

On the Gibbs states of the noncritical Potts model on \mathbb{Z}^2

Loren Coquille, Hugo Duminil-Copin, Dmitry Ioffe and Yvan Velenik

February 10, 2013

Abstract

We prove that all Gibbs states of the q -state nearest neighbor Potts model on \mathbb{Z}^2 below the critical temperature are convex combinations of the q pure phases; in particular, they are all translation invariant. To achieve this goal, we consider such models in large finite boxes with arbitrary boundary condition, and prove that the center of the box lies deeply inside a pure phase with high probability. Our estimate of the finite-volume error term is of essentially optimal order, which stems from the Brownian scaling of fluctuating interfaces. The results hold at any supercritical value of the inverse temperature $\beta > \beta_c(q) = \log(1 + \sqrt{q})$.

Keywords: Potts model, Gibbs states, DLR equation, Aizenman-Higuchi theorem, translation invariance, interface fluctuations, pure phases

MSC2010: 60K35, 82B20, 82B24

1 Introduction

1.1 History of the problem

Since the seminal works of Dobrushin and Lanford-Ruelle [16, 29], the equilibrium states of a lattice model of statistical mechanics in the thermodynamic limit — the so-called Gibbs states — are identified with the probability measures μ that are solutions of the DLR equation,

$$\mu(\cdot) = \int d\mu(\omega) \gamma_\Lambda(\cdot | \omega), \quad \text{for all finite subsets } \Lambda \text{ of the lattice,}$$

where the probability kernel γ_Λ is the Gibbsian specification associated to the system; see [19]. Under very weak assumptions (at least for bounded spins), it can be shown that the set \mathcal{G} of all Gibbs states is a non-empty simplex. The analysis of \mathcal{G} is thus reduced to determining its extremal elements. In general, this is a very hard problem which remains essentially completely open in dimensions 3 and higher, for any nontrivial model, even in perturbative regimes.

The problem of determining all extremal Gibbs states amounts to understanding all possible local behaviors of the system. Pirogov-Sinai's theory [32, 35] often allows, at very low temperatures, to determine the pure phases of the model, i.e., the extremal, translation invariant (or periodic) Gibbs states, as perturbations of the corresponding ground states. However, it might be the case that suitable boundary conditions induce interfaces resulting in the *local* coexistence of different thermodynamic phases. That such a phenomenon can occur was first proved for the nearest-neighbor ferromagnetic (n.n.f.) Ising model on \mathbb{Z}^3 by Dobrushin [17], by considering the model in a cubic box with $+$ spins on the top half boundary of the box and $-$ spins on the bottom half (the so-called Dobrushin boundary condition). He proved that, at low enough temperatures, the induced interface is rigid — it is given by a plane with local defects — and the corresponding Gibbs state is extremal.

In two dimensions, the situation is very different. Gallavotti [18] proved, by studying the fluctuations of the corresponding interface, that the Gibbs state of the (very low temperature) n.n.f. Ising model on \mathbb{Z}^2 obtained using the Dobrushin boundary condition is the mixture $\frac{1}{2}(\mu^+ + \mu^-)$, where μ^+ and μ^- are the two pure phases of the Ising model. This was refined by Higuchi [24], who proved

that the interface, after diffusive scaling, weakly converges to a Brownian bridge at sufficiently low temperatures. These two results were then pushed to all subcritical temperatures by, respectively, Messenger and Miracle-Sole [31] and Greenberg and Ioffe [22]. A weaker but very simple and general proof of the non-extremality of the state obtained using Dobrushin boundary condition can be found in [8].

The fact that the Dobrushin boundary condition gives rise to a translation invariant Gibbs state is a strong indication that all Gibbs states of the two-dimensional Ising model should be translation invariant: because of the large fluctuations of the interfaces, a small box deep inside the system should remain, with high probability, far away from any of the interfaces that are induced by the boundary condition. Thus the possible local behaviors of the system should correspond to the pure phases.

In the late 1970s, this phenomenology was established in the celebrated works of Aizenman [1] and Higuchi [25], based on important earlier work of Russo [34]. They proved that $\mathcal{G} = \{\alpha\mu^+ + (1-\alpha)\mu^- : 0 \leq \alpha \leq 1\}$ for the n.n.f. Ising model on \mathbb{Z}^2 . Their approaches relied on many specific properties of the Ising model (in particular, GKS, Lebowitz and FKG inequalities were used in the proof). A decade ago, Georgii and Higuchi [21] devised a variant of this proof with a number of advantages. In particular, their version only relies on the FKG inequality (and some lattice symmetries), which made it possible to obtain in the same way a complete description of Gibbs states in several other models: the n.n.f. Ising model on the triangular and hexagonal lattices, the antiferromagnetic Ising model in an homogeneous field and the hard-core lattice gas. It should be emphasized that all these works deal directly with the infinite-volume system, and have only very weak implications for large finite systems. In particular, the reasoning underlying these arguments (taking the form of a proof by contradiction) remains far from the heuristics of interfaces fluctuations.

A much more general result, restricted to very low temperatures, was established by Dobrushin and Shlosman [15]. They proved that, under suitable assumptions (finite single-spin space, bounded interactions, finite number of periodic ground states), all Gibbs states are periodic, and in particular are convex combinations of the pure phases corresponding to perturbations of the ground states of the model. Their approach deals with finite systems and is closer in spirit, if not fully in practice, to the above heuristics. Namely, even though interface fluctuations play a central role in the approach of [15], the authors resort to crude low temperature surgery estimates without developing a comprehensive fluctuation theory.

Very recently, a completely different approach to the Aizenman-Higuchi result was developed by two of us [13]. This new approach, although still restricted to the n.n.f. Ising model on \mathbb{Z}^2 , presents several advantages on the former ones. Unlike [15] it does not require a very low temperature assumption, and actually holds for all sub-critical temperatures. Furthermore, it provides a quantitative, finite-volume version of the Aizenman-Higuchi theorem, with the correct rate of relaxation. Another interesting feature of the proof is that it closely follows the outlined heuristics and, consequently, should be much more robust.

In the present work, we extend the approach of [13] to n.n.f. Potts models on \mathbb{Z}^2 . As will be seen below, two major factors make the proof substantially more difficult in this case. The first one is of a physical nature: In all previous non-perturbative studies, there were only two pure phases, and thus macroscopic interfaces were always line segments. In the Potts model with 3 or more states, there are more than two phases and, consequently, interfaces are more complicated objects, elementary macroscopic interfaces being trees rather than lines. The second difficulty is of a technical nature: The positive association of Ising spins, manifested through the FKG inequality, simplified many parts of the proof in [13]. Unfortunately, this property does not hold anymore in the context of general q -state Potts models. We will therefore avoid this difficulty by reformulating the problem in terms of the random-cluster representation.

1.2 Statement of the results

Let $\Omega = \{1, \dots, q\}^{\mathbb{Z}^2}$ be the space of configurations. Let Λ be a finite subset of \mathbb{Z}^2 , and $\Lambda^c = \mathbb{Z}^2 \setminus \Lambda$ be its complement. The *finite-volume Gibbs measure* in Λ for the q -state Potts model with boundary

conditions $\sigma \in \Omega$ and at inverse-temperature $\beta > 0$ is the probability measure on Ω (with the associated product σ -algebra) defined by

$$\mathbb{P}_{\beta, \Lambda}^{\sigma}(\eta) = \begin{cases} \frac{1}{Z_{\beta, \Lambda}^{\sigma}} e^{-\beta H_{\Lambda}(\eta)} & \text{if } \eta_i = \sigma_i \text{ for all } i \in \Lambda^c \\ 0 & \text{otherwise,} \end{cases}$$

where the normalization constant $Z_{\beta, \Lambda}^{\sigma}$ is the partition function. The Hamiltonian in Λ is given by

$$H_{\Lambda}(\eta) = - \sum_{\substack{i \sim j \\ \{i, j\} \cap \Lambda \neq \emptyset}} \delta_{\eta_i, \eta_j}$$

where $i \sim j$ if i and j are nearest neighbors in \mathbb{Z}^2 . In the case of pure boundary condition $i \in \{1, \dots, q\}$, meaning that $\sigma_x = i$ for every $x \in \Lambda^c$, we denote the measure by $\mathbb{P}_{\beta, \Lambda}^{(i)}$.

For an arbitrary subset A of \mathbb{Z}^2 , let \mathcal{F}_A be the sigma-algebra generated by spins in Λ . A probability measure \mathbb{P} on Ω is an *infinite-volume Gibbs measure* for the q -state Potts model at inverse temperature β if and only if it satisfies the following DLR condition:

$$\mathbb{P}(\cdot | \mathcal{F}_{\Lambda^c})(\sigma) = \mathbb{P}_{\beta, \Lambda}^{\sigma} \quad \text{for } \mathbb{P}\text{-a.e. } \sigma, \text{ and all finite subsets } \Lambda \text{ of } \mathbb{Z}^2.$$

Let $\mathcal{G}_{q, \beta}$ be the space of infinite-volume q -state Potts measures.

Non-emptiness of $\mathcal{G}_{q, \beta}$ can be proved constructively in this model. For $i \in \{1, \dots, q\}$, $(\mathbb{P}_{\beta, \Lambda}^{(i)})_{\Lambda}$ converges when $\Lambda \nearrow \mathbb{Z}^2$ (in particular, the limit does not depend on the sequence of boxes chosen); this follows easily, e.g., from the random cluster representation. We denote by $\mathbb{P}_{\beta}^{(i)}$ the corresponding limit. It can be checked [20, Prop. 6.9] that the measures $\mathbb{P}_{\beta}^{(i)}$ ($i = 1, \dots, q$) belong to $\mathcal{G}_{q, \beta}$ and are translation invariant.

When β is less than the critical inverse temperature $\beta_c(q) = \log(1 + \sqrt{q})$ [6], it is known that there exists a unique infinite-volume Gibbs measure (in particular $\mathbb{P}_{\beta}^{(i)} = \mathbb{P}_{\beta}^{(j)}$ for every $i, j \in \{1, \dots, q\}$). The relevant values of β for a study of $\mathcal{G}_{\beta, q}$ are thus $\beta \geq \beta_c(q)$.

In the present work, we extend ideas of [13] in order to determine all infinite-volume Gibbs measures for the q -state Potts models at inverse temperature $\beta > \beta_c(q)$ on \mathbb{Z}^2 . More precisely, we show that every Gibbs state is a convex combination of infinite-volume measures with pure boundary condition:

Theorem 1.1. *For any $q \geq 2$ and $\beta > \beta_c(q)$,*

$$\mathcal{G}_{q, \beta} = \left\{ \sum_{i=1}^q \alpha_i \mathbb{P}_{\beta}^{(i)}, \text{ where } \alpha_i \geq 0, \forall i \in \{1, \dots, q\} \text{ and } \sum_{i=1}^q \alpha_i = 1 \right\}. \quad (1)$$

A straightforward yet important corollary of this theorem is the fact that any Gibbs state is invariant under translations.

Corollary 1.2. *For any $q \geq 2$ and $\beta > \beta_c(q)$, all elements of $\mathcal{G}_{q, \beta}$ are invariant under translations.*

A second important corollary is the fact that the extremal Gibbs measures (also called pure states) of the simplex $\mathcal{G}_{q, \beta}$ are the infinite-volume measures with pure boundary condition.

Corollary 1.3. *For any $q \geq 2$ and $\beta > \beta_c(q)$, the extremal elements of $\mathcal{G}_{q, \beta}$ are the $\mathbb{P}_{\beta}^{(i)}$, $i \in \{1, \dots, q\}$.*

This follows from Theorem 1.1: Define $\Delta(\beta)$ via $\mathbb{P}_{\beta}^{(i)}(\eta_0 \neq i) = (q-1)\Delta$, and observe that in the decomposition (1) of $\mathbb{P}_{\beta} \in \mathcal{G}_{q, \beta}$, the coefficient α_i equals to $\frac{\mathbb{P}_{\beta}(\eta_0=i) - \Delta}{\mathbb{P}_{\beta}^{(i)}(\eta_0=i) - \Delta}$.

Actually, our main result is stronger than Theorem 1.1. As in [13], we obtain a finite-volume, quantitative version of the latter theorem, which, together with its proof, fully vindicates the heuristics given above. For a measure μ and an integrable function f , we write $\mu[f] = \int f d\mu$.

Theorem 1.4. *Let $q \geq 2$ and $\beta > \beta_c(q)$, and set $\Lambda_n = \mathbb{Z}^2 \cap [-n, n]^2$. For any $\varepsilon > 0$ small enough, there exists $C_\varepsilon < \infty$ such that, for any boundary condition σ on $\partial\Lambda_n$, we can find $\alpha_1^n, \dots, \alpha_q^n \geq 0$ depending on (n, σ, β, q) only, such that*

$$\left| \mathbb{P}_{\Lambda_n, \beta}^\sigma[g] - \sum_{i=1}^q \alpha_i^n \mathbb{P}_\beta^{(i)}[g] \right| \leq C_\varepsilon \|g\|_\infty n^{-\frac{1}{2} + 14\varepsilon},$$

for any measurable function g of the spins in Λ_{n^ε} .

Note that the error term is essentially of the right order (which is $O(n^{-1/2})$); see [13] for a proof of this claim when $q = 2$.

The strategy of the proof is the following. We consider the conditioned random-cluster measure on Λ associated to the q -state Potts model with boundary condition σ . Boundary conditions for the Potts model get rephrased as absence of connections (in the random-cluster configuration) between specified parts of the boundary of Λ . In other words, boundary conditions for the Potts models correspond to conditioning on the existence of dual-clusters between some dual-sites on the boundary. Note that the conditioning can be very messy, since intricate boundary conditions correspond to microscopic conditioning on existence of dual-clusters. It will be seen that being a mixture of measures with pure boundary condition boils down to the fact that, with high probability, no dual-cluster connected to the boundary reaches a small box deep inside Λ (which, in particular, implies that the same is true for the Potts interfaces).

The techniques involved in the proof are two-fold. First, we use positivity of surface tension in the regime $\beta > \beta_c$, which was proved in [6], in order to get rid of the microscopic mess due to the conditioning and to show that, deep inside the box, the conditioning with respect to σ corresponds to the existence of macroscopic dual-clusters. The second part of the proof consists in proving that these clusters are very slim, and that they fluctuate in a diffusive way, so that the probability that they touch a small box centered at the origin is going to zero as the size of Λ goes to infinity. The crucial step here is the use of the Ornstein-Zernike theory of sub-critical FK clusters developed in [10].

1.3 Open problems

Before delving into the proof, let us formulate some important open problems related to the present study.

► *Critical 2d Potts models.* The behavior of two-dimensional q -state Potts models in the critical regime $\beta = \beta_c(q)$ is still widely open. It is conjectured that there is a unique Gibbs state when $q = 3$ and 4, but that, for $q \geq 5$, there is coexistence at β_c of $q + 1$ pure phases: the q low-temperature ordered pure phases and the high temperature disordered phase. This is known to be true when q is large enough [28, 30]. The extension of the latter result to every $q > 4$ remains a mathematical challenge.

► *Finite-range 2d models.* The extension of the present result, even in the Ising case $q = 2$, to general finite-range interactions still seems out of reach today. There are, at least, two main difficulties when dealing with such models: On the one hand, it is difficult to find a suitable non-perturbative definition of interfaces (the classical definitions used, e.g., in Pirogov-Sinaï theory become meaningless once the temperature is not very low); on the other hand, interfaces will not partition the system into (random) subsystems with pure boundary conditions anymore, which implies that it will be necessary to understand relaxation to pure phases from impure boundary conditions. Of course, the general *philosophy* of the approach we use should still apply.

► *The question of quasiperiodicity.* There is a general conjecture that two-dimensional models should always possess a finite number of extremal Gibbs states, all of which are periodic. In particular, this

would imply that all Gibbs states are periodic, and thus that a two-dimensional quasicrystal cannot exist (as an equilibrium state).

► *Models in higher dimensions.* Needless to say, the situation in higher dimensions is very different, due to the existence of translation non-invariant states. Even in the very low-temperature 3-dimensional n.n.f. Ising model, the set of extremal Gibbs states is not known. Note, however, that it has been proved, in the case of a d -dimensional Ising model for any $d \geq 3$, that all *translation invariant* Gibbs states are convex combinations of the two pure phases at all temperatures [7]. A similar result also holds for large enough values of q [30].

1.4 Notations

Each nearest-neighbor edge e of \mathbb{Z}^2 intersects a unique dual edge of $(\mathbb{Z}^2)^* = (\frac{1}{2}, \frac{1}{2}) + \mathbb{Z}^2$, that we denote by e^* . Consider a subgraph $G = (V, E)$ of \mathbb{Z}^2 , with vertex set V and edge set E . If E is a set of directed edges, then its dual is defined by $E^* = \{e^* : e \in E\}$. Furthermore, if G does not possess any isolated vertices, we can define the dual V^* as the endpoints of edges in E^* . Altogether, this defines a dual graph $G^* = (V^*, E^*)$.

Let Λ_n be the set of sites of $\mathbb{Z}^2 \cap [-n, n]^2$ and E_n be the set of all nearest-neighbor edges of Λ_n . The dual graph is denoted by (Λ_n^*, E_n^*) . For $m < n$, the annulus $\Lambda_n \setminus \Lambda_m$ is denoted by $A_{m,n}$. The *vertex-boundary* ∂V of a graph (V, E) is defined by $\partial V = \{x \in V : \exists y \sim x \text{ such that } y \notin V\}$. The *exterior vertex-boundary* $\partial^{\text{ext}} V$ of a graph (V, E) is defined by $\partial^{\text{ext}} V = \cup_{x \in V} \{y \notin V : y \sim x\}$. The *edge-boundary* ∂E of a graph (V, E) is the set of edges between two adjacent points of ∂V .

It will occasionally be convenient to think about ∂E_m as a closed contour in \mathbb{R}^2 or, more generally, to think about subsets of E (clusters, paths, etc) in terms of their embedding into \mathbb{R}^2 ; we shall do it without further comments in the sequel.

All constants in the sequel depend on β and q only. We shall use the notation $f = O(g)$ if there exists $C = C(\beta, q) > 0$ such that $|f| \leq C|g|$. We shall write $f = \Theta(g)$ if both $f = O(g)$ and $g = O(f)$.

2 From Potts model to random-cluster model

In this section, we relate Potts and random-cluster models. We will assume throughout this article that the reader is familiar with the basic properties of the Fortuin-Kasteleyn (FK) representation. A very concise and clear exposition including derivation of comparison inequalities could be found in [2]. Mixing properties of random cluster measures were studied in [3, 4]. There is an extensive review [20] and a book [23] on the subject. More recent results [10, 6] play an important role in our approach.

Let $G = (V(G), E(G))$ be a finite graph. An element $\omega \in \{0, 1\}^{E(G)}$ is called a configuration. An edge e is said to be open in ω if $\omega(e) = 1$ and closed if $\omega(e) = 0$. We shall work with two types of boundary conditions: **f**-free and **w**-wired. Recall that the random-cluster measure with edge-weight p and cluster-weight q on G with $*$ -boundary condition ($*$ = **f**, **w**) is given by

$$\mu_{G,p,q}^*(\omega) = \mu_G^*(\omega) = \frac{p^{\#\text{open edges}} (1-p)^{\#\text{closed edges}} q^{\#\text{* clusters}}}{Z_{G,p,q}^*},$$

where $Z_{G,p,q}^*$ is a normalizing constant and a cluster is a maximal connected component of the graph $(V(G), \{e \in E(G) : \omega(e) = 1\})$. The number $\#\text{f}$ clusters counts all the disjoint clusters, whereas the number $\#\text{w}$ clusters counts only those disjoint clusters which are not connected to the vertex boundary ∂V .

2.1 Coupling with a supercritical random-cluster model on $(\mathbb{Z}^2)^*$

We consider the q -state Potts model on the graph $(\mathbb{Z}^2)^*$ at inverse temperature $\beta > \beta_c(q)$. As the parameters β and q will always remain fixed, we drop them from the notation. Fix $\sigma \in \{1, \dots, q\}^{(\mathbb{Z}^2)^*}$.

For each n , we define the Potts measure $\mathbb{P}_{\Lambda_n^*}^\sigma$ on Λ_n^* with boundary condition σ on the vertex boundary $\partial\Lambda_n^*$.

It is a classical result (see, *e.g.*, [2, 20]) that the Potts model can be coupled with a random-cluster configuration in the following way. From a configuration of spins $\eta \in \{1, \dots, q\}^{V(\Lambda_n^*)}$, construct a percolation configuration $\omega^* \in \{0, 1\}^{E_n^*}$ by setting each edge in E_n^* to be

- closed if the two end-points have different spins,
- closed with probability $e^{-\beta}$ and open otherwise if the two end-points have the same spins.

The measure thus obtained is a random-cluster measure on $(\mathbb{Z}^2)^*$ with edge-weight $p^* = 1 - e^{-\beta}$, cluster-weight q and wired boundary condition on $\partial\Lambda_n^*$, conditioned on the following event, called $\text{Cond}_n[\sigma]$: writing $S_i = \{x \in \partial\Lambda_n^* : \sigma(x) = i\}$, the sets S_i and S_j are not connected by open edges in E_n^* , for every $i \neq j$ in $\{1, \dots, q\}$. We denote this measure by $\mu_{\Lambda_n^*}^{\text{w}}(\cdot \mid \text{Cond}_n[\sigma])$. When there is no conditioning, the random-cluster measure with wired (resp. free) boundary condition is denoted by $\mu_{\Lambda_n^*}^{\text{w}}$ (resp. $\mu_{\Lambda_n^*}^{\text{f}}$).

Reciprocally, the Potts measure can be obtained from $\mu_{\Lambda_n^*}^{\text{w}}(\cdot \mid \text{Cond}_n[\sigma])$ by assigning to every cluster a spin in $\{1, \dots, q\}$ according to the following rule:

- For every $i \in \{1, \dots, q\}$, sites connected to S_i receive the spin i ,
- The sites of a cluster which is not connected to S_i receive the same spin in $\{1, \dots, q\}$ chosen uniformly at random, independently of the spins of the other clusters.

Thanks to the connection between Potts measures and random-cluster measures, tools provided by the theory of random-cluster models can be used in this context. Note that the parameters of the corresponding random-cluster measure are supercritical ($p^* > p_c(q)$).

2.2 Coupling with the subcritical Random-Cluster model on \mathbb{Z}^2

Rather than working with the supercritical random-cluster measure on $(\mathbb{Z}^2)^*$, we will be working with its subcritical dual measure on \mathbb{Z}^2 (this is the reason for choosing to define the Potts model on $(\mathbb{Z}^2)^*$). There is a natural one-to-one mapping between $\{0, 1\}^{E_n^*}$ and $\{0, 1\}^{E_n}$. Namely, set $\omega(e) = 1 - \omega(e^*)$. In this way, both direct and dual FK configurations are defined on the same probability space. In the sequel, the same notation will be used for percolation events in direct and dual configurations. For instance, $\omega \in \text{Cond}_n[\sigma]$ means that $\omega^* \in \text{Cond}_n[\sigma]$. The corresponding direct FK measure is $\mu_{\Lambda_n}^{\text{f}}(\cdot \mid \text{Cond}_n[\sigma])$.

It is well-known [11] that this defines an FK measure with parameters q and p satisfying $pp^*/[(1-p)(1-p^*)] = q$.

Since we are working with the low temperature Potts model, the random-cluster model on $(\mathbb{Z}^2)^*$ corresponds to $p^* > p_c(q)$ so that the random-cluster model on \mathbb{Z}^2 is subcritical ($p < p_c(q)$). For this measure, $\text{Cond}_n[\sigma]$ is an increasing event which requires the existence of direct open paths disconnecting different dual S_i -s. This reduces the problem to the study of the stochastic geometry of subcritical clusters. In particular, this enables us to use known results on the subcritical model.

Let us recall the few properties we will be using in the next sections. First, there is a unique infinite-volume measure, denoted $\mu_{\mathbb{Z}^2}$. Second, there is exponential decay of connectivities in the random-cluster model with parameter $p < p_c(q)$. These two properties imply the following corollary.

Proposition 2.1. *There exists $c > 0$ such that, for n large enough and $2k \leq n \leq m$,*

$$\begin{aligned} \mu_{A_{k,n}}^{\text{w}}(\text{there exists a crossing of } A_{k,n}) &\leq e^{-cn}, \\ \mu_{\Lambda_n}^{\text{w}}(\text{there exists a cluster of cardinality } m \text{ in } \Lambda_{n/2}) &\leq e^{-cm}, \end{aligned}$$

where a crossing is a cluster of $A_{m,n}$ connecting the inner box to the outer box.

A cluster surrounding the inner box of $A_{m,n}$ inside the outer box of $A_{m,n}$ is said to be a *circuit*. Note that the existence of a dual circuit is a complementary event to the existence of a crossing between the inner and outer boxes.

Proposition 2.1 follows from the exponential decay of connectivities proved for any $p < p_c(q)$ in [6] together with the uniqueness of the infinite-volume measure (this is required to tackle wired boundary conditions, see [10, Appendix] for details). The result would not be true at criticality when q is very large, despite the fact that there is exponential decay for free boundary conditions.

Surface tension Surface tension in the supercritical dual model is the inverse correlation length in the primal sub-critical FK percolation. Let $p < p_c(q)$. The surface tension in direction x is defined by

$$\tau(x) = \tau_p(x) = - \lim_{k \rightarrow \infty} \frac{1}{k} \log \mu_{\mathbb{Z}^2}(0 \leftrightarrow [kx]),$$

where $y \leftrightarrow z$ means that y and z belong to the same connected component. We will also refer to it as the τ -distance. By Proposition 2.1, τ is equivalent to the usual Euclidean distance on \mathbb{R}^d . Furthermore, by [10] it is strictly convex, and the following sharp triangle inequality of [26, 33] holds: There exists $\rho = \rho(p) > 0$ such that

$$\tau(x) + \tau(y) - \tau(x + y) \geq \rho(|x| + |y| - |x + y|). \quad (2)$$

Define $d_\tau(A, B) = \sup_{a \in A} \inf_{b \in B} \tau(a - b)$ to be the τ -Hausdorff distance between two sets.

2.3 Reformulation of the problem in terms of the subcritical random-cluster model

Theorem 2.2. Fix $p < p_c(q)$ and let $\varepsilon \in (0, 1)$. Then, uniformly in all boundary conditions σ ,

$$\mu_{\Lambda_n}^f(\mathcal{C} \cap \Lambda_{n^\varepsilon} \neq \emptyset \mid \text{Cond}_n[\sigma]) = O(n^{-\frac{1}{2} + 14\varepsilon}) \quad (3)$$

where \mathcal{C} is the set of sites connected to the boundary $\partial\Lambda_n$.

The proof of this theorem will be the core of the paper. Before delving into the proof, let us show how it implies Theorem 1.4.

Lemma 2.3. Let $\beta > \beta_c(q)$. Then,

$$\mathbb{P}_{(\mathbb{Z}^2)^*}^f = \frac{1}{q} \sum_{i=1}^q \mathbb{P}_{(\mathbb{Z}^2)^*}^{(i)}. \quad (4)$$

Proof. Fix $\beta > \beta_c$. Note that $\mathbb{P}_{(\mathbb{Z}^2)^*}^{(i)}$ can be defined via the coupling with the random-cluster measure as follows. Let $\mu_{(\mathbb{Z}^2)^*}$ be the unique infinite-volume random-cluster measure on $(\mathbb{Z}^2)^*$. Since $p^* > p_c(q)$, this measure possesses a unique infinite cluster. The Potts measure $\mathbb{P}_{(\mathbb{Z}^2)^*}^{(i)}$ is constructed by assigning spin i to the infinite cluster, and a spin chosen uniformly at random for each finite cluster, independently of the spin of the other clusters. The Potts measure $\mathbb{P}_{(\mathbb{Z}^2)^*}^f$ can also be constructed from $\mu_{(\mathbb{Z}^2)^*}$ by assigning to each cluster (including the infinite one) a spin chosen uniformly at random, independently of the spin of the other clusters. We deduce (4) immediately.

Note that in general, $\mathbb{P}_{(\mathbb{Z}^2)^*}^{(i)}$ is constructed from the infinite-volume random-cluster measure $\mu_{(\mathbb{Z}^2)^*}^w$ while $\mathbb{P}_{(\mathbb{Z}^2)^*}^f$ is constructed from the infinite-volume random-cluster measure $\mu_{(\mathbb{Z}^2)^*}^f$. Therefore, if these two measures are different, (4) will not be valid. This is the case when $p = p_c(q)$ and q is large enough. \square




Figure 1: On the left (resp. center, right), the event \mathcal{E} (resp. $\mathcal{F}^f, \mathcal{F}^{(i)}$) is depicted.

Lemma 2.4. *There exists $c > 0$ such that, for any $n > 0$ and any subdomain Ω^* of $(\mathbb{Z}^2)^*$ containing Λ_{2n}^* ,*

$$\mathbb{P}_{\Omega^*}^f[g] = \mathbb{P}_{(\mathbb{Z}^2)^*}^f[g] + O(\|g\|_\infty e^{-cn}), \quad (5)$$

for any g depending only on spins in Λ_n^* . The same holds for pure boundary conditions $i \in \{1, \dots, q\}$.

Proof. We treat the case of the free boundary condition. The other cases follow from the same proof. Since $p < p_c(q)$, the random-cluster model on \mathbb{Z}^2 has exponential decay of connectivities. Therefore, [4, Theorem 1.7(ii)] implies the so-called ratio strong mixing property for the dual random-cluster model: If a percolation event A depends on edges from E_A and if B depends on edges from E_B , then,

$$\left| \frac{\mu_{(\mathbb{Z}^2)^*}^f(A \cap B)}{\mu_{(\mathbb{Z}^2)^*}^f(A) \mu_{(\mathbb{Z}^2)^*}^f(B)} - 1 \right| \leq \sum_{e_A \in E_A, e_B \in E_B} e^{-cd(e_A, e_B)}, \quad (6)$$

where $d(e_A, e_B)$ is a distance between edges e_A and e_B (for instance the distance between their mid-points).

Together with the observation that $\mu_{\Omega^*}^f = \mu_{(\mathbb{Z}^2)^*}^f(\cdot | \omega(e) = 0, \forall e \notin E(\Omega^*))$, this leads to

$$|\mu_{\Omega^*}^f[f] - \mu_{(\mathbb{Z}^2)^*}^f[f]| = O(e^{-cn} \mu_{(\mathbb{Z}^2)^*}^f[f]) \quad (7)$$

for any function f depending only on edges in $E_{3n/2}^*$. More generally, let F be the event that there does not exist an open crossing in the annulus $A_{n, 3n/2}$ (this corresponds to the existence of a dual circuit surrounding the origin). The complement F^c of this event has exponentially small probability by Proposition 2.1. Consider a function f depending a priori on every dual edges, but with the property that $f \mathbf{1}_F$ is measurable with respect to edges in $E_{3n/2}^*$. We immediately find that

$$\mu_{\Omega^*}^f[f] = \mu_{\Omega^*}^f[f \mathbf{1}_F] + O(\|f\|_\infty \mu_{\Omega^*}^f(F^c)) = \mu_{\Omega^*}^f[f \mathbf{1}_F] + O(\|f\|_\infty e^{-cn})$$

and similarly for $\mu_{(\mathbb{Z}^2)^*}^f[f]$, so that (7) is preserved for this class of functions.

Now, consider g depending only on spins in Λ_n^* . Via the coupling with the random-cluster model, $\mathbb{P}_{\Omega^*}^f[g]$ and $\mathbb{P}_{(\mathbb{Z}^2)^*}^f[g]$ can be seen as $\mu_{\Omega^*}^f[f]$ and $\mu_{(\mathbb{Z}^2)^*}^f[f]$ for a certain function f , depending a priori on every edge, but for which $f \mathbf{1}_F$ depends on edges in $E_{3n/2}^*$ only (on the event F , the dual connections between vertices of Λ_n^* are determined by edges in $E_{3n/2}^*$). We conclude that

$$|\mathbb{P}_{\Omega^*}^f[g] - \mathbb{P}_{(\mathbb{Z}^2)^*}^f[g]| = |\mu_{\Omega^*}^f[f] - \mu_{(\mathbb{Z}^2)^*}^f[f]| = O(\|f\|_\infty e^{-cn}).$$

□

Proof of Theorem 1.4. Fix $n > 0$ and a boundary condition σ on $\partial\Lambda_n$. Fix $\varepsilon > 0$ small.

We consider the coupling (η, ω) (the measure is denoted by \mathbf{P}) with marginals $\mathbb{P}_{\Lambda_n}^\sigma$ and $\mu_{\Lambda_n}^f(\cdot \mid \text{Cond}_n[\sigma])$ described in the previous section. Let \mathcal{E} be the event that ω contains an open crossing in $A_{2n^\varepsilon, n}$. Let \mathcal{F}^f be the event that ω contains an open circuit in $A_{2n^\varepsilon, n}$. Let $\mathcal{F}^{(i)}$ be the event that ω contains neither an open crossing nor an open circuit in $A_{2n^\varepsilon, n}$, and that $(\Lambda_{2n^\varepsilon})^*$ is connected in the dual configuration to S_i . Note that

$$\mathbf{P}(\mathcal{E}) = \mu_{\Lambda_n}^f(\mathcal{E} \mid \text{Cond}_n[\sigma]) = O(n^{-\frac{1}{2}+14\varepsilon}),$$

by applying Theorem 2.2.

- (conditioning on \mathcal{F}^f). Let Γ^* be the connected component of $\partial\Lambda_n^*$ in ω^* . Denote the connected component of $\Lambda_{2n^\varepsilon}^*$ in $\Lambda_n^* \setminus \Gamma^*$ by Ω^* . We have $\Lambda_{2n^\varepsilon}^* \subset \Omega^*$. Conditioning on Γ^* we infer, using (5) and (4) that

$$\begin{aligned} \mathbf{P}(g \mid \mathcal{F}^f) &= \mathbf{P}(\mathbb{P}_{\Omega^*}^f[g] \mid \mathcal{F}^f) = \mathbb{P}_{(\mathbb{Z}^2)^*}^f[g] + O(\|g\|_\infty e^{-cn^\varepsilon}) \\ &= \frac{1}{q} \sum_{i=1}^q \mathbb{P}_{(\mathbb{Z}^2)^*}^{(i)}[g] + O(\|g\|_\infty e^{-cn^\varepsilon}). \end{aligned}$$

- (conditioning on $\mathcal{F}^{(i)}$). In this case, let us condition on the connected cluster Γ of $\partial\Lambda_n$. We view Γ as the set of bonds. Define Ω^* as the connected component of $\Lambda_{2n^\varepsilon}^*$ in $(E_n \setminus \Gamma)^*$. By construction, $\Lambda_{2n^\varepsilon}^* \subset \Omega^*$ and $\Omega^* \cap S_i \neq \emptyset$. Consequently, using (5) once again, we obtain

$$\mathbf{P}(g \mid \mathcal{F}^{(i)}) = \mathbf{P}(\mathbb{P}_{\Omega^*}^{(i)}[g] \mid \mathcal{F}^{(i)}) = \mathbb{P}_{(\mathbb{Z}^2)^*}^{(i)}[g] + O(\|g\|_\infty e^{-cn^\varepsilon}).$$

By summing all these terms,

$$\begin{aligned} \mathbb{P}_{\Lambda_n}^\sigma[g] &= \mathbf{P}[g] = \mathbf{P}[g \mid \mathcal{E}] \mathbf{P}[\mathcal{E}] + \mathbf{P}[g \mid \mathcal{F}^f] \mathbf{P}[\mathcal{F}^f] + \sum_{i=1}^q \mathbf{P}[g \mid \mathcal{F}^{(i)}] \mathbf{P}[\mathcal{F}^{(i)}] \\ &= \sum_{i=1}^q \left(\frac{1}{q} \mathbf{P}[\mathcal{F}^f] + \mathbf{P}[\mathcal{F}^{(i)}] \right) \mathbb{P}_{(\mathbb{Z}^2)^*}^{(i)}[g] + O(\|g\|_\infty n^{-\frac{1}{2}+14\varepsilon}), \end{aligned}$$

which implies the claim readily. \square

3 Macroscopic flower domains

In the box Λ_n , the conditioning on $\text{Cond}_n[\sigma]$ can be very messy. Indeed, as we mentioned before, it forces the existence of open paths separating the sets S_i . For instance, the number of such paths forced by an alternating boundary condition $1, 2, \dots, q, 1, 2, \dots$ is necessarily of order n .

We first show that, no matter what the boundary condition σ is, with high probability only a bounded number of such interfaces is capable of reaching an inner box Λ_m , where m is a fraction of n . Furthermore, we shall argue that the number of sites in $\partial\Lambda_m$ which are connected to the original $\partial\Lambda_n$ is uniformly bounded. In terms of the original Potts model, this corresponds to the existence, with high probability, of a domain including the box Λ_m for which the boundary condition contains a uniformly bounded number of spin changes. This will be called a flower domain below.

3.1 Definition of flower domains

Let $m < n$. For a configuration ω , let $\mathbb{C}_{m,n} = \mathbb{C}_{m,n}(\omega)$ be the set of sites connected to $\partial\Lambda_n$ in $\omega \cap (E_n \setminus E_m)$. Define the set of *marked vertices* by

$$\mathbb{G}_{m,n} = \mathbb{G}_{m,n}(\omega) = \mathbb{C}_{m,n} \cap \partial\Lambda_m.$$




Figure 2: Description of a flower domain $\mathcal{D}_{m,n}$ (light grey area). The blue points are locations of spin changes (i.e. separation between sets S_i), the red points constitute $\mathbb{G}_{m,n}$, the solid black lines in the annulus $\Lambda_n \setminus \Lambda_m$ constitute $\mathcal{C}_{m,n}$.

The set $\mathbb{G}_{m,n} \cup (\Lambda_n \setminus \mathcal{C}_{m,n})$ may have several connected components, exactly one of them containing Λ_m . Let us call the latter the *flower domain* $\mathcal{D}_{m,n} = \mathcal{D}_{m,n}(\omega)$ rooted at m . Note that $\mathbb{G}_{m,n} = \partial \mathcal{D}_{m,n} \cap \partial \Lambda_m$, that is marked sites are unambiguously determined by the corresponding flower domains.

Fix a configuration ω . Let $\mathcal{C} = \mathcal{C}_{m,n}(\omega)$ and let $\mathcal{D} = \mathcal{D}_{m,n}(\omega)$ be the corresponding flower domain. Let also $\mathbb{G} = \mathbb{G}_{m,n}(\omega)$. By construction, the restriction of the conditional measure $\mu_{\Lambda_n}^f(\cdot | \mathcal{C}_{m,n} = \mathcal{C})$ to $\{0, 1\}^{\mathcal{E}_{\mathcal{D}}}$, where $\mathcal{E}_{\mathcal{D}}$ is the set of edges of \mathcal{D} , is the FK measure with free boundary conditions on $\partial \mathcal{D} \setminus \mathbb{G}$ and wiring between sites of \mathbb{G} inherited from connections in \mathcal{C} . We denote this restricted conditional measure as $\mu_{\mathcal{D}}^{\text{flower}}$. We also set $\mathcal{C}_{\mathbb{G}}$ for the connected component of \mathbb{G} in the restriction of ω to $\mathcal{E}_{\mathcal{D}}$.

3.2 Cardinality of $\mathbb{G}_{m,n}$

Flower domains have typically small sets $\mathbb{G}_{m,n}$, as the following proposition shows.

Proposition 3.1. *There exists $M > 0$ such that for any $\delta > 0$*

$$\mu_{\Lambda_n}^f \left(\exists m \in \left[\frac{\delta n}{3}, \delta n \right] : |\mathbb{G}_{m,n}| \leq M \mid \text{Cond}_n[\sigma] \right) \geq 1 - e^{-\delta n}, \quad (8)$$

uniformly in σ and n sufficiently large.

The notation M will now be reserved for an integer $M > 0$ satisfying the previous proposition. We shall prove this Proposition for $\delta = 1$; the general case follows by a straightforward adaptation.

Definition 3.2. *Let \mathcal{E}_r be the event that there exist r disjoint crossings of $A_{n/3, n/2}$.*

Lemma 3.3. *For all $r \geq 1$ and $n > 0$,*

$$\mu_{\Lambda_n}^f(\mathcal{E}_r) \leq e^{-crn},$$

where $c > 0$ is defined in Proposition 2.1.

Proof. We prove that for all $r \geq 1$ and $n > 0$,

$$\mu_{\Lambda_n}^f(\mathcal{E}_r) \leq \left(\mu_{A_{n/3, n/2}}^w(\mathcal{E}_1) \right)^r. \quad (9)$$

The conclusion will then follow easily, since Proposition 2.1 implies that $\mu_{A_{n/3, n/2}}^w(\mathcal{E}_1) \leq \exp(-cn)$.

In order to prove (9), we proceed by induction. First, note that $\mu_{\Lambda_n}^f$ restricted to $A_{n/3, n/2}$ is stochastically dominated by $\mu_{A_{n/3, n/2}}^w$.

Let $r \geq 1$ and consider $\mu_{\Lambda_n}^f(\mathcal{E}_{r+1}|\mathcal{E}_r)$. We number the vertices of $\partial\Lambda_n = \{x_1, \dots, x_{4n+4}\}$ in clockwise order, starting at the bottom right corner. Let k be the smallest number such that there are r crossings among the clusters containing x_1, \dots, x_k . Denote by \mathcal{S} the union of these clusters (which may contain isolated vertices). Observe that all edges in $A_{n/3, n/2} \setminus \mathcal{S}$ which are incident to vertices of \mathcal{S} are closed. Therefore, the conditional measure $\mu_{\Lambda_n}^f(\cdot|_{A_{n/3, n/2} \setminus \mathcal{S}}|\mathcal{S})$ is stochastically dominated by $\mu_{A_{n/3, n/2}}^w(\cdot|_{A_{n/3, n/2} \setminus \mathcal{S}})$. In both instances above, the symbol $\nu(\cdot|_B)$ means the restriction of ν to edges of the graph with the vertex set B . As a result, the probability, under $\mu_{\Lambda_n}^f(\cdot|_{A_{n/3, n/2} \setminus \mathcal{S}}|\mathcal{S})$, that there exists a crossing of $A_{n/3, n/2}$ is smaller than $\mu_{A_{n/3, n/2}}^w(\mathcal{E}_1)$. We obtain

$$\begin{aligned} \mu_{\Lambda_n}^f(\mathcal{E}_{r+1}) &= \mu_{\Lambda_n}^f(\mathcal{E}_{r+1}|\mathcal{E}_r)\mu_{\Lambda_n}^f(\mathcal{E}_r) = \mu_{\Lambda_n}^f[\mu_{\Lambda_n}^f(\mathcal{E}_{r+1}|\mathcal{S})]\mu_{\Lambda_n}^f(\mathcal{E}_r) \\ &\leq \mu_{A_{n/3, n/2}}^w(\mathcal{E}_1)\mu_{\Lambda_n}^f(\mathcal{E}_r) \leq \mu_{A_{n/3, n/2}}^w(\mathcal{E}_1)^{r+1}. \end{aligned}$$

□

Proof of Proposition 3.1. Obviously,

$$\mu_{\Lambda_n}^f(\forall m \in [\frac{n}{3}, \frac{n}{2}] : |\mathbb{G}_{m, n}| > M \mid \text{Cond}_n[\sigma]) \leq \frac{\mu_{\Lambda_n}^f(\forall m \in [\frac{n}{3}, \frac{n}{2}] : |\mathbb{G}_{m, n}| > M)}{\mu_{\Lambda_n}^f(\text{Cond}_n[\sigma])}. \quad (10)$$

Let us bound from below the denominator of (10). If all the edges of ∂E_n are open, then $\text{Cond}_n[\sigma]$ occurs. Moreover, the measure $\mu_{\Lambda_n}^f$ stochastically dominates independent Bernoulli edge percolation on $\{0, 1\}^{E_n}$ with $\tilde{p} = p/(p + (q-1)p)$, see [2, Theorem 4.1]. We deduce

$$\mu_{\Lambda_n}^f(\text{Cond}_n[\sigma]) \geq \mu_{\Lambda_n}^f(\text{all the edges in } \partial E_n \text{ are open}) \geq \tilde{p}^{8n}. \quad (11)$$

Let us now bound from above the numerator of (10). First,

$$\mu_{\Lambda_n}^f(\forall m \in [\frac{n}{3}, \frac{n}{2}] : |\mathbb{G}_{m, n}| > M) \leq \mu_{\Lambda_n}^f(|\mathbb{C}_{n/3, n} \cap A_{n/3, n/2}| \geq Mn/6).$$

Fix $R > 0$. If $|\mathbb{C}_{n/3, n} \cap A_{n/3, n/2}| \geq Mn/6$, either $A_{n/3, n/2}$ contains more than R crossings or one of the crossings has cardinality larger than $Mn/(6R)$. Proposition 2.1 implies that the probability of having clusters with size larger than $Mn/(6R)$ in $\Lambda_{n/2}$ is smaller than $\exp[-cMn/(6R)]$ for n large enough. Lemma 3.3 together with (10) implies that, for n large enough,

$$\mu_{\Lambda_n}^f(\forall m \in [\frac{n}{3}, \frac{n}{2}] : |\mathbb{G}_{m, n}| > M \mid \text{Cond}_n[\sigma]) \leq \tilde{p}^{-8n}[e^{-cRn} + e^{-cMn/(6R)}] \leq e^{-n},$$

provided that R and M be sufficiently large. □

3.3 Reduction to FK measures on flower domains with free boundary condition

We define

$$\mathcal{M}_n = \max\{m \leq n : |\mathbb{G}_{m, n}| \leq M\}, \quad (12)$$

where the maximum is set to be equal to ∞ if there is no $m \leq n$ such that $|\mathbb{G}_{m, n}| \leq M$. With this notation, we actually proved that $\mathcal{M}_n \in [\frac{n}{3}, n]$ with probability bounded below by $1 - e^{-n}$.

Let \mathcal{C} be a possible realization of $\mathbb{C}_{m, n}$ and $\mathcal{D} = \mathcal{D}_{m, n}$ be the corresponding flower domain. The restriction of $\mu_{\Lambda_n}^f(\cdot \mid \mathcal{M}_n = m; \mathbb{C}_{m, n} = \mathcal{C})$ to \mathcal{D} is $\mu_{\mathcal{D}}^{\text{flower}}$. Furthermore,

$$\text{Cond}_n[\sigma] \cap \{\mathcal{M}_n = m\} \cap \{\mathbb{C}_{m, n} = \mathcal{C}\}$$

is a product event $\Omega_{\sigma, \mathcal{C}} \times \{\mathcal{M}_n = m; \mathbb{C}_{m, n} = \mathcal{C}\}$, where $\Omega_{\sigma, \mathcal{C}} \subset \{0, 1\}^{\mathcal{D}}$. Then

$$\mu_{\Lambda_n}^f(\mathbb{C} \cap \Lambda_{n^\varepsilon} \neq \emptyset \mid \text{Cond}_n[\sigma]; \mathcal{M}_n = m; \mathbb{C}_{m, n} = \mathcal{C}) = \mu_{\mathcal{D}}^{\text{flower}}(\mathbb{C}_{\mathbb{G}} \cap \Lambda_{n^\varepsilon} \neq \emptyset \mid \Omega_{\sigma, \mathcal{C}}). \quad (13)$$

The event $\Omega_{\sigma, \mathcal{C}}$ has an obvious structure. It corresponds to the existence of certain connections between different sites of $\mathbb{G} = \mathcal{C} \cap \partial\Lambda_m = \mathcal{D} \cap \partial\Lambda_m$. More precisely, let $\mathcal{P}_{\mathbb{G}}$ be the collection of different partitions of \mathbb{G} . Elements of $\mathcal{P}_{\mathbb{G}}$ are of the form $\underline{\mathbb{G}} = (\mathbb{G}_1, \dots, \mathbb{G}_\ell)$. Define

$$\Omega_{\underline{\mathbb{G}}} = \bigcap_i \bigcap_{u, v \in \mathbb{G}_i} \{u \leftrightarrow v\} \subset \{0, 1\}^{\mathcal{E}_{\mathcal{D}}}.$$

Let us say that a partition $\underline{\mathbb{G}}$ is compatible with $\Omega_{\sigma, \mathcal{C}}$ if $\Omega_{\underline{\mathbb{G}}} \subseteq \Omega_{\sigma, \mathcal{C}}$. Note that we do not rule out that some elements \mathbb{G}_i of a partition $\underline{\mathbb{G}}$ are singletons. If \mathbb{G}_i is a singleton, then $\bigcap_{u, v \in \mathbb{G}_i} \{u \leftrightarrow v\}$ is, of course, a sure event, which could be dropped from the definition of $\Omega_{\underline{\mathbb{G}}}$. In other words, only non-singleton elements of $\underline{\mathbb{G}}$ are relevant for $\Omega_{\underline{\mathbb{G}}}$. Also note that the events $\Omega_{\underline{\mathbb{G}}}$ do not have to be disjoint. Still, for any σ ,

$$\Omega_{\sigma, \mathcal{C}} = \bigcup_{\underline{\mathbb{G}} \in \mathcal{P}'_{\mathbb{G}}} \Omega_{\underline{\mathbb{G}}},$$

where the set $\mathcal{P}'_{\mathbb{G}}$ corresponds to partitions which are compatible with the occurrence of the event $\Omega_{\sigma, \mathcal{C}}$, and which are *maximal* in the sense that one cannot find a finer partition which would be still compatible with $\Omega_{\sigma, \mathcal{C}}$.

The previous section implies the following reduction, which we will now consider for the rest of this work.

Proposition 3.4. *Fix $\delta > 0$. Then, writing B_k for the k^{th} Bell number, which counts the number of partitions of a set of k elements,*

$$\mu_{\Lambda_n}^f(\mathbb{C} \cap \Lambda_{n^\varepsilon} \neq \emptyset \mid \text{Cond}_n[\sigma]) \leq e^{-\delta n} + B_M q^M \max \mu_{\mathcal{D}}^f(\mathbb{C}_{\mathbb{G}} \cap \Lambda_{n^\varepsilon} \neq \emptyset \mid \Omega_{\underline{\mathbb{G}}}), \quad (14)$$

for all boundary conditions σ and n sufficiently large. The above maximum is over all flower domains \mathcal{D} rooted at $m \in [\frac{n}{3}, n]$ with at most $|\mathbb{G}| \leq M$ marked points, and over all partitions $\underline{\mathbb{G}} \in \mathcal{P}'_{\mathbb{G}}$.

Above the term q^M comes from the fact that the elements of \mathbb{G} are possibly wired together. It then bounds the Radon-Nikodym derivative between measures $\mu_{\mathcal{D}}^{\text{flower}}$ and $\mu_{\mathcal{D}}^f$. The quantity B_M bounds from above the number of sub-partitions of \mathbb{G} (the events $\Omega_{\underline{\mathbb{G}}}$ being not necessarily disjoint).

4 Macroscopic structure near the center of the box

This section studies the macroscopic structure of the set \mathbb{C} of sites connected to the boundary of Λ_n . Its main result, Proposition 4.2 below, implies that on a sufficiently small scale $\delta > 0$, the intersection $\mathbb{C} \cap \Lambda_k$ for boxes with $k \in [\frac{\delta n}{3}, \delta n]$ is with an overwhelming probability either empty, or close to a segment, or close to a tripod (three segments coming out from a point).

Before starting, note that Proposition 3.4 enables us to restrict attention to a flower domain $\mathcal{D} = \mathcal{D}_{m, n}$ with $m \in [\frac{n}{3}, n]$ and $|\mathbb{G}_{m, n}| \leq M$. We set $\mathbb{G} = \mathbb{G}_{m, n}$. We now fix this flower domain and work under $\mu_{\mathcal{D}}^f(\cdot \mid \Omega_{\underline{\mathbb{G}}})$ for some $\underline{\mathbb{G}} \in \mathcal{P}'_{\mathbb{G}}$. All constants in this section are independent of $\mathcal{D}_{m, n}$ and $\underline{\mathbb{G}}$ as long as $|\mathbb{G}_{m, n}| \leq M$. We will often recall this independence by using the expression “uniformly in $(\mathcal{D}, \underline{\mathbb{G}})$ with $|\mathbb{G}| \leq M$ ”.

Define $\mathbb{C}_{k, \mathbb{G}}$ to be the set of edges connected to \mathbb{G} in $\mathcal{D} \setminus \Lambda_k$ (which can consist of several connected components). Note that $\mathbb{G}_{k, n} = \mathbb{C}_{k, n} \cap \partial\Lambda_k = \mathbb{C}_{k, \mathbb{G}} \cap \partial\Lambda_k$. Given $v_1, v_2 \in \mathbb{R}^2$, we define $[v_1, v_2]$ to be the line segment with endpoints v_1 and v_2 , and $\angle(v_1, v_2)$ to be the angle between v_1 and v_2 , seen as vectors in the plane. We refer to Fig. 3 for an illustration of the following definitions.

Definition 4.1. *For $k < m$, $\nu > 0$ and $\ell = 1, 2, 3$, let us say that $E_{\nu, k}^\ell \subset \{0, 1\}^{\mathcal{E}_{\mathcal{D}}}$ occurs if S_ℓ below happens:*

- S1. $\mathbb{G}_{k, n} = \emptyset$.




Figure 3: Description of the events $E_{\nu,k}^{\ell}$, $\ell = 1, 2, 3$ from left to right. The set $\mathbb{G}_{k,n}$, partitioned into $\mathbb{V}_{k,n}^{\ell}$, $\ell = 1, 2, 3$, is indicated in red.

- S2. $\mathbb{G}_{k,n} = \mathbb{V}_{k,n}^1 \cup \mathbb{V}_{k,n}^2$, where $\mathbb{V}_{k,n}^1, \mathbb{V}_{k,n}^2$ are two disjoint sets of τ -diameter less than or equal to νk . Moreover,
- Each of the sets $\mathbb{V}_{k,n}^1$ and $\mathbb{V}_{k,n}^2$ is connected in $\mathbb{C}_{k,\mathbb{G}}$.
 - For any two vertices $v_i \in \mathbb{V}_{k,n}^i$; $i = 1, 2$, we have $[v_1, v_2] \cap \Lambda_{k/2} \neq \emptyset$.
- S3. $\mathbb{G}_{k,n} = \mathbb{V}_{k,n}^1 \cup \mathbb{V}_{k,n}^2 \cup \mathbb{V}_{k,n}^3$, where $\mathbb{V}_{k,n}^1, \mathbb{V}_{k,n}^2$ and $\mathbb{V}_{k,n}^3$ are disjoint sets with τ -diameter less than or equal to νk . Moreover,
- Each of the sets $\mathbb{V}_{k,n}^1, \mathbb{V}_{k,n}^2$ and $\mathbb{V}_{k,n}^3$ is connected in $\mathbb{C}_{k,\mathbb{G}}$,
 - For any choice of $v_i \in \mathbb{V}_{k,n}^i$; $i = 1, 2, 3$. there exists $x \in \Lambda_{k/2}$ such that $\mathcal{T} = \{v_1, v_2, v_3; x\}$ is a Steiner tripod (see Definition 4.7 below). In particular, as it follows from P2 of Proposition 4.3 below, $\angle(v_i - x, v_j - x) > \frac{\pi}{2} + \eta$ for every $i \neq j$.

We are now in a position to state the main proposition.

Proposition 4.2. *For any $\nu > 0$, there exist $\delta = \delta(\nu, M) > 0$ and $\kappa = \kappa(\nu, M) > 0$ such that*

$$\mu_{\mathcal{D}}^f \left(\bigcup_{k \geq \delta n} (E_{\nu,k}^1 \cup E_{\nu,k}^2 \cup E_{\nu,k}^3) \cap \{|\mathbb{G}_{k,n}| \leq M\} \mid \Omega_{\mathbb{G}} \right) \geq 1 - e^{-\kappa n}, \quad (15)$$

uniformly in $(\mathcal{D}, \mathbb{G})$ with $|\mathbb{G}| \leq M$.

The proof of Proposition 4.2 comprises two steps: First, we show that the implied geometric structure is characteristic of deterministic objects called *Steiner forests*. Then, we show that, with high $\mu_{\mathcal{D}}^f(\cdot \mid \Omega_{\mathbb{G}})$ -probability, the cluster $\mathbb{C}_{\mathbb{G}}$ sits in the vicinity of one such forest.

4.1 Steiner forests

Note that for every m the set \mathcal{K}_m of all compact subsets of Λ_m is a Polish space with respect to the d_{τ} -distance.

We now recall the concept of Steiner forest. Consider $E \subseteq \partial\Lambda_m$ with $|E| \leq M$. Let $\underline{E} = (E_1, \dots, E_i)$ be a partition of E and $\Omega_{\underline{E}}$ be the set of compact subsets of \mathbb{R}^2 such that E_j is included in one of their connected components for every $j \in \{1, \dots, i\}$. For the trivial partition $\underline{E} = \{E\}$, we shall write Ω_E .




Figure 4: A non trivial Steiner forest with a partition $\underline{E} = (E_1, E_2)$ with $E_1 = \{u_1, \dots, u_4\}$ and $E_2 = \{u_5, u_6\}$.

For a compact $\mathcal{S} \subset \mathbb{R}^2$, let $\tau(\mathcal{S})$ be the (one-dimensional) Hausdorff measure of \mathcal{S} in the τ -norm. Explicitly,

$$\tau(\mathcal{S}) = \lim_{\varepsilon \rightarrow 0} \inf \left\{ \sum \text{diam}_\tau(A_i) : \mathcal{S} \subseteq \cup A_i, \text{diam}_\tau(A_i) \leq \varepsilon \right\}, \quad (16)$$

where $\text{diam}_\tau(A) = \sup\{\tau(x - y) : x, y \in A\}$. Define the set of *Steiner forests* by

$$\Omega_{\underline{E}}^{\min} = \left\{ \mathcal{F} \in \Omega_{\underline{E}} : \tau(\mathcal{F}) = \min_{\mathcal{S} \in \Omega_{\underline{E}}} \tau(\mathcal{S}) \right\}.$$

We set

$$\tau_{\underline{E}} = \min_{\mathcal{S} \in \Omega_{\underline{E}}} \tau(\mathcal{S}) = \tau(\mathcal{F}),$$

for any Steiner forest $\mathcal{F} \in \Omega_{\underline{E}}^{\min}$.

In the sequel we shall work only with Steiner forests $\mathcal{F} \in \Omega_{\underline{E}}^{\min}$, when \underline{E} is a partition of a set $E \subset \partial\Lambda_m$ of cardinality $|E| \leq M$. Let $\Omega_{M,m}^{\min}$ be the collection of all such forests.

Proposition 4.3. *Fix $M > 0$. The following properties hold uniformly in m , in finite subsets $E \subseteq \partial\Lambda_m$ with $|E| \leq M$ and in partitions \underline{E} of E :*

- P1. **(Number of Steiner forests and compactness of $\Omega_{M,m}^{\min}$)** *There exists $k = k(M) < \infty$ such that $|\Omega_{\underline{E}}^{\min}| \leq k$. The set $\Omega_{M,m}^{\min}$ is a compact subset of (\mathcal{K}_m, d_τ) .*
- P2. **(Structure of Steiner forests)** *The sets $\mathcal{F} \in \Omega_{\underline{E}}^{\min}$ are forests (that is collections of disjoint trees). Each inner node (that is not belonging to E) of such \mathcal{F} has degree 3. Furthermore, there exists an $\eta > 0$ such that the angle between two edges incident to an inner node of \mathcal{F} is always larger than $\frac{\pi}{2} + \eta$.*
- P3. **(Well separateness of trees)** *There exists $\delta_1 = \delta_1(M) > 0$ such that any $\mathcal{F} \in \Omega_{\underline{E}}^{\min}$ satisfies:*
 - (a) *for any Steiner tree $\mathcal{T} \in \mathcal{F}$, two different nodes of \mathcal{T} in $\Lambda_{m/2}$ are at d_τ -distance at least $\delta_1 m$ of each other;*
 - (b) *if \mathcal{T}_1 and \mathcal{T}_2 are two disjoint trees of \mathcal{F} , then $d_\tau(\mathcal{T}_1 \cap \Lambda_{m/2}, \mathcal{T}_2 \cap \Lambda_{m/2}) \geq \delta_1 m$.*
- P4. **(Stability)** *For any $\delta_2 > 0$, there exists $\kappa_2 = \kappa_2(\delta_2, M) > 0$ such that, for any $|E| \leq M$, any partition \underline{E} of E and any $\mathcal{S} \in \Omega_{\underline{E}}$,*

$$\tau(\mathcal{S}) \leq \tau_{\underline{E}} + \kappa_2 m \quad \text{implies} \quad \min_{\mathcal{F} \in \Omega_{M,m}^{\min}} d_\tau(\mathcal{S}, \mathcal{F}) < \delta_2 m. \quad (17)$$

Proof. We shall be rather sketchy since the arguments are presumably well understood. We shall consider the case $m = 1$ (the general case follows by homogeneity).

Let us start with P4. The functional τ in (16) is lower semi-continuous on (\mathcal{K}_1, d_τ) and has compact level sets (meaning sets of the form $\{\mathcal{S} : \tau(\mathcal{S}) \leq R\}$). See, for instance, [14, Proposition 3.1], where these facts are explained for the inverse correlation length of sub-critical Bernoulli bond percolation.

Assume that P4 is wrong. Then there exists $\delta > 0$ and two sequences; \underline{E}_k and $\mathcal{S}_k \in \Omega_{\underline{E}_k}$, such that

$$\tau(\mathcal{S}_j) < \tau_{\underline{E}_k} + \frac{1}{k} \quad \text{but} \quad \min_{\mathcal{F} \in \Omega_{M,m}^{\min}} d_\tau(\mathcal{S}_j, \mathcal{F}) > \delta.$$

Since $|\underline{E}_k| \leq M$, the sequence $\tau_{\underline{E}_k}$ is bounded. Hence $\{\mathcal{S}_j\}$ is precompact. Possibly passing to subsequence we may assume that \underline{E}_k converges to \underline{E} (points might collapse, but this is irrelevant since this preserves $|\underline{E}| \leq M$), and that \mathcal{S}_k converges to $\mathcal{S} \in \Omega_{\underline{E}}$. Both convergence are, of course, in the sense of Hausdorff distance. By minimality it is evident that $\tau_{\underline{E}} = \lim \tau_{\underline{E}_k}$. By lower-semicontinuity $\tau(\mathcal{S}) \leq \liminf \tau(\mathcal{S}_k)$. Which means that $\mathcal{S} \in \Omega_{\underline{E}}^{\min}$. A contradiction.

A proof of the first assertion of P1 can be found in [12, Theorem 1]. Compactness of $\Omega_{M,1}^{\min}$ follows from compactness of level sets of τ and the fact that if $\mathcal{F}_k \in \Omega_{\underline{E}_k}^{\min}$ converges to $\mathcal{F} \in \Omega_{\underline{E}}$, then, as was already mentioned above, $\tau_{\underline{E}} = \lim \tau_{\underline{E}_k}$, and hence, by the lower-semicontinuity of τ , $\mathcal{F} \in \Omega_{\underline{E}}^{\min}$.

A proof of P2 can be found in [5].

Let us turn to the proof of P3. For trivial partitions, Steiner forests are trees. Now, assume that there exists a sequence of Steiner trees $\mathcal{T}_k \in \Omega_{E_k}^{\min}$ such that \mathcal{T}_k contains at least two inner nodes in $\Lambda_{1/2}$ at distance less or equal $\frac{1}{k}$. There is no loss of generality to assume that the sequence \mathcal{T}_k converges to some \mathcal{T}^* . As it follows from P4, $\mathcal{T}^* \in \Omega_{E^*}^{\min}$, where E^* is the corresponding limit of E_k . Obviously, $|E^*|$ is still less or equal to M , since boundary points can only collapse under the limiting procedure.

The total number of nodes of each of \mathcal{T}_k is uniformly bounded above. Hence by our assumption we can choose a number $\ell \geq 2$, a point $x \in \Lambda_{1/2}$, a radius $\varepsilon > 0$ and a sequence $\nu(k) \rightarrow 0$, so that

- (a) each of \mathcal{T}_k contains ℓ nodes in $\Lambda_{\nu(k)}(x) = x + \Lambda_{\nu(k)}$;
- (b) none of \mathcal{T}_k contains nodes in the annulus $A_{\nu(k),\varepsilon}(x)$.

Then the restriction of \mathcal{T}_k to $\Lambda_\varepsilon(x)$ is a Steiner tree, whereas the cardinality of the intersection $|\partial\Lambda_\varepsilon(x) \cap \mathcal{T}_k| = \ell + 2$. By the minimality of \mathcal{T}_k the points of $\partial\Lambda_\varepsilon(x) \cap \mathcal{T}_k$ are uniformly separated. Consequently, $|\partial\Lambda_\varepsilon(x) \cap \mathcal{T}^*| = \ell + 2 > 3$. We infer that the degree of x in the Steiner tree \mathcal{T}^* is $\ell + 2 > 3$, which is impossible by P2. This proves P3(a).

Consider now two disjoint Steiner trees $\mathcal{T}_1 \in \Omega_{E_1}^{\min}$ and $\mathcal{T}_2 \in \Omega_{E_2}^{\min}$, such that the forest $\{\mathcal{T}_1, \mathcal{T}_2\}$ belongs to $\Omega_{\{E_1, E_2\}}^{\min}$. By the strict convexity of τ , the trees are confined to their convex envelopes: $\mathcal{T}_i \in \text{co}(E_i)$ for $i = 1, 2$. Thus if both trees are disjoint and intersect $\Lambda_{1/2}$, it follows that $\text{co}(E_1) \cap \text{co}(E_2) = \emptyset$. Consequently, there exist $u_1, v_1 \in E_1$ and $u_2, v_2 \in E_2$, such that \mathcal{T}_1 lies below the interval $[u_1, v_1]$ and \mathcal{T}_2 lies above the interval $[u_2, v_2]$ (notions of above and below are with respect to the directions of normals). We are now facing two cases:

- \mathcal{T}_1 or \mathcal{T}_2 has an inner node in $\Lambda_{2/3}$. By P2, inner nodes are of degree three and angles between edges incident to inner nodes are at most $\pi - 2\eta$. This pushes inner nodes of \mathcal{T}_i away from $[u_i, v_i]$ uniformly in \mathcal{T}_1 and \mathcal{T}_2 . In such a case, P3 is satisfied.
- Both \mathcal{T}_1 and \mathcal{T}_2 do not contain nodes in $\Lambda_{2/3}$, but each contains an edge which crosses $\Lambda_{1/2}$. Having such edges close to each other (and hence running essentially in parallel across $\Lambda_{1/2}$) is easily seen to be incompatible with the minimality of \mathcal{F} .

This achieves the proof of P3(b). □

4.2 Forest skeleton of the cluster $C_{\mathbb{G}}$

Let $\underline{\mathbb{G}}$ be a partition of \mathbb{G} . We now aim to show that, under $\mu_{\mathcal{D}}^f(\cdot | \Omega_{\underline{\mathbb{G}}})$, the cluster $C_{\underline{\mathbb{G}}}$ stays typically close to one of the Steiner forests from $\Omega_{M,m}^{\min}$. In order to do that, we introduce the notion of forest

skeleton of the cluster. This notion is a modification of the coarse-graining procedure developed in Section 2.2 of [10].

Let \mathbf{U}_τ be the unit ball in τ -norm. Fix a large number $c > 0$ and consider K such that $c \log K < K$. For any $y \in \mathbb{Z}^2$, set

$$\mathbf{B}_K(y) = (y + K \cdot \mathbf{U}_\tau) \cap \mathbb{Z}^2 \quad \text{and} \quad \hat{\mathbf{B}}_K(y) = \mathbf{B}_{K+c \log K}(y).$$

If $x \in A \subset \mathbb{Z}^2$ and $y \in A \cup \partial^{\text{ext}} A$, we shall use $\{x \xrightarrow{A} y\}$ to denote the event that x and y are connected by an open path from x to y whose vertices belong to A , with the possible exception of the terminal point y itself.

Let us construct the forest skeleton \mathcal{F}_K of the cluster $\mathbb{C}_\mathbb{G}$ (see Figure 5). Here and below, vertices in \mathbb{Z}^2 are ordered using the lexicographical ordering. In the following construction, we will often refer to the *minimal* vertex having some property.

STEP 1. Set $r = 1, i = 1$. Set $x_0^1 = u_{i_1}$ be the minimal vertex of \mathbb{G} . Set $V = \{x_0^1\}$ and $\mathcal{C} = \hat{\mathbf{B}}_K(x_0^1)$. Go to STEP 2.

STEP 2. If there exists $x \in V$ and $u \in \mathbb{G} \setminus V$ such that $u \in \hat{\mathbf{B}}_{2K}(x)$, then choose $u^* \in \mathbb{G} \setminus V$ to be the minimal such vertex. Set $x_i^r = u^*, A_i^r = \mathbf{B}_K(x_i^r)$ and go to STEP 3. Otherwise, go to STEP 4.

STEP 3. Update $V \rightarrow V \cup \{x_i^r\}, \mathcal{C} \rightarrow \mathcal{C} \cup \hat{\mathbf{B}}_K(x_i^r)$ and $i \rightarrow i + 1$. Go to STEP 2.

STEP 4. If there is at least one vertex $y \in \partial^{\text{ext}} \mathcal{C}$ such that

$$y \xrightarrow{\mathbb{C}_\mathbb{G} \setminus \mathcal{C}} \partial^{\text{ext}} \mathbf{B}_K(y) \setminus \mathcal{C},$$

then choose y^* to be the minimal such vertex, set $x_i^r = y^*, A_i^r = \mathbf{B}_K(x_i^r) \setminus \mathcal{C}$, and go to STEP 3. Otherwise, go to STEP 5.

STEP 5. If $\mathbb{G} \subset V$, then terminate the construction. Otherwise, choose u^* to be the minimal vertex of $\mathbb{G} \setminus V$. Update $r \rightarrow r + 1$ and set $x_0^r = u^*$. Update $V \rightarrow V \cup \{x_0^r\}$ and $i = 1$. Go to STEP 2.

Definition 4.4. *The above procedure produces r disjoint sets of vertices $V^1 = \{x_0^1, x_1^1, \dots\}, V^2 = \{x_0^2, x_1^2, \dots\}, \dots, V^r = \{x_0^r, x_1^r, \dots\}$. The vertices x_i^j constructed on STEP 4 are equipped with sets $A_i^j, j = 1 \dots r$. Exit paths through such A_i^j -s contribute multiplicative factors e^{-K} each. Sets A_i^j for vertices x_i^j constructed on STEP 2 play no role and are introduced for notational convenience only (see (20) below). By construction, there are at most M such vertices.*

The edges within each group $\ell = 1, \dots, r$ are constructed as follows: x_i^ℓ is connected to the vertex of

$$\{x_j^\ell : j < i \text{ and } x_i \in \hat{\mathbf{B}}_{2K}(x_j)\}$$

which has smallest index j .

This produces a graph which is a set of r trees $\mathcal{T}_K^1, \dots, \mathcal{T}_K^r$. The union of the trees is called the forest skeleton $\mathcal{F}_K = \cup_\ell \mathcal{T}_K^\ell$.

Note that we consider these graphs as compact subsets of \mathbb{R}^2 . An example of forest skeleton is drawn on Figure 5. The following result follows trivially from the construction of the forest skeleton.

Proposition 4.5. *Let \mathcal{F}_K be the forest skeleton of $\mathbb{C}_\mathbb{G}$, then*

1. \mathbb{G} is included in the vertices of \mathcal{F}_K .
2. Two vertices $u, v \in \mathbb{G}$ which were connected in $\mathbb{C}_\mathbb{G}$ are also connected in \mathcal{F}_K .
3. $\mathbb{C}_\mathbb{G} \subseteq \cup_{\ell, i} \hat{\mathbf{B}}_{2K}(x_i^\ell)$.




Figure 5: Construction of the forest skeleton $\mathcal{F}_K = \{\mathcal{T}_K^1, \mathcal{T}_K^2\}$ of the cluster C_G (in black), consisting of the trees $\mathcal{T}_K^i = \{t^i, \mathfrak{B}^i\}$, $i = 1, 2$. The Steiner forest corresponding to the partition $\underline{\mathcal{G}} = (\{u_1, u_2, u_5\}, \{u_3, u_4\})$ is drawn in dashed green.

4.3 Distance between $\mathbb{C}_{\mathbb{G}}$ and Steiner forests

Proposition 4.6. *For every $\delta_3 > 0$, there exists $\kappa_3 = \kappa_3(M) > 0$ such that for n large enough,*

$$\mu_{\mathcal{D}}^f \left(\min_{\mathcal{F} \in \Omega_{M,m}^{\min}} d_{\tau}(\mathbb{C}_{\mathbb{G}}, \mathcal{F}) > \delta_3 n \mid \Omega_{\mathbb{G}} \right) \leq e^{-\kappa_3 n},$$

uniformly in $(\mathcal{D}, \mathbb{G})$ with $|\mathbb{G}| \leq M$.

Proof. Let \mathcal{F}_K be the forest skeleton of $\mathbb{C}_{\mathbb{G}}$ at scale K (K will be chosen later). By the third item of Proposition 4.5,

$$d_{\tau}(\mathbb{C}_{\mathbb{G}}, \mathcal{F}_K) \leq 2K + c \log 2K.$$

The proposition thus reduces to the following claim: for any $\delta_3 > 0$, there exist $K = K(M) > 0$ and $\kappa_3 = \kappa_3(M) > 0$ such that

$$\mu_{\mathcal{D}}^f \left(\min_{\mathcal{F} \in \Omega_{M,m}^{\min}} d_{\tau}(\mathcal{F}_K, \mathcal{F}) > \delta_3 n \mid \Omega_{\mathbb{G}} \right) \leq e^{-\kappa_3 n},$$

uniformly in $(\mathcal{D}, \mathbb{G})$ with $|\mathbb{G}| \leq M$. We now prove this statement.

Writing $E := \{\min_{\mathcal{F} \in \Omega_{M,m}^{\min}} d_{\tau}(\mathcal{F}_K, \mathcal{F}) > \delta_3 n\}$, we have

$$\mu_{\mathcal{D}}^f(E \mid \Omega_{\mathbb{G}}) = \frac{\mu_{\mathcal{D}}^f(E \cap \Omega_{\mathbb{G}})}{\mu_{\mathcal{D}}^f(\Omega_{\mathbb{G}})} \leq \frac{\mu_{\mathcal{D}}^f(\tau(\mathcal{F}_K) \geq \tau_{\mathbb{G}} + \kappa_2 n)}{\mu_{\mathcal{D}}^f(\Omega_{\mathbb{G}})} \quad (18)$$

where in the last inequality we used Property P4 of Proposition 4.3, applied with $\delta_2 = \delta_3$.

Let \mathcal{F} be a Steiner forest in $\Omega_{\mathbb{G}}^{\min}$ and \mathcal{F}' be the forest obtained by replacing each inner node of \mathcal{F} by the closest vertex of \mathbb{Z}^2 . Now, by the FKG inequality, we can lower bound the denominator

$$\begin{aligned} \mu_{\mathcal{D}}^f(\Omega_{\mathbb{G}}) &\geq \mu_{\mathcal{D}}^f \left(\bigcap_{\{x,y\} \in \mathcal{E}(\mathcal{F}')} \{x \leftrightarrow y\} \right) \geq \prod_{\{x,y\} \in \mathcal{E}(\mathcal{F}')} \mu_{\mathcal{D}}^f(x \leftrightarrow y) \\ &\geq \prod_{\{x,y\} \in \mathcal{E}(\mathcal{F}')} e^{-\tau(y-x)(1+o_{|y-x|}(1))} = e^{-\tau_{\mathbb{G}}(1+o_n(1))}. \end{aligned} \quad (19)$$

where $\lim_{k \rightarrow \infty} o_k(1) = 0$ by definition, and the product is taken over the set $\mathcal{E}(\mathcal{F}')$ of all the inner edges of the approximate Steiner forest \mathcal{F}' .

To obtain an upper bound on the numerator, we follow [10, Section 2]. Let $|V(\mathcal{F}_K)| = \sum_{\ell=1}^r |V^{\ell}|$ be the total number of vertices of the forest skeleton \mathcal{F} , then

$$\begin{aligned} e^{-2KM} \mu_{\mathcal{D}}^f(\mathcal{F}_K = \mathcal{F}) &\leq \mu_{\mathcal{D}}^f \left(\bigcap_{\ell=1}^r \bigcap_{i=0}^{|V^{\ell}|} x_i^{\ell} \overset{A_i^{\ell}}{\leftrightarrow} \partial^{\text{ext}} \mathbf{B}_K(x_i^{\ell}) \right) \leq \prod_{\ell=1}^r \prod_{i=1}^{|V^{\ell}|} \mu_{\mathbf{A}_K^i}^w(x_i^{\ell} \overset{A_i^{\ell}}{\leftrightarrow} \partial^{\text{ext}} \mathbf{B}_K(x_i^{\ell})) \\ &\leq (e^{-K(1-o_K(1))})^{\sum_{\ell=1}^r |V^{\ell}|} = e^{-K|V(\mathcal{F})|(1-o_K(1))} \\ &\leq e^{-\tau(\mathcal{F})(1-o_K(1)-o_n(1))}, \end{aligned} \quad (20)$$

where in the first inequality the term e^{-2KM} compensates (by the FKG inequality) the inclusion of events $x_i^{\ell} \overset{A_i^{\ell}}{\leftrightarrow} \partial^{\text{ext}} \mathbf{B}_K(x_i^{\ell})$ for points $x_i^{\ell} \in \mathbb{G}$, whereas in the second inequality we expand the probability of the intersection as a product of conditional expectations and then use the FKG inequality to compare this conditional expectations with the probability with wired boundary conditions, and in the second line we use that $\mu_{\mathbf{B}_K(x)}^w(x \leftrightarrow \partial^{\text{ext}} \mathbf{B}_K(x)) = e^{-K(1-o_K(1))}$ (this follows from [10, Corollary 1.1], which is now known to be valid up to $p_c(q)$ thanks to Proposition 2.1). If we now upper bound crudely the

number of forest K -skeletons rooted at \mathbb{G} with $\tau(\mathcal{F}) = T$ (and so with less than $C_1 T/K$ vertices) by $(C_2 K)^{C_3 T/K}$, we get

$$\begin{aligned} \mu_{\mathcal{D}}^f(\tau(\mathcal{F}_K) \geq \tau_{\mathbb{G}} + \kappa_2 n) &= \sum_{\mathcal{F}: \tau(\mathcal{F}) \geq \tau_{\mathbb{G}} + \kappa_2 n} \mu_{\mathcal{D}}^f(\mathcal{F}_K = \mathcal{F}) = \sum_{T \geq \tau_{\mathbb{G}} + \kappa_2 n} \sum_{\substack{\mathcal{F}: \\ \tau(\mathcal{F}) = T}} \mu_{\mathcal{D}}^f(\mathcal{F}_K = \mathcal{F}) \\ &= \sum_{T \geq \tau_{\mathbb{G}} + \kappa_2 n} e^{(C_1 T/K) \log(C_2 K) - T(1 - o_K(1) - o_n(1))} \\ &\leq C_4 \cdot e^{-(\tau_{\mathbb{G}} + \kappa_2 n)(1 + o_K(1) + o_n(1))}, \end{aligned} \quad (21)$$

where we used (20) in the second line. The result follows by comparison with (19):

$$\mu_{\mathcal{D}}^f(E|\Omega_{\mathbb{G}}) \leq e^{-(\tau_{\mathbb{G}} + \kappa_2 n)(1 - o_K(1) - o_n(1)) + \tau_{\mathbb{G}}(1 + o_n(1))} \leq e^{-n\kappa_3}.$$

Note that $\tau_{\mathbb{G}} o_n(1) = o(n)$ since $\tau_{\mathbb{G}} = O(n)$. The latter follows from the fact that $\tau_{\mathbb{G}}$ is bounded by the τ -length of the forest obtained by opening all the edges of ∂E_m (recall that τ is an equivalent norm on \mathbb{R}^2). \square

Proof of Proposition 4.2.

Fix $\mathcal{D} = \mathcal{D}_{m,n}$ and $\mathbb{G} = \mathbb{G}_{m,n}$ with $m \geq \frac{n}{3}$ and $|\mathbb{G}| \leq M$. Let $\nu > 0$. Fix an arbitrary $0 < \delta \ll 1$ such that $\Lambda_{250\delta n} \subset \delta_1 n \mathbf{U}_{\tau}$, where δ_1 is given by P3. By definition of δ , we know that for any forest $\mathcal{F} \in \Omega_{M,m}^{\min}$, $\mathcal{F} \cap \Lambda_{250\delta n}$ is connected and contains at most one node. Therefore, we have three cases: either $\mathcal{F} \cap \Lambda_{2\delta n} = \emptyset$, or $\mathcal{F} \cap \Lambda_{2\delta n} \neq \emptyset$ but $\mathcal{F} \cap \Lambda_{20\delta n}$ contains only one edge, or $\mathcal{F} \cap \Lambda_{20\delta n}$ contains more than one edge. In the later case, the fact that edges incident to a node make an angle larger or equal to $\frac{\pi}{2} + \eta$ implies that $\mathcal{F} \cap \Lambda_{40\delta n}$ contains a node.

Also set $\delta_3 < \min\{\nu, \delta\}$. Proposition 4.6 implies that

$$\min_{\mathcal{F} \in \Omega_{M,m}^{\min}} d_{\tau}(\mathbb{C}_{\mathbb{G}}, \mathcal{F}) \leq \delta n, \quad (22)$$

with probability larger than $1 - e^{-\kappa_3 n}$ for n large enough. We now assume that this inequality is indeed satisfied. Since, by P1 of Proposition 4.3 the set $\Omega_{M,m}^{\min}$ is compact, and since we are after an upper bound which vanishes with n , it will be enough to fix a Steiner forest $\mathcal{F} \in \Omega_{M,m}^{\min}$ and to assume that

$$d_{\tau}(\mathbb{C}_{\mathbb{G}}, \mathcal{F}) \leq \delta n, \quad (23)$$

Let us treat the three previous cases separately.

- C1. $\mathcal{F} \cap \Lambda_{2\delta n} = \emptyset$. In such case, (23) shows that $\mathbb{C}_{\mathbb{G}} \cap \Lambda_{\delta n} = \emptyset$. Thus, $E_{\nu, \delta n}^1$ holds true and $\mathbb{G}_{\delta n, n} = \emptyset$.
- C2. $\mathcal{F} \cap \Lambda_{20\delta n} = [u_1, u_2]$ with u_1 and u_2 on $\partial \Lambda_{20\delta n}$ and $[u_1, u_2] \cap \Lambda_{2\delta n} \neq \emptyset$. In such case, (23) shows that $\mathbb{C}_{\mathbb{G}}$ intersects $\Lambda_{3\delta n}$ which in turns implies that $E_{\nu, k}^2$ holds for every $k \in [6\delta n, 18\delta n]$. Proposition 3.1 implies the existence of $k \in [6\delta n, 18\delta n]$ with $|\mathbb{G}_{k,n}| \leq M$ on an event of probability larger than $1 - e^{-18\delta n}$.
- C3. There exists a node $x \in \Lambda_{40\delta n}$ and therefore $\mathcal{F} \cap \Lambda_{250\delta n} = [u_1, x] \cup [u_2, x] \cup [u_3, x]$ with u_1, u_2, u_3 on $\partial \Lambda_{250\delta n}$ such that $\angle(u_i - x, u_j - x) > \frac{\pi}{2} + \eta$ for every $i \neq j$. In such case, (23) shows that $E_{\nu, k}^3$ holds for every $k \in [82\delta n, 246\delta n]$. Proposition 3.1 implies the existence of $k \in [82\delta n, 246\delta n]$ with $|\mathbb{G}_{k,n}| \leq M$ on an event of probability larger than $1 - e^{-246\delta n}$.

Altogether, we obtain the claim. \square

For later use, let us introduce the following definition:

Definition 4.7. For u_1, u_2, u_3 in general position the function $\phi(y) := \sum_{i=1}^3 \tau(u_i - y)$ is strictly convex and quadratic around its minimum point; see [9, Lemma 3]. Let x be the unique minimizer of ϕ . In this way the notation $\mathcal{T}(u_1, u_2, u_3; x)$ is reserved for the minimal Steiner forest (in this case it is a tree) which contains u_1, u_2, u_3 . It might happen, of course, that x coincides with one of the u_i -s. When, however, this is not the case, we shall refer to $\mathcal{T}(u_1, u_2, u_3; x)$ as a Steiner tripod.

5 Fluctuation theory and proof of Theorem 2.2

We are now in a position to prove Theorem 2.2. Let $\nu > 0$ small enough to be fixed later. By (14) and (15), we can assume that there exist $\delta = \delta(\nu) > 0$ and $k \geq \delta n$ such that $|\mathbb{G}_{k,n}| \leq M$ and $E_{\nu,k}^\ell$ holds true for some $\ell \in \{1, 2, 3\}$. Let

$$\mathcal{R}_n = \max \{k \geq \delta n : |\mathbb{G}_{k,n}| \leq M \text{ and } E_{\nu,k}^1 \cup E_{\nu,k}^2 \cup E_{\nu,k}^3\} \in [\delta n, n].$$

Let \mathcal{C} be a possible realization of $\mathbb{C}_{k,n}$ and $\mathcal{D} = \mathcal{D}_{k,n}$ be the corresponding flower domain. We also set $\mathbb{G} = \mathbb{G}_{k,n}$. The restriction of $\mu_{\Lambda_n}^f(\cdot | \mathcal{R}_n = k; \mathbb{C}_{k,n} = \mathcal{C})$ to \mathcal{D} is $\mu_{\mathcal{D}}^{\text{flower}}$. Exactly as in Section 3.3,

$$\text{Cond}_n[\sigma] \cap \{\mathcal{R}_n = k\} \cap \{\mathbb{C}_{k,n} = \mathcal{C}\} = \Omega_{\sigma, \mathcal{C}} \times \{\mathcal{R}_n = k; \mathbb{C}_{k,n} = \mathcal{C}\},$$

where $\Omega_{\sigma, \mathcal{C}} = \cup_{\underline{\mathbb{G}} \in \mathcal{P}'_{\mathbb{G}}} \Omega_{\underline{\mathbb{G}}}$ is defined as in Section 3.3. This reduction shows that it is sufficient to prove that

$$\mu_{\mathcal{D}}^f(\mathbb{C}_{\mathbb{G}} \cap \Lambda_{n^\varepsilon} \neq \emptyset | \Omega_{\sigma, \mathcal{C}}) = O(n^{\varepsilon-1/2}),$$

uniformly in the possible realizations of \mathcal{D} , \mathcal{C} and \mathbb{G} .

From now on, we fix $k \geq \delta n$ such that $|\mathbb{G}_{k,n}| \leq M$ and $E_{\nu,k}^\ell$ holds true for some $\ell \in \{1, 2, 3\}$. We set $\mathcal{D} = \mathcal{D}_{k,n}$, $\mathcal{C} = \mathcal{C}_{k,n}$ and $\mathbb{G} = \mathbb{G}_{k,n}$.

Since each set \mathbb{V}^i is already assumed to be connected outside of \mathcal{D} (since $E_{\nu,k}^1 \cup E_{\nu,k}^2 \cup E_{\nu,k}^3$ occurs), partitions $\underline{\mathbb{G}} \in \mathcal{P}'_{\mathbb{G}}$ can be of four different types (recall that they are maximal in the sense defined in the previous section): singletons only, singletons together with one pair of elements in two different \mathbb{V}^i (this cannot occur in $E_{\nu,k}^1$), singletons together with one triplet of elements in three different \mathbb{V}^i (this can occur only in $E_{\nu,k}^3$), singletons together with two pairs (u, v) and (u', w) , where u and u' belong to the same \mathbb{V}^i , and v and w belong to the other \mathbb{V}^j (this can occur only in $E_{\nu,k}^3$). Let $\mathcal{P}_{\mathbb{G}}^*$ be the set of partitions in $\mathcal{P}'_{\mathbb{G}}$ of one of the first three types. If the configuration is in $\Omega_{\sigma, \mathcal{C}} \setminus \cup_{\underline{\mathbb{G}} \in \mathcal{P}_{\mathbb{G}}^*} \Omega_{\underline{\mathbb{G}}}$, there are two different clusters connecting two pairs of vertices (u, v) and (u', w) satisfying the conditions described above. By choosing $\nu > 0$ small enough, the assumption that $E_{\nu,k}^3$ holds implies that $\tau(u - v) + \tau(u' - w) \geq (1 + \varepsilon)\tau_{\mathbb{G}}$ (where $\varepsilon = \varepsilon(\delta_3, \nu) > 0$) uniformly in the possible pairs (u, v) and (u', w) . As in the proof of Proposition 4.6, one obtains after a small computation that

$$\mu_{\mathcal{D}}^f(\Omega_{\sigma, \mathcal{C}} \setminus \cup_{\underline{\mathbb{G}} \in \mathcal{P}_{\mathbb{G}}^*} \Omega_{\underline{\mathbb{G}}} | \Omega_{\sigma, \mathcal{C}}) = O(e^{-ck}),$$

for some constant $c > 0$. Hence, a reduction in the spirit of Proposition 3.4 shows that Theorem 2.2 would follow from the bound

$$\mu_{\mathcal{D}}^f(\mathbb{C}_{\mathbb{G}} \cap \Lambda_{n^\varepsilon} \neq \emptyset | \Omega_{\underline{\mathbb{G}}}) = O(n^{\varepsilon-1/2}), \quad (24)$$

where the right-hand side is uniform in the possible realizations of \mathcal{D} and in the $\underline{\mathbb{G}} \in \mathcal{P}_{\mathbb{G}}^*$. We decompose the proof of (24) into three cases, depending on the type of $\underline{\mathbb{G}}$.

Scenario S1: No imposed crossing. This occurs in the following two cases (cf. Definition (4.1)): (i) $E_{\nu,k}^1$ occurs; (ii) $E_{\nu,k}^2 \cup E_{\nu,k}^3$ occurs and the partition $\underline{\mathbb{G}}$ is composed of singletons only. In this case, the measure $\mu_{\mathcal{D}_{k,n}}^f(\cdot | \Omega_{\underline{\mathbb{G}}})$ is unconditioned (i.e. $\Omega_{\underline{\mathbb{G}}} = \Omega$). Proposition 2.1 then implies that $\mu_{\mathcal{D}}^f(\mathbb{C}_{\mathbb{G}} \cap \Lambda_{n^\varepsilon} \neq \emptyset | \Omega_{\underline{\mathbb{G}}})$ decays exponentially with n .

Scenario S2: One imposed crossing. This occurs when $E_{\nu,k}^2 \cup E_{\nu,k}^3$ occurs and $\underline{\mathbb{G}}$ is composed of singletons together with a unique pair (u, v) , where $u \in \mathbb{V}^i, v \in \mathbb{V}^j$ with $i \neq j$. In other words, $\Omega_{\underline{\mathbb{G}}} = \{u \leftrightarrow v\}$. In this case, the cluster $\mathcal{C}_{\underline{\mathbb{G}}} \subseteq \mathcal{D}$ may contain several connected components, but, up to exponentially small (in k) $\mu_{\mathcal{D}}^f(\cdot | u \leftrightarrow v)$ -conditional probabilities, only one of them, namely the connected cluster $\mathbb{C}(u, v)$ of $\{u, v\}$ is capable of reaching Λ_{n^ε} . However, the law of the cluster connecting u and v converges to the law of a Brownian bridge. In fact, one obtains the following stronger result:

$$\mu_{\mathcal{D}}^f(x \in \mathbb{C}(u, v) | u \leftrightarrow v) \leq \frac{C}{\sqrt{|u - v|}} \exp\left(-\kappa \frac{d_\tau(x, [u, v])^2}{|u - v|}\right), \quad (25)$$

where κ and C are constants depending on p only, and $[u, v]$ denotes the segment between u and v . In the case of Ising interfaces, such bound was obtained in [22, (3.31)]. The proof relies on the positive curvature of the surface tension and on the effective random walk with exponentially decaying step distribution representation of the interface. The theory developed in [10] enables a literal adaptation to the case of sub-critical FK-clusters, see Theorems C and E and Subsections 4.4 and 4.5 in [10]. Consequently,

$$\mu_{\mathcal{D}}^f(\mathbb{C}_{\mathbb{G}} \cap \Lambda_{n^\varepsilon} \neq \emptyset \mid \Omega_{\mathbb{G}}) = O(n^{2\varepsilon-1/2}). \quad (26)$$

Scenario S3: One tripod. This can only happen when $E_{\nu,k}^3$ occurs and \mathbb{G} is composed of singletons together with one triplet (u_1, u_2, u_3) with $u_1 \in \mathbb{V}^1, u_2 \in \mathbb{V}^2, u_3 \in \mathbb{V}^3$. Thus, in this case $\Omega_{\mathbb{G}} = \{\mathbb{C}(u_1, u_2, u_3) \neq \emptyset\}$, where $\mathbb{C}(u_1, u_2, u_3)$ is the joint connected cluster of $\{u_1, u_2, u_3\}$. Again, $\mathbb{C} = \mathbb{C}_{\mathbb{G}} \subseteq \mathcal{D}$ may contain several connected components, but, up to exponentially small (in k) $\mu_{\mathcal{D}}^f(\cdot \mid \mathbb{C}(u_1, u_2, u_3) \neq \emptyset)$ -conditional probabilities, only one of them, namely $\mathbb{C}(u_1, u_2, u_3)$ itself, is capable of reaching Λ_{n^ε} . By definition, there exists a unique $x = x(u_1, u_2, u_3) \in \Lambda_{k/2}$ (see Definition 4.7) such that $\mathcal{T}_x = \{u_1, u_2, u_3; x\}$ is a Steiner tripod. To lighten the notation, we set

$$E(u_1, u_2, u_3, x) = \{u_1, u_2, u_3 \text{ are connected and } d_\tau(\mathbb{C}_{\mathbb{G}}, \mathcal{T}_x) \leq \nu k\}$$

and redefine $\mathbb{C} = \mathbb{C}(u_1, u_2, u_3)$. Thanks to Propositions 4.2 and 4.6, we now aim at proving the bound

$$\mu_{\mathcal{D}}^f(\mathbb{C} \cap \Lambda_{n^\varepsilon} \neq \emptyset \mid E(u_1, u_2, u_3, x)) = O(n^{\varepsilon-1/2}). \quad (27)$$

This bound will imply Theorem 2.2.

Proving (27) is more complicated than proving (26). Nevertheless, the idea remains the same: The tripod has Gaussian fluctuations, therefore it intersects a small box with probability going to 0. In the case of percolation, fluctuations of tripods on the level of local limit results were studied in [9]. We are not after a full local limit picture here, and merely explain how techniques from [10] allow to derive (27). Let us write $\Lambda_r(x)$ for $x + \Lambda_r$.

Definition 5.1. (Cones $\mathcal{Y}_1, \mathcal{Y}_2, \mathcal{Y}_3$)

Since, by Property P2 of the Steiner forests, for every $i \neq j$,

$$\angle(u_i - x, u_j - x) \geq \frac{\pi}{2} + \eta,$$

there exist disjoint cones $\mathcal{Y}_1, \mathcal{Y}_2$ and \mathcal{Y}_3 such that each \mathcal{Y}_i contains exactly one lattice direction in its interior (i.e., one of the four vectors $(1, 0)$, $(0, 1)$, $(-1, 0)$ and $(0, -1)$, denoted by \mathbf{f}_i), and there exists $\varepsilon_1 > 0$ such that $u_i \in \text{int}(y + \mathcal{Y}_i)$ for every $y \in \Lambda_{\varepsilon_1 k}(x)$ and every $i \in \{1, 2, 3\}$, and $u_i \in \text{int}(u_j - \mathcal{Y}_i)$ for every $i \neq j$.

Definition 5.2. (Event $\mathcal{S}(t, y)$)

Given $y \in \Lambda_{\varepsilon_1 k}(x)$ and $t \in \mathbb{N}$, let $\mathcal{S}(t, y)$ be the event that the following three conditions occur:

R1. u_1, u_2 and u_3 are pairwise disconnected in $\mathbb{C} \setminus \Lambda_t(y)$,

R2. \mathbb{C} intersects $\partial\Lambda_t(y)$ in exactly three vertices.

For $i = 1, 2, 3$, let $\mathbb{C}_i(t, y)$ be the connected component of $\mathbb{C} \setminus \Lambda_t(y)$ containing u_i , and $v_i(t, y) = \mathbb{C}_i(t, y) \cap \partial\Lambda_t(y)$. Define $\mathbb{C}_0(t, y) = \mathbb{C} \setminus (\mathbb{C}_1(t, y) \cup \mathbb{C}_2(t, y) \cup \mathbb{C}_3(t, y))$. We will drop the reference to t and y when no confusion is possible.

R3. \mathbb{C}_0 is contained in $\bigcap_{i=1}^3 (v_i - \mathcal{Y}_i)$ and $\mathbb{C}_i \subset (v_i + \mathcal{Y}_i)$ for $i \in \{1, 2, 3\}$.




Figure 6: Description of the event $S(t, y)$, namely the cones and the decomposition of the cluster C into $C_i(t)$ and $v_i(t)$, $i = 1, 2, 3$.

Lemma 5.3. *Fix $\varepsilon_1 > 0$ and let $\varepsilon > 0$ be sufficiently small. There exists $C > 0$ such that*

$$\mu_{\mathcal{D}}^f \left(\bigcup_{t \leq Ck^\varepsilon} \bigcup_{y \in \Lambda_{\varepsilon_1 k}(x)} \mathcal{S}(t, y) \mid E(u_1, u_2, u_3, x) \right) \geq 1 - O(e^{-k^\varepsilon}). \quad (28)$$

Proof of Lemma 5.3. First of all, we notice that coarse-graining on the k^ε -scale enables a reduction to particularly simple geometric structures. Consider a forest skeleton of the cluster C at scale k^ε . Note that, conditionally on $E(u_1, u_2, u_3, x)$, this forest is in fact a tree $\mathcal{T}_{k^\varepsilon}$.

We define the *trunk* \mathfrak{t}_ε of $\mathcal{T}_{k^\varepsilon}$ as the minimal subtree of $\mathcal{T}_{k^\varepsilon}$ which spans $\{u_1, u_2, u_3\}$.

We define the *branches* of $\mathcal{T}_{k^\varepsilon}$ as $\mathfrak{B}_\varepsilon = \mathcal{T}_{k^\varepsilon} \setminus \mathfrak{t}_\varepsilon$. In this case, we obtain the following reduced geometry of typical $\mathcal{T}_{k^\varepsilon}$, which holds uniformly in all situations in question, up to probabilities which are exponentially small in k^ε :

T1. $\mathcal{T}_{k^\varepsilon}$ does not have branches. This means that the tree $\mathcal{T}_{k^\varepsilon}$ consists only of a trunk which is a tripod, i.e. with one vertex of degree 3 and all other vertices of degree at most 2. We will write x_ε for the only triple point of $\mathcal{T}_{k^\varepsilon}$, and $\mathcal{T}_{k^\varepsilon}^i = \{u_{i,\varepsilon}^{n_i}, \dots, u_{i,\varepsilon}^1 = x_\varepsilon\}$, $i = 1, 2, 3$, for the three legs of $\mathcal{T}_{k^\varepsilon}$. Note that $u_i \in \hat{\mathbf{B}}_{2k^\varepsilon}(u_{i,\varepsilon}^{n_i})$.

T2. Fix $\kappa > 0$ small. For every $\varepsilon > 0$ and each $\varepsilon' \in (0, \varepsilon/2)$, the skeletons $\mathcal{T}_{k^{\varepsilon'}}^i \setminus \Lambda_{k^\varepsilon}(x_{\varepsilon'}) \subseteq x_{\varepsilon'} + \mathcal{Y}_{i,2\kappa}$ as soon as k becomes sufficiently large, where cones $\mathcal{Y}_{i,2\kappa}$ are defined via

$$\mathcal{Y}_{i,r} = \{z : \angle(z, u_i - x_{\varepsilon'}) \leq r\}. \quad (29)$$

That is, the vertices of each of the three branches of $\mathcal{T}_{k^{\varepsilon'}}$ outside the box Λ_{k^ε} are confined to the respective cones $x_{\varepsilon'} + \mathcal{Y}_{i,2\kappa}$.

Before proving Properties T1 and T2, let us describe how they can be used to prove the lemma. First of all, note that, by Proposition 4.6, we may assume that $|x_{\varepsilon'} - x| \leq \delta_1 k$ with $\delta_1 > 0$ fixed as small as we wish. In particular, we may assume that $u_i \in \text{int}(x_{\varepsilon'} + \mathcal{Y}_i)$ (see Definition 5.1) and, consequently,

that $\mathcal{Y}_{i,2\kappa} \subset \mathcal{Y}_i$.

By Proposition 4.5, the connected cluster \mathbf{C} is included in $\mathcal{T}_{k^{\varepsilon'}} + 2k^{\varepsilon'}\mathbf{U}_\tau$. Therefore, Properties T1 and T2 imply that $\mathbf{C} \setminus \Lambda_{k^\varepsilon}(x_{\varepsilon'}) = \tilde{\mathbf{C}}_1 \cup \tilde{\mathbf{C}}_2 \cup \tilde{\mathbf{C}}_3$, where $\tilde{\mathbf{C}}_1$, $\tilde{\mathbf{C}}_2$ and $\tilde{\mathbf{C}}_3$ are the clusters (in $\mathbf{C} \setminus \Lambda_{k^\varepsilon}(x_{\varepsilon'})$) of u_1 , u_2 and u_3 respectively. Note that, by T2, clusters \mathbf{C}_i are confined to the sets (actually truncated cones) $(x_{\varepsilon'} + \mathcal{Y}_{i,2\kappa} + 2k^{\varepsilon'}\mathbf{U}_\tau) \setminus \Lambda_{k^\varepsilon}(x_{\varepsilon'})$, which are well separated on the k^ε -scale. Consequently coarse-graining estimates developed in [10, Section 2] apply to each of \mathbf{C}_i separately. As a result, the claim of Lemma 5.3 follows by a straightforward adaptation of the mass-gap arguments of [10, Section 2] applied separately to each of the three disjoint clusters $\tilde{\mathbf{C}}_1$, $\tilde{\mathbf{C}}_2$ and $\tilde{\mathbf{C}}_3$. For instance, one can show the following: Fix r large enough so that $\Lambda_{k^{\varepsilon'}}(x_{\varepsilon'}) \subset v - \mathcal{Y}_i$ for any $v \in (x_{\varepsilon'} + \mathcal{Y}_{i,2\kappa}) \cap (\Lambda_{2rk^\varepsilon}(x_{\varepsilon'}) \setminus \Lambda_{rk^\varepsilon}(x_{\varepsilon'}))$ and $i = 1, 2, 3$. Then, up to probabilities which are exponentially small in k^ε , there exists $t \in [rk^\varepsilon, 2rk^\varepsilon]$ such that each of the clusters $\tilde{\mathbf{C}}_i$ contains a \mathcal{Y}_i -break point on $\partial\Lambda_t(x_{\varepsilon'})$. That is,

- for $i = 1, 2, 3$, the intersection $v_i = \tilde{\mathbf{C}}_i \cap \partial\Lambda_t(x_{\varepsilon'})$ is a singleton;
- for $i = 1, 2, 3$ the cluster $\tilde{\mathbf{C}}_i \subset (v_i + \mathcal{Y}_i) \cup (v_i - \mathcal{Y}_i)$.

This ensures $\mathcal{S}(t, y)$ for some $y \in \Lambda_{\varepsilon_1 k}$ and (28) follows. \square

For the proof of Property T1, we refer to [10, Lemma 2.1 and 2.2].

Proof of Property T2. Let us start with a lower bound on $\mu_{\mathcal{D}}^f(E(u_1, u_2, u_3, x))$ which will be used later as a test threshold quantity for ruling out improbable events. Let y be a lattice approximation of x . By the FKG inequality,

$$\mu_{\mathcal{D}}^f(E(u_1, u_2, u_3, x)) \geq \mu_{\mathcal{D}}^f\left(\bigcap_{i=1}^3 \{y \stackrel{\mathcal{D}}{\leftrightarrow} u_i\}\right) \geq \prod_{i=1}^3 \mu_{\mathcal{D}}^f(y \stackrel{\mathcal{D}}{\leftrightarrow} u_i).$$

Theorem A in [10] gives sharp asymptotics of quantities $\mu^f(y \leftrightarrow v_i)$. These sharp asymptotics are built upon an effective random walk representation of events $\{y \leftrightarrow u\}$ as described in Subsection 4.1 of the paper. Steps of this random walk have effective drift from u_i towards y , and, since $\Lambda_k \subset \mathcal{D}$, it is easy to adjust the arguments therein in order to show that

$$\mu_{\mathcal{D}}^f(y \stackrel{\mathcal{D}}{\leftrightarrow} u) \geq \frac{C_0}{\sqrt{k}} e^{-\tau(u-y)},$$

uniformly in $y \in \Lambda_{\frac{k}{2}}$ and $v \in \partial\Lambda_k$, where C_0 (and, similarly, C_1, C_2, \dots below) is a universal constant, in the sense that (30) applies uniformly in all the situations in question as soon as k is sufficiently large. Consequently,

$$\mu_{\mathcal{D}}^f(E(u_1, u_2, u_3, x)) \geq \exp\left(-\sum_{i=1}^3 \tau(u_i - x) - C_1 \log k\right), \quad (30)$$

also uniformly in all the situations in question as soon as k is sufficiently large.

Next, let us say that $w \in \mathcal{T}_{k^{\varepsilon'}}^i$ is a 2κ -cone point of $\mathcal{T}_{k^{\varepsilon'}}^i$ if $\mathcal{T}_{k^{\varepsilon'}}^i \subset (w - \mathcal{Y}_{i,2\kappa}) \cup (w + \mathcal{Y}_{i,2\kappa})$. In our notation,

$$\tau(\mathcal{T}_{k^{\varepsilon'}}^i) = \sum_{i=1}^3 \tau(\mathcal{T}_{k^{\varepsilon'}}^i)$$

Since τ is a strictly convex norm ([10, Subsection 1.3.2]),

$$\tau(\mathcal{T}_{k^{\varepsilon'}}^i) \geq \tau(u_i - x_{\varepsilon'}) (1 + \delta(\kappa)) \geq \tau(u_i - x_{\varepsilon'}) + C_2 k, \quad (31)$$

whenever $\mathcal{T}_{k^{\varepsilon'}}^i$ does not contain 2κ -cone points at all. This is essentially Lemma 2.4 of [10]. In view of (20), and in view of the lower bound (30), we are entitled to ignore the situation when any of the $\mathcal{T}_{k^{\varepsilon'}}^i$ does not have 2κ -cone points at all.

In the sequel, we use w_i^* to denote the first 2κ -cone point of $\mathcal{T}_{k^{\varepsilon'}}^i$ (starting at $x_{\varepsilon'}$) and N_i to denote its serial number; that is, $w_i^* = u_{i,\varepsilon'}^{N_i}$. Define $\mathcal{T}_{k^{\varepsilon'}}^{i,*} = \{u_{i,\varepsilon'}^1, \dots, u_{i,\varepsilon'}^{N_i} = w_i^*\}$ as the portion of $\mathcal{T}_{k^{\varepsilon'}}^i$ up to w_i^* . Given y and $\underline{w} = (w_1, w_2, w_3)$, define the percolation event $E_{\varepsilon'}(y, \underline{w}) \subset E(u_1, u_2, u_3, x)$ as

$$E_{\varepsilon'}(y, \underline{w}) = \{x_{\varepsilon'} = y; w_i^* = w_i \text{ for } i = 1, 2, 3\}.$$

In view of (30), Property T2 will follow as soon as we shall have checked that

$$\mu_{\mathcal{D}}^f(E_{\varepsilon'}(y, \underline{w})) \leq e^{-\sum_i \tau(u_i - y) - C_3 k^\varepsilon}, \quad (32)$$

uniformly in k , tripods \mathcal{T}_x , y and $\underline{w} \notin \Lambda_{k^\varepsilon}(y)$. For fixed realizations $\mathcal{T}^{i,*}$ of $\mathcal{T}_{k^{\varepsilon'}}^{i,*}$ we have

$$\mu_{\mathcal{D}}^f(E_{\varepsilon'}(y, \underline{w}); \mathcal{T}_{k^{\varepsilon'}}^{i,*} = \mathcal{T}^{i,*} \text{ for } i = 1, 2, 3) \leq \exp\left\{-\sum_{i=1}^3 \{\tau(u_i - w_i) + \tau(\mathcal{T}^{i,*})(1 - o_{k^{\varepsilon'}}(1))\} + C_4 k^{2\varepsilon'}\right\}.$$

This follows from (20) and from the finite energy property (applied for configurations on $\Lambda_{C_3 k^{\varepsilon'}}(w_i)$) of the FK measures. Indeed, the finite energy property and the exponential ratio mixing property (6) enable to decouple between the event $\bigcap_i \{\mathcal{T}_{k^{\varepsilon'}}^{i,*} = \mathcal{T}^{i,*}\}$ and the events $\bigcap_i \{w_i \xleftrightarrow{\mathcal{Y}_{i,2\kappa}} u_i\}$.

Assume, for instance, that $w_1 \notin \Lambda_{k^\varepsilon}(y)$. There are two cases to consider:

CASE 1: $w_1 \in y + \mathcal{Y}_{1,\kappa}$. Then, exactly as in (31), $\tau(\mathcal{T}^{1,*}) \geq \tau(w_1 - y) + C_6 |w_1 - y|$. As in (21) the entropic factor related to the number of possible compatible realizations $\mathcal{T}^{1,*}$ is suppressed, and (32) follows as soon as we choose $\varepsilon > 2\varepsilon'$.

CASE 2: $w_1 \in (y + \mathcal{Y}_{1,2\kappa}) \setminus (y + \mathcal{Y}_{1,\kappa})$. By construction, $\tau(\mathcal{T}^{1,*}) \geq \tau(w_1 - y)$. However, by the sharp triangle inequality (2),

$$\tau(w_1 - y) + \tau(u_1 - w_1) - \tau(u_1 - y) \geq C_7 |w_1 - y|,$$

uniformly in w_1 under consideration. Again, since the entropic factor is suppressed, (32) follows. \square

Lemma 5.4. *Let $S_\varepsilon(y) = \bigcup_{t \leq Ck^\varepsilon} S(t, y)$. There exist two universal constants $\kappa > 0$ and $C < \infty$ such that*

$$\mu_{\mathcal{D}}^f(S_\varepsilon(y) \mid E(u_1, u_2, u_3, x)) = O\left(k^{12\varepsilon-1} \exp\left(-\kappa \frac{|y-x|^2}{k}\right)\right), \quad (33)$$

uniformly in $y \in \Lambda_{\varepsilon_1 k}(x)$.

Proof. Decompose

$$S(t, y) = \bigcup_W S^W(t, y)$$

according to the triple $W = \{v_1 - y, v_2 - y, v_3 - y\} \subset \partial\Lambda_t$ which shows up in the definition. From now on, we set $w_1 = v_1 - y$, $w_2 = v_2 - y$ and $w_3 = v_3 - y$.

Since, under the event $S^W(t, y)$, we have that $C_0 \subseteq \bigcap_{i=1}^3 (v_i - \mathcal{Y}_i)$ and that the points u_i lie deep in the interior of the corresponding cones $v_i + \mathcal{Y}_i$, with $v_i \in \partial\Lambda_t(y)$ and $t \leq Ck^\varepsilon$, the Ornstein-Zernike asymptotics of [10, Theorem A] imply that

$$\mu_{\mathcal{D}}^f\left(\bigcap_{i=1}^3 \{C_i \subset v_i + \mathcal{Y}_i\} \mid C_0(t, y)\right) = \Theta\left(k^{-3/2} e^{-\sum_{i=1}^3 \tau(u_i - v_i)}\right), \quad (34)$$

uniformly in any possible realization C_0 of $C_0(t, y)$ compatible with $S^W(t, y)$. Note that if $C_0(t, y)$ is compatible with $S^W(t, y)$, then shifts $C_0^u \triangleq C_0 + u$ are compatible with shifted events $S^W(t, y + u)$.

Recall Definition 4.7 of $\phi(y)$. Given a triple $W = \{w_1, w_2, w_3\} \subset \Lambda_{Ck^\varepsilon}$, let us define

$$\phi_W(y) = \sum_{i=1}^3 \tau(u_i - w_i - y).$$

Together with (34), we obtain

$$\begin{aligned} \frac{\mu_{\mathcal{D}}^f(S^W(t, y))}{\mu_{\mathcal{D}}^f(S^W(t, z))} &= \frac{\sum_{\mathcal{C}_0} \mu_{\mathcal{D}}^f\left(\bigcap_{i=1}^3 \{\mathcal{C}_i \subset y + w_i + \mathcal{Y}_i\} \mid \mathcal{C}_0\right) \mu_{\mathcal{D}}^f(\mathcal{C}_0(t, y) = \mathcal{C}_0)}{\sum_{\mathcal{C}_0} \mu_{\mathcal{D}}^f\left(\bigcap_{i=1}^3 \{\mathcal{C}_i \subset z + w_i + \mathcal{Y}_i\} \mid \mathcal{C}_0\right) \mu_{\mathcal{D}}^f(\mathcal{C}_0(t, z) = \mathcal{C}_0^{z-y})} \\ &= \Theta(e^{\phi_W(y) - \phi_W(z)}), \end{aligned} \quad (35)$$

uniformly in $t \leq Ck^\varepsilon$, $W = \{w_1, w_2, w_3\} \subset \partial\Lambda_t$ and $y, z \in \Lambda_{\varepsilon_1 k}$, where the sum is over \mathcal{C}_0 compatible with $S^W(t, y)$ and where in the last line we used classical ratio-mixing properties of subcritical random-cluster measures [3, Theorem 3.4] and (6) to compare $\mu_{\mathcal{D}}^f(\mathcal{C}_0(t, y) = \mathcal{C}_0)$ and $\mu_{\mathcal{D}}^f(\mathcal{C}_0(t, z) = \mathcal{C}_0^{z-y})$.

The function ϕ_W has a non-degenerate quadratic minimum at some $x_{\min}(W)$ (see [9, Lemma 3]). In view of the homogeneity of τ , a quadratic expansion around x_{\min} yields

$$\phi_W(y) - \phi_W(x_{\min}) = \Theta\left(\frac{|y - x_{\min}|^2}{k}\right), \quad (36)$$

uniformly in all situations in question. Since $|\phi(y) - \phi_W(y)| = O(k^\varepsilon)$, its minimizers $x_{\min}(W)$ solve

$$F(x, W) \triangleq \nabla_x \phi_W(x) = 0.$$

Since $\text{Hess}(\phi)$ is non-degenerate at x , the implicit function theorem applies. As a result, $|x_{\min}(W) - x| = O(\sum_i |w_i|) = O(k^\varepsilon)$ uniformly in all W in question. This fact, together with (36), yields

$$\phi_W(y) - \phi_W(x) = \Theta\left(\frac{|y - x|^2}{k} + k^{2\varepsilon-1}\right). \quad (37)$$

Since there are at most $O(k^{3\varepsilon})$ possible choices for W and $O(k^\varepsilon)$ possible choices for t , we deduce from (35) and (37) that

$$\frac{1}{O(k^{4\varepsilon})} \exp(-C_1 \frac{|y - x|^2}{k}) \leq \frac{\mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(y))}{\mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(x))} \leq O(k^{4\varepsilon}) \exp(-C_2 \frac{|y - x|^2}{k}). \quad (38)$$

Above $\mathcal{S}_\varepsilon(x)$ means in fact $\mathcal{S}_\varepsilon(\lfloor x \rfloor)$. We can now compute

$$\begin{aligned} \mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(y) \mid E(u_1, u_2, u_3, x)) &= \frac{\mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(y))}{\mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(x))} \cdot \frac{\mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(x))}{\mu_{\mathcal{D}}^f(E(u_1, u_2, u_3, x))} \\ &\leq O(k^{4\varepsilon}) \exp\left(-C_2 \frac{|y - x|^2}{k}\right) \frac{\mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(x))}{\mu_{\mathcal{D}}^f(E(u_1, u_2, u_3, x))} \end{aligned} \quad (39)$$

where we used the second inequality in (38). In order to see that the rightmost term in (39) is of the right order, observe that $|y - x| \leq k^{1/2-\varepsilon}$ implies that $e^{-C_2|y-x|^2/k}$ is of order 1, and therefore the ratio in (38) is smaller than $O(k^{4\varepsilon})$. Therefore, by looking at the $k^{1-2\varepsilon}$ sites which are at distance at most $k^{1/2-\varepsilon}$ from x , we deduce, using the first inequality in (38), that

$$\mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(x)) \leq O(k^{-1+6\varepsilon}) \sum_{y \in \Lambda_{k^{1/2-\varepsilon}}(x)} \mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(y)) \leq O(k^{-1+8\varepsilon}) \mu_{\mathcal{D}}^f(E(u_1, u_2, u_3, x)),$$

where in the second inequality, we used the fact that in a given configuration there are at most $O(k^{2\varepsilon})$ sites y such that the corresponding events $\mathcal{S}_\varepsilon(y)$ occur. This implies that

$$\frac{\mu_{\mathcal{D}}^f(\mathcal{S}_\varepsilon(x))}{\mu_{\mathcal{D}}^f(E(u_1, u_2, u_3, x))} \leq O(k^{8\varepsilon-1}).$$

Together with (39), we obtain (33). \square

Lemmas 5.3 and 5.4 imply that

$$\mu_{\mathcal{D}}^f(\mathbb{C} \cap \Lambda_{k^\varepsilon} \neq \emptyset \mid E(u_1, u_2, u_3, x)) \leq O(k^{12\varepsilon-1}) \sum_{y \in \Lambda_{\varepsilon_1 k}(x)} e^{-\kappa|y-x|^2/k} \mu_{\mathcal{D}}^f(\mathbb{C} \cap \Lambda_{k^\varepsilon} \neq \emptyset \mid \mathcal{S}_\varepsilon(y)) + O(e^{-k^\varepsilon}). \quad (40)$$

It remains to provide an upper bound on $\mu_{\mathcal{D}}^f(\mathbb{C} \cap \Lambda_{k^\varepsilon} \neq \emptyset \mid \mathcal{S}_\varepsilon(y))$. There are two cases to consider:

CASE 1: $y \in \Lambda_{2Ck^\varepsilon}$. In this case, we simply use

$$\mu_{\mathcal{D}}^f(\mathbb{C} \cap \Lambda_{k^\varepsilon} \neq \emptyset \mid \mathcal{S}_\varepsilon(y)) \leq 1. \quad (41)$$

The total contribution to the right-hand side of (40) is then bounded by $O(k^{14\varepsilon-1})$, which is negligible with respect to our target estimate (27).

CASE 2: $y \notin \Lambda_{2Ck^\varepsilon}$. In this case, Λ_{k^ε} can intersect at most one of the $\Lambda_{Ck^\varepsilon}(y) + \mathcal{Y}_i$, and therefore can be hit by only one cluster \mathbb{C}_i . Without loss of generality, let us assume that $\mathbb{C}_i = \mathbb{C}_1$. Conditioning on the smallest t such that $\mathcal{S}(t, y)$ occurs as well as on $\mathbb{C}_0, \mathbb{C}_2$ and \mathbb{C}_3 , the cluster \mathbb{C}_1 obeys, as was explained after (25), a diffusive scaling. In particular,

$$\mu_{\mathcal{D}}^f(z \in \mathbb{C}_1 \mid \mathcal{S}(t, y), \mathbb{C}_0, \mathbb{C}_2, \mathbb{C}_3) = O\left(\frac{1}{\sqrt{|v_1 - z|}} \exp\left[-\kappa' \frac{d_\tau(z, [v_1, u_1])^2}{|v_1 - z|}\right]\right).$$

In the previous inequality, $v_1 = v_1(t, y)$. We find:

$$\begin{aligned} \mu_{\mathcal{D}}^f(\mathbb{C} \cap \Lambda_{k^\varepsilon} \neq \emptyset \mid \mathcal{S}(t, y), \mathbb{C}_0, \mathbb{C}_2, \mathbb{C}_3) &\leq \sum_{z \in \partial\Lambda_{k^\varepsilon}} \mu_{\mathcal{D}}^f(z \in \mathbb{C}_1 \mid \mathcal{S}(t, y), \mathbb{C}_0, \mathbb{C}_2, \mathbb{C}_3) \\ &\leq \sum_{z \in \partial\Lambda_{k^\varepsilon}} O\left(\frac{1}{\sqrt{|v_1 - z|}} \exp\left[-\kappa' \frac{d_\tau(z, [v_1, u_1])^2}{|v_1 - z|}\right]\right) \\ &= O\left(\frac{k^\varepsilon}{\sqrt{|y|}} \exp\left\{-\kappa'' \frac{d_\tau(0, [y, u_1])^2}{|y|}\right\}\right). \end{aligned} \quad (42)$$

In the last line, we used the fact that $y, v_1 \notin \Lambda_{Ck^\varepsilon}$ and $|v_1 - y| \leq Ck^\varepsilon$. Let us substitute (42) into the sum on the right-hand side of (40) to obtain

$$\begin{aligned} \mu_{\mathcal{D}}^f(\mathbb{C} \cap \Lambda_{k^\varepsilon} \neq \emptyset \mid E(u_1, u_2, u_3, x)) &\leq O(e^{-k^\varepsilon}) + O(k^{14\varepsilon-1}) \\ &\quad + O(k^{13\varepsilon-1}) \sum_{y \in \Lambda_{\varepsilon_1 k}(x) \setminus \Lambda_{2Ck^\varepsilon}} \frac{1}{\sqrt{|y|}} \exp\left[-\kappa \frac{|y-x|^2}{k} - \kappa'' \frac{d_\tau(0, [y, u_1])^2}{|y|}\right]. \end{aligned} \quad (43)$$

After a simple estimate, one sees easily that the sum on the right is bounded above as

$$2 \sum_{y \in \Lambda_{k^{1/2+\varepsilon}}(x) \setminus \Lambda_{2Ck^\varepsilon}} \frac{1}{\sqrt{|y|}} \exp\left[-\kappa'' \frac{d_\tau(0, [y, u_1])^2}{|y|}\right] \leq \sum_{\ell=\max\{2Ck^\varepsilon, |x|-k^{1/2+\varepsilon}\}}^{|x|+k^{1/2+\varepsilon}} \frac{O(\sqrt{\ell})}{\sqrt{\ell}} = O(k^{\frac{1}{2}+\varepsilon}),$$

uniformly in x . In order to obtain the first inequality, we used the fact that $\exp(-\kappa|y-x|^2/k)$ is very small for sites outside of $\Lambda_{k^{1/2+\varepsilon}}(x)$. For the second, observe that sites y at distance ℓ contributing substantially to this sum must satisfy the condition that 0 is at distance $O(\sqrt{\ell})$ of $[y, u_1]$. There are $O(\sqrt{\ell})$ of them. This concludes the proof.

Acknowledgements. H. D.-C. was supported by the EU Marie-Curie RTN CODY, the ERC AG CONFRA, as well as by the Swiss National Science Foundation. The research of D.I. was supported by the Israeli Science Foundation (grant No 817/09). L.C. and Y.V. were partially supported by the Swiss National Science Foundation.

References

- [1] M. Aizenman. Translation invariance and instability of phase coexistence in the two-dimensional Ising system. *Comm. Math. Phys.*, 73(1):83–94, 1980.
- [2] M. Aizenman, J. T. Chayes, L. Chayes and C. M. Newman. Discontinuity of the magnetization in one-dimensional $1/|x - y|^2$ Ising and Potts models. *J. Stat. Phys.*, 50(1): 1–40. 1988.
- [3] K. S. Alexander. On weak mixing in lattice models. *Probab. Theory Relat. Fields*, 110(4):441–471, 1998.
- [4] K. S. Alexander. Mixing properties and exponential decay for lattice systems in finite volumes *Ann. Probab.* 32(1A):441–487, 2004.
- [5] M. Alfaro, M. Conger, K. Hodges, A. Levy, R. Kochar, L. Kuklinski, Z. Mahmood and K. Von Haam. The structure of singularities in Φ -minimizing networks in \mathbb{R}^2 , *Pacific J. of Math.* 149(2):201–210, 1991.
- [6] 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 Relat. Fields*, 153(3), 511–542, 2012.
- [7] T. Bodineau. Translation invariant Gibbs states for the Ising model. *Probab. Theory Relat. Fields*, 135(2):153–168, 2006.
- [8] J. Bricmont, J. L. Lebowitz and C. E. Pfister. On the equivalence of boundary conditions. *J. Statist. Phys.*, 21(5):573–582, 1979.
- [9] M. Campanino and M. Gianfelice. A local limit theorem for triple connections in subcritical Bernoulli percolation. *Probab. Theory Relat. Fields*, 143(3-4):353–378, 2009.
- [10] M. Campanino, D. Ioffe and Y. Velenik. Fluctuation theory of connectivities for subcritical random cluster models. *Ann. Probab.*, 36(4):1287–1321, 2008.
- [11] J. T. Chayes, L. Chayes and R. H. Schonmann. Exponential decay of connectivities in the two-dimensional Ising model. *J. Statist. Phys.*, 49(3-4):433–445, 1987.
- [12] E. J. Cockayne. On the Steiner problem, *Canad. Math. Bull.* 10:431–450, 1967.
- [13] L. Coquille and Y. Velenik. A finite-volume version of Aizenman–Higuchi theorem for the 2d Ising model. *Probab. Theory Relat. Fields*, 153(1-2):25–44, 2012.
- [14] O. Couronné. A large deviation result for the subcritical Bernoulli percolation. *Annales de la Faculté des Sciences de Toulouse*, 14(2), 201–204, 2005.
- [15] R. L. Dobrushin and S. B. Shlosman. The problem of translation invariance of Gibbs states at low temperatures. In *Mathematical physics reviews, Vol. 5*, volume 5 of *Soviet Sci. Rev. Sect. C Math. Phys. Rev.*, pages 53–195. Harwood Academic Publ., Chur, 1985.
- [16] R. L. Dobrušin. Gibbsian random fields for lattice systems with pairwise interactions. In *Funkcional. Anal. i Priložen.* 4(2):31–43, 1968.
- [17] R. L. Dobrušin. The Gibbs state that describes the coexistence of phases for a three-dimensional Ising model. *Teor. Verojatnost. i Primenen.*, 17:619–639, 1972.
- [18] G. Gallavotti. The phase separation line in the two-dimensional Ising model. *Comm. Math. Phys.*, 27:103–136, 1972.
- [19] H.-O. Georgii. Gibbs measures and phase transitions, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, 1988.

- [20] H.-O. Georgii, O. Häggström and C. Maes. The random geometry of equilibrium phases. In *Phase transitions and critical phenomena*, Phase Transit. Crit. Phenom., vol. 18, pp. 1–142, Academic Press, San Diego, 2001.
- [21] H.-O. Georgii and Y. Higuchi. Percolation and number of phases in the two-dimensional Ising model. *J. Math. Phys.*, 41(3):1153–1169, 2000.
- [22] L. Greenberg and D. Ioffe. On an invariance principle for phase separation lines. *Ann. Inst. H. Poincaré Probab. Statist.*, 41(5):871–885, 2005.
- [23] G. R. Grimmett. The random-cluster model, vol. 333 of Grundlehren der Mathematischen Wissenschaften, Springer-Verlag, Berlin, 2006.
- [24] Y. Higuchi. On some limit theorems related to the phase separation line in the two-dimensional Ising model. *Z. Wahrsch. Verw. Gebiete*, 50(3):287–315, 1979.
- [25] Y. Higuchi. On the absence of non-translation invariant Gibbs states for the two-dimensional Ising model. In *Random fields, Vol. I, II (Esztergom, 1979)*, volume 27 of *Colloq. Math. Soc. János Bolyai*, pages 517–534. North-Holland, Amsterdam, 1981.
- [26] D. Ioffe. Large deviations for the 2D Ising model: a lower bound without cluster expansions. *J. Statist. Phys.*, 74(1-2):411–432, 1994.
- [27] L. Laanait, A. Messenger, S. Miracle-Solé, J. Ruiz and S. Shlosman. Interfaces in the Potts model. I. Pirogov-Sinai theory of the Fortuin-Kasteleyn representation. *Comm. Math. Phys.*, 140(1):81–91, 1991.
- [28] L. Laanait, A. Messenger and J. Ruiz. Phases coexistence and surface tensions for the Potts model. *Comm. Math. Phys.*, 105(4):527–545, 1986.
- [29] O. E. Lanford and D. Ruelle. Observables at infinity and states with short range correlations in statistical mechanics. *Comm. Math. Phys.*, 13:194–215, 1969.
- [30] D. H. Martirosian. Translation invariant Gibbs states in the q -state Potts model. *Comm. Math. Phys.*, 105(2):281–290, 1986.
- [31] A. Messenger and S. Miracle-Sole. Correlation functions and boundary conditions in the Ising ferromagnet. *J. Statist. Phys.*, 17(4):245–262, 1977.
- [32] S. A. Pirogov and Ja. G. Sinaï. Phase diagrams of classical lattice systems. *Teoret. Mat. Fiz.*, 25(3):358–369, 1975.
- [33] C-E. Pfister and Y. Velenik. Interface, Surface Tension and Reentrant Pinning Transition in the 2D Ising Model. *Comm. Math. Phys.* 204(2), 269–312, 1998.
- [34] L. Russo. The infinite cluster method in the two-dimensional Ising model. *Comm. Math. Phys.*, 67(3):251–266, 1979.
- [35] M. Zahradník. An alternate version of Pirogov-Sinaï theory. *Comm. Math. Phys.*, 93(4):559–581, 1984.

DÉPARTEMENT DE MATHÉMATIQUES
 UNIVERSITÉ DE GENÈVE
 GENÈVE, SWITZERLAND
 E-MAIL: loren.coquille@unige.ch; hugo.duminil@unige.ch; yvan.velenik@unige.ch
 FACULTY OF IE&M
 TECHNION
 HAIFA, ISRAEL
 E-MAIL: ieioffe@ie.technion.ac.il