THE WORK OF HUGO DUMINIL-COPIN

MARTIN HAIRER

ABSTRACT

The past decade has seen tremendous progress in our understanding of the behavior of many probabilistic models at or near their "critical point." On the 5th of July 2022, Hugo Duminil-Copin was awarded the Fields medal for the crucial role he played in many of these developments. In this short review article, we will try to put his work into context and present a small selection of his results.

1. INTRODUCTION

Hugo Duminil-Copin was awarded the Fields medal in Helsinki during the opening ceremony of the 2022 virtual ICM. In this short note, I will try to put his work into context and to give the reader a glimpse of why the questions it addresses are not only very interesting from a purely mathematical perspective, but also contribute to further our understanding of nature at a fundamental level. I should start first of all with a disclaimer. Hugo Duminil-Copin is an astounding problem solver and, while his interest falls squarely into the general area of probability theory and in particular the type of probabilistic problems that arise when studying microscopic models for statistical mechanics, I will not be able to do justice to the breadth of his contributions. Furthermore, my own area of expertise is somewhat tangential to that of Duminil-Copin, so this note should be taken as the point of view of an interested outsider. In particular, any misrepresentations of his results and/or techniques will be entirely due to my own ignorance.

In its broadest form, classical statistical mechanics can be thought of as the study of the global behavior of "large" systems (of "size" $N \gg 1$) that comprise many identical "small" subsystems interacting with each other. One typically indexes the subsystems by a discrete set Λ_N with $\lim_{N \to \infty} |\Lambda_N| = \infty$ and one is interested in quantities that are stable as $N \to \infty$. In many cases of interest, one has $\Lambda_N \subset \Lambda$ for Λ a discrete subset of Euclidean space (typically, a regular lattice) and its elements are interpreted as a physical location of the corresponding subsystem; the interaction between subsystems may then depend on their locations. (In most models they actually depend only on their relative positions, a notion that generalizes very well to locations taking values in more general symmetric spaces.)

Let us write S for the state space of one single such subsystem, so that the state space for the full system is $\mathcal{S}_N \overset{\text{def}}{=} S^{\Lambda_N}$. In *equilibrium statistical mechanics*, we furthermore assume that S is equipped with a "reference" probability measure μ (think of μ as being normalized counting measure if S is a finite set, normalized volume measure if it is a compact manifold, etc.) and that our system is described by an *energy function* $H^{(N)} \colon \mathcal{S}_N \to \mathbf{R}$, which typically comprises a contribution for each subsystem, as well as additional interaction terms. In full generality, one would have something like

$$H^{(N)}(\sigma) = \sum_{A \subset \mathcal{S}_N} H_A(\sigma_A), \tag{1.1}$$

where σ_A denotes the restriction of $\sigma \in S^{\Lambda_N}$ to S^A and the function H_A typically only depends on the "shape" of the subset A, so satisfies natural invariance properties under translations and possibly reflections and/or discrete rotations. In many classical models, the only nonvanishing terms in (1.1) are those with $|A| \leq 2$.

Given such an energy function, we obtain a probability measure $\mu_{\beta,N}$ on $\mathcal{S}_N$ by setting

$$\mu_{\beta,N}(d\sigma) = Z_{\beta,N}^{-1} \exp\big(-\beta H^{(N)}(\sigma)\big) \prod_{u \in \Lambda_N} \mu(d\sigma_u), \tag{1.2}$$

where $Z_{\beta,N}$ is chosen in such a way that $\mu_{\beta,N}(\mathcal{S}_N) = 1$. Physically, the parameter $\beta > 0$ appearing in this expression is the inverse of the temperature of the system. To a large extent,

(equilibrium) statistical mechanics is the study of $\mu_{\beta,N}$ as $N \to \infty$ with a particular emphasis on the behavior under $\mu_{\beta,N}$ of observables that take a "macroscopic" (of the order of the size of the domain Λ_N) or "mesoscopic" (tending to infinity as $N \to \infty$ but much smaller than $|\Lambda_N|$) number of components of σ into account.

1.1. Bernoulli percolation

The simplest such example is that of $S = \{-1, 1\}$, $H_N = 0$, and $\mu(\{-1\}) = \mu(\{1\}) = \frac{1}{2}$. Regarding the index set Λ_N, we consider the case of the even elements of a large box in $\mathbf{Z}^2$, namely $\Lambda_N = \{u \in \{-N, \ldots, N\}^2 : u_1 + u_2 \text{ even}\}$. (The reason why we make this strange choice rather than simply taking all elements of $\{-N, \ldots, N\}^2$ will soon become clear.)

One of the simplest kind of "global" observables for this system is given by the following kind of linear statistics. Given a smooth function $\phi \colon [-1, 1]^2 \to \mathbf{R}$, we define $I_\phi^N \colon S_N \to \mathbf{R}$ by

$$I_\phi^N(\sigma) = N^{-\alpha} \sum_{u \in \Lambda_N} \sigma_u \phi(u/N). \tag{1.3}$$

Note that this is exhaustive: for any fixed N, if we know $I_\phi^N(\sigma)$ for every smooth function ϕ, then we can in principle recover the argument σ itself. A version of the central limit theorem then immediately yields the following result:

Theorem 1.1. *Setting $\alpha = 1$, the joint distribution of $I_\phi^N(\sigma)$ for any finite collection of test functions ϕ as above converges as $N \to \infty$ to the law of a collection of jointly centered Gaussian random variables I_ϕ such that*

$$\mathbf{E} I_\phi \, I_\psi = \frac{1}{2} \int_{[-1,1]^2} \phi(x)\psi(x)\,dx.$$

(The factor $\frac{1}{2}$ appearing here comes from the fact that the local density of Λ_N in $\mathbf{Z}^2$ is $\frac{1}{2}$.)

A much more interesting kind of global observables is given by the connectivity properties of σ, which were first studied by Broadbent and Hammersley [12]. These are however *much* harder to analyze and, even though the model just described appears at first sight to be somewhat trivial, most of its results already lead us squarely into the 21st century mathematics. In order to describe what we mean by "connectivity" in this context, instead of interpreting elements $u \in \Lambda_N$ as points in $\mathbf{Z}^2$, we interpret them as nearest-neighbor edges of a suitable sublattice of $\mathbf{Z}^2$ by associating to u the unique edge e_u of $\mathbf{Z}_{\text{even}} \times \mathbf{Z}_{\text{odd}}$ with midpoint u. We will also write e_u^* for the edge of $\mathbf{Z}_{\text{odd}} \times \mathbf{Z}_{\text{even}}$ with midpoint u. In other words, we set

$$e_u = \begin{cases} (u_\downarrow, u_\uparrow) & \text{if } u_1 \text{ is even,} \\ (u_\leftarrow, u_\rightarrow) & \text{if } u_1 \text{ is odd,} \end{cases} \qquad e_u^* = \begin{cases} (u_\leftarrow, u_\rightarrow) & \text{if } u_1 \text{ is even,} \\ (u_\downarrow, u_\uparrow) & \text{if } u_1 \text{ is odd.} \end{cases}$$

Here, given $u = (u_1, u_2) \in \mathbf{Z}^2$, we write $u_\leftarrow = (u_1 - 1, u_2)$, etc. The endpoints of these edges do belong to the stated sublattices of $\mathbf{Z}^2$ since $u_1 + u_2$ is even, so either both u_1 and u_2 are even or both are odd.
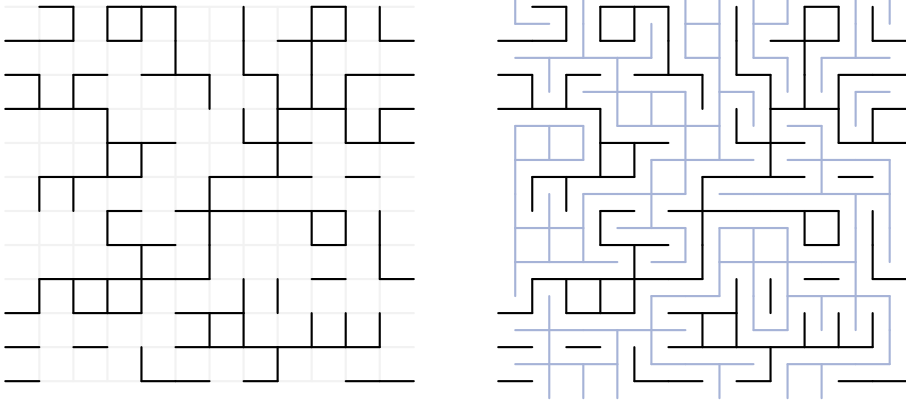
FIGURE 1

On the left, we draw a typical percolation configuration for $N = 11$. On the right, the same configuration is drawn together with its dual configuration in light blue.

Given a configuration $\sigma \in \mathcal{S}_N$, we interpret edges e_u with $\sigma_u = -1$ as "open" and draw them in black, while the remaining edges are considered "closed" and are drawn in light grey. This yields a picture like shown on the left in Figure 1. We can then ask, for example, what is the probability p_N that it is possible to go from the left boundary of the light gray graph to the right boundary (the "boundary" here consists of the ends of the dangling edges) while only traversing black edges. It turns out that this probability does take nontrivial values even for large values for N. As a matter of fact, it is independent of N as the following classical result (see, for example, **[36, LEMMA 11.21]**) shows.

Theorem 1.2. *One has $p_N = \frac{1}{2}$ for every N.*

Proof. The trick is to observe that given a configuration $\sigma \in \mathcal{S}_N$, if we draw the dual configuration $\sigma^* \in \mathcal{S}_N$ defined by $\sigma_u^* = -\sigma_u$ by coloring (in blue, say) the edges e_u^* with $\sigma_u^* = -1$, then we obtain a drawing with the property that blue edges never intersect black edges. As a consequence, it is possible to cross the square from left to right by traversing only black edges if and only if it is *not* possible to cross it from top to bottom by traversing only blue edges. (See Figure 1.) On the other hand, the law of the collection of blue edges is the same as that of the collection of black edges, only rotated by 90°, so that we must have $p_N = 1 - p_N$ as claimed. ∎

Remark 1.3. If, instead of choosing edges to be open with probability $\frac{1}{2}$, we choose them to be open with some probability p, then we have $p_N \to 1$ for $p > \frac{1}{2}$ and $p_N \to 0$ for $p < \frac{1}{2}$. This is an example of *phase transition*: an abrupt change in the behavior of some global observables as a parameter of the model is varied continuously. In this specific example, we were able to determine the *critical value $p_c = \frac{1}{2}$* explicitly by exploiting an exact duality.

It is similarly possible to obtain a large collection of interesting global observables by taking a shape $\mathcal{U} \subset [-1, 1]^2$ diffeomorphic to a square and considering the analogous event $A_{\mathcal{U}}^{(N)} \subset \mathcal{S}_N$ asking whether it is possible to connect the left and right edges of $N\mathcal{U}$ (without ever leaving $N\mathcal{U}$) by a path following only open edges of a given configuration $\sigma \in \mathcal{S}_N$. Again, the knowledge of these events is an exhaustive statistics for any given fixed value of N. It is furthermore known that for any finite number of such shapes $\{\mathcal{U}_i\}_{i \in I}$ (for I some finite index set) the random variables $\{[A_{\mathcal{U}_i}^{(N)}]\}_{i \in I}$ converge in law to a nondegenerate limit $\{[A_{\mathcal{U}_i}]\}_{i \in I}$ as $N \to \infty$ [56]. (Here, we write $[A]$ for the indicator function of an event A.) An amazing fact is that this scaling limit is conformally invariant: if $\phi \colon D \to D'$ is a conformal map between two smooth simply connected domains $D, D' \subset \mathbf{C}$ such that $[-1, 1]^2 \subset D$ and such that $\mathcal{V}_i \stackrel{\text{def}}{=} \phi(\mathcal{U}_i) \subset [-1, 1]^2$, then the joint law of the random variables $\{[A_{\mathcal{V}_i}]\}_{i \in I}$ is *the same* as that of $\{[A_{\mathcal{U}_i}]\}_{i \in I}$.

This conformal invariance turns out to be a crucial feature of the scaling limits of many equilibrium statistical mechanics models in two dimensions. It provides a link to conformal field theory which, at a purely mathematical level, can be thought of as the study of irreducible representations of the Virasoro algebra. In particular, it strongly suggests that the possible large-scale behaviors one can see for two-dimensional equilibrium models come in a one-parameter family of "universality classes" parametrized by the central charge of the corresponding conformal field theory. (In the case of percolation, it turns out that this central charge is given by $c = 0$.)

1.2. The Ising model

The next-"simplest" model of statistical mechanics falling into the category of equilibrium models described above is the Ising model [41, 43]. (See also the review article [16] in these proceedings which contains a more detailed account of the various developments spawned by this model.) In this case, the index set is given by $\Lambda_N = \{-N, \ldots, N\}^d$ for some $d \geq 1$, the reference measure μ and local state space S are as above, but this time one has $H_A = 0$ unless $A = \{u, v\}$ with $u, v \in \mathbf{Z}^d$ such that $|u - v| = 1$, in which case one sets $H_A(\sigma) = -\sigma_u \sigma_v$. This time, the model has a nontrivial dependence on the parameter β appearing in (1.2), which plays a role somewhat similar to the parameter p that appeared in Remark 1.3.

At a very qualitative level, the situation is somewhat similar to what happened in the case for percolation: in every dimension $d \geq 2$, there exists a critical (dimension-dependent) value β_c which delineates two different regimes. At "high temperature," namely for $\beta < \beta_c$, the *spontaneous magnetization*, namely the random quantity $N^{-d} \sum_{i \in \Lambda_N} \sigma_i$, 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 $\beta_c = \log \sqrt{1 + \sqrt{2}}$ [50]. (There is no phase transition at all in dimension 1 and the spontaneous magnetization always vanishes, so in some sense $\beta_c = +\infty$ there.)

It is again possible to ask the same questions as in the case of Bernoulli percolation. This time, however, even the analogue of Theorem 1.1, which was an essentially trivial

consequence of the central limit theorem (or at least a version thereof), is already highly nontrivial. It was shown in a recent series of works [13, 14] that if one chooses $\beta = \beta_c$ and $\alpha = 15/8$ in the expression (1.3) in dimension $d = 2$, then it converges in law to nontrivial limiting random variables, jointly for any fixed number of test functions. This time, however, the limiting distributions are not Gaussian (they actually exhibit an even faster decaying tail behavior). Note that the exponent α is closely related to the behavior of $\mathbf{E}_c \sigma_u \sigma_v$ (where $\mathbf{E}_c$ denotes the expectation under the Gibbs measure (1.2) for the critical value of the inverse temperature β) since, assuming that $\mathbf{E}_c \sigma_u \sigma_v \approx |u - v|^{-\delta}$, one finds that

$$\mathbf{E}_c \left(I_\phi^N(\sigma) \right)^2 = N^{-2\alpha} \sum_{u,v} \phi(u/N) \phi(v/N) \mathbf{E}_c \sigma_u \sigma_v \lesssim N^{-2\alpha} \sum_{u,v} |u - v|^{-\delta} \approx N^{2d - (\delta \wedge d) - 2\alpha},$$

so that one expects the relation $\alpha = d - (\delta \wedge d)/2$, which (correctly) leads to the prediction $\delta = \frac{1}{4}$. This and a number of other properties of the Ising model at criticality allow associating it to the conformal field theory with central charge $c = \frac{1}{2}$.

The picture in higher dimensions is much less clear, however. For $d \geq 5$, it was shown in [1, 2, 29] that the correct scaling exponent to use in (1.3) 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 Neuman 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 level with the development of the "conformal bootstrap" [23, 24]. 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 [3] where they show that any subsequential limit for expressions of the form (1.3) as $N \to \infty$ (and $\beta \to \beta_c$) must necessarily be Gaussian.

In fact, some of the results just mentioned are shown for the "lattice Φ^4 model" which is the equilibrium model with $S = \mathbf{R}$, as well as

$$H_{\{u\}}(\sigma) = V(\sigma_u) \stackrel{\text{def}}{=} \sigma_u^4 - \alpha \sigma_u^2, \quad H_{\{u,v\}}(\sigma) = \frac{1}{2}(\sigma_u - \sigma_v)^2,$$

again provided that u and v are nearest-neighbors, and with c an additional parameter. While this appears to be very different from the Ising model at first sight, we can see that it is actually a generalization of it: if the constant α is large, then the potential V has two very deep wells with minima located at $\pm\sqrt{\alpha}$, so its effect is to impose that $\sigma_u \approx \pm\sqrt{\alpha}$ with high probability. The main contribution then comes from the cross-term of the square in the two-body term which is the same as for the Ising model. These kind of considerations lead one to expect that, since these models exhibit long-range correlations at the critical temperature (in the sense that the correlation $\mathbf{E}\sigma_x \sigma_y$ decays slowly in $|x - y|$ as already pointed out earlier) which should furthermore lead to some form of self-averaging, the Ising model and the Φ^4 model exhibit the same behavior at criticality.

1.3. A general picture

The general picture that should by now be emerging from our discussion can be summarized as follows:

(1) Many of the simplest local equilibrium systems do exhibit a phase transition, namely there exists a critical value β_c at which the qualitative large scale behavior 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 behavior of the system displays qualitatively different behavior. 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 $2d$ conformal field theories (and the associated probabilistic objects like SLE [55], QLE [45], etc.) provides a hugely powerful machinery to predict – and in a number of cases also rigorously prove – their behavior.

(3) The universe of local statistical mechanics models can be subdivided into broad classes of models that exhibit a shared large-scale behavior at criticality. These are called "universality classes" and, in the $2d$ equilibrium case, they are expected to come in families parametrized 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.)

(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 behavior at criticality is "trivial" in the sense that it exhibits Gaussian behavior. (Typically, with correlation function given by the Green's function of the Laplacian.) 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 Hugo Duminil-Copin's many contributions to this vast programme.

---

1   See, however, the recent breakthrough made in the approximation of the critical exponents of the $3d$ Ising model using the "conformal bootstrap" [23, 24] already mentioned above.

This represents, of course, a mere sliver of his work and completely ignores very substantial chunks of it. By presenting not just a long laundry list of results that he proved and conjectures that he settled but instead an overview of the strategy of proof for a few select results, I hope to be able to convey one of the features of Duminil-Copin's body of work, namely that he has a knack for finding just the right way of looking at a problem that had hitherto been overlooked. In many cases, this only provides small cracks in the problem's armor that still require tremendous technical skill to be wedged open, but in some cases it results in surprisingly simple but ingenious proofs. Either way, I am very much looking forward to learning more from Duminil-Copin's insights for many years to come.

2. (DIS)CONTINUITY OF PHASE TRANSITIONS

One very natural question in this area is whether one can take the limit $N \to \infty$ in (1.2). At this stage, we note that the definition of $H^{(N)}$ given in (1.1) is not necessarily the most natural one since it restricts the sum over those clusters A that are constrained to *entirely* lie in S_N. Another possibility that appears just as natural would be to restrict the sum over clusters that merely intersect S_N, but to specify some fixed "boundary condition" $\bar{\sigma} \in S^\Lambda$ that is used to compute the values of the H_A with A intersecting both Λ_N and $\Lambda \setminus \Lambda_N$ in the sense that we interpret σ_A in (1.1) as $\sigma_{A,x} = \sigma_x$ for $x \in A \cap \Lambda_N$ and $\sigma_{A,x} = \bar{\sigma}_x$ otherwise.

In many examples of interest (including the case of the Ising model, but *not* the case of percolation), 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
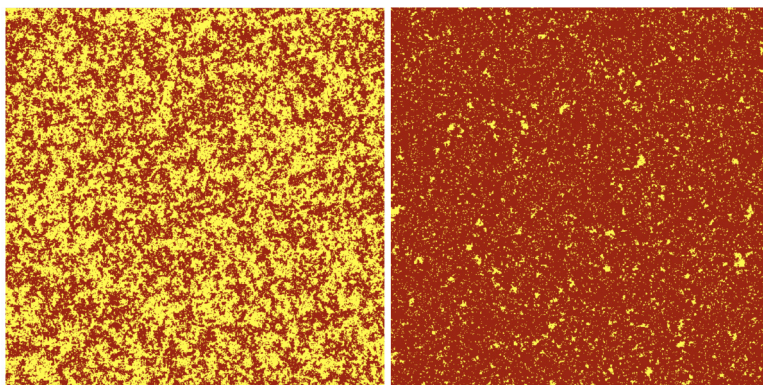
FIGURE 2
Typical Ising configurations for $\beta < \beta_c$ (left) and $\beta > \beta_c$ (right).

would have obtained a sample with the opposite behavior, which illustrates the nonunique-ness 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 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^-.

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 [5,60] 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 $d = 1$ [41] where no phase transition is present, but it was a major breakthrough when Onsager obtained his exact solution for $d = 2$ [50]). 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 using a number of powerful techniques, including, for example, the *lace expansion* [39,54].

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 [4], Duminil-Copin et al. gave the first rigorous proof that this is indeed the case. The proof relies on the introduction of the quantity

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

where μ_β^0 denotes the infinite volume limit obtained from using "free" conditions, as well as three main steps. First, they rely on results of [30, 31] to argue that the Fourier transform of $x \mapsto \int \sigma_0 \sigma_x \, \mu_\beta^0(d\sigma)$ 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" [35] to show that the quantity $\int \sigma_0 \sigma_x \, \mu_\beta^+(d\sigma) - \int \sigma_0 \sigma_x \, \mu_\beta^0(d\sigma)$ 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 magnetization 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.

To illustrate the fact that continuity of the phase transition, whatever the dimension, is a rather nontrivial property that is not necessarily expected in general, a good example is that of the Potts model [53]. This is defined similarly to the Ising model, but this time the local state space S is given by $S = \{1, \ldots, q\}$ for some $q \geq 2$ endowed again with the normalized counting measure as its reference measure. As in the Ising model, one sets $H_A = 0$ unless $A = \{u, v\}$ with $u, v \in \mathbf{Z}^d$ such that $|u - v| = 1$, in which case one sets $H_A(\sigma) = \mathbf{1}_{\sigma_u = \sigma_v}$. For $q = 2$, this is equivalent to the Ising model since their energy functionals only differ by a constant. Let us also remark that there is an essentially equivalent model called the random cluster model (or sometimes the FK model after Fortuin and Kasteleyn who introduced it in [28]) in which one directly considers partitions of $\mathbf{Z}^d$ into connected "clusters" (which one should think of as the edge-connected components of the sets $\{u : \sigma_u = i\}$ for $i \in S$ and a given configuration σ of the Potts model) and which makes sense also for noninteger values of $q \geq 1$. (In the case $q = 1$, the FK model actually reduces to regular Bernoulli percolation.) See (4.1) below for a more precise definition of this model.

It was conjectured by Baxter in the 1970s [8,9] that the Potts model on $\mathbf{Z}^2$ exhibits a continuous phase transition if and only if $q \leq 4$. The pair of articles [17,21] 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 5 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–Lieb [59] in a restricted context and then by Baxter et al. [10] in more generality, between the FK model on $\mathbf{Z}^2$ and the so-called six-vertex model Configurations of the latter can be visualized as jigsaws where one assigns to each vertex of $\mathbf{Z}^2$ (or a subset thereof) one of the six (oriented) tiles
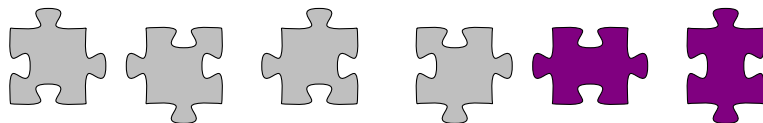
and one enforces the admissibility constraint that the tiles fit together seamlessly. One further postulates that the probability of seeing a given admissible configuration is proportional to $c^{\#p}$, where $\#p$ denotes the number of purple tiles in the configuration and c is some fixed constant. The relation between the six-vertex model and the critical FK model holds for the specific choice $c = \sqrt{2 + \sqrt{q}}$. The advantage gained from this relation is that the six-vertex model is "solvable" in a certain sense using the transfer matrix formalism. This does not 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 [21] 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 behavior 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 deducing the discontinuity of the phase transition.

To conclude this section, I would like to mention the beautiful article [20] which, although not quite dealing with the question of continuity of the phase transition, does have a related flavor. 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 $f : S \to \mathbf{R}$) 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 will 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 a duality argument [11], but it is not known for more general situations. The main result of [20] 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 generalization of the OSSS inequality [49]. The context here is that of increasing random variables $f : \{0, 1\}^E \to [0, 1]$ (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 $\mathbf{P}(w_e = 1 \mid \mathcal{F}_F)$ are increasing functions. (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 $w \in \{0, 1\}^E$ 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 been obtained yet at that point.) The algorithm stops once the revealed values are sufficient to determine the value of $f(w)$, thus yielding a random set $\hat{E} \subset E$ of revealed values. The result of [20] is then that one has the inequality

$$\mathrm{Var}(f) \leq \sum_{e \in E} \mathbf{P}(e \in \hat{E})\,\mathrm{Cov}(f, w_e), \qquad (2.1)$$

which looks formally the same as the result of [49], 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 $\mathbf{P}(w_e = 1 \mid \mathcal{F}_F)$ is constant), the results of [49] follow as a special case.

Using this result, the authors of [20] then obtain the following dichotomy which yields the desired sharpness statement.

Theorem 2.1. *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 (2.1) is known, the proof 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 (2.1). 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 showing that one has the bound

$$\theta_n'(\beta) \gtrsim \sum_{e \in E} \mathrm{Cov}_\beta\left(\mathbf{1}_{0\leftrightarrow\partial\Lambda_n}, w_e\right) \geq \frac{n}{8\Sigma_n(\beta)} \theta_n(\beta)\left(1 - \theta_n(\beta)\right), \tag{2.2}$$

where $\Sigma_n = \sum_{k=0}^{n-1} \theta_n$. 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 (2.2) necessarily satisfies a dichotomy of the type appearing in the statement of Theorem 2.1. Since we are not interested in the regime where θ_n is large, we can rewrite (2.2) 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.

3. TRIVIALITY OF Φ_4^4

It has been known since the groundbreaking work of Osterwalder and Schrader [51, 52] that, at least in some cases, the construction of a (bosonic) quantum field theory satisfying the Wightman axioms is equivalent to the construction of a probability measure on the space of distributions satisfying a number of natural properties. One of the pinnacles of that line of enquiry was the construction in the seventies of the Φ_2^4 and Φ_3^4 measures [22, 25, 27, 33, 34, 47, 48, 57], which corresponds to the simplest case of an interacting theory in two or three space-time dimensions with one type of boson.

At a heuristic level, the Φ_d^4 measure is the measure $\mu^{(d)}$ on the space of Schwartz distributions $\mathcal{S}'(\mathbf{R}^d)$ (or on the d-dimensional torus) given by

$$\mu^{(d)}(d\Phi) = Z^{-1} \exp\left(-\frac{1}{2}\int \left(|\nabla\Phi(x)|^2 - C\Phi^2(x) + \Phi^4(x)\right) dx\right) d\Phi,$$

where "$d\Phi$" denotes the infinite-dimensional Lebesgue measure on $\mathcal{S}'(\mathbf{R}^d)$. This expression is, of course, problematic at many levels: infinite-dimensional Lebesgue measure does not exist, distributions cannot be squared, etc. If it were only for the term $|\nabla\Phi|^2$, one could reasonably interpret this expression as the Gaussian measure μ_0 with covariance operator given by the Green's function of the Laplacian, which is a well-defined probability measure (modulo technicalities arising from the constant mode which can easily be fixed). The measure μ_0 is called the Gaussian Free Field (GFF) since it corresponds to a quantum field theory in which particles are free, i.e., do not interact with each other at all.

This suggests that a more refined interpretation of the Φ_d^4 measure could be given by

$$\mu^{(d)}(d\Phi) = Z^{-1} \exp\left(-\frac{1}{2}\int \Phi^4(x)\,dx\right)\mu_0(d\Phi). \tag{3.1}$$

This is still ill-defined since the GFF is supported on distributions rather than functions for any dimension $d \geq 2$. However, setting $\Phi_\varepsilon = \rho_\varepsilon \star \Phi$, the *Wick power*

$$:\Phi^4: = \lim_{\varepsilon\to 0}\left(\Phi_\varepsilon^4 - 3\Phi_\varepsilon^2 \mathbf{E}\Phi_\varepsilon^2\right), \tag{3.2}$$

turns out to be a well-defined random Schwartz distribution (i.e., the limit exists and is independent of the choice of ρ_ε) in dimensions $d < 4$. In dimension 2, Nelson showed in [47] that the Radon–Nikodym factor appearing in (3.1) with Φ^4 replaced by $:\Phi^4:$ yields an integrable random variable, thus leading to a definition of $\mu^{(2)}$. In particular, the Φ_2^4 measure is equivalent to the GFF. In dimension 3, this turns out not to be the case, but it is still possible to show that the measure

$$\mu^{(3)}(d\Phi) = \lim_{\varepsilon\to 0} Z_\varepsilon^{-1}\exp\left(-\frac{1}{2}\int \Phi_\varepsilon^4(x) - C_\varepsilon\Phi_\varepsilon^2(x)\,dx\right)\mu_0(d\Phi) \tag{3.3}$$

is well-defined for a suitable choice of the constant C_ε which differs from the choice $3\mathbf{E}\Phi_\varepsilon^2 \sim \varepsilon^{-1}$ suggested by (3.2) by a logarithmically divergent term. (An alternative construction of this measure was recently obtained by completely different techniques in [37, 38, 46].)

This discussion begs the question of what happens for $d \geq 4$ and especially when $d = 4$ which is the physically most interesting case from the QFT perspective (remember that dimension here corresponds to space-time). Regarding the case $d > 4$, it was already shown in the eighties by Aizenman and Fröhlich [1, 2, 29] that pretty much any "reasonable" definition of the Φ_d^4 measure actually coincides with the GFF. This still left the case $d = 4$ which has always been expected to be the hard case since it is "critical" in the sense that, at least at a formal level, the terms Φ^4 and $|\nabla\Phi|^2$ scale in the same way in the following sense. Writing $\mathcal{S}_\lambda$ for the transformation $(\mathcal{S}_\lambda F)(x) = F(\lambda x)$, the GFF has the self-similarity property $\mathcal{S}_\lambda \Psi \overset{\text{law}}{=} \lambda^{\frac{2-d}{2}}\Psi$ for Ψ drawn from μ_0. Pretending that Ψ behaves like a function (even though it really is a random distribution), we deduce that

$$\mathcal{S}_\lambda|\nabla\Psi|^2 = \lambda^{-2}|\nabla\mathcal{S}_\lambda\Psi|^2 \overset{\text{law}}{=} \lambda^{-d}|\nabla\Psi|^2, \quad \mathcal{S}_\lambda(\Psi^4) = (\mathcal{S}_\lambda\Psi)^4 \overset{\text{law}}{=} \lambda^{4-2d}\Psi^4.$$

These exponents are indeed equal if and only if $d = 4$. A heuristic calculation actually suggests that, at higher order, the term $|\nabla\Psi|^2$ dominates the term Ψ^4 at large scales. Variants of this observation have been made rigorous in a number of works [26, 32, 40], including most

recently in an impressive series of works by Bauerschmidt–Brydges–Slade (see [6,7] and the references therein).

One way of formulating one of their main results is the framework given in our introduction with $S = \mathbf{R}$, μ being Lebesgue measure, $H_{\{x\}}(\phi) = \frac{g}{4}\phi_x^4 + \frac{\nu}{2}\phi_x^2$, $H_{\{x,y\}}(\phi) = |\phi_x - \phi_y|^2$ when x and y are neighboring lattice sites in $\mathbf{Z}^4$, and $H_A = 0$ otherwise. This model behaves in a way that is very similar to the Ising model, to which it degenerates in the regime $g \to \infty$ and $\nu = -g$. Traditionally, one considers the Φ^4 model with $\beta = 1$, since one can always reduce oneself to this case by adjusting g and ν, and possibly rescaling the ϕ_x's by a factor. One typically also considers g fixed, it is therefore the parameter ν that is tuneable and plays the role of a "temperature" in this model. Just like the Ising model, it exhibits a phase transition at some value $\nu_c \in \mathbf{R}$: for $\nu > \nu_c$, there exists a unique infinite volume measure which is symmetric under $\phi \mapsto -\phi$. For $\nu < \nu_c$, on the other hand, one finds two distinct infinite-volume measures (as well as their convex combinations) depending on the boundary conditions one chooses.

A state $\phi \in S^{\Lambda_N}$ with $\Lambda_N = \{-N, \ldots, N\}^4$ is viewed as a distribution $\iota\phi$ on the torus (of size 2) by setting, for every smooth test function $f : \mathbf{T}^4 \to \mathbf{R}$,

$$(\iota\phi)(f) = \sum_{x \in \Lambda_N} \sigma_N \phi_x f(x/N),$$

for a sequence of values σ_N chosen in such a way that $\mathbf{E}((\iota\phi)(1)^2) = 1$. It is then shown in [6] that if g is sufficiently small and ν is chosen in a suitable way (close but not quite equal to the critical value ν_c), then $\iota\phi$ converges to a massive GFF, namely the Gaussian field with covariance given by $(m^2 - \Delta)^{-1}$ for some $m \in \mathbf{R}$ (which depends on the specific way in which ν is being tuned to approach ν_c as $N \to \infty$).

While this result strongly suggests that there exists no nontrivial Φ_4^4 measure, it does not rule out the possibility of having a nontrivial scaling limit for the discrete field we just described at (or near) criticality when the constant g is sufficiently large (in other words, "at strong coupling"). The technique of proof of [6] was to implement a rigorous version of the "renormalization group technique," which relies on a subtle analysis of the behavior of the renormalization map near the fixed point given by the GFF. This is unfortunately perturbative in nature and so has little hope of being able to deal with arbitrary g. In the recent work [3], however, Aizenman and Duminil-Copin finally succeeded in showing the following result.

Theorem 3.1. *For every way of adjusting $g = g_N$ and $\nu = \nu_N$ as $N \to \infty$ such that $\nu_N \geq \nu_{c,N}$, every $M_N \to \infty$ with $1 \ll M_N \ll N$, and every smooth compactly supported test function f, the law of $\xi_N^f = \sum_{x \in \Lambda_N} \phi_x f(x/M_N)$, normalized so that its variance is one, converges to a normal distribution.*

Remark 3.2. The condition $\nu_N \geq \nu_{c,N}$ can actually be slightly relaxed but not too much. This is because, in the "low temperature" regime and with free (or periodic) boundary conditions, one would expect the law of ξ_N^f to converge to a Bernoulli random variable rather than a Gaussian.

At a high level, the main reason why [3] can deal with arbitrary couplings is that one can think of their setting as being more akin to "perturbing around $g = \infty$" rather than around $g = 0$. In the setting of the introduction, they start by considering the Ising model as described there (i.e., with $\mu = \frac{1}{2}(\delta_1 + \delta_{-1})$), but then expand their class of models to allow for each site to represent a collection of spins with arbitrary ferromagnetic interactions within a site, instead of a single spin. This has the effect of replacing μ by any measure that can be obtained as the law of $\delta \sum_{i=1}^{K} s_i$ for some $\delta > 0$ and $K \in \mathbf{N}$, and where the $s_i \in \{-1, 1\}$ are random variables with a joint distribution proportional to $\exp(-\sum_{ij} a_{ij} s_i s_j)$ for some arbitrary but *positive* coefficients a_{ij}. As was shown already in the 1970s [58, THEOREM 1], all probability measures on $\mathbf{R}$ of the type $Z^{-1} \exp(cx^2 - gx^4) \, dx$ can be obtained as limits of such measures, so that the discrete Φ_4^4 model can be viewed as a limit of block-spin models.

Recall that to show that a collection $\{X_a\}_{a \in A}$ of real-valued random variables is jointly Gaussian it suffices to show that all joint fourth cumulants $\mathbf{E}_c\{X_{a_1}, \ldots, X_{a_4}\}$ with $a_i \in A$ vanish. It is therefore not surprising that fourth cumulants of the spin variables play an important role in any proof of Gaussianity for Ising-type models. In dimension $d \geq 5$, the proof in [2] relied on two very important facts. First, writing $C(x, y) = \mathbf{E}(\sigma_x \sigma_y)$ for the spin correlation function, one shows that for any temperature any any ferromagnetic interaction, one has the bound

$$\left| \mathbf{E}_c\{\sigma_{x_1}, \ldots, \sigma_{x_4}\} \right| \leq 2 \sum_{y \in \mathbf{Z}^d} C(x_1, y) \cdots C(x_4, y). \tag{3.4}$$

One then observes that *at the critical temperature*, the function C is bounded by

$$C(x, y) \lesssim |x - y|^{2-d}. \tag{3.5}$$

Consider now four smooth compactly supported test functions f_i and define

$$X_i = \sum_{x \in \mathbf{Z}^d} \sigma_x f_i(x/M).$$

In particular, the sum ranges over $\mathcal{O}(M^d)$ sites. If one assumes that (3.5) is sharp, then one expects to have $\mathbf{E}X_i^2 \approx M^{d+2}$, so that the "correct" normalization for the X_i's to have unit variance is expected to be $\xi_i = M^{-\frac{d+2}{2}} X_i$. On the other hand, combining the covariance bound with the bound on the fourth cumulant, a powercounting argument shows that $\mathbf{E}_c\{\xi_1, \ldots, \xi_4\} \lesssim M^{-2(d+2)} M^{d+8} = M^{4-d}$, which does indeed converge to 0 as $M \to \infty$ when $d > 4$, thus showing that the ξ_i's are jointly Gaussian in the limit.

Clearly, this calculation does not allow us to conclude anything when $d = 4$. The main contribution of [3] is to show that (3.4) can actually be improved to a bound of the type

$$\left| \mathbf{E}_c\{\sigma_{x_1}, \ldots, \sigma_{x_4}\} \right| \lesssim \frac{\sum_{y \in \mathbf{Z}^d} C(x_1, y) \cdots C(x_4, y)}{\left(\sum_{|x| \leq M} C(0, x)^2 \right)^c}, \tag{3.6}$$

for some (possibly very small) $c > 0$. Here, one assumes that the x_i's are all at distances at least M of each other.

Remark 3.3. If one believes that the bound (3.5) represents the correct behavior of C at criticality, then the denominator appearing in (3.6) is of order $(\log M)^c$ in dimension 4.

This, however, is not known and is also not used by [3], whether for deriving (3.6) or for deducing Theorem 3.1 from it.

The proof of (3.6) relies on the "random current" representation of the Ising model in which the configuration space consists of "currents", namely maps $\mathbf{n}\colon E \to \mathbf{N}$ where E denotes the set of (unoriented) nearest-neighbor pairs in our lattice. The Ising measure then naturally leads to a weight w on currents defined by $w(\mathbf{n}) = \prod_{e \in E} \frac{\beta^{\mathbf{n}(e)}}{\mathbf{n}(e)!}$ as well as the notion of "source" of a current given by

$$\partial \mathbf{n} \stackrel{\text{def}}{=} \left\{ x \; : \; \sum_{e \ni x} \mathbf{n}(e) \text{ is odd} \right\}.$$

The link between currents and the Ising model is the following formula. Given any finite set $A \subset \mathbf{Z}^d$, one has

$$\mathbf{E} \prod_{a \in A} \sigma_a = \frac{\sum_{\mathbf{n} : \partial \mathbf{n} = A} w(\mathbf{n})}{\sum_{\mathbf{n} : \partial \mathbf{n} = \emptyset} w(\mathbf{n})}.$$

A natural notion then is that of a "random current with source A" for which the probability of seeing a given current $\mathbf{n}$ is nonvanishing only when $\partial \mathbf{n} = A$ in which case it is proportional to $w(\mathbf{n})$. When $A = \{x, y\}$, a current $\mathbf{n}$ with source A can be interpreted (not uniquely!) as the occupation measure of a collection of loops in $\mathbf{Z}^d$, together with a non-self-intersecting path joining x and y. In particular, the restriction of $\mathbf{n}$ to the collection of loops connected (either directly or indirectly through other loops) to the path joining x and y can be thought of as the occupation measure of one single random path joining x to y.

The bound (3.6) can then be reformulated in terms of intersection properties of such random paths. From a heuristic perspective, one gets a lot of mileage from thinking of these random paths as simple random walk trajectories. Note that dimension 4 is critical for the question whether the traces of two random walk trajectories intersect or not: in $d < 4$, the trajectories of two independent random walks with any two starting points will intersect almost surely. In $d > 4$, on the other hand, they only intersect with positive probability (going to 0 as the two starting points are taken far from each other) and, if they do, they only have a finite number of intersection points. In dimension $d = 4$, the probability that two random walks starting at distance of order M from each other do intersect decays like $1/\log M$, but the *expected* number of intersection times remains of order one as $M \to \infty$. This shows that if they do intersect, then the number of intersection points is typically quite large, of order $\log M$.

The bulk of the hard work performed in [3] is to show that the random paths arising in the random cluster representation of the Ising model at criticality exhibit a similar behavior, but with $\log M$ replaced by some quantity of size at least $(\log M)^c$ for some $c > 0$. The argument is a masterpiece combining a delicate multiscale analysis, topological arguments, and probabilistic reasoning. One of the main problem the authors have to overcome is the fact that these random paths are *very* far from being simple random walks and only satisfy some spatial version of the Markov property.

4. ROTATIONAL INVARIANCE FOR THE CRITICAL FK MODELS

As already mentioned a number of times, a crucial feature of $2d$ equilibrium statistical mechanics is the fact that most models are expected to obey a form of conformal invariance (or equivariance) when considering large-scale observables at the critical temperature. This expectation and the resulting link to the well understood world of $2d$ conformal field theories allows to generate a plethora of conjectures regarding the large-scale behavior of these models, but these are in many cases extremely hard to prove. Consider for example the N-step $2d$ self-avoiding random walk which is simply the uniform measure on all functions $h: \{0, \ldots, N\} \to \mathbf{Z}^2$ such that $h(0) = 0$ and such that $|h(i + 1) - h(i)| = 1$ for all $i < N$. Exploiting the expected conformal invariance of its suitably rescaled large-N limit, 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 $\mathrm{SLE}_{8/3}$ [42]. Rigorously, almost *nothing* nontrivial is known: 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 [44]. 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 [18]. 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 [15] and Smirnov's proof of conformal invariance for critical percolation [56]. The two-dimensional FK model with $q \leq 4$ already mentioned in Section 2 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 [19], Duminil-Copin et al. show that the large-scale behavior of these models is indeed rotationally invariant.

To define the notion of "large-scale behavior," 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 does not 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 postulates 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 metrizes 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

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

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 the 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.

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) 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 $\Omega \to \mathbf{R}^2$. Given an angle $\theta \in \mathbf{R}$, we also 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*

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

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 can be deduced from it without too much effort. In fact, the authors of [19] show a type of universality statement for the FK model on rectangular lattices, but its formulation requires some preparation. We start 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 biinfinite 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

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

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 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 centers 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-neighbor edges under ι_α.
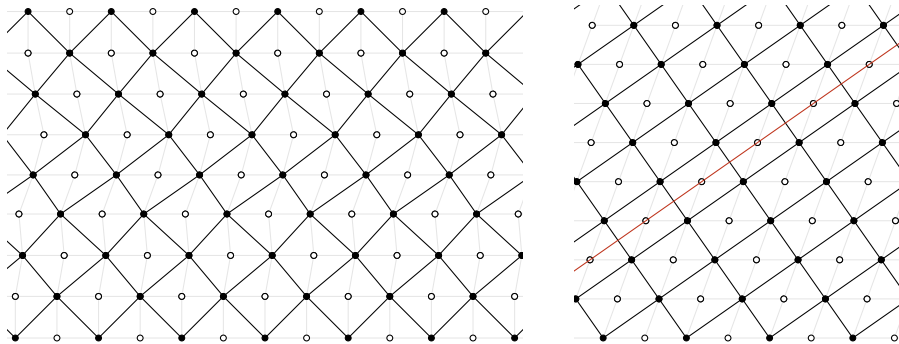
FIGURE 3

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.

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

$$\left(\prod_{e \in \omega} p_e \right) \left(\prod_{e \in E \setminus \omega} (1 - p_e) \right) q^{k(\omega)}, \tag{4.1}$$

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 $\alpha_k^{(i)} = \alpha \mathbf{1}_{k \geq N}$ and then swapping components in such a way as to move some of the nonzero components

down until one ends up with $\alpha_k^{(f)} = \alpha(\mathbf{1}_{|k| \leq N} + \mathbf{1}_{k > 3N})$. 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

$$d_H\left(\mathbf{P}_\varepsilon, R_\alpha^* \mathbf{P}_\varepsilon\right) \leq d_H(\mathbf{P}_{\varepsilon, L(0)}, \mathbf{P}_{\varepsilon, L(\alpha)}) + d_H\left(\mathbf{P}_{\varepsilon, L(\alpha)}, R_\alpha^* \mathbf{P}_{\varepsilon, L(0)}\right) = 2d_H(\mathbf{P}_{\varepsilon, L(0)}, \mathbf{P}_{\varepsilon, L(\alpha)}),$$

and the claim follows.

FUNDING

REFERENCES

[1]   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.

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

[3]   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.

[4]   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.

[5]   M. Aizenman and R. Fernández, On the critical behavior of the magnetization in high-dimensional Ising models. *J. Stat. Phys.* **44** (1986), no. 3–4, 393–454.

[6]   R. Bauerschmidt, D. C. Brydges, and G. Slade, Scaling limits and critical behaviour of the 4-dimensional n-component $|\phi|^4$ spin model. *J. Stat. Phys.* **157** (2014), no. 4–5, 692–742.

[7]   R. Bauerschmidt, D. C. Brydges, and G. Slade, A renormalisation group method. III. Perturbative analysis. *J. Stat. Phys.* **159** (2015), no. 3, 492–529.

[8]   R. J. Baxter, Generalized ferroelectric model on a square lattice. *Stud. Appl. Math.* **50** (1971), 51–69.

[9]   R. J. Baxter, Potts model at the critical temperature. *J. Phys. C, Solid State Phys.* **6** (1973), no. 23, L445–L448.

[10] R. J. Baxter, S. B. Kelland, and F. Y. Wu, Equivalence of the Potts model or Whitney polynomial with an ice-type model. *J. Phys. A: Math. Gen.* **9** (1976), no. 3, 397–406.

[11] V. Beffara and H. Duminil-Copin, The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probab. Theory Related Fields* **153** (2012), no. 3–4, 511–542.

[12] S. R. Broadbent and J. M. Hammersley, Percolation processes: I. Crystals and mazes. *Math. Proc. Cambridge Philos. Soc.* **53** (1957), no. 3, 629–641.

[13] 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.

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

[15] 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.

[16] H. Duminil-Copin, 100 years of the (critical) Ising model on the hypercubic lattice. In *Proceedings of the International Congress of Mathematicians, Vol. 1*, pp. 164–210, EMS Press, 2022.

[17] 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 $q > 4$. *Ann. Sci. Éc. Norm. Supér. (4)* **54** (2021), no. 6, 1363–1413.

[18] H. Duminil-Copin and A. Hammond, Self-avoiding walk is sub-ballistic. *Comm. Math. Phys.* **324** (2013), no. 2, 401–423.

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

[20] 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.

[21] 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.

[22] J.-P. Eckmann and K. Osterwalder, On the uniqueness of the Hamiltionian and of the representation of the CCR for the quartic boson interaction in three dimensions. *Helv. Phys. Acta* **44** (1971), 884–909.

[23] 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.

[24] 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.

[25] J. Feldman, The $\lambda\varphi_3^4$ field theory in a finite volume. *Comm. Math. Phys.* **37** (1974), 93–120.

[26] J. Feldman, J. Magnen, V. Rivasseau, and R. Sénéor, Construction and Borel summability of infrared Φ_4^4 by a phase space expansion. *Comm. Math. Phys.* **109** (1987), no. 3, 437–480.

[27] J. S. Feldman and K. Osterwalder, The Wightman axioms and the mass gap for weakly coupled $(\Phi^4)_3$ quantum field theories. *Ann. Physics* **97** (1976), no. 1, 80–135.

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

[29] J. Fröhlich, On the triviality of $\lambda\varphi_d^4$ theories and the approach to the critical point in $d > 4$ dimensions. *Nuclear Phys. B* **200** (1982), no. 2, 281–296.

[30] 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.

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

[32] K. Gawȩdzki and A. Kupiainen, Massless lattice φ_4^4 theory: a nonperturbative control of a renormalizable model. *Phys. Rev. Lett.* **54** (1985), no. 2, 92–94.

[33] J. Glimm, Boson fields with the :Φ^4: interaction in three dimensions. *Comm. Math. Phys.* **10** (1968), 1–47.

[34] J. Glimm and A. Jaffe, Positivity of the ϕ_3^4 Hamiltonian. *Fortschr. Phys.* **21** (1973), 327–376.

[35] R. B. Griffiths, C. A. Hurst, and S. Sherman, Concavity of magnetization of an Ising ferromagnet in a positive external field. *J. Math. Phys.* **11** (1970), 790–795.

[36] G. Grimmett, *Percolation*. 2nd edn., Grundlehren Math. Wiss. 321, Springer, Berlin, 1999.

[37] M. Hairer, A theory of regularity structures. *Invent. Math.* **198** (2014), no. 2, 269–504.

[38] M. Hairer and J. Mattingly, The strong Feller property for singular stochastic PDEs. *Ann. Inst. Henri Poincaré Probab. Stat.* **54** (2018), no. 3, 1314–1340.

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

[40] T. Hara and H. Tasaki, A rigorous control of logarithmic corrections in four-dimensional ϕ^4 spin systems. II. Critical behavior of susceptibility and correlation length. *J. Stat. Phys.* **47** (1987), no. 1–2, 99–121.

[41] E. Ising, Beitrag zur Theorie des Ferromagnetismus. *Z. Phys.* **31** (1925), no. 1, 253–258.

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

[43] W. Lenz, Beitrag zum Verständnis der magnetischen Erscheinungen in festen Körpern. *Z. Phys.* **21** (1920), 613–615.

[44] 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.

[45] J. Miller and S. Sheffield, Quantum Loewner evolution. *Duke Math. J.* **165** (2016), no. 17, 3241–3378.

[46] A. Moinat and H. Weber, Space-time localisation for the dynamic Φ_3^4 model. *Comm. Pure Appl. Math.* **73** (2020), no. 12, 2519–2555.

[47] E. Nelson, A quartic interaction in two dimensions. In *Mathematical theory of elementary particles (Proc. Conf., Dedham, Mass., 1965)*, pp. 69–73, MIT Press, Cambridge, MA, 1966.

[48] E. Nelson, Construction of quantum fields from Markoff fields. *J. Funct. Anal.* **12** (1973), 97–112.

[49] 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)*, pp. 31–39, IEEE, Pittsburg, PA, 2005.

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

[51] K. Osterwalder and R. Schrader, Axioms for Euclidean Green's functions. *Comm. Math. Phys.* **31** (1973), 83–112.

[52] K. Osterwalder and R. Schrader, Axioms for Euclidean Green's functions II. *Comm. Math. Phys.* **42** (1973), 281–305.

[53] R. B. Potts, Some generalized order-disorder transformations. *Proc. Camb. Philos. Soc.* **48** (1952), 106–109.

[54] A. Sakai, Lace expansion for the Ising model. *Comm. Math. Phys.* **272** (2007), no. 2, 283–344.

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

[56] O. Schramm and S. Smirnov, On the scaling limits of planar percolation. *Ann. Probab.* **39** (2011), no. 5, 1768–1814.

[57] B. Simon, *The $P(\phi)_2$ Euclidean (quantum) field theory*. Princeton Ser. Phys., Princeton University Press, Princeton, NJ, 1974.

[58] B. Simon and R. B. Griffiths, The $(\phi^4)_2$ field theory as a classical Ising model. *Comm. Math. Phys.* **33** (1973), 145–164.

[59] 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. R. Soc. Lond. Ser. A* **322** (1971), no. 1549, 251–280.

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

MARTIN HAIRER

Imperial College London, UK, m.hairer@imperial.ac.uk