if one chooses $\alpha = 15/8$ in dimension $d = 2$, then this is indeed the case. Actually even more was shown there; namely, one can consider the joint distribution of finitely many quantities of the form

$$I_\phi^N(\sigma) = N^{-\alpha} \sum_{x\in\Lambda_N} \phi(x/N)\sigma_x \;, \tag{1.1}$$

for ϕ a smooth test function supported on $[-1, 1]^d$, and these all converge. One way of interpreting this is that there exists a random distribution ζ on the hypercube such that the quantities $I_\phi^N(\sigma)$ all converge jointly in law to the quantities $\zeta(\phi)$.

This time however, unlike in the central limit theorem, the limiting distributions are not Gaussian (the random variables $\zeta(\phi)$ actually exhibit an even faster decaying tail behaviour with $\log \mathbf{P}(\zeta(\phi) > K) \eqsim -K^{16}$ rather than $-K^2$) and no nice closed form expression exists for them. However, there *does* exist a closed form expression for their joint moments, which was first derived heuristically in the physics literature in [BPZ84, Car84, BG93] and was recently made rigorous in [CHI15]. Note that the exponent α appearing in (1.1) is closely related to the behaviour of $\mathbf{E}_c\sigma_u\sigma_v$ (where $\mathbf{E}_c$ denotes the expectation under μ_{β_c}) since, assuming that $\mathbf{E}_c\sigma_u\sigma_v \approx |u - v|^{-2\delta}$, one finds that

$$\begin{aligned}\mathbf{E}_c\big(I_\phi^N(\sigma)\big)^2 &= N^{-2\alpha}\sum_{u,v}\phi(u/N)\phi(v/N)\mathbf{E}_c\sigma_u\sigma_v \\ &\lesssim N^{-2\alpha}\sum_{u,v}|u-v|^{-2\delta} \approx N^{2d-(2\delta\wedge d)-2\alpha}\;,\end{aligned}$$

so that one expects the relation $\alpha = d - (\delta \wedge d/2)$, which (correctly) leads to the prediction $\delta = \frac{1}{8}$. Interestingly, the limiting random distribution ζ exhibits a form of covariance under the action of the conformal groupoid in the following sense. Given any smooth simply connected domain $D \subset \mathbf{R}^2$, one can consider expressions like (1.1), but this time with $\Lambda_N = ND \cap \mathbf{Z}^2$. It turns out that these do again converge, this time to a random distribution ζ_D on the domain D. Given two such domains D and $\bar D$ and a bijective conformal map $\psi\colon D \to \bar D$, the pushforward η of $\zeta_{\bar D}$ to D given by

$$\eta(\phi) = \zeta_{\bar D}(\phi \circ \psi^{-1}) \tag{1.2}$$

is equal in law to the random distribution $\bar\eta$ given by

$$\bar\eta(\phi) = \zeta_D(|\psi'|^{15/8}\phi) = \zeta_D(|\psi'|^{\alpha}\phi)\;, \tag{1.3}$$

where $\alpha = 2 - \delta$ is as above. This and a number of other properties of the Ising model at criticality allows us to associate it to the conformal field theory with central charge $c = \frac{1}{2}$.

The picture in dimensions greater than 2 is less clear. For $d \geq 5$, it was shown in [Aiz81, Aiz82, Frö82] that the correct scaling exponent to use in (1.1) at $\beta = \beta_c$ is $\alpha = 1 + \frac{d}{2}$ and that the limit is a Gaussian free field, namely the Gaussian random distribution with correlation function given by the Green's function of the Laplacian (with Neumann boundary conditions on the square). In dimension $d = 3$, virtually nothing is known rigorously about the critical Ising model, not even the value of its scaling exponents, although much progress has been made at a nonrigorous (but very well supported) level with the development of the *conformal bootstrap* [ESPP+12, ESPP+14]. Regarding the case $d = 4$, it was somewhat unclear until very recently whether the Ising model at criticality should be *trivial* (i.e., described by Gaussian distributions) or not. This was eventually settled by Aizenman and Duminil-Copin in the work [ADC21] where they show that any subsequential limit for expressions of the form (1.1) as $N \to \infty$ (and $\beta \to \beta_c$) must necessarily be Gaussian.

1.1. **A general picture.** The general picture that has been emerging over the past half century or so regarding the behaviour of many statistical mechanics systems can be summarised as follows.

(1) Many of the simplest local equilibrium systems in dimension 2 or higher do exhibit a phase transition, namely there exists a critical value β_c at which the qualitative large scale behaviour of the system changes abruptly. In general, a system may depend on additional parameters in which case one may see a more complicated *phase diagram* with several regions in parameter space where the global behaviour of the system displays qualitatively different behaviour. In any case, the "high temperature / small β phase" is expected to behave in such a way that what happens in well separated regions of space is very close to independent.

(2) In dimension 2, many of these systems appear to exhibit a form of conformal invariance at criticality, even though no rotation symmetry is built a priori into their description. When this happens, the link to two-dimensional conformal field theories (and the associated probabilistic objects like Schramm–Loewner evolution (SLE) [Sch00], Quantum Loewner evolution (QLE) [MS16], etc.) provides a hugely powerful machinery to predict—and in a number of cases also rigorously prove—their behaviour. In the case of the Potts model (see Section 2 for its definition), these links are on a strong rigorous footing for $q \in \{0, 2\}$, but much needs to be done for other values of q.

(3) The universe of local statistical mechanics models can be subdivided into broad classes of models that exhibit a shared large-scale behaviour at criticality. These are called *universality classes* and, in the two-dimensional equilibrium case, they are expected to come in families parametrised by a real parameter, the central charge. (For certain values of the central charge, one expects to have several *subclasses*, but we will not discuss this kind of subtlety here.) In particular, the large-scale behaviour of such models is expected to be very stable under changes in the details of their microscopic description, such as the shape of the grid on which they are defined, the range of their interactions, etc.

(4) Although one still expects conformal invariance at criticality in higher dimensions, this is a much smaller symmetry there and therefore appears to provide somewhat less insight.[1] One also expects the situation there to be more rigid than in two dimensions, with fewer universality classes. (Possibly only a discrete family.)

(5) Models that have *obvious* variants in every dimension typically have a critical dimension above which their behaviour at criticality is trivial in the sense that it exhibits Gaussian behaviour. (Typically, their correlation function is given by the Green's function of the Laplacian when one considers models with a clear underlying spatial structure.) In the case of the Ising universality class, this critical dimension is 4, while in the case of Bernoulli percolation it is 6.

One important branch of modern probability theory aims to put this general picture onto rigorous mathematical footing. The remainder of this article is devoted to a short overview of some of the recent contributions to this vast programme, mainly focusing around the example of the critical Potts model where much recent progress was made by Hugo Duminil-Copin and his collaborators.

2. The Potts and random cluster models

The Potts model is a natural generalisation of the Ising model: this time the configuration space is given by $\Omega = \{1, \ldots, q\}^G$, and the corresponding energy functional is given by $E_q(\sigma) = -\sum_{x\sim y} \mathbf{1}\{\sigma_x = \sigma_y\}$. This is often visualised by interpreting the q values as "colours", so that the effect of the energy functional is to favour configurations where neighbouring vertices tend to have the same colour. We denote by

$$\mathbf{P}_{\beta,q}(\sigma) \propto \exp\bigl(-\beta E_q(\sigma)\bigr) ,$$

[1] See however the recent breakthrough made in the approximation of the critical exponents of the $3d$ Ising model using the conformal bootstrap [ESPP+12, ESPP+14] already mentioned above.

the corresponding Gibbs measure. Note that the case $q = 2$ yields the Ising model, modulo a recentering of the energy (which doesn't affect the measures $\mu_\beta = \mathbf{P}_{\beta,2}$ since they are normalised to be probability measures). For $q \neq 2$, the Potts model does not exhibit the kind of exact solvability that the Ising model does in two dimensions (as discovered by Onsager [Ons44] in his famous computation of its partition function), so that it is one of the simplest possible models of statistical mechanics that isn't known to be exactly solvable.

One important feature of the Potts model is that it is very closely related to a different model, the random cluster or FK model, introduced by Fortuin and Kasteleyn [FK72], which however makes sense for all $q > 0$, not just integer values. This model is usually interpreted as a percolation model, i.e., its state space is given by $\bar\Omega = \{0,1\}^E$, where E denotes the set of edges of the graph G and, given a configuration $\omega \in \bar\Omega$, we say that the edge e is *open* if $\omega_e = 1$ and *closed* otherwise. Given two fixed parameters $p \in (0,1)$ and $q > 0$, the probability of a configuration ω is then proportional to

$$\mathbf{Q}_{p,q}(\omega) \propto p^{|\omega|}(1-p)^{|1-\omega|}q^{|K_\omega|} ,$$

where $|\omega| = \sum_{e\in E} \omega_e$ and K_ω denotes the set of connected components (also called *clusters* in this context) of the subgraph G_ω of G given by replacing the edge set E with the set $E_\omega = \{e : \omega_e = 1\}$ of open edges. (Here an isolated vertex counts as a connected component.)

It turns out (see for example [Gri06, Thm. 1.13]) that given any finite graph G and provided that β and p are related by the identity

$$p = 1 - e^{-\beta} \tag{2.1}$$

(and that q is an integer), one can find a probability measure $\mathbf{P}$ on $\Omega \times \bar\Omega$ with the following properties.

- The marginal of $\mathbf{P}$ on Ω coincides with the Potts model, namely $\mathbf{P}(A\times\bar\Omega) = \mathbf{P}_{\beta,q}(A)$.
- The marginal of $\mathbf{P}$ on $\bar\Omega$ coincides with the random cluster model, namely $\mathbf{P}(\Omega \times A) = \mathbf{Q}_{p,q}(A)$.
- Under $\mathbf{P}$, almost every configuration (σ,ω) is such that for every open edge xy (i.e., such that $\omega_{xy} = 1$), one has $\sigma_x = \sigma_y$.
- Conditional on a configuration σ, the law of ω under $\mathbf{P}$ is obtained by setting the values $\{\omega_{xy} : \sigma_x = \sigma_y\}$ to be i.i.d. Bernoulli random variables with parameter p.
- Conditional on a configuration ω, the law of σ under $\mathbf{P}$ is obtained by assigning to every cluster $A \in K_\omega$ independently a colour $\sigma_A \in \{1,\dots,q\}$ and then setting $\sigma_x = \sigma_A$ for all $x \in A$.

The advantage of the random cluster model is that it exhibits a nice duality in the case when G is a connected planar graph (for example a chunk of the two-dimensional lattice). In that case, one can define a dual graph (G^*, E^*) whose vertex set G^* consists of the faces of the original graph G and such that there is a bijection between E and E^* mapping any edge $e \in E$ to the edge e^* connecting the two faces adjacent to e. (This may generate self-loops since G is allowed to have vertices of degree 1.)

Every configuration ω on E then determines a dual configuration ω^* on E^* by setting $\omega^*_{e^*} = 1 - \omega_e$, where e and e^* are related as just described. See Figure 1 for an example of a configuration ω on a chunk of the square lattice, as well as the

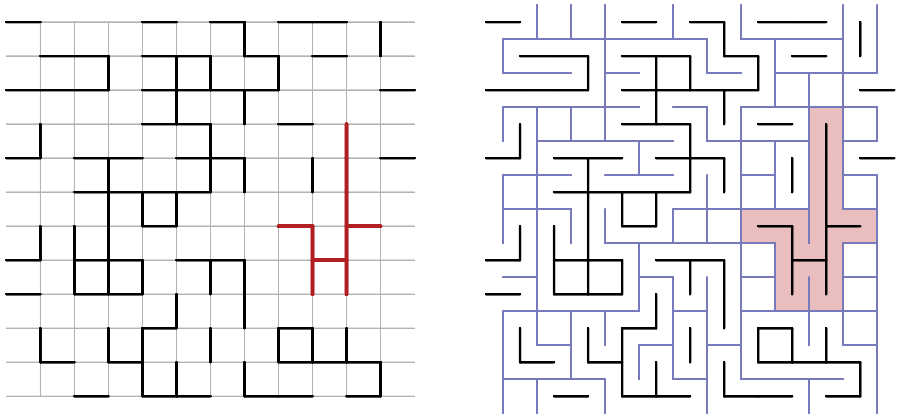

FIGURE 1. On the left, we draw a configuration ω for the random cluster model with $N = 11$, with one of the clusters highlighted in red. On the right, the same configuration is drawn together with its dual configuration in light blue. The face of the dual configuration corresponding to the cluster is shaded in light red.

corresponding dual configuration. Write $\mathbf{Q}_{p,q}^*$ for the pushforward of the measure $\mathbf{Q}_{p,q}$ under the map $\omega \mapsto \omega^*$. One then has the following result.

Proposition 2.1. *The measure* $\mathbf{Q}_{p,q}^*$ *coincides with the random cluster model on* G^* *with parameters* (p^*, q)*, where* p^* *is given by*

$$p^* = \frac{q - pq}{p + q - pq} \;.$$

Proof. Recall that, given any configuration ω, G_ω is the (planar) subgraph of G obtained by only retaining the *open edges* $E_\omega = \{e : \omega_e = 1\}$. The proof is then based on two remarks. First, writing F_ω for the set of faces of G_ω (with the usual convention that there is an infinite outer face) and K_ω for the set of its connected components, we note that one has the identity

$$|G| + |F_\omega| = 1 + |E_\omega| + |K_\omega| \;.$$

(This variant of the Euler characteristic formula is true for any planar graph and can easily be shown by induction over the number of vertices and edges. The reason why we have G appearing there is to emphasise that the vertex set of the graph G_ω is independent of the configuration ω, which will be important in what follows.) The second remark relates the graph G_ω to the subgraph $G^*_{\omega^*}$ of G^* generated by the configuration dual to ω.[2] One can see that connected components of G_ω are then in one-to-one correspondence with faces of $G^*_{\omega^*}$; see Figure 1 for an illustration of this fact. In other words, one has the identity $|K_\omega| = |F^*_{\omega^*}|$.

[2]Note that $G^*_{\omega^*}$ is very different from the dual graph of G_ω.

Using this correspondence and the fact that $|E_\omega| + |E^*_{\omega^*}| = |E|$ by definition of the dual configuration, it then follows that

$$\begin{aligned}\mathbf{Q}^*_{p,q}(\omega^*) &\propto p^{|\omega|}(1-p)^{|1-\omega|}q^{k(\omega)} \propto \big(p/(1-p)\big)^{|E_\omega|}q^{|K_\omega|}\\ &= \big(p/(1-p)\big)^{|E|-|E^*_{\omega^*}|}q^{|F^*_{\omega^*}|} \propto \big((1-p)/p\big)^{|E^*_{\omega^*}|}q^{1+|K^*_{\omega^*}|+|E^*_{\omega^*}|-|G^*|}\\ &\propto \big(q(1-p)/p\big)^{|E^*_{\omega^*}|}q^{|K^*_{\omega^*}|} = \big(p^*/(1-p^*)\big)^{|E^*_{\omega^*}|}q^{|K^*_{\omega^*}|} \propto \mathbf{Q}_{p^*,q}\ ,\end{aligned}$$

which is precisely the desired claim. □

Since the square lattice is self-dual, this leads to the natural conjecture that the critical value of p for the random cluster model on $\mathbf{Z}^2$ is given by the (unique) value $p_c \in (0,1)$ such that $p_c^* = p_c$, namely

$$p_c^2(q-1) - 2p_c q + q = 0 \qquad \Rightarrow \qquad p_c = \frac{q-\sqrt{q}}{q-1} = 1 - \frac{1}{1+\sqrt{q}}\ .$$

Thanks to (2.1) and the close link between the random cluster model and the Potts model, this motivates the following recent result [BDC12].

Theorem 2.2. *The critical inverse temperature for the q-colour Potts model is given by* $\beta_c = \log(1+\sqrt{q})$.

Of course the analogous result also holds for the random cluster model. In the remainder of this article, we describe several recent results for the random cluster and Potts models at criticality. Our main focus is on the two-dimensional case, but we will see that one important result is the continuity of the phase transition in dimension 3.

3. (Dis)continuity of phase transitions

One very natural question in statistical mechanics is whether one can take the limit $N \to \infty$ for the finite volume Gibbs measures. At this stage, we note that there are actually several inequivalent natural ways in which one can define the Ising or Potts model in a region of size N of $\mathbf{Z}^d$. One possibility is to simply consider $\Lambda_N = \{-N,\dots,N\}^d$ as a subgraph of the lattice $\mathbf{Z}^d$, as we have done so far. However, one could also extend configurations $\sigma \in \{1,\dots,q\}^{\Lambda_N}$ to all of $\{1,\dots,q\}^{\mathbf{Z}^d}$ by fixing a reference configuration $\bar\sigma \in \{1,\dots,q\}^{\mathbf{Z}^d}$ and postulating that $\sigma_x = \bar\sigma_x$ for $x \notin \Lambda_N$. (A natural choice is to take $\bar\sigma$ constant, and we will mainly consider such a situation here.) Finally, one could identify $-N$ with N in Λ_N and consider the Potts model on larger and larger discrete tori. In this way, we have different choices of *boundary conditions* yielding different definitions for the finite volume measures $\mu_{\beta,N}$.

In many examples of interest (including the case of the Potts models), the measure $\mu_\beta = \lim_{N\to\infty}\mu_{\beta,N}$ is well defined (i.e., independent of the choice of boundary condition) for $\beta < \beta_c$ while one can obtain several distinct limits in the case $\beta > \beta_c$. Figure 2 shows typical samples drawn from μ_β for the Ising model with $\bar\sigma \equiv 1$. In the case $\beta > \beta_c$, the resulting sample clearly "remembers" the bias introduced by $\bar\sigma$ in the sense that a typical configuration consists of a "sea" of spins taking the dominant value $+1$ (brown) with small "islands" of spins taking the value -1 (yellow). Had we set $\bar\sigma \equiv -1$, we would have obtained a sample with the opposite behaviour, which illustrates the nonuniqueness of the infinite-volume measure μ_β in this case. In the case $\beta < \beta_c$ on the other hand, each one of the two possible spin values is

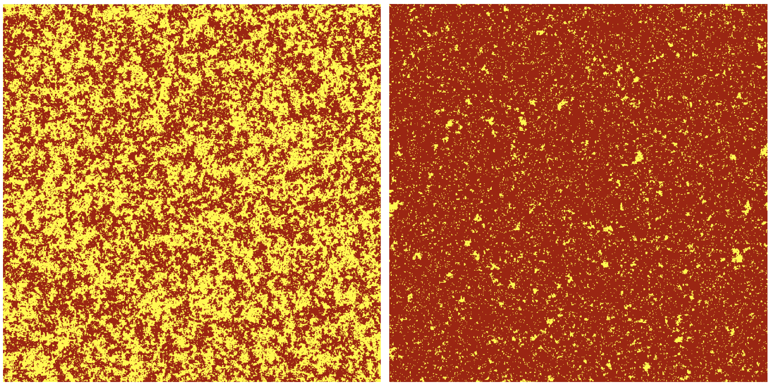

FIGURE 2. Typical Ising configurations for $\beta < \beta_c$ (left) and $\beta > \beta_c$ (right).

about equally represented, and the measure is symmetric under the substitution $1 \leftrightarrow -1$, which illustrates the uniqueness of μ_β. It is in fact a theorem in the case of the Ising model that for $\beta > \beta_c$ there exist exactly two translation invariant infinite volume measures $\mu_\beta^\pm$ corresponding to boundary conditions $\bar\sigma \equiv \pm 1$, and that every accumulation point of $\mu_{\beta,N}$ for any sufficiently homogeneous boundary condition as $N \to \infty$ is a convex combination of μ_β^+ and μ_β^-. (In fact a similar statement holds for the Potts model with q states, where one has exactly q distinct infinite volume Gibbs measures when $\beta > \beta_c$.)

This raises the question of the uniqueness of μ_β at $\beta = \beta_c$. If it is, then we say that the phase transition is *continuous*, otherwise it is said to be *discontinuous*. The reason for this terminology is that continuity in this sense turns out to be equivalent to the continuity of the maps $\beta \mapsto \mu_\beta^\pm$ at $\beta = \beta_c$. It has been known for quite some time [Yan52, AF86] that the phase transition for the Ising model is continuous in dimensions $d = 1, 2$ as well as $d \geq 4$. The reason why dimensions 1 and 2 are typically much better understood is that the Ising model is *solvable* in these dimensions in the sense that explicit expressions can be derived for the expectation of a large number of observables under $\mu_{\beta,N}$. (This solution is straightforward in $d = 1$ [Isi25] where no phase transition is present, but it was a major breakthrough when Onsager obtained his exact solution for $d = 2$ [Ons44].) Dimension $d = 4$ on the other hand is the *upper critical dimension* beyond which the model is expected to be trivial (i.e., described by Gaussian random variables in the scaling limit) which allows us to use a number of powerful techniques, including for example the *lace expansion* [HS94, Sak07].

This leaves the case $d = 3$ which is of course the physically most interesting one since the Ising model is a toy model of ferromagnetism, and its dimensions represent the usual spatial dimensions. Heuristic considerations suggest that the phase transition is also continuous there, and this is consistent with physical experiments, assuming that the Ising model belongs to the same universality class as that of a genuine physical magnet. In the article [ADCS15], Duminil-Copin et al. gave the first rigorous proof that this is indeed the case. The proof relies on the introduction of the quantity

$$M(\beta) = \inf_{B \subset \mathbf{Z}^3} \frac{1}{|B|^2} \sum_{x,y \in B} \int \sigma_x \sigma_y \, \mu^0_\beta(d\sigma) ,$$

where μ^0_β denotes the infinite volume limit obtained from using *free* conditions, as well as three main steps. First, they rely on results of [FSS76, FILS78] to argue that the Fourier transform of $x \mapsto \int \sigma_0 \sigma_x \, \mu^0_\beta(d\sigma)$ belongs to L^1 at $\beta = \beta_c$, which implies that $M(\beta_c) = 0$. Then, and this is the main step, they show that having $M(\beta) = 0$ implies that a certain percolation model with long-range correlations constructed from the Ising model admits no infinite clusters. Finally, they use a variant of the *switching lemma* [GHS70] to show that the quantity $\int \sigma_0 \sigma_x \, \mu^+_\beta(d\sigma) - \int \sigma_0 \sigma_x \, \mu^0_\beta(d\sigma)$ is dominated by an explicit function times the probability of the origin belonging to an infinite cluster in the above mentioned model and therefore has to vanish at $\beta = \beta_c$. Once this is known, it is not too difficult to show that the spontaneous magnetisation of the Ising model at criticality must vanish (namely one has $\int \sigma_0 \, \mu^+_{\beta_c}(d\sigma) = 0$), which in turn yields the desired uniqueness statement.

Considering more general values of q for the Potts model illustrates in a rather striking way the fact that continuity of the phase transition, whatever the dimension, is a rather nontrivial property that isn't necessarily expected in general. Indeed, it was conjectured by Baxter in the 1970s [Bax71, Bax73] that the Potts model on $\mathbf{Z}^2$ exhibits a continuous phase transition if and only if $q \leq 4$. The pair of articles [DCST17, DCGH+21] by Duminil-Copin et al. provides proofs of both directions of this conjecture. For the sake of brevity we will not comment on the proofs in any detail, but we note that the proof of continuity of the phase transition for $q \leq 4$ is almost completely disjoint from that in the case of the $3d$ Ising model. A milestone is again to show that the model at criticality with boundary condition set to one fixed element of S admits no infinite cluster. However both the proof of this fact (exploiting a form of discrete holomorphicity of certain cleverly chosen observables) and the proof of its equivalence with the uniqueness of the infinite-volume measure at criticality (actually they show equivalence of a list of five quite distinct properties which are of independent interest for the study of the critical Potts model) are completely different.

Regarding the proof of *dis*continuity when $q > 4$, the main tool is a close relation, first discovered by Temperley and Lieb [TL71] in a restricted context and then by Baxter et al. [BKW76] in more generality, between the critical FK model on $\mathbf{Z}^2$ and the so-called six-vertex model. Configurations of the latter can be visualised as jigsaw pieces where one assigns to each vertex of $\mathbf{Z}^2$ (or a subset thereof) one of the following six (oriented) tiles.

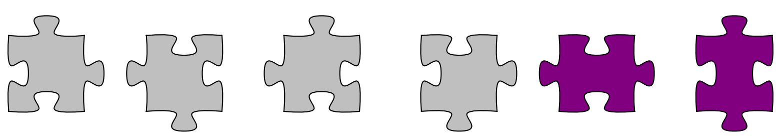

Further, one enforces the admissibility constraint that the tiles fit together seamlessly. One also postulates that the probability of seeing a given admissible configuration is proportional to $c^{\#p}$, where $\#p$ denotes the number of purple tiles in the configuration and c is some fixed constant. The relation between the six-vertex model and the critical FK model holds for the specific choice $c = \sqrt{2 + \sqrt{q}}$. Similar to the link between the FK and Potts models, the relation between the two models is not quite deterministic, but takes the form of a coupling between the FK model on a square lattice and the six-vertex model on a smaller rotated square lattice (the six-vertex model on $\mathbf{Z}^2$ with nearest-neighbour connections is coupled to the cluster model on the even sublattice of $\mathbf{Z}^2$ with nearest-neighbour connections) such that the corresponding conditional probabilities are straightforward to describe by a sequence of i.i.d. choices.

The advantage gained from this relation is that the six-vertex model is solvable in a certain sense using the transfer matrix formalism. This doesn't get one out of the woods since the transfer matrices V_N involved are very large: they act on a vector space of dimension 2^N, but are block diagonal with each block $V_N^{[n]}$ acting on a subspace of dimension $\binom{n}{N}$. Each of these blocks is irreducible with positive entries and therefore admits a Perron–Frobenius vector. The main technical result of [DCST17] is a very sharp asymptotic for the Perron–Frobenius eigenvalues of $V_N^{[N/2-r]}$ for fixed r as $N \to \infty$. Interestingly, the authors are able to prove that the ratios between these values converge to finite (and explicit, at least as explicit convergent series) limits as $N \to \infty$ and that the values themselves diverge exponentially in N with known exponent, but the common lower-order behaviour of that divergence is not known. This asymptotic is however sufficient to obtain good control over the partition function of the six-vertex model and to exploit it to compute an explicit expression for the inverse correlation length of the critical Potts model with free boundary conditions when $q > 4$. The finiteness of that expression finally allows us to deduce the discontinuity of the phase transition. In fact it allows us to prove more than that, namely it shows that under the measure with free boundary conditions, the probability that the cluster containing the origin touches the boundary of a large ball decays exponentially in the radius of the ball, with an exponent that can be deduced from the above asymptotic.

To conclude this section, I would like to mention the beautiful article [DCRT19] which, although not quite dealing with the question of continuity of the phase transition, does have a related flavour. The question there is that of the *sharpness* of the phase transition which in this particular case is couched as the question whether it is really true that the measure μ_β has exponentially decaying correlations (in the sense that the covariance between $f(\sigma_0)$ and $f(\sigma_x)$ decays exponentially fast as $|x| \to \infty$ for any *nice enough* function $f\colon S \to \mathbf{R}$) for *every* $\beta < \beta_c$ and not just for small enough values where a perturbation argument around $\beta = 0$ (where $f(\sigma_0)$ and $f(\sigma_x)$ are independent under μ_0 as soon as $x \neq 0$) may apply. One difficulty with this type of statement is that one does in general not know any closed-form expression for β_c: in the case of the FK model on the square lattice such an expression can be derived by the duality argument described in Section 2, but it is not known for more general situations. The main result of [DCRT19] is that the phase transition of the FK model on *any* vertex-transitive infinite graph is sharp.

The main tool in their proof is a novel and far-reaching generalisation of the OSSS inequality [OSSS05]. The context here is that of increasing random variables $f\colon \{0,1\}^E \to [0,1]$ (for a finite set E and for the natural coordinate-wise partial order on $\{0,1\}^E$), where $\{0,1\}^E$ is furthermore equipped with a probability measure $\mathbf{P}$ that is itself *monotonic* in the sense that for every $F \subset E$ and every $e \in E \setminus F$, the conditional probabilities $\mathbf{P}(w_e = 1 \,|\, \mathcal{F}_F)$ are increasing functions on $\{0,1\}^F$.

(Here $\mathcal{F}_F$ denotes the σ-algebra generated by the evaluations $w \mapsto w_e$ for $e \in F$.) One then considers *any* algorithm that reveals one by one the values of an input $w \in \{0,1\}^E$ to f in such a way that the coordinate to be revealed next depends in a deterministic way on the information gleaned from the revealment up to that point. (In particular, the first coordinate to be revealed is always the same since no information has yet been obtained.) The algorithm stops once the revealed values are sufficient to determine the value $f(w)$ with certainty, thus yielding a random set $\hat{E} \subset E$ of revealed values. The result of [DCRT19] is then that one has the inequality

$$\mathrm{Var}(f) \leq \sum_{e \in E} \mathbf{P}(e \in \hat{E})\ \mathrm{Cov}(f, w_e)\ , \tag{3.1}$$

which looks formally the same as the result of [OSSS05], but the assumption there was that the measure $\mathbf{P}$ is simply the uniform measure. Since the latter is clearly monotonic (it is such that $\mathbf{P}(w_e = 1 \,|\, \mathcal{F}_F)$ is constant), the results of [OSSS05] follow as a special case.

Using this result, [DCRT19] then obtain the following dichotomy which yields the desired sharpness statement.

Theorem 3.1. *Let G be any transitive graph and let $\mathbf{P}_{\beta,n}$ be the FK measure on the ball Λ_n of radius n in G. Then there exists $\beta_c \in \mathbf{R}$ such that, for every $\beta < \beta_c$ there exists $c_\beta > 0$ such that $\mathbf{P}_{\beta,n}(0 \leftrightarrow \partial\Lambda_n) \lesssim e^{-c_\beta n}$, uniformly in n. For $\beta > \beta_c$ on the other hand, there exists $c > 0$ such that $\mathbf{P}_{\beta,n}(0 \leftrightarrow \partial\Lambda_n) \geq c \min\{1, \beta - \beta_c\}$.*

Once (3.1) is known, the proof of Theorem 3.1 is surprisingly simple and relies on two ingredients. First, one can show that the measures $\mathbf{P}_{\beta,n}$ and the function $\mathbf{1}_{0\leftrightarrow\partial\Lambda_n}$ satisfy the assumptions of (3.1). Setting $\theta_n(\beta) = \mathbf{P}_{\beta,n}(0 \leftrightarrow \partial\Lambda_n)$, a clever choice of search algorithm for the (potential) cluster connecting the origin 0 to $\partial\Lambda_n$ then allows us to show that one has the bound

$$\theta_n'(\beta) \gtrsim \sum_{e\in E} \mathrm{Cov}_\beta(\mathbf{1}_{0\leftrightarrow\partial\Lambda_n}, w_e) \geq \frac{n}{8\Sigma_n(\beta)}\theta_n(\beta)(1-\theta_n(\beta))\ , \tag{3.2}$$

where $\Sigma_n = \sum_{k=0}^{n-1} \theta_k$. The fact that the first inequality holds is known and can be checked in an elementary way. The second fact is that *any* sequence of functions $\beta \mapsto \theta_n(\beta)$ satisfying a differential inequality of the form (3.2) necessarily satisfies a dichotomy of the type appearing in the statement of Theorem 3.1. Since we are not interested in the regime where θ_n is large, we can rewrite (3.2) as $\theta_n' \geq \frac{cn}{\Sigma_n}\theta_n$. The fact that the θ_n then should satisfy such a dichotomy is quite clear: if β is such that they converge to a nonvanishing limit θ, then $\Sigma_n/n \sim \theta$ and one must have $\theta' \geq c$. If on the other hand they converge to 0 on a whole interval $[a,b]$, then that convergence must take place sufficiently fast so that $\Sigma_n/n \gg \theta_n$ (since otherwise the previous argument applies). Since $\Sigma_n/n \sim \theta_n$ for $\theta_n \sim n^{-\alpha}$ as soon as $\alpha < 1$, it is then plausible that for any $c < b$ one has $\theta_n \ll n^{-1/2}$ (say), implying $\theta_n' \gtrsim \sqrt{n}\theta_n$ and therefore $\theta_n \lesssim e^{-\sqrt{n}(c-\beta)}$ for $\beta < c$. This shows that Σ_n is bounded for $\beta < c$, leading to $\theta_n' \gtrsim n\theta_n$ and therefore an exponentially (in n) small bound as claimed.

4. Rotational invariance for the critical FK models

As already mentioned a number of times, a crucial feature of two dimensional equilibrium statistical mechanics is that *most* models (at least those with sufficiently *local* interactions) are expected to obey a form of conformal invariance, or rather equivariance as in (1.2), when considering large-scale observables (crossing probabilities, averages, etc.) at the critical temperature. This expectation and the resulting link to the well understood world of two-dimensional conformal quantum field theories allows us to generate a plethora of conjectures regarding the large-scale behaviour of these models, but these are in many cases extremely hard to prove. Consider for example the N-step two-dimensional self-avoiding random walk, which is simply the uniform measure on all *injective* functions $h\colon \{0,\ldots,N\} \to \mathbf{Z}^2$ such that $h(0) = 0$ and such that $|h(i+1) - h(i)| = 1$ for all $i < N$. If the injectivity condition on h is dropped, one recovers the simple random walk which is well known [Don51] to converge to Brownian motion for large N, provided that it is rescaled by $\sqrt{N}$.

Back to the self-avoiding random walk, exploiting the expected conformal invariance of its suitably rescaled large-N limit and the known properties of the Schramm–Loewner evolutions, one expects the size of $h(N)$ to be of order $N^{3/4}$ and its rescaling by $N^{3/4}$ to converge to a specific continuous random curve, namely $\mathrm{SLE}_{8/3}$ [LSW04]. Rigorously however, almost *nothing* nontrivial is known about this model: although the diameter of the range of h trivially has to be at least $\sqrt{N/\pi}$, the current best lower bound on the endpoint does not even match that! Instead, one only knows the bound $(\mathbf{E}|h(N)|^p)^{1/p} \geq \frac{1}{6}N^{p/(2p+2)}$ that was recently obtained by Madras [Mad14]. Similarly, while one trivially has $|h(N)| \leq N$, the best nontrivial upper bound is pretty much the weakest possible improvement, namely that for every $p \geq 1$ one has $\lim_{N\to\infty} N^{-1}(\mathbf{E}|h(N)|^p)^{1/p} = 0$, obtained around the same time by Duminil-Copin and Hammond [DCH13]. One main obstruction is that there is at the moment no proof showing that the self-avoiding random walk is conformally invariant at large scales.

While this illustrates the importance of showing that statistical models are conformally invariant (or at least rotationally invariant as a crucial first step) at criticality, the strategy of proof for such claims has so far mostly relied on finding a large enough collection of observables that already satisfy a discrete analogue of conformal invariance, typically by solving a discrete analogue of the Cauchy–Riemann equations. See for example Chelkak and Smirnov's proof of conformal invariance for the Ising model on isoradial graphs [CS12] and Smirnov's proof of conformal invariance for critical percolation [SS11]. The two-dimensional FK model with $q \leq 4$ already mentioned in Section 3 is one of the simplest models where conformal invariance at criticality is expected, but where it is not known how to obtain this from a suitable discrete conformal invariance. In the recent work [DCKK+20], Duminil-Copin et al. succeeded in taking the first step towards conformal invariance by showing that the large-scale behaviour of these models is indeed rotationally invariant.

To define the notion of *large-scale behaviour*, we recall that the configuration space of the FK model is the same as that for regular percolation; see Figure 1. Such a configuration can alternatively be described as a collection of non-self-intersecting loops separating the percolation clusters from the clusters of the dual configuration. (Actually it naturally yields *two* collections of loops, depending on whether the loop encloses a percolation cluster of the primary or of the dual configuration, but we will ignore this detail for the sake of our exposition.) Given two collections $\mathcal{F}$ and $\bar{\mathcal{F}}$ of non-self-intersecting loops in the plane, one then defines a distance between them in the following way. Given (small) $\eta > 0$, write $\mathcal{B}_\eta \subset \mathbf{R}^2$ for a large chunk of a fine lattice in $\mathbf{R}^2$, for example $\mathcal{B}_\eta = \eta\mathbf{Z}^2 \cap [-\eta^{-1}, \eta^{-1}]^2$. Given a loop γ and assuming that its image doesn't intersect the set $\mathcal{B}_\eta$, one then denotes by $[\eta]_\gamma$ its homotopy class in $\mathbf{R}^2 \setminus \mathcal{B}_\eta$. One then obtains a distance by postulating that $d_H(\mathcal{F}, \bar{\mathcal{F}}) \leq \eta$ if and only if, for every $\gamma \in \mathcal{F}$ that encloses at least two elements of $\mathcal{B}_\eta$ but not all of it, there exists $\bar{\gamma} \in \bar{\mathcal{F}}$ such that $[\gamma]_\eta = [\bar{\gamma}]_\eta$ and vice versa. (The H here stands for *homotopy*.)

Given a metric space (M, d), the metric d lifts naturally to a metric on the space of probability measures on M which metrises the topology of weak convergence (at least when M is nice, for example Polish). This is done by considering the Wasserstein (also sometimes called Kantorovich–Rubinstein or Monge–Kantorovich) distance

$$d(\mu, \nu) = \inf_{\mathbf{P} \in \mathcal{C}(\mu,\nu)} \int d(x, y)\, \mathbf{P}(dx, dy) \;, \tag{4.1}$$

where $\mathcal{C}(\mu_1, \mu_2)$ denotes the set of all couplings between μ_1 and μ_2, that is, probability measures on M^2 with ith marginal equal to μ_i. Note that with this definition, the map that assigns to x the probability measure δ_x concentrated at x is an isometry.[3]

Fix now once and for all $q \in [1, 4]$ and consider a smooth bounded simply connected domain $\Omega \subset \mathbf{R}^2$. For $\varepsilon > 0$, write $\mathbf{P}_{\varepsilon,\Omega}$ for the critical FK measure (viewed as a measure on collections of loops separating clusters and dual clusters) on $\varepsilon\mathbf{Z}^2 \cap \Omega$ with free boundary conditions. We also write $\mathbf{P}_\varepsilon$ for the limit of $\mathbf{P}_{\varepsilon,\Omega}$ as $\Omega \to \mathbf{R}^2$ and, given an angle $\theta \in \mathbf{R}$, we write R_θ for the rotation by θ, which naturally acts on loops in $\mathbf{R}^2$. The large-scale rotational invariance of the critical FK model can then be formulated as follows.

Theorem 4.1. *For every domain $\Omega \subset \mathbf{R}^2$ as above and every angle θ one has*

$$\lim_{\varepsilon \to 0} d_H\big(R_\theta^* \mathbf{P}_{\varepsilon,\Omega}, \mathbf{P}_{\varepsilon, R_\theta \Omega}\big) = 0 \;.$$

Furthermore, one has $\lim_{\varepsilon \to 0} d_H(R_\theta^* \mathbf{P}_\varepsilon, \mathbf{P}_\varepsilon) = 0$.

We only focus on the second statement since it turns out that the first one can be deduced from it without too much effort. In fact, the authors of [DCKK+20] show a type of universality statement for the FK model on rectangular lattices, but its formulation requires some preparation. One starts by defining a specific class of isoradial embeddings of the two-dimensional square lattice into the plane. Recall that a planar graph embedded in the plane is isoradial if, for each face f, there exists a circle of radius 1 containing all the vertices of f. (For example, the canonical embedding of the square lattice is isoradial.)

Given a bi-infinite sequence $\alpha\colon \mathbf{Z} \to (-\frac{\pi}{2}, \frac{\pi}{2})$, we consider the map $\iota_\alpha\colon \mathbf{Z}^2 \to \mathbf{R}^2$ given by

$$\iota_\alpha\colon (x, y) \mapsto \big(x + s_y, c_y\big) \;, \qquad s_y = \sum_{k \in (0,y]} \sin(\alpha_k) \;, \quad c_y = \sum_{k \in (0,y]} \cos(\alpha_k) \;,$$

with the convention that for $y < 0$, $\sum_{(0,y]} = -\sum_{(y,0]}$. This defines an isoradial graph $L(\alpha)$ by considering the embedding of $\{(x, y) : x + y \text{ even}\}$ (joined by diagonal edges) under ι_α (see Figure 3). The dual graph $L^*(\alpha)$ of $L(\alpha)$ is then given

[3] In fact, (4.1) in general only makes sense between measures such that the distance to some arbitrary fixed point has a finite first moment. One can always either restrict oneself to such measures or replace d by $1 \wedge d$ which generates the same topology as d.

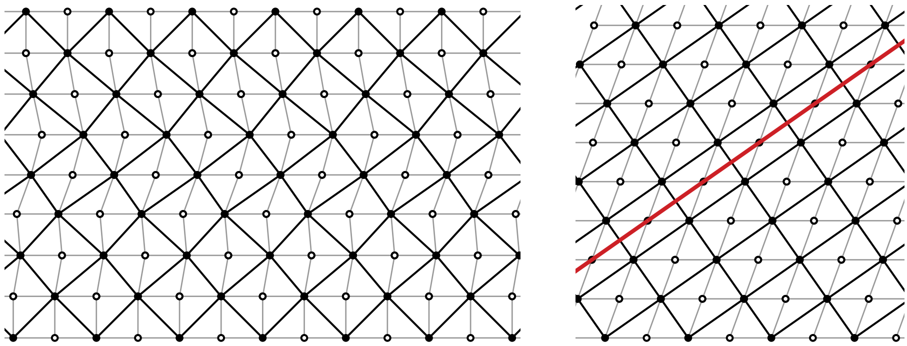

Figure 3. Examples of graphs $L(\alpha)$. On the left is a generic α while on the right α is constant but nonzero. The graph itself is drawn in black, the vertices of its dual graph are drawn in white, and the associated diamond graph is light gray. In red, we draw one of the symmetry axes of the second graph.

by the embedding of $\{(x, y) : x + y \text{ odd}\}$. The associated *diamond graph* has as its vertices both the vertices of $L(\alpha)$ and the centres of its faces, and its edges are given by all pairs (v, f) with v a vertex and f a face such that $v \in f$. The diamond graph is simply given by the embedding of the usual lattice $\mathbf{Z}^2$ with nearest-neighbour edges under ι_α.

It is crucial at this stage to note that the critical FK model on $L(\alpha)$ is *not* given by simply pushing forward the critical FK model on $\mathbf{Z}^2$ under the map ι_α. Instead, one reweighs each edge of the graph in a very specific way that depends on the length of the edge. More specifically, viewing a configuration of the FK model as a subset $\omega \subset E$ of the set of edges of the (finite) graph on which the model is considered, the probability of seeing a given configuration ω is proportional to

$$\Big(\prod_{e\in\omega} p_e\Big)\Big(\prod_{e\in E\setminus\omega} (1-p_e)\Big) q^{k(\omega)}\;, \tag{4.2}$$

where $k(\omega)$ denotes the number of connected components of the subgraph ω. The formula for p_e as a function of q and the length of the edge e is explicit but not relevant for the sake of this discussion.

The most important step in the proof is to show that the large-scale connectivity properties of the critical FK model on $L(\alpha)$ are very close to those of the model on $L(T_j\alpha)$, where T_j swaps the jth and $(j+1)$-th component:

$$(T_j\alpha)_k = \begin{cases} \alpha_{j+1} & \text{if } k = j,\\ \alpha_j & \text{if } k = j+1,\\ \alpha_k & \text{otherwise.}\end{cases}$$

Furthermore, there exists a natural coupling between the FK measures on the two lattices which implements this *closedness*. This part of the proof exploits the link to the six-vertex model and its solvability using the transfer matrix formalism. One then deduces from this that the model on the standard lattice $L(0)$ is very close to that on a rotated rectangular lattice $L(\alpha)$ with $k \mapsto \alpha_k$ constant (see the right half of Figure 3). This works by fixing some large $N > 0$ (which is then eventually sent to infinity) and starting from $\alpha_k^{(i)} = \alpha \mathbf{1}_{k\ge N}$ and then swapping components in such a way as to move some of the nonzero components down until one ends up with $\alpha_k^{(f)} = \alpha(\mathbf{1}_{|k|\le N} + \mathbf{1}_{k>3N})$. Since one has $L(0) \approx L(\alpha^{(i)})$ and $L(\alpha) \approx L(\alpha^{(f)})$, the desired statement follows if one can control the error made at each step of the argument. This turns out to be extremely delicate, and one has to exploit subtle stochastic cancellations along the way. One trick is to allow the vertices of the set $\mathcal{B}_\eta$ around which the homotopy classes are computed to move a little bit with each application of a swapping operator T_j and to show that this motion ends up being diffusive (and therefore *slow*) rather than ballistic.

Once one knows that $\lim_{\varepsilon\to 0} d_H(\mathbf{P}_{\varepsilon,L(0)}, \mathbf{P}_{\varepsilon,L(\alpha)}) = 0$, the second part of Theorem 4.1 follows at once. The idea is simply to note that $L(\alpha)$ is invariant under reflection along a line with angle $\frac{\pi}{4} - \frac{\alpha}{2}$, but that the effect of this reflection on $L(0)$ is the same as that of a rotation by angle α (since it is itself invariant under reflection along a line with angle $\frac{\pi}{4}$), so that

$$\begin{aligned} d_H(\mathbf{P}_\varepsilon, R_\alpha^*\mathbf{P}_\varepsilon) &\le d_H(\mathbf{P}_{\varepsilon,L(0)}, \mathbf{P}_{\varepsilon,L(\alpha)}) + d_H(\mathbf{P}_{\varepsilon,L(\alpha)}, R_\alpha^*\mathbf{P}_{\varepsilon,L(0)}) \\ &= 2d_H(\mathbf{P}_{\varepsilon,L(0)}, \mathbf{P}_{\varepsilon,L(\alpha)})\;, \end{aligned}$$

and the claim follows.

About the author

Martin Hairer works in the general area of probability theory with a main focus on the analysis of stochastic partial differential equations. His work has been distinguished with a 2014 Fields Medal and the 2021 Breakthrough Prize in Mathematics. He currently holds positions at École Polytechnique Fédérale de Lausanne (Switzerland) and Imperial College London (UK).

References

[ADC21] M. Aizenman and H. Duminil-Copin, *Marginal triviality of the scaling limits of critical 4D Ising and ϕ_4^4 models*, Ann. of Math. (2) **194** (2021), no. 1, 163–235, DOI 10.4007/annals.2021.194.1.3. MR4276286

[ADCS15] M. Aizenman, H. Duminil-Copin, and V. Sidoravicius, *Random currents and continuity of Ising model's spontaneous magnetization*, Comm. Math. Phys. **334** (2015), no. 2, 719–742, DOI 10.1007/s00220-014-2093-y. MR3306602

[AF86] M. Aizenman and R. Fernández, *On the critical behavior of the magnetization in high-dimensional Ising models*, J. Statist. Phys. **44** (1986), no. 3-4, 393–454, DOI 10.1007/BF01011304. MR857063

[Aiz81] M. Aizenman, *Proof of the triviality of φ_d^4 field theory and some mean-field features of Ising models for $d > 4$*, Phys. Rev. Lett. **47** (1981), no. 1, 1–4, DOI 10.1103/PhysRevLett.47.1. MR620135

[Aiz82] M. Aizenman, *Geometric analysis of φ^4 fields and Ising models. I, II*, Comm. Math. Phys. **86** (1982), no. 1, 1–48. MR678000

[Bax71] R. J. Baxter, *Generalized ferroelectric model on a square lattice*, Studies in Appl. Math. **50** (1971), 51–69, DOI 10.1002/sapm197150151. MR307646

[Bax73] R. J. Baxter, *Potts model at the critical temperature*, Journal of Physics C: Solid State Physics **6** (1973), no. 23, L445–L448. doi:10.1088/0022-3719/6/23/005.

[BDC12] V. Beffara and H. Duminil-Copin, *The self-dual point of the two-dimensional random-cluster model is critical for $q \ge 1$*, Probab. Theory Related Fields **153** (2012), no. 3-4, 511–542, DOI 10.1007/s00440-011-0353-8. MR2948685

[BG93] T. W. Burkhardt and I. Guim, *Conformal theory of the two-dimensional Ising model with homogeneous boundary conditions and with disordred boundary fields*, Phys. Rev. B **47** (1993), 14306–14311. doi:10.1103/PhysRevB.47.14306.

[BKW76] R. J. Baxter, S. B. Kelland, and F. Y. Wu, *Equivalence of the Potts model or Whitney polynomial with an ice-type model*, Journal of Physics A: Mathematical and General **9** (1976), no. 3, 397–406. doi:10.1088/0305-4470/9/3/009.

[BPZ84] A. A. Belavin, A. M. Polyakov, and A. B. Zamolodchikov, *Infinite conformal symmetry in two-dimensional quantum field theory*, Nuclear Phys. B **241** (1984), no. 2, 333–380, DOI 10.1016/0550-3213(84)90052-X. MR757857

[Car84] J. L. Cardy, *Conformal invariance and surface critical behavior*, Nuclear Physics B **240** (1984), no. 4, 514–532. doi:https://doi.org/10.1016/0550-3213(84)90241-4.

[CGN15] F. Camia, C. Garban, and C. M. Newman, *Planar Ising magnetization field I. Uniqueness of the critical scaling limit*, Ann. Probab. **43** (2015), no. 2, 528–571, DOI 10.1214/13-AOP881. MR3305999

[CGN16] F. Camia, C. Garban, and C. M. Newman, *Planar Ising magnetization field II. Properties of the critical and near-critical scaling limits* (English, with English and French summaries), Ann. Inst. Henri Poincaré Probab. Stat. **52** (2016), no. 1, 146–161, DOI 10.1214/14-AIHP643. MR3449298

[CHI15] D. Chelkak, C. Hongler, and K. Izyurov, *Conformal invariance of spin correlations in the planar Ising model*, Ann. of Math. (2) **181** (2015), no. 3, 1087–1138, DOI 10.4007/annals.2015.181.3.5. MR3296821

[CS12] D. Chelkak and S. Smirnov, *Universality in the 2D Ising model and conformal invariance of fermionic observables*, Invent. Math. **189** (2012), no. 3, 515–580, DOI 10.1007/s00222-011-0371-2. MR2957303

[DCGH+21] H. Duminil-Copin, M. Gagnebin, M. Harel, I. Manolescu, and V. Tassion, *Discontinuity of the phase transition for the planar random-cluster and Potts models with* $q > 4$ (English, with English and French summaries), Ann. Sci. Éc. Norm. Supér. (4) **54** (2021), no. 6, 1363–1413, DOI 10.24033/asens.2485. MR4371394

[DCH13] H. Duminil-Copin and A. Hammond, *Self-avoiding walk is sub-ballistic*, Comm. Math. Phys. **324** (2013), no. 2, 401–423, DOI 10.1007/s00220-013-1811-1. MR3117515

[DCKK+20] H. Duminil-Copin, K. K. Kozlowski, D. Krachun, I. Manolescu, and M. Oulamara, *Rotational invariance in critical planar lattice models*, `arXiv:2012.11672`, 2020.

[DCRT19] H. Duminil-Copin, A. Raoufi, and V. Tassion, *Sharp phase transition for the random-cluster and Potts models via decision trees*, Ann. of Math. (2) **189** (2019), no. 1, 75–99, DOI 10.4007/annals.2019.189.1.2. MR3898174

[DCST17] H. Duminil-Copin, V. Sidoravicius, and V. Tassion, *Continuity of the phase transition for planar random-cluster and Potts models with* $1 \leq q \leq 4$, Comm. Math. Phys. **349** (2017), no. 1, 47–107, DOI 10.1007/s00220-016-2759-8. MR3592746

[Don51] M. D. Donsker, *An invariance principle for certain probability limit theorems*, Mem. Amer. Math. Soc. **6** (1951), 12. MR40613

[ESPP+12] S. El-Showk, M. F. Paulos, D. Poland, S. Rychkov, D. Simmons-Duffin, and A. Vichi, *Solving the 3d Ising model with the conformal bootstrap*, Phys. Rev. D **86**, (2012), 025022. `arXiv:1203.6064`, 2012.

[ESPP+14] S. El-Showk, M. F. Paulos, D. Poland, S. Rychkov, D. Simmons-Duffin, and A. Vichi, *Solving the 3d Ising model with the conformal bootstrap II. c-minimization and precise critical exponents*, J. Stat. Phys. **157** (2014), no. 4-5, 869–914, DOI 10.1007/s10955-014-1042-7. MR3269694

[FILS78] J. Fröhlich, R. Israel, E. H. Lieb, and B. Simon, *Phase transitions and reflection positivity. I. General theory and long range lattice models*, Comm. Math. Phys. **62** (1978), no. 1, 1–34. MR506363

[FK72] C. M. Fortuin and P. W. Kasteleyn, *On the random-cluster model. I. Introduction and relation to other models*, Physica **57** (1972), 536–564. MR359655

[Frö82] J. Fröhlich, *On the triviality of* $\lambda\varphi_d^4$ *theories and the approach to the critical point in* $d \geq 4$ *dimensions*, Nuclear Phys. B **200** (1982), no. 2, 281–296, DOI 10.1016/0550-3213(82)90088-8. MR643591

[FSS76] J. Fröhlich, B. Simon, and T. Spencer, *Infrared bounds, phase transitions and continuous symmetry breaking*, Comm. Math. Phys. **50** (1976), no. 1, 79–95. MR421531

[GHS70] R. B. Griffiths, C. A. Hurst, and S. Sherman, *Concavity of magnetization of an Ising ferromagnet in a positive external field*, J. Mathematical Phys. **11** (1970), 790–795, DOI 10.1063/1.1665211. MR266507

[Gri06] G. Grimmett, *The random-cluster model*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 333, Springer-Verlag, Berlin, 2006, DOI 10.1007/978-3-540-32891-9. MR2243761

[HS94] T. Hara and G. Slade, *Mean-field behaviour and the lace expansion*, Probability and phase transition (Cambridge, 1993), NATO Adv. Sci. Inst. Ser. C: Math. Phys. Sci., vol. 420, Kluwer Acad. Publ., Dordrecht, 1994, pp. 87–122. MR1283177

[Isi25] E. Ising, *Beitrag zur Theorie des Ferromagnetismus*, Zeitschrift fur Physik **31** (1925), no. 1, 253–258. doi:10.1007/BF02980577.

[LSW04] G. F. Lawler, O. Schramm, and W. Werner, *On the scaling limit of planar self-avoiding walk*, Fractal geometry and applications: a jubilee of Benoît Mandelbrot, Part 2, Proc. Sympos. Pure Math., vol. 72, Amer. Math. Soc., Providence, RI, 2004, pp. 339–364, DOI 10.1090/pspum/072.2/2112127. MR2112127

[Mad14] N. Madras, *A lower bound for the end-to-end distance of the self-avoiding walk*, Canad. Math. Bull. **57** (2014), no. 1, 113–118, DOI 10.4153/CMB-2012-022-6. MR3150723

[MS16] J. Miller and S. Sheffield, *Quantum Loewner evolution*, Duke Math. J. **165** (2016), no. 17, 3241–3378, DOI 10.1215/00127094-3627096. MR3572845

[Ons44] L. Onsager, *Crystal statistics. I. A two-dimensional model with an order-disorder transition*, Phys. Rev. (2) **65** (1944), 117–149. MR10315

[OSSS05] R. O'Donnell, M. Saks, O. Schramm, and R. Servedio, *Every decision tree has an influential variable*, in *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, 31–39. 2005. `arXiv:cs/0508071`. doi:10.1109/SFCS.2005.34.

[Sak07] A. Sakai, *Lace expansion for the Ising model*, Comm. Math. Phys. **272** (2007), no. 2, 283–344, DOI 10.1007/s00220-007-0227-1. MR2300246

[Sch00] O. Schramm, *Scaling limits of loop-erased random walks and uniform spanning trees*, Israel J. Math. **118** (2000), 221–288, DOI 10.1007/BF02803524. MR1776084

[SS11] O. Schramm and S. Smirnov, *On the scaling limits of planar percolation*, Ann. Probab. **39** (2011), no. 5, 1768–1814, DOI 10.1214/11-AOP659. With an appendix by Christophe Garban. MR2884873

[TL71] H. N. V. Temperley and E. H. Lieb, *Relations between the "percolation" and "colouring" problem and other graph-theoretical problems associated with regular planar lattices: some exact results for the "percolation" problem*, Proc. Roy. Soc. London Ser. A **322** (1971), no. 1549, 251–280, DOI 10.1098/rspa.1971.0067. MR498284

[Yan52] C. N. Yang, *The spontaneous magnetization of a two-dimensional Ising model*, Phys. Rev. (2) **85** (1952), 808–816. MR51740

École Polytechnique Fédérale de Lausanne, Switzerland; and Imperial College London, United Kingdom

Email address: `martin.hairer@epfl.ch`

Email address: `m.hairer@imperial.ac.uk`