

Gaussian free field convergence of the six-vertex model

with $-1 \leq \Delta \leq -\frac{1}{2}$

Hugo Duminil-Copin^{*†}, Karol Kajetan Kozłowski[‡], Piet Lammers[§] and Ioan Manolescu[¶]

6 March 2026

Abstract

We study the isotropic six-vertex model on \mathbb{Z}^2 with spectral parameter $\Delta \in [-1, -1/2]$, that is, with weights $\mathbf{a} = \mathbf{b} = 1$ and $\mathbf{c} \in [\sqrt{3}, 2]$. We show that the associated height function converges, in the scaling limit, to a properly scaled full-plane Gaussian free field. The result extends to anisotropic weights $\mathbf{a} \neq \mathbf{b}$ upon using a suitable embedding of the lattice.

A	Introduction	2
1	Motivation	2
2	Statement of our main result	7
3	First applications of our result	11
4	Overview of the proof and ingredients	14
5	Organisation of the paper	21
B	Proof of the main results	21
6	Sub-sequential GFF limits for the two-point function	22
7	Sub-sequential GFF limits for multi-point functions	38
8	Sub-sequential GFF limits	42
9	Proof of the main result (Theorem 2.8)	44
10	Proof for the anisotropic case (Theorem 3.3)	45
C	Ingredient 4: Spectral representation of correlation functions	48
11	Transfer matrix formalism	48
12	Spectral representation of general observables (Theorem 4.15)	50
13	Spectral representation for the two-point function (Theorem 4.12)	52
D	Ingredient 3: Regularity estimates and qualitative behaviour	55
14	Spin representation and RSW theory	55
15	FKG inequality of the spin representation	61
16	Percolation estimates for the spin configuration	65
17	Full-plane spin representation (incl. Theorem 4.10)	70
18	Torus and cylinder spin representation (Theorem 2.2)	72
19	Bounds on crossing counts (Theorem 4.9)	75
20	Level line tree and branching function	80
21	Regularity estimate in full plane (Theorem 4.5)	85
22	Regularity estimate (cylinder, Corollary 4.7)	89
23	Mixing estimate (Theorem 4.8)	91
E	Ingredient 2: Glimpse of scale invariance	97
24	Free energy second derivative (Theorem 4.4)	97
25	Proof of Theorem 4.3	101
26	Proof of the splitting (Proposition 25.3)	110
	References	123

^{*}Institut des Hautes Études Scientifiques, duminil@ihes.fr
[†]Université de Genève, hugo.duminil@unige.ch
[‡]ENSL, CNRS, Laboratoire de Physique, F-69342 Lyon, France, karol.kozlowski@ens-lyon.fr
[§]CNRS, Sorbonne Université, LPSM, piet.lammers@cnrs.fr
[¶]Université de Fribourg, ioan.manolescu@unifr.ch

Part A

Introduction

1 Motivation

1.1 Phase transitions in a nutshell

Physical systems undergoing a continuous phase transition can often be understood mathematically through lattice models, in which the microscopic degrees of freedom are encoded by variables attached to the sites, edges or faces of a graph. These models serve as effective descriptions of the underlying physical interactions while remaining amenable to rigorous analysis. In this framework, k -point correlation functions correspond to limits of expectations of products of *local* operators $\langle \prod_i O_{u_i}^{(i)} \rangle_\delta$. In this context, $\langle \cdot \rangle_\delta$ encodes the correlations of the model on a lattice of mesh size δ , u_1, \dots, u_k are points in space, the $O^{(i)}$ are (potentially different) observables measuring local properties of the system near the origin 0, and $O_{u_i}^{(i)}$ are the translates of these observables to the points u_i so that $O_{u_i}^{(i)}$ is measurable in terms of the behaviour in a small neighbourhood of u_i . In the sequel, we suppose each observable to be centred.

Away from the phase transition, correlations are extremely weak: for any collection of distinct points, the correlation functions decay exponentially fast in the separation distance (here of order $1/\delta$) between the points u_1, \dots, u_k . In other words, individual observables decorrelate fast and become asymptotically independent.

At a continuous phase transition, however, asymptotic independence still occurs, but correlations are expected to be much stronger. The k -point functions exhibiting a power-law decay with the distance:

$$\lim_{\delta \rightarrow 0} \delta^{-(\alpha_1 + \dots + \alpha_k)} \langle \prod_{i=1}^k O_{u_i}^{(i)} \rangle_\delta = \mathcal{C}_{O^{(1)}, \dots, O^{(k)}}(u_1, \dots, u_k) \quad (1)$$

for certain non-trivial functions $\mathcal{C}_{O^{(1)}, \dots, O^{(k)}}$ and *critical exponents* $\alpha_i = \alpha_i(O^{(i)})$ determining the rate of algebraic decay of the correlators at large distances.

The functions $\mathcal{C}_{O^{(1)}, \dots, O^{(k)}}$ are predicted to be *invariant under dilations* and *rotations*, as well as *universal* in the sense that they do not depend on the specific local interaction potential appearing in the definition of the lattice model. One purpose of statistical mechanics is to group models undergoing continuous phase transitions into *universality classes* with matching critical exponents α_i and limiting functions $\mathcal{C}_{O^{(1)}, \dots, O^{(k)}}$.

The very ideas on the behaviour of systems at continuous phase transitions date back to the foundational works in theoretical physics from the mid-1960s to the early 1970s, notably those of Fisher [Fis67; Fis66; Fis64], Kadanoff [Kad66], and Widom [Wid65b; Wid65a], culminating in Wilson's renormalisation-group theory [Wil71a; Wil71b]. The predictions have been tested extensively against perturbative calculations in a wide range of settings, effective models, and experimental data, as well as through numerical simulations [Kad+67; EF63; DS57; Bak61; El+12; KPS14].

1.2 The CFT conjectural limit of lattice models

It is not a priori clear what the hypothetical universality classes should be. One line of thought, originating in the works of Patashinskii and Pokrovskii [PP64] and later Polyakov [Pol68; Pol70b; Pol70a], is that universality classes should be captured by quantum

field theories that are invariant under not only scaling and rotations, but also *conformal transformations*. Such quantum fields theories are called Conformal Field Theories (CFTs). This led to the following loosely stated conjecture.

Conjecture 1. *The scaling limits in the sense of Equation (1) of two-dimensional statistical mechanics models undergoing a continuous phase transition, are given by the correlation functions of a CFT.*

In general dimension, the additional requirement of being invariant under conformal maps is not so restrictive, since the local conformal group is finite-dimensional and thus imposes only limited constraints. However, as observed in the groundbreaking work of Belavin, Polyakov, and Zamolodchikov [BPZ84], the situation improves dramatically in two dimensions: there, the conformal group is infinite-dimensional and the resulting local conformal symmetry leads to the integrability of numerous CFTs in the plane; we refer to the mathematically-oriented textbook [Gaw] for details.

Two-dimensional CFTs are classified by their central charge $c \in \mathbb{R}_{>0}$. After extensive investigations in the 1980s and 1990s, these theories turned out to be amenable to a substantial degree of analysis on physical grounds [Tes19; DMS97]. On the mathematical side, apart from the numerous early developments (cf. [Gaw; Seg88]), there has recently been a renewed interest in these CFTs within the probability community, driven by the development of the Schramm–Loewner Evolution [Sch00; Law14] (and the associated models constructed from it [SW12; MS16; She09; MS20; GRV16]) as well as by rigorous techniques based on [Dav+16; RV14; GRV19].

In many cases, one can provide a rather detailed description of the CFTs correlation functions. The understanding of the CFT with $c = 1$ – the free boson model, known to mathematicians as the Gaussian free field – is straightforward thanks to its Gaussian nature. In particular, Wick’s rule expresses k -point correlations in terms of products of 2-point correlations, which significantly streamlines their analysis. We refer to [She07; BP25] for details on this specific case. The rational CFTs, notably the minimal models corresponding to $c = 1 - 6\frac{(p-q)^2}{pq}$ with $p, q \geq 2$ coprime integers, are also very well-understood thanks to algebraic structures specific to these central charges [DGT24; DW25; GZ26; Fje+06; FRS02; FRS04a; FRS04b; FRS05]. For $c \geq 25$, the theory is well developed, as in [Tes95; Tes01], and more recently in the mathematics literature in [DS11; BP25; DS11; DS11; Gui+24]. Fully determining the correlations of CFTs still is the a subject of intense study both on the physical and mathematical sides, see e.g. [Ang+21; KRV20; Gui+24; CRV23; GZ26; DW25].

Today’s understanding of CFTs, even though not fully exhaustive, is very impressive. Still, even if one were given a complete and rigorous description of all two-dimensional CFTs, a thorough understanding of the continuous phase transition exhibited by a given lattice model still requires the following two challenging steps:

- (i) identifying which CFT (that is, which value of the central charge c and which representation of the corresponding CFT algebraic structure), if any, captures the relevant limit at criticality;
- (ii) identifying which correlation functions of that CFT describe the limits of the k -point functions of the chosen local observables. Concretely, in the setting of Equation (1), this amounts to identifying the functions $\mathcal{C}_{O^{(1)}, \dots, O^{(k)}}$ with appropriate correlation functions of the CFT in question.

1.3 From discrete to continuum

In order to go in this direction, two principal strategies have been developed over the years.

1.3.1 Extracting information from the transfer matrix formalism

For models that admit a transfer-matrix formulation, analysing the leading eigenvalues of the transfer matrix T_L of a system of size L with periodic boundary conditions, yields valuable information about the limiting behaviour. It is well known that the exponential growth rate of the largest eigenvalue of T_L determines the free energy of the model. Moreover, as argued heuristically in [Aff86; BCN86], the finite-size correction to the free energy – equivalently, the subleading correction to the principal eigenvalue – produces a constant that is interpreted, conjecturally, as the central charge of the conformal field theory expected to govern the model’s behaviour at criticality.

Additional information can be extracted from the large- L asymptotics of the subleading eigenvalues of the transfer matrix and the so-called translation operator (both of which will appear in this paper). In particular, conjectural equations relate the spacing of these eigenvalues to the possible values of critical exponents associated with local observables. The above physical picture is thus rather comprehensive, although it does not specify which critical exponent corresponds to a given lattice correlation function. This last step is typically achieved, at a physical level of rigour, through symmetry considerations. We refer to [Car86; Car08] for reviews.

While constructing a transfer matrix from a local Hamiltonian is often straightforward, extracting the large-volume asymptotics is an arduous task. For a generic model, this problem is hopeless – even at a heuristic level. However, the situation improves dramatically for integrable models of two-dimensional statistical mechanics [Bax82].

The simplest example in this direction is the two-dimensional Ising model in vanishing external field, originally solved by Onsager [Ons44] and later revisited in many different ways (see [MW14; Dum23] for historical accounts). Kaufman [Kau49] and later Lieb-Schultz-Mattis [LSM61] observed that the Ising model becomes equivalent, through a simple algebraic transformation, to a model of non-interacting (or *free*) fermions. This free-fermionic structure explains, to a large extent, the particularly simple exact solvability of the model. In particular, it leads to a closed and fully explicit formula for the eigenvalues of its transfer matrix, see for example [Bax82] for a modern treatment. This analysis provides evidence that the scaling limit of the model at criticality is governed by the conformal field theory with central charge $c = 1/2$.

The situation is considerably more involved for other integrable models of two-dimensional statistical mechanics undergoing continuous phase transitions. These models are genuinely interacting, meaning that no simple mapping reduces them to free fermions. They remain solvable in the sense that their transfer matrices can be diagonalised in a relatively explicit manner, but doing so requires solving a system of Bethe Ansatz equations, determining which solutions correspond to the dominant and and to the tower of sub-dominant eigenvalues, and analysing their behaviour as L tends to infinity.

In the mid-1980s and 1990s, substantial effort produced, on heuristic grounds, expansions of the top eigenvalues of the transfer matrix for a wide range of integrable two-dimensional models at criticality [DV95; dW85; IKR89; KB90; KBP91; KWZ93]. Although non-rigorous, this body of work yielded extensive conjectures for the central charges and critical exponents governing the scaling limits of many integrable lattice models.

A main obstacle to full mathematical rigour lies in the difficulty of proving the *condensation* of Bethe roots, a property originally conjectured by Hülten [Hul38]. Significant progress was made in [DS09; Dum+21; Dum+22; Gus80; Gus85; Koz18], culminating in the first rigorous derivations of expansions of top eigenvalues for the transfer matrix of the six-vertex model [Koz18] and for the staggered six-vertex model [FGK25], which arises in the analysis of the XXZ spin-1/2 chain at finite temperature.

1.3.2 Discrete holomorphicity and conformal invariant scaling limits

A second strategy roots in the development of *discrete holomorphicity*. The development of boundary CFTs, which incorporate the effect of boundary conditions, has significantly deepened the physical analysis of conformal field theories. The intuition that studying physical systems in planar domains can reveal additional structure has led, within the mathematical literature, to major breakthroughs in our understanding of the scaling limits of two-dimensional models. The idea is to harvest the fact that conformally invariant (or rather covariant) families of correlations in the continuum are often harmonic or holomorphic solutions of certain boundary value problems. It is therefore natural to expect that in a certain sense, discrete ancestors to these correlations are discrete harmonic or holomorphic solutions to the corresponding discrete boundary value problem. This observation proved successful for a few models during the first decade of the millennium.

One of the early rigorous proofs of conformal invariance concerns domino tilings, which possess an underlying free-fermionic structure. Kenyon [Ken00; CKP01] established conformal invariance in the scaling limit for the height-function distribution of domino tilings on bipartite Temperleyan planar graphs by showing that certain observables are discrete holomorphic and satisfy a Dirichlet boundary value problem. He further proved the convergence of the height function (seen as a distribution) to the Gaussian free field, corresponding to the CFT with central charge $c = 1$ (the free boson). This analysis was extended to periodic isoradial graphs in [Til07], a class preserving exact solvability and discrete holomorphic structure.

In 2001, Smirnov [Smi01] proved the celebrated Cardy formula [Car92] for Bernoulli site percolation on the triangular lattice. Smirnov's argument relies on a certain approximately discrete holomorphic observables satisfying certain boundary conditions, which can be shown to converge to the solution of the continuum analogue of this boundary value problem. This major achievement led to a very precise description of the critical regime, including the determination of critical exponents and links to $c = 0$ CFT [LSW02; Smi01; CN07; CN06; CN26; CF24].

In 2004, Lawler, Schramm, and Werner [LSW04] derived the conformal invariance of the loop-erased random walk and the associated Uniform Spanning Tree (UST) by exploiting properties of the discrete Green function, making the approach close in spirit to strategies above. Note that the UST is closely connected to the dimer model and also enjoys a free-fermionic structure.

In 2008, Smirnov [Smi10] and Chelkak-Smirnov [CS12] developed a discrete-holomorphic framework for fermionic observables of the two-dimensional Ising model, showing that they converge in the scaling limit to holomorphic solutions of appropriate Riemann-Hilbert boundary-value problems. Extensions of this method to isoradial graphs yielded robust convergence results for fermionic observables and demonstrated universality and conformal symmetry in the scaling limit. Further developments led to a multitude of scaling-limit results for interfaces [Che+14], the energy-density observable [HS13], its n -point generalisations [Hon], spin correlators [CHI15], and eventually correlation functions corresponding to primary CFT operators [CHI22]. These works collectively yield the link to the $c = \frac{1}{2}$ CFT.

1.4 Beyond the free fermion point

Except for Bernoulli site percolation [Smi01] and a reverse-engineered model known as the harmonic explorer [SS05], all rigorous progress on scaling limits achieved so far concern models that reduce to *non-interacting* (free) fermions. Yet the main appeal of the CFT universality conjecture (Conjecture 1) for continuous phase transitions lies precisely in its

breadth: it is intended to encompass genuinely *interacting* models, far removed from any structure reminiscent of free theories. However, once interactions are present, the analysis becomes considerably more intricate.

Rigorous progress has been achieved for small perturbations of free-fermionic models, such as weakly interacting dimers. These results rely on *constructive renormalisation*, the rigorous embodiment of Wilson’s renormalisation-group program. Initiated in [PS], this method has been further developed to establish scaling properties for a variety of perturbative models [Mas04; GGM12; BFM14a; BFM14b; GMT17; GT19]. While these works represent a significant advance, they remain confined to perturbative regimes: the very nature of constructive renormalisation makes it extremely difficult to reach non-perturbative settings or scaling limits far from free fermionic ones.

In this light, the genuinely interacting *integrable* models of two-dimensional statistical mechanics form a natural testing ground for Conjecture 1. These models are not equivalent to free fermions except, possibly, at isolated parameter values, and they possess a rich algebraic structure arising from the representation theory of quantum groups. This structure manifests in key identities such as the star–triangle relation [Ken99; DLM18] and in the commutativity of families of transfer matrices [Bax72; STF79], providing powerful tools that are unavailable for generic models.

A particularly prominent example of a genuinely interacting integrable model is the *six-vertex model*; see, for instance, [LW80; Bax82].

This paper provides the first scaling-limit result for the six-vertex model across a substantial range of parameters. Specifically, for $\Delta \in [-1, -\frac{1}{2}]$, we prove that the height function of the full-plane model converges to the Gaussian free field as the mesh size tends to zero. Although our result is currently restricted to the full-plane setting, it constitutes an important first step towards a more general framework for establishing scaling limits of planar models, as it applies to a broad class of genuinely interacting models. As such, it belongs to a larger program aimed at determining the behaviour of critical lattice models in two dimensions deep in the interaction regime.

It is of course premature to discuss the proof (especially since we did not yet introduce the model nor the result properly) in full detail, but let us briefly indicate that the argument draws on ideas inspired by both the transfer-matrix formalism and discrete holomorphicity. In this sense, it synthesises elements of the two historical approaches outlined above. Roughly speaking, we exploit the properties of a certain spectral measure that encodes the averaged behaviour of the eigenvalues of the transfer matrix and shift operator, and use this to show that the k -point correlations of the six-vertex model are harmonic in the limit. To achieve this, we combine the rotational invariance of the model obtained in [Ave+a] with a novel analysis of spectral properties. This constitutes the main innovation of the paper. The harmonicity, combined with an analysis of the behaviour near singular points, allows us to identify these limiting k -point correlations, which can then be bootstrapped to obtain stronger modes of convergence. To be more precise, our proof begins by establishing convergence only along certain suitable sub-sequences. It ends up being sufficient to deduce the full result, which in turn retroactively ensures that passing to sub-sequences was unnecessary.

The ability to extract such sub-sequential limits is in fact another main innovation of the paper. The framework in which we work permits the use of qualitative estimates for the model – reminiscent of the RSW-type theory familiar to percolation specialists (we shall discuss RSW theory in depth) – to obtain compactness and extract convergent sub-sequences of the spectral measure mentioned above. This strategy resonates with

what made the proof of Cardy’s formula for Bernoulli site percolation on the triangular lattice possible. In [Smi01], Smirnov goes around the problem of having observables that are only approximately holomorphic by showing that RSW-theory enables one to extract sub-sequential limits for these discrete observables, which end up being holomorphic. A leitmotif emerges: in order to move beyond the ‘ultra-integrable’ cases of models possessing a free-fermion structure in which exactly discrete harmonic or holomorphic observables can be found, one likely needs to work in a setting where the absence of exact discrete harmonicity or holomorphicity can be compensated by *a priori* estimates that ensure the existence of sub-sequential limits. The objects introduced in this paper illustrates the advantage and the potential of such a perspective.

2 Statement of our main result

This section formally introduces the model and states the main results.

2.1 Definition of the six-vertex model

In this paper, the six-vertex model is defined on graphs which locally look like the square lattice graph. We start with a definition on tori. For $M, L \in \mathbb{Z}_{\geq 4}$, let $\mathbb{T}_{M,L} = (V(\mathbb{T}_{M,L}), E(\mathbb{T}_{M,L}))$ denote the toroidal square grid on the vertex set $V(\mathbb{T}_{M,L}) := (\mathbb{Z}/M\mathbb{Z}) \times (\mathbb{Z}/L\mathbb{Z})$, with edges placed between vertices at Euclidean distance 1 from each other.

An *arrow configuration* ω is an assignment of an orientation to each edge. An arrow configuration is said to satisfy the *ice rule* (or be a *six-vertex configuration*) if every vertex has exactly two incoming and two outgoing edges. As a result, there are six possible arrangements of incoming and outgoing edges around each vertex, labelled according to Figure 1.

Definition 2.1 (Six-vertex model on the torus). For parameters $\mathbf{a}_1, \mathbf{a}_2, \mathbf{b}_1, \mathbf{b}_2, \mathbf{c}_1, \mathbf{c}_2 > 0$, the *weight* of an arrow configuration ω on $\mathbb{T}_{M,L}$ is given by

$$W_{6V}(\omega) = \mathbb{1}[\omega \text{ satisfies the ice rule}] \cdot \mathbf{a}_1^{n_1} \mathbf{a}_2^{n_2} \mathbf{b}_1^{n_3} \mathbf{b}_2^{n_4} \mathbf{c}_1^{n_5} \mathbf{c}_2^{n_6}, \quad (2)$$

where n_i denotes the number of vertices in $V(\mathbb{T}_{M,L})$ of type i in ω . The Gibbs measure $\mathbb{P}_{\mathbb{T}_{M,L}}$ on arrow configurations ω is given by

$$\mathbb{P}_{\mathbb{T}_{M,L}}[\{\omega\}] := \frac{1}{Z_{\mathbb{T}_{M,L}}} \cdot W_{6V}(\omega), \quad (3)$$

where $Z_{\mathbb{T}_{M,L}}$ is the unique constant, called the *partition function*, rendering $\mathbb{P}_{\mathbb{T}_{M,L}}$ a probability measure.

The weights were taken to be positive reals for a probabilistic interpretation. In this work, we further specialise to the case where

$$\mathbf{a}_1 = \mathbf{a}_2 =: \mathbf{a}; \quad \mathbf{b}_1 = \mathbf{b}_2 =: \mathbf{b}; \quad \mathbf{c}_1 = \mathbf{c}_2 =: \mathbf{c}, \quad (4)$$




Figure 1: The six vertex configurations with labels and weights.

which renders $W_{6V}(\omega)$ invariant under flipping all orientations in ω . The parameters \mathbf{a} , \mathbf{b} , and \mathbf{c} are always fixed in this article, which is why they do not appear in notations. It is standard to introduce the *spectral parameter* defined by the formula

$$\Delta = \Delta(\mathbf{a}, \mathbf{b}, \mathbf{c}) = \frac{\mathbf{a}^2 + \mathbf{b}^2 - \mathbf{c}^2}{2\mathbf{a}\mathbf{b}}. \quad (5)$$

The model may be extended to infinite volume in the following fashion. Any configuration on $\mathbb{T}_{M,L}$ that obeys the ice rule has the same number of left-arrows on each vertical column – we call this *preservation of horizontal arrows*. A configuration is called *balanced* if, for every vertical column of horizontal arrows, the number of left-arrows equals the number of right-arrows. Write $\{\text{balanced}\}$ for the collection of balanced arrow configurations satisfying the ice rule. From now on, L is always even, so that balanced configurations exist.

For $\Delta < -1$, the six-vertex model is known to be in a localised regime, implying trivial limiting behaviour [Dum+21; RS22; GP23]. Therefore, we focus in the whole paper on the case $\Delta \geq -1$, which corresponds, when $\mathbf{a} = \mathbf{b} = 1$, to $\mathbf{c} \leq 2$. We shall derive the following known result *en passant* in Lemma 18.7).

Theorem 2.2 (Infinite-volume six-vertex model). *Fix $\mathbf{a} = \mathbf{b} = 1$ and $\mathbf{c} \in [1, 2]$. The weak limit of the measures $\mathbb{P}_{\mathbb{T}_{M,L}}[\cdot | \{\text{balanced}\}]$ exists when the limits are taken in the following order: first M tends to infinity, and then L tends to infinity. We denote it $\mathbb{P}_{\mathbb{Z}^2}$ and call it the six-vertex measure in the plane with slope zero. It is invariant under the automorphism group of \mathbb{Z}^2 .*

This result asserts that $\mathbb{P}_{\mathbb{Z}^2}$ is the unique probability measure on arrow configurations of the square lattice graph \mathbb{Z}^2 such that

$$\lim_{L \rightarrow \infty} \lim_{M \rightarrow \infty} \mathbb{P}_{\mathbb{T}_{M,L}}[A | \{\text{balanced}\}] = \mathbb{P}_{\mathbb{Z}^2}[A] \quad (6)$$

for any event A that is measurable in terms of the orientation of finitely many edges of the square lattice graph \mathbb{Z}^2 . The measure $\mathbb{P}_{\mathbb{Z}^2}$ may be characterised in several other ways. In addition to the above description, it is also the weak limit of $\mathbb{P}_{\mathbb{T}_{M,L}}$ as M, L tend to infinity in arbitrary fashion, it is the unique ergodic Gibbs measure which is invariant under a global arrow flip, and it is the unique minimiser of a free energy functional related to the six-vertex model. The last two equivalent statements do not play a role in this work.

Theorem 2.2 and the above equivalences follow from the general analysis of height functions in [She05] combined with *delocalisation* of the height function. Delocalisation was first derived at $\mathbf{c} = 1$ [Cha+21], $\mathbf{c} = 2$ [GP23], and $\mathbf{c} \in [(2 + 2^{1/2})^{1/2}, 2]$ [Lis21], before the full range $\mathbf{c} \in [1, 2]$ was covered in [Dum+24] via a Bethe Ansatz argument, and later in [GL25] using a percolation approach.

2.2 Height function of the six-vertex model

A *height function* is a function $h : F(\mathbb{Z}^2) \rightarrow \mathbb{Z}$ on the faces $F(\mathbb{Z}^2)$ of the square lattice which differs by exactly ± 1 between any two adjacent faces. We also require the face on the north-east of the origin to have an *even* height.

We will consider height functions up to addition of an even constant. Formally, consider two height functions h' and h *equivalent* if there exists some constant $a \in 2\mathbb{Z}$ such that $h'(u) = h(u) + a$ for all $u \in F(\mathbb{Z}^2)$. Gradients of height functions are simply the equivalence classes of height functions for this equivalence relation. For all practical purposes, we identify gradients of height functions with any representative of the equivalence class. Finally, we shall write also h for the piecewise constant function

$$h : \mathbb{R}^2 \rightarrow \mathbb{Z}, (x, y) \mapsto h(\text{the face whose bottom-left corner is } ([x], [y])), \quad (7)$$




Figure 2: The six-vertex configuration and gradient height function are related such that the higher height is on the left of each arrow.

and, for any $\delta > 0$, we define the scaled height function $h^{(\delta)} : \mathbb{R}^2 \rightarrow \mathbb{Z}$, $u \mapsto h(u/\delta)$.

Definition 2.3 (Height function of a six-vertex configuration). Full-plane six-vertex configurations are in bijection with gradient height functions. More precisely, we associate any six-vertex configuration ω with the height functions for which the height of the face on the *left* of each arrow is one unit higher than the height of that on its *right*; see Figure 2.

The multi-point correlation functions of gradient height-functions will be the core observables in our study of the six-vertex model's height function.

Definition 2.4 (Six-vertex multi-point correlation functions). The k -point correlation function of the six-vertex model assigns to the vector $\mathbf{u} = (u_1, u'_1, \dots, u_k, u'_k) \in (\mathbb{R}^2)^{2k}$ whose $2k$ coordinates are built out of points $u_1, u'_1, \dots, u_k, u'_k \in \mathbb{R}^2$, the quantity

$$\Phi_k(\mathbf{u}) := \mathbb{E}_{\mathbb{Z}^2} \left[\prod_{i=1}^k (h(u'_i) - h(u_i)) \right], \quad (8)$$

where h denotes any height function associated with the six-vertex configuration sampled according to $\mathbb{P}_{\mathbb{Z}^2}$. Moreover, for $\delta > 0$, the *scaled k -point correlation function* is defined as

$$\Phi_k^{(\delta)}(\mathbf{u}) := \Phi_k(u_1/\delta, u'_1/\delta, \dots, u_k/\delta, u'_k/\delta). \quad (9)$$

The k -point correlation function is well-defined, as the integrand in the expectation depends only on the gradient of the height function. We immediately recognise a few basic properties of Φ_k , which we often use without further mention:

- $\Phi_k \equiv \Phi_k^{(\delta)} \equiv 0$ for k odd since h and $-h$ have the same distribution,
- Φ_k and $\Phi_k^{(\delta)}$ are antisymmetric under swapping u_i and u'_i ,
- Φ_k and $\Phi_k^{(\delta)}$ are invariant under permuting the pairs $\{u_i, u'_i\}$,
- Φ_k and $\Phi_k^{(\delta)}$ satisfy the following *additivity property* for any u_1, u'_1, u''_1 and for any fixed $\mathbf{v} = (u_2, u'_2, \dots, u_k, u'_k)$:

$$\Phi_k(u_1, u'_1, \mathbf{v}) + \Phi_k(u'_1, u''_1, \mathbf{v}) = \Phi_k(u_1, u''_1, \mathbf{v}). \quad (10)$$

Although one primary object of interest is the k -point correlation, we will also treat the height function h as a random distribution. To that end, we introduce the distribution defined by integrating h against test functions.

Definition 2.5 (Six-vertex test functions). A *generalised test function* is a finite, compactly supported, signed measure φ on \mathbb{R}^2 with $\varphi(\mathbb{R}^2) = 0$. For any generalised test function φ , define

$$\langle h^{(\delta)}, \varphi \rangle := \int h^{(\delta)}(x) d\varphi(x) = \int h(x/\delta) d\varphi(x). \quad (11)$$

Observe that $\langle h^{(\delta)}, \varphi \rangle$ is a gradient measurable random variable because $\varphi(\mathbb{R}^2) = 0$. We may therefore interpret it as a random variable.

2.3 Background on the Gaussian free field

With the model in place, we next describe the scaling limit that will ultimately arise. Write $|\cdot|$ for the Euclidean norm on \mathbb{R}^2 . Define the full-plane Green function $G_{\mathbb{R}^2}$ as

$$G_{\mathbb{R}^2} : \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow (-\infty, \infty], (x, y) \mapsto -\frac{1}{2\pi} \log |y - x|. \quad (12)$$

Our main result below says that as δ tends to zero, the random height function $h^{(\delta)}$ obtained from $\mathbb{P}_{\mathbb{Z}^2}$ converges (up to scaling by a constant $\sigma = \sigma(\Delta)$) to the *Gaussian free field* (GFF), which we denote Γ . There are various mathematical ways to view the GFF depending on the desired regularity of the object; here, we need the following three (which resonate with Definitions 2.4 and 2.5). For more details on GFF, see [BP25].

Definition 2.6 (GFF definitions). We consider three perspectives on the GFF.

(i) **Multi-point correlation functions.** For any $k \geq 1$, define

$$\mathcal{D}_k := \{(u_i, u'_i)_{i=1, \dots, k} \in (\mathbb{R}^2 \times \mathbb{R}^2)^k : \forall i \neq j, \{u_i, u'_i\} \cap \{u_j, u'_j\} = \emptyset\}, \quad (13)$$

and define the functions

$$\Psi_k^{\text{GFF}} : \mathcal{D}_k \rightarrow \mathbb{R}, \mathbf{u} \mapsto \sum_{a, \pi} (-1)^{\varepsilon(a)} \prod_{ij \in \pi} G_{\mathbb{R}^2}(a_i, a_j), \quad (14)$$

where:

- The sum over a runs over maps $\{1, \dots, k\} \rightarrow \{u_1, u'_1, \dots, u_k, u'_k\}$ with $a_i \in \{u_i, u'_i\}$,
- $\varepsilon(a) \in \mathbb{Z}/2\mathbb{Z}$ is the parity of the number of indices i such that u_i is chosen,
- The sum over π runs over all *pairings* of $\{1, \dots, k\}$, that is, all partitions of $\{1, \dots, k\}$ into *pairs* (i.e. sets containing two points).

These are precisely the correlation functions corresponding to a Gaussian process with covariance $G_{\mathbb{R}^2}$. Notice that $\Psi_k^{\text{GFF}} \equiv 0$ when k is odd since π is then empty.

(ii) **Finite-dimensional marginals.** We say that a generalised test function φ has *finite Dirichlet energy* if $\int G_{\mathbb{R}^2}(u, v) d\varphi(u) d\varphi(v) < \infty$. For any finite family $\boldsymbol{\varphi} = (\varphi_1, \dots, \varphi_n)$ of finite Dirichlet energy generalised test functions, introduce the associated $n \times n$ covariance matrix $\Sigma(\boldsymbol{\varphi})$ defined via

$$\Sigma(\boldsymbol{\varphi})_{ij} := \int G_{\mathbb{R}^2}(u, v) d\varphi_i(u) d\varphi_j(v). \quad (15)$$

We then think of $(\langle \Gamma, \varphi_i \rangle)_i$ as a random variable having the law $\mathcal{N}(0, \Sigma(\boldsymbol{\varphi}))$.

(iii) **Random element of a negative regularity Hölder space.** Fix $\alpha \in (-1, 0)$ and a bounded open set $U \subset \mathbb{R}^2$. Recall that the *Hölder space* $\mathcal{C}^\alpha(U)$ of regularity α on U is the completion of $C_c^\infty(\mathbb{R}^2)$ with respect to the semi-norm

$$\|\cdot\|_{\mathcal{C}^\alpha(U)} : C_c^\infty(\mathbb{R}^2) \rightarrow \mathbb{R}, f \mapsto \sup_{\substack{(\varepsilon, \varphi) \in (0, 1] \times \mathcal{T}_1((-1, 1)^2), \\ \text{support}(\varphi) \subset U/\varepsilon}} \varepsilon^{-\alpha} \int f(x/\varepsilon) d\varphi(x), \quad (16)$$

where $\mathcal{T}(U)$ denotes the set of generalised test functions whose support is included in U and $\mathcal{T}_1(U) \subset \mathcal{T}(U)$ the set of such test functions whose density is 1-Lipschitz. We view Γ as a random element in $\mathcal{C}^\alpha(U)$ such that, for any finite $\boldsymbol{\varphi} = (\varphi_1, \dots, \varphi_k)$ with $\varphi_i \in \mathcal{T}(U)$ for every i , the law of $(\langle \Gamma, \varphi_i \rangle)_i$ is $\mathcal{N}(0, \Sigma(\boldsymbol{\varphi}))$.

2.4 Statement of the main result

Three modes of convergence are considered.

Definition 2.7 (GFF convergence). Consider some fixed value $\sigma \in \mathbb{R}_{\geq 0}$ as well as a random gradient height function $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ in some probability measure \mathbb{P} . We say that the *scaling limit* of h is $\sigma\Gamma$ or $\sigma \cdot \text{GFF}$ if *all* of the statements below hold true.

- (i) **Convergence of multi-point correlation functions.** For any $k \geq 1$, $\Phi_k^{(\delta)}$ converges uniformly on compact subsets of \mathcal{D}_k to $\sigma^k \Psi_k^{\text{GFF}}$ as δ tends to zero.
- (ii) **Convergence of finite-dimensional marginals.** For any family $\varphi = (\varphi_1, \dots, \varphi_n)$ of finite Dirichlet energy generalised test functions, the law of $(\langle h^{(\delta)}, \varphi_i \rangle)_i$ converges weakly to $\mathcal{N}(0, \sigma^2 \Sigma(\varphi))$ as δ tends to zero.
- (iii) **Convergence in law in a negative regularity Hölder space.** For any $\alpha \in (-1, 0)$ and any open bounded $U \subset \mathbb{R}^2$, the law of $h^{(\delta)}$ converges to that of $\sigma\Gamma$ in $\mathcal{C}^\alpha(U)$ as δ tends to zero. Finally, we also require convergence in the classical Besov spaces $\mathcal{B}_{p,q}^\alpha(U)$ and Sobolev spaces $W^{\alpha,p}(U)$ for any $p \in [1, \infty)$ and $q \in [1, \infty]$ (see Section 8 for details).

We are now in a position to state our main result.

Theorem 2.8 (Scaling limit of the six vertex model with isotropic weights). *The scaling limit of the six-vertex model's height function on \mathbb{Z}^2 with $\mathbf{a} = \mathbf{b} = 1$ and $\sqrt{3} \leq \mathbf{c} \leq 2$ is $\sigma \cdot \text{GFF}$, where*

$$\sigma^2 = \frac{2}{\arccos \Delta} = \frac{1}{\arcsin \frac{\mathbf{c}}{2}}. \quad (17)$$

Let us comment on our requirement on the parameters. In terms of Δ , the previous theorem covers the regime $-1 \leq \Delta \leq -\frac{1}{2}$. The restriction on the isotropic case (i.e., with weights $\mathbf{a} = \mathbf{b} = 1$) is lifted in Theorem 3.3 below. Recall that the isotropic six-vertex model with $\mathbf{c} > 2$ ($\Delta < -1$) is known to be in a localised regime. Theorem 2.8 is expected to hold true for all $\mathbf{c} \in (0, 2]$ (i.e., $\Delta \in [-1, 1)$) and was previously obtained at the free fermion point $\mathbf{c} = \sqrt{2}$ ($\Delta = 0$) in [Ken00] and for \mathbf{c} close to $\sqrt{2}$ in [GMT17].

Two aspects of our proof require lower bounds on \mathbf{c} . First, the Fortuin–Kasteleyn–Ginibre (FKG) property for several representations of the six-vertex model requires that $\mathbf{c} \geq 1$. This property is used to obtain regularity of our objects at various stages of the proof. Second, our proof relies on the asymptotic rotational invariance of the multi-point correlation functions Φ_k , which was obtained in [Ave+a] using the corresponding random-cluster model [Dum+20]. For that result to apply, the random-cluster model must exhibit the FKG property, which requires $\mathbf{c} \geq \sqrt{3}$.

An extension of this asymptotic rotational invariance result to $\mathbf{c} \geq 1$ is conceivable, for instance by using the above-mentioned representations of the six-vertex model. This would immediately allow the extension of Theorem 2.8 to the interval $1 \leq \mathbf{c} \leq 2$ (corresponding to $-1 \leq \Delta \leq 1/2$). However, circumventing the absence of the FKG inequality for the representations of the six-vertex model appears to be a formidable challenge, placing the case $0 < \mathbf{c} < 1$ beyond the reach of current methods.

3 First applications of our result

The six-vertex model in the regime $-1 \leq \Delta < 1$ lies at the crossroads of a large family of two-dimensional lattice models. It is closely related to the dimer model, the Ising and Potts models, the critical random-cluster model, loop $O(n)$ models, Ashkin–Teller models, random

permutations, and quantum spin chains [FK72; Nie82; GP23; Lis22; MW68]. While it is not yet clear how much information can ultimately be extracted from the GFF convergence established here, several significant applications are already available.

3.1 Critical exponents of other two-dimensional lattice models

Thanks to the Baxter–Kelland–Wu (BKW) correspondence [BKW76], the six-vertex model is intimately connected, at its continuous phase transition, to the critical *random-cluster model* – also called the *Fortuin–Kasteleyn (FK) percolation* – introduced in [FK72]. Our main result enables the following consequences for random-cluster models.

- **One-arm exponent** α_1 . For the random-cluster model with cluster-weight $q \in [1, 4]$, [Che+b] obtains the existence and value of the one-arm critical exponent α_1 , describing the decay of the probability that a vertex connects to distance n . In turn, this derives the classical exponents η , ζ , and δ governing the behaviour of the two-point function, the cluster-size tail at criticality, and the ghost-field connectivity; see [Che+b] for details.
- **Two-arm exponent** α_2 . For the random-cluster model with $q \in [1, 4]$, [ADH] obtains the existence and value of the two-arm exponent α_2 , describing the probability that a vertex lies on a primal/dual interface extending to distance n . As a consequence, one deduces the fractal dimension of any sub-sequential scaling limit of critical interfaces, in agreement with the predicted $\text{CLE}(\kappa)$; see [ADH] for details.
- **Energy exponent** ι . For the random-cluster model with $q \in [4 - \varepsilon, 4]$ (for some small value of ε), the articles [Che+a; Ave+b] determine the so-called *influence exponent* ι , which controls the covariance of observables at criticality. Combined with the scaling relations established in [DM22], this yields the thermodynamic critical exponents α , β , γ , and ν , governing respectively the behaviour of the free energy, the spontaneous magnetisation, the susceptibility, and the correlation length.

Finally, since the random-cluster model is coupled to the Potts model, all the above critical exponents transfer directly to the two, three and four state Potts models. These results were already known for two-state Potts case, better known as the Ising model, but are new for the three- and four-state Potts models.

We expect that more results can be obtained in this direction, both for the random-cluster model and other models.

Remark 3.1. We stress that the critical exponents obtained here do not rely on evaluating the top eigenvalues of the transfer matrix. Instead, they harvest Theorem 2.8 which is based on an analysis of the average behaviour of certain eigenvalues, well-separated from the spectral edge. In this sense, our approach circumvents the major difficulty of providing a rigorous justification for computations of the leading eigenvalues.

3.2 Scaling limit of specific random-cluster observables

Harvesting the BKW correspondence [BKW76] allows one to derive the scaling limit of certain random-cluster observables at criticality. In particular, the characteristic function of six-vertex test functions $\langle h^{(\delta)}, \varphi \rangle$ can be expressed as the expectation – under the random-cluster model – of a product of suitably twisted weights associated with the loops of a percolation configuration. More precisely, if $\mu = \frac{1}{2\pi} \arccos(\sqrt{q}/2)$ and $\cos_\mu(\cdot) := \cos(\cdot + 2\pi\mu)/\cos(2\pi\mu)$, then

$$\mathbb{E}_{\mathbb{Z}^2}^{\text{GV}} \left[e^{i\langle h^{(\delta)}, \varphi \rangle} \right] = \phi_{\delta\mathbb{Z}^2, q} \left[\prod_{\ell \in \mathcal{L}} \cos_\mu(\varphi(\text{int}(\ell))) \right], \quad (18)$$

where $\phi_{\delta\mathbb{Z}^2, q}$ denotes the law of the ensemble \mathcal{L} of loops on $\delta\mathbb{Z}^2$ arising from the loop representation of the critical random-cluster measure with cluster-weight q , and where $\text{int}(\ell)$ is the interior of the loop ℓ , i.e., the bounded connected component of $\mathbb{R}^2 \setminus \ell$ when ℓ is viewed as a continuous simple path. Variants of this identity have appeared repeatedly in the literature; see, for instance, [Dub11]. We refer to [Ave+a] for more details on this formula.

The convergence to the Gaussian free field yields the following corollary, which underpins the derivation of the critical exponents presented in the previous section.

Corollary 3.2. *Fix $q \in [1, 4]$. For every finite Dirichlet energy generalized test function φ ,*

$$\lim_{\delta \rightarrow 0} \phi_{\delta\mathbb{Z}^2, q} \left[\prod_{\ell \in \mathcal{L}} \cos_{\mu}(\varphi(\text{int}(\ell))) \right] = \exp \left(-\frac{1}{2} \sigma^2 \iint G_{\mathbb{R}^2}(x, y) \varphi(x) \varphi(y) dx dy \right), \quad (19)$$

where $\sigma^2 = 2 / \arccos(-\frac{\sqrt{q}}{2})$.

3.3 Applications to the anisotropic six-vertex model

The universality of the the random-cluster model derived in [Dum+20] enables us to transfer our main result to the anisotropic six-vertex model (weights $\mathbf{a} \neq \mathbf{b}$). It is customary to parametrize $(\mathbf{a}, \mathbf{b}, \mathbf{c})$ in the following way: let $\zeta = \arccos(-\Delta)$ and $\theta \in (0, \pi)$ be the unique angle such that, if $1 > \Delta > -1$

$$\mathbf{a} \sin \frac{\zeta}{2} = \sin \left(\left(1 - \frac{\theta}{\pi}\right) \zeta \right), \quad \mathbf{b} \sin \frac{\zeta}{2} = \sin \frac{\theta \zeta}{\pi}, \quad \mathbf{c} = 2 \cos \frac{\zeta}{2}. \quad (20)$$

and if $\Delta = -1$,

$$\mathbf{a} = 2 \frac{\pi - \theta}{\pi}, \quad \mathbf{b} = 2 \frac{\theta}{\pi}, \quad \mathbf{c} = 2. \quad (21)$$

The parameter θ encodes the natural embedding of the square lattice; $\theta = \pi/2$ encodes the isotropic case. More precisely, let

$$\mathbf{L}_{\theta} : \mathbb{R}^2 \rightarrow \mathbb{R}^2, (x, y) \mapsto (x + \cos(\theta)y, \sin(\theta)y). \quad (22)$$

The combination of the universality result of [Dum+20], the consequences of the Baxter–Kelland–Wu coupling obtained in [Ave+a], and Theorem 2.8 implies the following result.

Theorem 3.3 (Scaling limit of the six vertex model with general weights). *The height function of the six-vertex model on $\mathbf{L}_{\theta}\mathbb{Z}^2$ with $\mathbf{a}, \mathbf{b}, \mathbf{c} > 0$ such that $\Delta \in [-1, -1/2]$ is converging in the sense of Definition 2.7, Items (i) and (ii) to $\sigma \cdot \text{GFF}$, where*

$$\sigma^2 = \frac{2}{\arccos \Delta} = \frac{1}{\arcsin \frac{\zeta}{2}} = \frac{2}{(\pi - \zeta)}. \quad (23)$$

Note that we do not claim convergence in the sense of Definition 2.7(iii), i.e., in Hölder spaces. This omission is purely technical. In order to keep the paper to a reasonable length, we chose to rely as much as possible on the so-called spin representation of the six-vertex height function (discussed at length below). Establishing the necessary RSW theory of this spin representation is relatively direct in the isotropic case but is not available the anisotropic setting. While [Ave+a] works directly in the anisotropic regime, it does not provide the regularity estimate required here (see Remark 10.1 for further discussion).

4 Overview of the proof and ingredients

We give below a roadmap to the proof of our main result, Theorem 2.8. Theorem 3.3 will be derived from Theorem 2.8 in Section 10; outside of that section, we only consider the isotropic case $\mathbf{a} = \mathbf{b} = 1$.

4.1 Overview of the proof structure

The proof of the convergence result in the isotropic case (Theorem 2.8) consists of four steps performed in Part B, informally described as follows. We say that h converges to $\sigma \cdot \text{GFF}$ along some sequence $(\delta_n)_n$ tending to zero if each of the convergences in Definition 2.7 holds true along that sequence of scales.

1. **Theorem 6.1** asserts that the two-point function $\Phi_2^{(\delta)}$ satisfies a dichotomy: either $\Phi_2^{(\delta)}$ converges to $\sigma^2 \cdot \Psi_2^{\text{GFF}}$ for some $\sigma \geq 0$, or such a convergence holds true along two sub-sequences with two distinct values of σ .
2. **Theorem 7.1** enables one to pass from two-point to multi-point correlation functions in the following sense: if along some sub-sequence $(\delta_n)_n$, $\Phi_2^{(\delta_n)}$ tends to $\sigma^2 \cdot \Psi_2^{\text{GFF}}$, then $\Phi_k^{(\delta_n)}$ tends to $\sigma^k \cdot \Psi_k^{\text{GFF}}$ for every k . This extends the dichotomy of Theorem 6.1 to all multi-point correlation functions.
3. **Theorem 8.1** states that if along some sub-sequence $(\delta_n)_n$, all correlation functions converge to those of $\sigma \cdot \text{GFF}$, then the limit of $h^{(\delta_n)}$ is $\sigma \cdot \text{GFF}$ (in the sense of Definition 2.7). This extends the dichotomy of Theorem 6.1 to all modes of GFF convergence.
4. **Theorems 4.3 and 4.4** jointly imply that if for some $(\delta_n)_n$, $h^{(\delta_n)}$ converges to $\sigma \cdot \text{GFF}$, then $\sigma^2 = 2/\arccos \Delta$. This makes the dichotomy collapse to a single case, and completes the proof of the main result.

To compactly state the proofs of these steps in Part B, we rely on four proof ingredients which are used as “black boxes” in Part B (see Table 1). These “black boxes” are developed in Parts C–E. In short, they are described as follows.

1. **Rotational invariance.** The correlation functions are asymptotically rotationally invariant. More precisely, any *sub-sequential* scaling limit of the correlation functions, is rotationally invariant.
2. **Glimpse of scale invariance.** A suitably chosen “observable” converges in the scaling limit, and we can calculate its limit explicitly.
3. **Regularity estimates and qualitative behaviour.** We establish suitable bounds on the correlation functions which hold true at all scales.
4. **Spectral representation of correlation functions.** Correlation functions may be expressed in terms of the spectra of two commuting transfer matrices.

The first two ingredients (rotation invariance and a weak form of scale invariance) echo

	Ingredient 1: <i>Rotation invariance</i>	Ingredient 2: <i>Scale invariance</i>	Ingredient 3: <i>Regularity of correlations</i>	Ingredient 4: <i>Spectral representation</i>
Step 1: <i>Two-point</i>	X		X	X
Step 2: <i>Multi-point</i>	X		X	X
Step 3: <i>Test functions</i>			X	
Step 4: <i>Finding σ</i>		X	X	

Table 1: Use of the proof ingredients in the global proof steps

the discussion at the beginning of the paper and are consistent with physics predictions, especially those arising from the renormalization-group formalism. We emphasize, however, that these properties are *not* obtained through a rigorous renormalization-group analysis. The third ingredient (regularity estimates) corresponds to qualitative bounds expected for generic continuous phase transitions.

The fourth ingredient is more mysterious. We interpret it as follows. Since the work of Polyakov [Pol68; Pol70b; Pol70a], it has been predicted that conformal invariance should follow from rotation, scale, and translation invariance (the latter being trivial for our model), provided the theory also satisfies a suitable *locality* principle. While our spectral representation does not directly yield locality of the observables, it allows us to relate the effect of applying the Laplacian to correlation functions at different spatial positions. Even if *a priori* of a different kind than locality, this remarkable feature provides the additional structure needed to carry out our analysis.

The subsections below formally describe these ingredients so that they can be used as black boxes in Part B, and proved in later parts.

4.2 Ingredient 1: Rotation invariance

The following result serves as a key external input for the paper. It states that k -point correlations functions are invariant under rotations in the limit as δ tends to zero.

Theorem 4.1 (Rotation invariance of k -point correlations [Ave+a]). *Fix $\sqrt{3} \leq \mathbf{c} \leq 2$. Then, for any $k \in 2\mathbb{Z}_{\geq 1}$ and any compact set $K \subset \mathcal{D}_k$,*

$$\lim_{\delta \rightarrow 0} \sup_{\mathbf{u} \in K} \sup_I |\Phi_k^{(\delta)}(\mathbf{u}) - \Phi_k^{(\delta)}(I\mathbf{u})| = 0, \quad (24)$$

where the second supremum is taken over all isometries $I : \mathbb{R}^2 \rightarrow \mathbb{R}^2$.

This is the origin of the restriction $\mathbf{c} \geq \sqrt{3}$. Indeed, [Dum+20] proves asymptotic rotational invariance for the critical random-cluster model with cluster weight $q \in [1, 4]$. Through the Baxter–Kelland–Wu (BKW) correspondence [BKW76], this range of q matches the regime $\mathbf{c} \in [\sqrt{3}, 2]$ of the six-vertex model. In [Ave+a], this correspondence is used to transfer the asymptotic rotational invariance from the critical random-cluster model to the six-vertex height function.

We expect Theorem 4.1 to hold for all $\mathbf{c} \in (0, 2]$. For $\mathbf{c} \in [1, 2]$, we believe that there may exist a proof following [Dum+20], but working directly with the six-vertex model rather than its FK-percolation representation. The arguments of the present paper would extend verbatim to $\mathbf{c} \in [1, 2]$ if Theorem 4.1 were available in that parameter range.

For $\mathbf{c} < 1$, although convergence to the GFF is still expected, the six-vertex model lacks positive association, and several steps of the proof (related to Ingredient 3) fail (in their present form) without this positive association.

4.3 Ingredient 2: A glimpse of scale invariance

It may be natural to expect that, in addition to rotational invariance, scale invariance is an important ingredient in identifying the scaling limit. While having such a property would simplify considerably our argument, it seems currently out of reach of direct techniques. Still, a glimpse of scale invariance is provided by the fact that the *free energy* or *surface tension* of the six-vertex model with a slope is twice differentiable at zero slope. Indeed, its second derivative will be identified as the limit of a certain quantity as the scale δ tends to zero. The convergence of said quantity will act as our indicator of scale invariance.

Below, we make the previous claim explicit. Let us start by recalling the definition of the free energy.

Definition 4.2 (Free energy). Fix $\mathbf{c} > 0$. Define the *free energy at slope* $s \in [-1, 1]$ via

$$f(s) = \lim_{L \rightarrow \infty} \lim_{M \rightarrow \infty} \frac{1}{ML} \log \left(Z_{\mathbb{T}_{M,L}} \mathbb{P}_{\mathbb{T}_{M,L}} \left[\left\{ \frac{I}{ML} = \frac{\lfloor Ls/2 \rfloor}{L/2} \right\} \right] \right), \quad (25)$$

where I is the number of left arrows minus the number of right arrows on any given vertical column of horizontal edges.

While the height function may not be defined for unbalanced six-vertex configurations on the torus, I/ML should be interpreted as its average slope in the vertical direction. Note that $f : [-1, 1] \rightarrow \mathbb{R}$ is an even function thanks to the symmetry by flipping all arrows.

Twice differentiability of the free energy was proved in [Dum+22] for the six-vertex model with $\mathbf{c} \in (0, 2]$ using Bethe Ansatz techniques. It was used in [Dum+24] to prove the delocalisation of the zero-slope six-vertex model for $\mathbf{c} \in [1, 2]$. By further harnessing [Dum+22], we explicitly compute the second derivative of f at 0 – see Section 24 for $\mathbf{c} \in [1, 2]$.

We use this result as a starting point for a two-step analysis, summarised in the following two results. The first step shows that the second derivative is indeed related to scale invariance of our actual six-vertex model: it determines the amplitude of the GFF limit (assuming such a limit exists). In the second step, we turn the characterisation of [Dum+22] into an explicit computation.

Theorem 4.3 (GFF-LDP correspondence). Fix $\mathbf{c} \in [1, 2]$. Assume that the six-vertex model has a sub-sequential scaling limit of the form $\sigma \cdot \text{GFF}$. Then,

$$\sigma^2 = -\frac{1}{f''(0)}. \quad (26)$$

Theorem 4.4 (Computation of $f''(0)$). Fix $\mathbf{c} \in [1, 2]$. Then, f is symmetric and twice differentiable at $s = 0$, with

$$f''(0) = -\frac{1}{2} \arccos \Delta = -\arcsin(\mathbf{c}/2), \quad (27)$$

where the relation between \mathbf{c} and Δ is given in Equation (5).

Jointly, the two theorems prove that any sub-sequential scaling limit that is a multiple of the GFF must have the explicitly computed variance.

Let us briefly comment on Theorem 4.3. The formula $\sigma^2 = -1/f''(0)$ would boil down to a “back of the envelope” calculation if the topology of the sub-sequential convergence towards σ GFF were strong enough to include convergence of probabilities of large deviation events. Unfortunately, the topology is not compatible with events whose probability decays exponentially fast to zero, and we must therefore obtain the formula by different means. We shall derive the formula by essentially expressing the probability of a large deviation event as the product of many probabilities of GFF events that *are* compatible with the topology of the sub-sequential convergence.

4.4 Ingredient 3: Regularity estimates and qualitative behaviour

The recently developed Russo–Seymour–Welsh (RSW) theory for the six-vertex model implies a *circuit estimate* for a suitable percolation representation of the model. This

RSW theory was first developed in [Dum+24] using the Bethe Ansatz. The representation employed here first appeared in [Lis21; Lis22]. The corresponding circuit estimate follows from [Dum+24] and was later obtained independently in [GL25] through a different approach not relying on the Bethe Ansatz. We will later return to these aspects in more detail.

For the purpose of this introduction, we adopt the following principle: this paragraph records the consequences of the representation and the circuit estimate without describing the representation of the circuit estimate explicitly, so that they may be treated as black boxes in Part B. All statements below will be proved in Part D.

First, define the *scale separation functions* for any $\{a, a'\}, \{b, b'\} \subset \mathbb{R}^2$:

$$S_{\mathbb{R}^2}(\{a, a'\}, \{b, b'\}) := \log \frac{\text{dist}(\{a, a'\}, \{b, b'\})}{\min\{|a' - a|, |b' - b|\}}; \quad (28)$$

$$S'_{\mathbb{R}^2}(\{a, a'\}, \{b, b'\}) := \log \frac{1 \vee \text{dist}(\{a, a'\}, \{b, b'\})}{\min\{|a' - a|, |b' - b|\}}. \quad (29)$$

Recall that $|\cdot|$ denotes Euclidean distance in these formulas; dist denotes the Euclidean distance between the two sets (that is, the minimum distance between any point in the first set and any point in the second set). The first function is truly invariant under scaling; the second function is more adapted to the discrete setting as it allows the pairs of points to overlap.

The following estimate bears a resemblance to Equation (14) and forms the basis of our qualitative analysis of correlation functions.

Theorem 4.5 (Regularity estimate). *For any $k \in 2\mathbb{Z}_{\geq 1}$, there exist constants $\alpha_k > 0$ and $C_k < \infty$ such that for any $\mathbf{c} \in [1, 2]$ and $\mathbf{u} \in (\mathbb{Z}^2)^{2k}$,*

$$|\Phi_k(\mathbf{u})| \leq C_k \sum_{\pi} \prod_{ij \in \pi} \begin{cases} e^{-\alpha_k S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\})} & \text{if } S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\}) \geq 20k^2, \\ 1 \vee -S'_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\}) & \text{if } S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\}) < 20k^2, \end{cases} \quad (30)$$

where π runs over pairings of $\{1, \dots, k\}$.

The previous theorem has the following important corollary, which follows immediately from the additivity property in Equation (10).

Corollary 4.6 (Precompactness of correlation functions). *Fix $\mathbf{c} \in [1, 2]$ and $k \in 2\mathbb{Z}_{\geq 1}$. Let (δ_n) denote any sequence tending to zero. Then the following two are equivalent:*

- (i) $\Phi_k^{(\delta_n)}$ converges to Ψ_k pointwise on a countable dense subset of \mathcal{D}_k ;
- (ii) $\Phi_k^{(\delta_n)}$ converges to Ψ_k uniformly on compact subsets of \mathcal{D}_k .

In this case, we say simply that $\Phi_k^{(\delta_n)}$ converges to Ψ_k on \mathcal{D}_k .

Moreover, $\sup_{\delta} |\Phi_k^{(\delta)}|$ is finite on \mathcal{D}_k , and therefore the family $(\Phi_k^{(\delta)})_{\delta}$ is precompact in this topology. Finally, any sub-sequential limit is continuous on \mathcal{D}_k .

Proof. The proof follows by carefully manipulating the regularity estimate (Theorem 4.5), the definition of the scale separation functions (Equations (28) and (29)), and additivity (Equation (10)).

We first claim that, for any neighbourhood $\mathcal{N} \subset \mathcal{D}_k$ of some compact set $\mathcal{K} \subset \mathcal{D}_k$, we can find some $\delta_0 > 0$ such that

$$\sup_{\mathbf{v} \in \mathcal{K}, \delta \in (0, \delta_0)} |\Phi_k^{(\delta)}(\mathbf{v})| \leq \sup_{\mathbf{u} \in \mathcal{N}} B(\mathbf{u}) \quad (31)$$

where $B(\mathbf{u})$ is defined as

$$C_k \sum_{\pi} \prod_{ij \in \pi} \begin{cases} e^{-\alpha_k S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\})} & \text{if } S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\}) \geq 20k^2, \\ 1 \vee -S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\}) & \text{if } S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\}) < 20k^2. \end{cases} \quad (32)$$

This claim is convenient: it does not take into account the integer restriction in Theorem 4.5, and it does not involve the function $S'_{\mathbb{R}^2}$.

To prove the claim, we first want to find a $\delta_0 > 0$ such that:

- For any $\mathbf{v} \in \mathcal{K}$ and $\delta \in (0, \delta_0)$, the point $\lfloor \mathbf{v}/\delta \rfloor \in (\mathbb{Z}^2)^{2k}$ with all coordinates rounded down still lies in $(\delta^{-1})\mathcal{N} \subset \mathcal{D}_k$,
- For any $\delta \in (0, \delta_0)$ and $\mathbf{u} \in (\delta^{-1})\mathcal{N}$, Equations (28) and (29) coincide for any $i \neq j$.

The first item holds true by basic topological considerations. For the second item, it is easy to see that (28) and (29) coincide on $(\delta^{-1})\mathcal{N}$ for small enough δ by compactness of \mathcal{N} , simply because $\text{dist}(\{u_i/\delta, u'_i/\delta\}, \{u_j/\delta, u'_j/\delta\}) \geq 1$.

It is easy to derive the claim from the two above items. Notice that B inherits scale-invariance from $S_{\mathbb{R}^2}$, and therefore the left side of (31) is bounded by $\sup_{\mathbf{u} \in \mathcal{N}, \delta \in (0, \delta_0)} B(\mathbf{u}/\delta) = \sup_{\mathbf{u} \in \mathcal{N}} B(\mathbf{u})$. This proves the claim (Equation (31)). We will now derive the statements in the corollary from this claim.

First, it follows immediately that $\sup_{\delta} |\Phi_k^{(\delta)}|$ is bounded on \mathcal{D}_k . For the other two statements (locally uniform convergence towards a continuous function), it suffices to prove that for any point $\mathbf{u} \in \mathcal{D}_k$ and $\varepsilon > 0$, there exists a neighbourhood $\mathcal{N} \subset \mathcal{D}_k$ of \mathbf{u} and some $\delta_0 > 0$ such that

$$\sup_{\mathbf{v} \in \mathcal{N}, \delta \in (0, \delta_0)} |\Phi_k^{(\delta)}(\mathbf{u}) - \Phi_k^{(\delta)}(\mathbf{v})| \leq \varepsilon. \quad (33)$$

This difference can be written as a telescopic sum of $2k$ terms, where each term is of the form $\Phi_k^{(\delta)}(\mathbf{v}') - \Phi_k^{(\delta)}(\mathbf{v}'')$ for some \mathbf{v}' and \mathbf{v}'' that differ in only one entry. In that case, we can apply additivity (Equation (10)) to write this difference as a single correlation function $\Phi_k^{(\delta)}(\mathbf{w})$ where \mathbf{w} has the property that two points are very close (since \mathcal{N} is a tiny neighbourhood of \mathbf{u}) and the other points are fixed. Equation (31) then tells us that $\Phi_k^{(\delta)}(\mathbf{w})$ can be made as small as desired by shrinking \mathcal{N} . This implies the desired statement. \square

The following cylinder estimates are straightforward adaptations of Theorem 4.5. Below, $\Phi_{\text{Cyl}, 2}$ is the two-point correlation function on the cylinder $\mathbb{Z} \times (\mathbb{Z}/L\mathbb{Z})$, see Subsection 4.5 below for details.

Corollary 4.7 (Regularity estimate for the cylinder). *There exist constants $C, c \in (0, \infty)$ such that for every $\mathbf{c} \in [1, 2]$, $L \in 2\mathbb{Z}_{\geq 1}$ and $k \in \mathbb{Z}_{\geq 1}$,*

$$|\Phi_{\text{Cyl}, 2}((0, 0), (k, 0), (2k, 0), (3k, 0))| \leq C, \quad (34)$$

$$|\Phi_{\text{Cyl}, 2}((0, 0), (0, \ell), (k, 0), (k, \ell))| \leq C(\ell/k)^c \quad \text{for } 0 < \ell \leq \min\{8k, L/2\}. \quad (35)$$

A more subtle manifestation of the RSW theory takes the form of a mixing estimate. It is standard that RSW estimates induce polynomial mixing estimates between scales. Here, we state a non-optimised version in terms of the multi-point correlation functions.

Theorem 4.8 (Mixing estimate). *Fix $\mathbf{c} \in [1, 2]$, $k \in 2\mathbb{Z}_{\geq 1}$. In this theorem, we consider $\mathbf{u} = (\mathbf{u}^{(1)}, \mathbf{u}^{(2)}) \in \mathcal{D}_k$ with $\mathbf{u}^{(1)} = (u_1, u'_1, u_2, u'_2)$ and $\mathbf{u}^{(2)} = (u_3, \dots, u'_k)$; we consider each of the $2k$ points fixed, except for the first point u_1 which is variable. Then there exists some constant $C = C(u'_1, u_2, u'_2, \dots) < \infty$ such that*

$$\limsup_{u_1 \rightarrow u_2} \limsup_{\delta \rightarrow 0} \left| \Phi_k^{(\delta)}(\mathbf{u}) - \Phi_2^{(\delta)}(\mathbf{u}^{(1)}) \Phi_{k-2}^{(\delta)}(\mathbf{u}^{(2)}) \right| \leq C. \quad (36)$$

Next, we state two intermediate results which are useful in Part E, where we identify the variance of the limiting GFF. These intermediate results are stated in terms of (random) subsets of $F(\mathbb{Z}^2)$. We endow such subsets with nearest-neighbour connectivity: faces are neighbours if and only if they share an edge. We identify a path of such faces with the union of the line segments connecting the centres of the faces, so that we may view such paths as subsets of \mathbb{R}^2 .

Theorem 4.9 (Arm exponents). *Fix $\mathbf{c} \in [1, 2]$. There exists a constant $c > 0$ such that for any $k \in \mathbb{Z}_{\geq 0}$ and $R, r \in \mathbb{Z}_{\geq 4}$ satisfying $R \geq 2r$,*

$$\mathbb{P}_{\mathbb{Z}^2} \left[\left\{ \exists a \in 2\mathbb{Z} \text{ such that the two sets of faces } \{h + a \leq 0\} \text{ and } \{h + a \geq k\} \text{ both contain paths from } [-r, r]^2 \text{ to } \partial[-R, R]^2 \right\} \right] \leq (r/R)^{ck^2}. \quad (37)$$

The last intermediate result follows directly from the representation (and not the RSW theory). Informally, *flip domination* says that if h is below some fixed $m \in 2\mathbb{Z}$ on a closed circuit of faces, then, on the faces surrounded by this circuit, $h - m$ is stochastically dominated by $m - h$. The formal statement is slightly more involved, owing to the gradient nature of the six-vertex height function.

Theorem 4.10 (Flip domination). *Fix $\mathbf{c} \in [1, 2]$. Consider the following setup:*

- γ is an arbitrary self-avoiding $F(\mathbb{Z}^2)$ -circuit,
- $F_\gamma \subset F(\mathbb{Z}^2)$ denotes the faces strictly surrounded by γ (not those visited by γ),
- $m \in 2\mathbb{Z}$ denotes the maximum of $2\lceil \frac{1}{2}h|_{F_\gamma} \rceil$, so that $h - m$ is gradient measurable,
- \mathcal{E} is any gradient event measurable in terms of $h|_{F(\mathbb{Z}^2) \setminus F_\gamma}$ with positive probability.

Then, for the measure $\mathbb{P}_{\mathbb{Z}^2}[\cdot | \mathcal{E}]$, the height function $(h - m)|_{F_\gamma}$ is stochastically dominated by $(m - h)|_{F_\gamma}$. More precisely, for any bounded increasing function X , we have

$$\mathbb{E}_{\mathbb{Z}^2} [X((h - m)|_{F_\gamma}) | \mathcal{E}] \leq \mathbb{E}_{\mathbb{Z}^2} [X((m - h)|_{F_\gamma}) | \mathcal{E}]. \quad (38)$$

4.5 Ingredient 4: Spectral representation of correlation functions

The purpose of the next few paragraphs is to state a spectral representation (which is a consequence of the transfer matrix formalism) as it is used in the core of the proof, without going into detail on how it is obtained. The spectral representation does not rely on the Yang–Baxter equations or the Bethe Ansatz. We start with a definition of the six-vertex model on the cylinder.

4.5.1 Six-vertex model on the cylinder

The transfer matrix can be used to derive identities in the six-vertex model on a cylinder. Recall that $\mathbb{P}_{\mathbb{T}_{M,L}}$ is the six-vertex measure on the torus $\mathbb{T}_{M,L}$. For L even, let Cyl_L denote the graph on the vertex set $\mathbb{Z} \times (\mathbb{Z}/L\mathbb{Z})$ with nearest-neighbour connectivity. The definition of a balanced six-vertex configuration extends to the cylinder: a six-vertex configuration is called *balanced* if, in each column of horizontal edges, there are exactly $L/2$ arrows pointing to the right and $L/2$ arrows pointing to the left. The event of balanced six-vertex configurations is denoted $\{\text{balanced}\}$ like before. The following lemma forms the starting point of our spectral representation (it is proved later on in Equation (195)).

Lemma 4.11 (Cylinder measure). *Fix $\mathbf{c} > 0$. The weak limit of $\mathbb{P}_{\mathbb{T}_{M,L}}[\cdot | \{\text{balanced}\}]$ as $M \rightarrow \infty$ exists and is denoted $\mathbb{P}_{\text{Cyl}_L}$.*

Recall that Theorem 2.2 implies that when $\mathbf{c} \in [1, 2]$, $\mathbb{P}_{\text{Cyl}_L}$ converges weakly to $\mathbb{P}_{\mathbb{Z}^2}$ as L tends to infinity.

Balanced six-vertex configurations ω on Cyl_L are in bijection with gradient height functions $h : F(\text{Cyl}_L) \rightarrow \mathbb{Z}$. To see that this is true, we remark that any oriented loop on the dual graph of Cyl_L intersects the same number of left- and right-pointing arrows (relative to the orientation of the loop), thanks to the ice rule and the balanced condition. We may extend the domain of such gradient height functions to $\mathbb{R} \times (\mathbb{R}/L\mathbb{Z})$ (via an analogue of Equation (7)) and to \mathbb{R}^2 by a simple lift.

Recall Definition 2.4. For any $\mathbf{u} = (u_1, u'_1, \dots, u_k, u'_k) \in (\mathbb{R}^2)^{2k}$, we shall write

$$\Phi_{\text{Cyl}_L, k}(\mathbf{u}) := \mathbb{E}_{\text{Cyl}_L} \left[\prod_{i=1}^k (h(u'_i) - h(u_i)) \right]. \quad (39)$$

4.5.2 Spectral representation of the two-point function

Let us introduce some more notation. For $\mathbf{u} = (u_1, u'_1, \dots, u_{2k}, u'_{2k}) \in (\mathbb{R}^2)^{2k}$, define x_i, y_i, x'_i , and y'_i such that

$$u'_i - u_i = (x_i, y_i) \quad \text{and} \quad u_{i+1} - u'_i = (x'_i, y'_i). \quad (40)$$

We say that a sequence $(u_1, u'_1, \dots, u_{2k}, u'_{2k})$ is *horizontally ordered* whenever $x_i, x'_i \geq 0$ for all i , and *horizontally strictly ordered* whenever $x_i \geq 0$ and $x'_i > 0$ for all i .

Theorem 4.12 (Spectral representation of the two-point function). *Fix $\mathbf{c} > 0$ and $L \in 2\mathbb{Z}_{>0}$. Then, there exists a finite positive measure μ_L on $\mathbb{R}_{>0} \times \mathbb{R}$ such that*

$$\Phi_{\text{Cyl}_L, 2}(\mathbf{u}) = \int \left((1-a)^{x_2} e^{-ib y_2} - 1 \right) (1-a)^{x'_1} e^{-ib y'_1} \left(1 - (1-a)^{x_1} e^{-ib y_1} \right) d\mu_L(a, b) \quad (41)$$

for any horizontally ordered $\mathbf{u} = (u_1, u'_1, u_2, u'_2) \in \mathbb{Z}^2$, and which is supported on the set $(0, 2] \times [-\pi, \pi]$ and invariant under the map $(a, b) \mapsto (a, -b)$. Furthermore, $\mu_L(\{|b| \in (0, 2\pi/L)\}) = 0$.

The measure μ_L is called *positive* to distinguish it from *signed* or *complex* measures. We shall later see that μ_L is a finite sum of Dirac masses induced by the spectrum of the transfer matrix of the six-vertex model. The measure is therefore referred to as a *spectral measure*. The heart of the proof of our main theorem will be to study the full-plane two-point function Φ_2 via an appropriate limit (as L tends to infinity) of the measures μ_L .

4.5.3 Spectral representation of general observables

We shall derive an expression for general observables similar to (41), but less explicit. At the heart of this expression is a crucial symmetry under *reflection* and the associated *reflection positivity*.

To describe it, let $\mathcal{R} := \{\frac{1}{2}\} \times \mathbb{R} \subset \mathbb{R}^2$ denote the reflection line. The definition is chosen such that \mathcal{R} traverses face centres (and not vertices). Let $\dagger : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ denote the reflection with respect to \mathcal{R} . Let $F^-(\text{Cyl}_L)$ denote the cylinder faces on the left of \mathcal{R} , and let $F^+(\text{Cyl}_L)$ denote the cylinder faces on the right of \mathcal{R} ; the faces whose centres lie on \mathcal{R} are included in both sets. The function \dagger is also interpreted as an involution on $F(\text{Cyl}_L)$, and acts on height functions via $h^\dagger := h \circ \dagger$.

Remark 4.13. It is important that \dagger is applied to the height function, not the arrows.

Let \mathfrak{A} denote the set of *local observables*, that is, real-valued random variables \mathbf{X} which are measurable in terms of the restriction of the gradient of h to *finitely many* faces (such observables are necessarily bounded). Let

$$\mathfrak{A}^\pm := \{\mathbf{X} \in \mathfrak{A} : \mathbf{X} \text{ is measurable in terms of } h|_{F^\pm(\text{Cyl}_L)}\}. \quad (42)$$

Notice that the reflection \dagger may be interpreted as a bijection from \mathfrak{A}^- to \mathfrak{A}^+ , via

$$\mathbf{X}^\dagger : h \mapsto \mathbf{X}(h^\dagger). \quad (43)$$

Finally, for any $v \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}$, define the (translation) map $\tau_v : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2$, $u \mapsto u + v$, which is extended to \mathfrak{A}^+ via

$$\tau_v : \mathfrak{A}^+ \rightarrow \mathfrak{A}^+, \mathbf{X} \mapsto (h \mapsto \mathbf{X}(h \circ \tau_v)). \quad (44)$$

Remark 4.14. To better grasp the definition, consider the following example. Let $(u_1, u'_1, \dots, u_k, u'_k) \subset (\mathbb{Z}_{\geq 0} + \frac{1}{2}) \times (\mathbb{Z} + \frac{1}{2})$ denote a family of centres of faces in $F^+(\text{Cyl}_L)$. Then:

- The observable \mathbf{X} defined via $\mathbf{X} : h \mapsto \prod_i (h(u'_i) - h(u_i))$ belongs to \mathfrak{A}^+ ,
- For any $v \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}$, we have $\tau_v(\mathbf{X}) \in \mathfrak{A}^+$ with $\tau_v(\mathbf{X})(h) = \prod_i (h(u'_i + v) - h(u_i + v))$,
- We have $\mathbf{X}^\dagger \in \mathfrak{A}^-$ with $\mathbf{X}^\dagger(h) = \prod_i (h((u'_i)^\dagger) - h(u_i^\dagger))$.

We now state the main theorem about the spectral representation of general observables.

Theorem 4.15 (Spectral representation of general observables). *Fix $\mathbf{c} > 0$. For any triple $(\mathbf{X}, \mathbf{Y}, L) \in \mathfrak{A}^- \times \mathfrak{A}^+ \times 2\mathbb{Z}_{\geq 1}$, there exists a finite complex-valued measure $\mu_{\mathbf{X}, \mathbf{Y}, L}$ supported on $[0, 2)$ such that:*

(i) For any $k \geq 0$,

$$\mathbb{E}_{\text{Cyl}_L}[\mathbf{X} \cdot \tau_{(k,0)}(\mathbf{Y})] = \int (1-a)^k d\mu_{\mathbf{X}, \mathbf{Y}, L}(a), \quad (45)$$

(ii) (**Cauchy–Schwarz inequality**) $\mu_{\mathbf{X}, \mathbf{X}^\dagger, L}$ and $\mu_{\mathbf{Y}^\dagger, \mathbf{Y}, L}$ are positive measures and

$$|\mathbb{E}_{\text{Cyl}_L}[\mathbf{X}\mathbf{Y}]|^2 \leq \|\mu_{\mathbf{X}, \mathbf{Y}, L}\|^2 \leq \|\mu_{\mathbf{X}, \mathbf{X}^\dagger, L}\| \cdot \|\mu_{\mathbf{Y}^\dagger, \mathbf{Y}, L}\| = \mathbb{E}_{\text{Cyl}_L}[\mathbf{X}\mathbf{X}^\dagger] \mathbb{E}_{\text{Cyl}_L}[\mathbf{Y}^\dagger \mathbf{Y}] \quad (46)$$

where $\|\cdot\|$ denotes the total variation metric.

Note that (45) holds true and is non-trivial for $k = 0$. The second property can be understood as a reflection positivity property, see e.g. [Bis09] and references therein.

5 Organisation of the paper

The paper is organised into four parts, according to Table 2. Recall from Section 4 that Part B contains the main proofs and depends on all of the four ingredients. The other parts are independent of one another, except that Ingredient 2 (Part E) relies on some ideas developed in Ingredient 3 (Part D). Ingredient 1 (the rotational invariance of Theorem 4.1) was proved in prior work, and no part is dedicated to it.

Part B

Proof of the main results

This part implements the main proofs as outlined in Section 4. The ingredients stated formally in the introduction are used as external inputs. The main results of Sections 6,

Part	Content	Techniques	External inputs
B	Main proofs (Theorems 2.8 and 3.3)	measure theory; complex analysis	Ingredients 1–4
C	Ingredient 4 Spectral representation	elementary; linear algebra	
D	Ingredient 3 Regularity estimates	percolation	[Dum+24] (or [GL25])
E	Ingredient 2 Glimpse of scale invariance	percolation; complex analysis	[Dum+22] and Ingredient 3

Table 2: Organisation of the paper: content, techniques, and dependencies of each part.

7, and 8 are Theorems 6.1, 7.1, and 8.1, respectively, and they correspond to Steps 1–3 outlined in Section 4. Sections 6–8 may be read independently of one another. Section 9 then combines these results to prove the main result in the isotropic case (Theorem 2.8). Section 10 presents the proof of the anisotropic case (Theorem 3.3). The value of $\mathbf{c} \in [\sqrt{3}, 2]$ does not play a role in the proofs and is omitted from notations.

6 Sub-sequential GFF limits for the two-point function

6.1 Statement of the theorem and outline of the proof

Section 6 is dedicated to proving the following result.

Theorem 6.1 (Dichotomy for the two-point function). *One of the following two properties holds true:*

- *There exists some $\sigma \in \mathbb{R}_{\geq 0}$ such that $\Phi_2^{(\delta)}$ converges to $\sigma^2 \Psi_2^{\text{GFF}}$ uniformly on every compact subset of \mathcal{D}_2 as δ tends to zero,*
- *There exist two distinct $\sigma, \sigma' \in \mathbb{R}_{\geq 0}$ and two sequences $(\delta_n)_n, (\delta'_n)_n$ tending to zero such that $\Phi_2^{(\delta_n)}$ and $\Phi_2^{(\delta'_n)}$ converge to $\sigma^2 \Psi_2^{\text{GFF}}$ and $(\sigma')^2 \Psi_2^{\text{GFF}}$ respectively, uniformly on every compact subset of \mathcal{D}_2 as n tends to infinity.*

Remark 6.2. The theorem does not assert that all sub-sequential limits are multiples of Ψ_2^{GFF} . We believe that its proof allows for the possibility of other sub-sequential scaling limits, such as, for example, $\sigma^2 \Psi_2^{\text{GFF}} + (\sigma')^2 \Psi'_2$, where Ψ' is the two-point correlation function of a *massive* GFF.

Let us sketch the proof. The first step (Subsection 6.2) consists in taking the limit as L tends to infinity and taking a scaling limit along a sub-sequence in the spectral representation formula from Theorem 4.12 for the two-point correlation function of the six-vertex model on the cylinder. By compactness arguments, this yields a limiting two-point function Ψ_2 and a positive measure μ on $\mathbb{R}_{>0} \times \mathbb{R}$, invariant under $(a, b) \mapsto (a, -b)$, such that for every $\mathbf{u} = (u_1, u'_1, u_2, u'_2) \in (\mathbb{R}^2)^4$ satisfying some simple geometric constraints,

$$\Psi_2(\mathbf{u}) = \int \left(e^{-ax_2 - iby_2} - 1 \right) e^{-ax'_1 - iby'_1} \left(1 - e^{-ax_1 - iby_1} \right) d\mu(a, b) \quad (47)$$

(recall Equation (40) for a definition of $(x_1, y_1, x'_1, y'_1, x_2, y_2)$ associated with such a \mathbf{u}). The measure μ fully encodes the two-point function Ψ_2 , and our aim is to analyse its structure.

Assuming we may differentiate under the integral, applying the Laplacian at *any* argument of Ψ_2 produces a factor $a^2 + (ib)^2$ in the integrand, which hints that harmonicity of Ψ_2 in each argument on \mathcal{D}_2 is equivalent to the concentration property $\mu[\{b^2 \neq a^2\}] = 0$. In fact, it is a simple exercise (see also the proof of Theorem 9.1) to check from the expression above that $\Psi_2 = \sigma^2 \cdot \Psi_2^{\text{GFF}}$ for some $\sigma > 0$ if and only if $\mu[\{b^2 \neq a^2\}] = 0$ and the density of the first marginal of μ is $\frac{\sigma^2}{2\pi a} da$.

We will not be able to establish these two properties of μ directly. The key ingredient we do exploit, however, is the rotational invariance of Ψ_2 (Theorem 4.1), which provides a collection of identities relating values of Ψ_2 at different points. These identities translate into constraints on the joint distribution of a and b under μ – see (78) and (87). By combining them with certain analyticity properties (of the function I_F defined in Subsection 6.3), we will derive in Subsection 6.4 that

$$\mu[\{a^2 < b^2\}] = 0. \quad (48)$$

One might hope that Theorem 4.1 would also yield the opposite bound $\mu[\{a^2 > b^2\}] = 0$, thereby implying $\mu[\{b^2 \neq a^2\}] = 0$. The density of the first marginal would then follow readily from the rotational invariance constraints, thus determining μ up to multiplicative constant. Unfortunately this is not the case: as pointed out in Remark 6.2, one may construct an entire family of scaling limits consistent with rotational invariance and all regularity assumptions, but for which $\mu[\{b^2 \neq a^2\}] \neq 0$.

Although (48) is insufficient to uniquely determine Ψ_2 , it is nonetheless enough to show (see Subsection 6.5) that the large-scale and small-scale behaviour of Ψ_2 is compatible with that of two-point correlation functions of the GFF. In particular, if Ψ_2 is not a multiple of Ψ_2^{GFF} , then its sub-sequential scaling limits at large and small scales must be distinct multiples of Ψ_2^{GFF} . This yields the dichotomy of Theorem 6.1.

6.2 Compactness of the spectral representation

Here and below, we keep the convention from Equation (40) for $(x_1, y_1, x'_1, y'_1, x_2, y_2)$ introduced in the introduction for every horizontally ordered \mathbf{u} . Recall from Theorem 4.12 that

$$\Phi_{\text{Cyl}_L, 2}(\mathbf{u}) = \int \chi_{\mathbf{u}}^{\text{discr}}(a, b) d\mu_L(a, b) \quad (49)$$

where

$$\chi_{\mathbf{u}}^{\text{discr}}(a, b) := \left((1-a)^{x_2} e^{-iby_2} - 1 \right) (1-a)^{x'_1} e^{-iby'_1} \left(1 - (1-a)^{x_1} e^{-iby_1} \right) \quad (50)$$

for any horizontally ordered $\mathbf{u} = (u_1, u'_1, u_2, u'_2) \in (\mathbb{Z}^2)^4$. This subsection proves that the family $(\mu_L)_L$ belongs to a compact space of measures, which allows us to derive a spectral representation for the sub-sequential scaling limit of the two-point function.

6.2.1 A compact space of measures

In this section, we introduce a compact space \mathcal{M} of measures satisfying some scale-invariant qualitative bounds, and then show that the cylinder measures μ_L belong to this space.

Definition 6.3 (The compact space $\mathcal{M}_{c,C}$). For $c, C \in (0, \infty)$, let $\mathcal{M}_{c,C}$ denote the set of positive measures ν on $\mathbb{R}_{>0} \times \mathbb{R}$, invariant under $(a, b) \mapsto (a, -b)$ and satisfying the following bounds:

- (i) $\nu[\{a \in (\alpha, 2\alpha)\}] \leq C$, for any $\alpha > 0$,

(ii) $\nu[\{(a, |b|) \in (0, \alpha) \times (\beta, 2\beta)\}] \leq C(\alpha/\beta)^c$, for any $\beta \geq \alpha > 0$.

We endow $\mathcal{M}_{c,C}$ with the *weak* (or *vague*) topology, that is, the topology making the map $\nu \mapsto \nu[f]$ continuous for any continuous function $f : \mathbb{R}_{>0} \times \mathbb{R} \rightarrow \mathbb{C}$ whose support is a compact subset of $\mathbb{R}_{>0} \times \mathbb{R}$.

We now show that the cylinder measures μ_L belong to $\mathcal{M}_{c,C}$ for suitable constants $c, C \in (0, \infty)$.

Lemma 6.4. *There exist constants $c, C \in (0, \infty)$ such that, for every L , the measure μ_L belongs to $\mathcal{M}_{c,C}$.*

From this point on, we fix such a pair $c, C \in (0, \infty)$ once and for all, and we write $\mathcal{M} := \mathcal{M}_{c,C}$ in what follows.

Proof. Reflection symmetry suggests that some $\mathbf{u} \in (\mathbb{Z}^2)^4$ are special: namely those \mathbf{u} where the first two points are chosen on the left of some vertical reflection line, and where the last two points are the reflections of the first two points. This leads to $\chi_{\mathbf{u}}^{\text{discr}}$ being real and of constant sign, which is crucial for the proof. We shall fix L throughout.

Step 1: Checking (i) in the definition of \mathcal{M} . For $k \in \mathbb{Z}_{\geq 0}$ and $\ell \in 2\mathbb{Z}_{\geq 0}$, set

$$\mathbf{u} = \begin{pmatrix} u_1 \\ u'_1 \\ u_2 \\ u'_2 \end{pmatrix} := \begin{pmatrix} 0 & 0 \\ k & 0 \\ k + \ell & 0 \\ 2k + \ell & 0 \end{pmatrix} \quad (51)$$

so that

$$\chi_{\mathbf{u}}^{\text{discr}}(a, b) = -(1-a)^\ell (1-(1-a)^k)^2 \leq 0. \quad (52)$$

We now split the argument in two cases depending on the value of α . First, consider the case $\alpha \leq 1/16$. Let $k = \ell \in [\frac{1}{8\alpha}, \frac{1}{4\alpha}] \cap 2\mathbb{Z}$. Then, for any $(a, b) \in [\alpha, 2\alpha] \times \mathbb{R}$, a short computation gives that $|\chi_{\mathbf{u}}^{\text{discr}}(a, b)| \geq 1/1000$. Also, Corollary 4.7 implies that $|\Phi_{\text{Cyl}_L, 2}(\mathbf{u})| \leq C$ uniformly in L and α . Combining the claims of the two last sentences and $\chi_{\mathbf{u}}^{\text{discr}}(a, b) < 0$ leads to

$$\mu_L[\{a \in [\alpha, 2\alpha]\}] \leq -1000 \int \chi_{\mathbf{u}}^{\text{discr}}(a, b) d\mu_L(a, b) = 1000 |\Phi_{\text{Cyl}_L, 2}(\mathbf{u})| \leq 1000C. \quad (53)$$

For the case $\alpha \geq 1/16$, simply set $k = 1$ and $\ell = 0$ and use a similar strategy to get

$$\mu_L[\{a \in [1/16, \infty)\}] \leq 256 |\Phi_{\text{Cyl}_L, 2}(\mathbf{u})| \leq 256. \quad (54)$$

On the right we used that height differences are bounded by 1, and thus $|\Phi_{\text{Cyl}_L, 2}(\mathbf{u})| \leq 1$.

Step 2: Checking (ii) in the definition of \mathcal{M} . We will prove the existence of constants $c, C > 0$ such that

$$\mu_L[\{(a, |b|) \in (0, \alpha) \times (\beta, 2\beta)\}] \leq C\left(\frac{\alpha}{\beta}\right)^c \quad \text{for } \alpha < \frac{1}{16} \text{ and } \beta = \frac{\pi}{2\ell} \quad (55)$$

with $1 \leq \ell \leq L/2$ integer. Condition (ii) may be deduced from (55), together with point (i) above and the fact that μ_L is supported on $\{(a, b) : 0 < a \leq 2 \text{ and } 2\pi/L \leq |b| \leq \pi\}$ (see Theorem 4.12), by simple algebraic manipulations.

Fix values of α and β as in (55) and choose some even integer $k \in [\frac{1}{8\alpha}, \frac{1}{4\alpha}]$. Set

$$\mathbf{u} = \begin{pmatrix} u_1 \\ u'_1 \\ u_2 \\ u'_2 \end{pmatrix} := \begin{pmatrix} 0 & 0 \\ 0 & \ell \\ k & 0 \\ k & \ell \end{pmatrix} \quad (56)$$

so that $x_1 = x_2 = 0$ and $x'_1 = k$ and

$$\chi_{\mathbf{u}}^{\text{discr}}(a, b) = 2(1 - a)^k(1 - \cos \ell b) \geq 0. \quad (57)$$

Then, for any $(a, b) \in [0, \alpha] \times [\beta, 2\beta]$, a short computation gives that $\chi_{\mathbf{u}}^{\text{discr}}(a, b) \geq 1$. Also, Corollary 4.7 implies that $|\Phi_{\text{Cyl}, 2}(\mathbf{u})| \leq C'(\alpha/\beta)^{c'}$ for some constants $C', c' \in (0, \infty)$ that are uniform in L, α , and β . The conclusion follows by the same argument as in point (i). \square

We now start manipulating the measures in \mathcal{M} . By construction, these measures have the property that

$$\sup_{\nu \in \mathcal{M}} \int (a \wedge \frac{1}{a}) d\nu(a, b) < \infty. \quad (58)$$

This means that the function $a \wedge \frac{1}{a}$ is a good domination function for applying the dominated convergence theorem. In the following lemma, we collect a few more properties of \mathcal{M} . We leave it as a straightforward exercise to the reader.

Lemma 6.5 (Properties of \mathcal{M}). *We have the following properties:*

- (i) \mathcal{M} is a compact topological space,
- (ii) Suppose that $f : \mathbb{R}_{>0} \times \mathbb{R} \rightarrow \mathbb{C}$ is a continuous function such that

$$\sup_{a, b} (a \vee \frac{1}{a}) |f(a, b)| < \infty, \quad (59)$$

then the function $\mathcal{M} \rightarrow \mathbb{C}, \nu \mapsto \nu[f]$ is continuous,

- (iii) Let $(\nu_n)_n \subset \mathcal{M}$ denote a sequence of measures, and let $(f_n)_n$ denote a sequence of continuous functions $f_n : \mathbb{R}_{>0} \times \mathbb{R} \rightarrow \mathbb{C}$ such that $\sup_{n, a, b} (a \vee \frac{1}{a}) |f_n(a, b)| < \infty$. If $(\nu_n)_n$ converges to ν and $(f_n)_n$ converges uniformly on every compact subset of $\mathbb{R}_{>0} \times \mathbb{R}$ to a function f , then

$$\lim_{n \rightarrow \infty} \nu_n[f_n] \rightarrow \nu[f]. \quad (60)$$

6.2.2 Passage to the full-plane limit of cylinder measures

In this section, we analyse the limit of the cylinder measures as L tends to infinity. Since the measures μ_L belong to \mathcal{M} , we may fix, once and for all, a sub-sequential limit $\mu_\infty \in \mathcal{M}$ of the sequence $(\mu_L)_L$. The full sequence $(\mu_L)_L$ may actually be shown to converge — for instance using (61) and the uniqueness of the $L \rightarrow \infty$ limit of the two-point function stemming from the existence of the full-plane measure — but this fact will not play a role in what follows. Therefore we simply work with a fixed sub-sequential limit.

The full-plane two-point correlation function may then be expressed in terms of μ_∞ as in the case of the cylinder.

Lemma 6.6. *For any horizontally ordered sequence $\mathbf{u} \subset \mathbb{Z}^2$ with $\{y_1, y_2\} \ni 0$, we have*

$$\Phi_2(\mathbf{u}) = \int \chi_{\mathbf{u}}^{\text{discr}}(a, b) d\mu_\infty(a, b). \quad (61)$$

Proof. Suppose that \mathbf{u} satisfies the assumptions of the lemma. Using the full-plane limit of the six-vertex model (Theorem 2.2), we obtain

$$\Phi_2(\mathbf{u}) = \lim_{L \rightarrow \infty} \Phi_{\text{Cyl}_{L,2}}(\mathbf{u}) \stackrel{(49)}{=} \lim_{L \rightarrow \infty} \mu_L[\chi_{\mathbf{u}}^{\text{discr}}]. \quad (62)$$

Since μ_L is supported on $(0, 2] \times \mathbb{R}$, we may insert the indicator $\mathbb{1}[a \leq 2]$ and write

$$\Phi_2(\mathbf{u}) = \lim_{L \rightarrow \infty} \mu_L[\mathbb{1}[a \leq 2] \cdot \chi_{\mathbf{u}}^{\text{discr}}]. \quad (63)$$

Because $y_1 = 0$ or $y_2 = 0$, we have $\mathbb{1}[a \leq 2] \cdot \chi_{\mathbf{u}}^{\text{discr}} = O(a \wedge \frac{1}{a})$ (see Equation (50)). Lemma 6.5 therefore implies that

$$\Phi_2(\mathbf{u}) = \mu_{\infty}[\mathbb{1}[a \leq 2] \cdot \chi_{\mathbf{u}}^{\text{discr}}]. \quad (64)$$

The proof is completed by discarding the indicator $\mathbb{1}[a \leq 2]$, which is justified by the fact that $\mu_{\infty}[|a| > 2] = 0$. This property is inherited from the corresponding bound for the measures μ_L , by the definition of \mathcal{M} and a second application of Lemma 6.5. \square




Figure 3: The spectral representation passes to the scaling limit when the points $\mathbf{u} \in (\mathbb{R}^2)^4$ satisfies the geometric constraints of the figure: the pairs of points cannot overlap “horizontally” (in the figure, $x_1 \geq 0$ and $x'_1 > 0$), and one pair must line on a horizontal line (in the figure, $y_2 = 0$).

6.2.3 Passage to the scaling limit

We now pass to the scaling limit by letting the mesh size of the lattice tend to zero. For any $\nu \in \mathcal{M}$ and $\delta > 0$, denote by $\nu^{(\delta)} \in \mathcal{M}$ the measure defined by

$$\nu^{(\delta)}(U) := \nu(\delta U) \quad (65)$$

for any measurable subset $U \subset \mathbb{R}_{>0} \times \mathbb{R}$. By construction, for every integrable function f ,

$$\int f(\delta a, \delta b) d\nu^{(\delta)}(a, b) = \int f(a, b) d\nu(a, b). \quad (66)$$

For $u = (x, y) \in \mathbb{R}^2$, set $u/\delta := (\lfloor x/\delta \rfloor, \lfloor y/\delta \rfloor) \in \mathbb{Z}^2$, and for any $\mathbf{u} = (u_1, u'_1, \dots, u_k, u'_k) \in (\mathbb{R}^2)^{2k}$ for some k , write $\mathbf{u}/\delta := (u_1/\delta, u'_1/\delta, \dots, u_k/\delta, u'_k/\delta) \in (\mathbb{Z}^2)^{2k}$.

The following notion will be used throughout.

Definition 6.7 (Convergence sequence). A *convergence sequence* is a sequence of positive reals $(\delta_n)_n$ tending to zero such that $\mu_{\infty}^{(\delta_n)}$ converges in the compact topological space \mathcal{M} . We denote by μ its limit. The dependence on $(\delta_n)_n$ will always be clear from context.

Convergence sequences exist by compactness of \mathcal{M} . Moreover, from any sequence of scales $(\delta'_n)_n$ tending to zero, one may extract a convergence sub-sequence $(\delta_n)_n$.

To state the next result, we introduce a variant χ of the function χ^{discr} from Equation (50): for any $\mathbf{u} = (u_1, u'_1, u_2, u'_2) \in \mathbb{R}^2$, define

$$\chi_{\mathbf{u}}(a, b) := (e^{-ax_2 - iby_2} - 1)e^{-ax'_1 - iby'_1}(1 - e^{-ax_1 - iby_1}). \quad (67)$$

Lemma 6.8 (Spectral representation of the sub-sequential scaling limit). *Let $(\delta_n)_n$ be a convergence sequence. Then, for any horizontally strictly ordered $\mathbf{u} = (u_1, u'_1, u_2, u'_2) \in (\mathbb{R}^2)^4$ satisfying $\{y_1, y_2\} \ni 0$*

$$\Psi_2(\mathbf{u}) := \lim_{n \rightarrow \infty} \Phi_2^{(\delta_n)}(\mathbf{u}) = \int \chi_{\mathbf{u}}(a, b) d\mu(a, b). \quad (68)$$

Remark 6.9. Compared to (61), we require above that the points be horizontally *strictly* ordered. This additional strictness is necessary because, under scaling, one must also control the behaviour of $\chi_{\mathbf{u}}$ for *large* values of a , and this geometric requirement is precisely what ensures that such behaviour is properly handled.

Proof. For notational simplicity we assume that the points \mathbf{u} belong to $(\mathbb{Z}^2)^{2k} \subset (\mathbb{R}^2)^{2k}$, that $1/\delta_n \in \mathbb{Z}$ for all n , and that $y_2 = 0$. Since \mathbf{u} is horizontally strictly ordered and due to the smoothness of $\mathbf{u} \mapsto \chi_{\mathbf{u}}(a, b)$, the bounds in the definition of \mathcal{M} and the Hölder nature of the correlators, these assumptions are harmless.

The proof would be easy if μ_{∞} was supported on $\{a \leq 1\}$ (meaning that all eigenvalues of the transfer matrix were non-negative). The negative eigenvalues, corresponding to $a \in (1, 2]$, make the proof slightly more technical.

For any n , split $\Phi_2^{(\delta_n)}(\mathbf{u}) = P_n + N_n$ depending on the contributions of positive and negative eigenvalues respectively:

$$P_n := \int_{\{a \leq 1\}} \chi_{\mathbf{u}/\delta_n}^{\text{discr}}(a, b) d\mu_{\infty}(a, b); \quad (69)$$

$$N_n := \int_{\{1 < a \leq 2\}} \chi_{\mathbf{u}/\delta_n}^{\text{discr}}(a, b) d\mu_{\infty}(a, b). \quad (70)$$

To conclude, it suffices to show that P_n tends to $\int \chi_{\mathbf{u}}(a, b) d\mu(a, b)$ (as $n \rightarrow \infty$), and N_n to zero.

Step 1: Limit of N_n . Recall from the definition of \mathcal{M} that $\mu_{\infty}[\{1 < a \leq 2\}] < \infty$. By writing out $\chi_{\mathbf{u}/\delta_n}^{\text{discr}}$ explicitly, it is easy to see that

$$|\chi_{\mathbf{u}/\delta_n}^{\text{discr}}(a, b)| \leq 4(1 - a)^{x'_1/\delta_n}. \quad (71)$$

For any fixed $a \in (1, 2)$, the bound on the right converges to zero as n tends to infinity. Thus, using the dominated convergence theorem, the only contribution to the limit can come from an atom at $a = 2$. However, $\mu_{\infty}[\{a = 2\}] = 0$. Indeed, writing $e_1 = (1, 0)$, one has, as k tends to infinity over the odd integers,

$$\Phi_2(0, e_1, (k+1)e_1, (k+2)e_1) = - \int a^2(1 - a)^k d\mu_{\infty}(a, b) \rightarrow 4\mu_{\infty}[\{a = 2\}]. \quad (72)$$

But we know that the correlation functions on the left tend to zero, thanks to our regularity estimate (Theorem 4.5). This concludes Step 1.

Step 2: Limit of P_n . We may write

$$P_n = \int_{\{a \leq 1/\delta_n\}} \chi_{\mathbf{u}/\delta_n}^{\text{discr}}(\delta_n a, \delta_n b) d\mu_{\infty}^{(\delta_n)}(a, b) = \int f_n(a, b) d\mu_{\infty}^{(\delta_n)}(a, b), \quad (73)$$

where $f_n(a, b) := \mathbb{1}[a \leq 1/\delta_n] \cdot \chi_{\mathbf{u}/\delta_n}^{\text{discr}}(\delta_n a, \delta_n b)$. By working out an explicit expression for f_n , it is straightforward to see that f_n converges on every compact subset of $\mathbb{R}_{>0} \times \mathbb{R}$ to $\chi_{\mathbf{u}}$ as n tends to infinity. Also, the fact that $\{y_1, y_2\} \ni 0$ implies that $\sup_{n,a,b} (a \vee \frac{1}{a}) |f_n(a, b)| < \infty$. Since $\mu_{\infty}^{(\delta_n)}$ tends to μ (by hypothesis), the convergence of P_n now follows from Lemma 6.5. \square

6.3 Consequences of rotational invariance

Let us briefly recap what we did so far. We introduced the compact set of measures \mathcal{M} described in Definition 6.3 and studied some of its properties in Lemma 6.5. This enabled us to prove that any sequence of mesh sizes tending to zero has a ‘‘convergence sub-sequence’’ $(\delta_n)_n$ along which, for some $\mu \in \mathcal{M}$,

$$\Psi_2(\mathbf{u}) := \lim_{n \rightarrow \infty} \Phi_2^{(\delta_n)}(\mathbf{u}) = \int \chi_{\mathbf{u}}(a, b) d\mu(a, b) \quad (74)$$

for any horizontally strictly ordered $\mathbf{u} = (u_1, u'_1, u_2, u'_2) \in \mathcal{D}_2$ with $\{y_1, y_2\} \ni 0$ (see Equation (67), Lemma 6.8, and Figure 3).

In this section we study the properties of the sub-sequential scaling limits μ of the spectral measures μ_{∞} , in particular the consequences of the rotational invariance of Theorem 4.1. These are specifically manifested in (78), which will eventually lead to Theorem 6.1. Rotational invariance will be used again in the proof of Theorem 7.1, but we consider (78) to be its main consequence.

Set $\mathbb{C}_+ := \{z \in \mathbb{C} : \text{Re}(z) > 0\}$.

Definition 6.10. To a convergence sequence $(\delta_n)_n$, associate the following functions:

$$F : \mathbb{C}_+ \times \mathbb{R} \rightarrow \mathbb{C}, (x, y) \mapsto \int a e^{-ax - iby} d\mu(a, b), \quad (75)$$

$$I_F : \mathbb{R}_{>0} \rightarrow \mathbb{R}, s \mapsto - \int_1^s F(x, 0) dx. \quad (76)$$

The bounds in the definition of \mathcal{M} ensure that the above functions are indeed well-defined. In the definition of I_F , the integral is taken with a sign, so that $I'_F(s) = -F(s, 0)$ for all $s > 0$.

We gather a few properties of F first.

Lemma 6.11. *For every convergence sequence $(\delta_n)_n$, the function F satisfies the following properties:*

(i) *For any horizontally strictly ordered $(u_1, u'_1, u_2, u'_2) \in (\mathbb{R}^2)^4$, we have*

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \Psi_2(u_1, u'_1, u_2, u_2 + (\varepsilon, 0)) = F(x_1 + x'_1, y_1 + y'_1) - F(x'_1, y'_1). \quad (77)$$

(ii) *The function F is continuous on $\mathbb{C}_+ \times \mathbb{R}$.*

(iii) *For any fixed $y \in \mathbb{R}$, the function $F(\cdot, y)$ is holomorphic on \mathbb{C}_+ ,*

(iv) *For any (x, y) , we have $|F(x, y)| \leq F(\text{Re}(x), 0)$, and $F(x, 0)$ is decreasing in $x \in \mathbb{R}_{>0}$,*

(v) *We have $\lim_{x \rightarrow \infty} F(x, 0) = 0$.*

Proof. All properties follow from (68), the bounds in the definition of \mathcal{M} and the dominated convergence theorem with the dominating function $a \wedge \frac{1}{a}$. \square

We now turn to properties of I_F . Recall that $|\cdot|$ is the Euclidean norm on \mathbb{R}^2 .

Lemma 6.12. *For every convergence sequence $(\delta_n)_n$, the functions F and I_F satisfy the following properties.*

- (i) *The function I_F is analytic.*
- (ii) *For every $\mathbf{u} = (u_1, u'_1, u_2, u'_2) \in \mathcal{D}_2$, the limit $\Psi_2(\mathbf{u}) := \lim_{n \rightarrow \infty} \Phi_2^{(\delta_n)}(\mathbf{u})$ is well-defined and equals*

$$\Psi_2(\mathbf{u}) = I_F(|u_2 - u_1|) + I_F(|u'_2 - u'_1|) - I_F(|u_2 - u'_1|) - I_F(|u'_2 - u_1|). \quad (78)$$

- (iii) *For any $(x, y) \in \mathbb{R}_{>0} \times \mathbb{R}$, we have*

$$F(x, y) = -\partial_x(I_F(|(x, y)|)). \quad (79)$$

Remark 6.13. Note the similarity between (78) and the expression of the GFF correlation function (14). That $\Psi_2(\mathbf{u})$ may be decomposed as a sum of four terms depending on the pairwise differences of \mathbf{u} is a simple consequence of (68) – at least for \mathbf{u} to which (68) applies. That the terms only depend on the Euclidean norm of the differences is a crucial fact, which encodes the rotational invariance of the scaling limits (Theorem 4.1).

Proof. Analyticity of I_F follows from the properties of F . We turn to proofs of (ii) and (iii).

Define $\mathbb{R}_\pm^2 := \{(x, y) \in \mathbb{R}^2 : \pm x > 0\}$. Recall that Ψ_2 is well-defined and expressed in terms of μ for all horizontally strictly ordered \mathbf{u} with $y_1 = 0$ or $y_2 = 0$ – see Lemma 6.8. The proof has four steps. First, we prove that for fixed $(u_1, u'_1, u_2) \in (\mathbb{R}_-^2)^2 \times \mathbb{R}_+^2$, the function $\Psi_2(u_1, u'_1, u_2, \cdot)$ is well-defined and C^1 on \mathbb{R}_+^2 . Second, we prove Equation (79). Third, we prove Equation (78) on $\mathbb{R}_-^2 \times \mathbb{R}_-^2 \times \mathbb{R}_+^2 \times \mathbb{R}_+^2$ when $y_2 = 0$. Finally, we extend Equation (78) to all of \mathcal{D}_2 .

Step 1: Definition of $\Psi_2(u_1, u'_1, u_2, \cdot)$. Readers may help themselves with Figure 4. Fix $(u_1, u'_1) \in (\mathbb{R}_-^2)^2$. Let us first prove well-definedness of $\Psi_2(u_1, u'_1, \cdot, \cdot)$ on $(\mathbb{R}_+^2)^2$. Fix two points $u_2, u'_2 \in (\mathbb{R}_+^2)^2$. Let U denote an open set containing the line segment between u_2 and u'_2 , and such that its closure is a compact subset of \mathbb{R}_+^2 (see Figure 4). Let $\theta : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ denote the rotation by a sufficiently small angle so that $\theta u_1, \theta u'_1 \in \mathbb{R}_-^2$ and $\theta U \subset \mathbb{R}_+^2$. We may now find a path from u_2 to u'_2 consisting of finitely many straight line segments $S \subset U$, which each have the property that either S is horizontal or θS is horizontal (see Figure 4). Write S_- and S_+ for the starting point and endpoint of S , respectively. By additivity of $\Phi_2^{(\delta_n)}$, we get

$$\Phi_2^{(\delta_n)}(\mathbf{u}) = \sum_S \Phi_2\left(\frac{u_1}{\delta_n}, \frac{u'_1}{\delta_n}, \frac{S_-}{\delta_n}, \frac{S_+}{\delta_n}\right) \quad (80)$$

As $n \rightarrow \infty$, all terms in the finite sum converge: if S is horizontal, then the term converges by Lemma 6.8; if θS is horizontal, then the term converges by applying Lemma 6.8 to the rotated system and using Theorem 4.1.

We now prove that $\Psi_2(u_1, u'_1, u_2, \cdot)$ is C^1 on U . By Lemma 6.11(i), the horizontal partial derivative exists and is continuous (we also use the additivity property of Ψ_2). By applying the same reasoning to the rotated system, we find that the partial derivative in the direction $\theta^{-1}e_1$ (with $e_1 = (1, 0)$) also exists and is continuous. Since the vectors $\{e_1, \theta^{-1}e_1\}$ span the tangent space of \mathbb{R}^2 , we conclude that $\Psi_2(u_1, u'_1, u_2, \cdot)$ is C^1 on U .

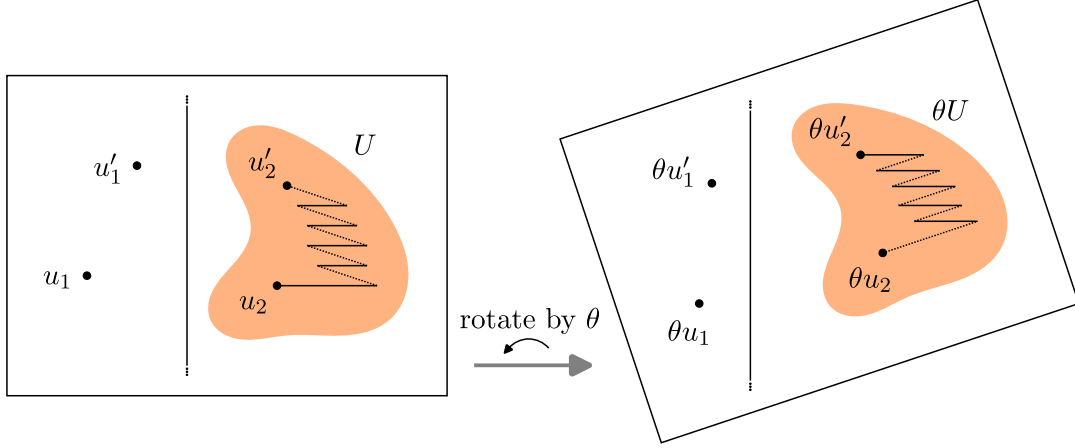


Figure 4: In Step 1 in the proof of Lemma 6.12, we prove that $\Psi_2(u_1, u'_1, \cdot, \cdot)$ is well-defined and C^1 on $(\mathbb{R}_+^2)^2$ by using Lemma 6.11 and rotation invariance of Ψ_2 .

Step 2: Proof of Equation (79). Readers may help themselves with Figure 5. Set $u'_1 = (0, 0)$, $u_2 \in \mathbb{R}_+^2$, and $u_1 = -(n-1) \cdot u_2$ for some large $n > 1$. As before, let $(x'_1, y'_1) = u_2 - u'_1$. Recall Equation (77) from Lemma 6.11. There are two observations to make.

First, the dominated convergence theorem implies that $\lim_{n \rightarrow \infty} F(nx'_1, ny'_1) = 0$, and therefore

$$\lim_{n \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \Psi_2(u_1, u'_1, u_2, u_2 + (\varepsilon, 0)) = \lim_{n \rightarrow \infty} F(nx'_1, ny'_1) - F(x'_1, y'_1) = -F(x'_1, y'_1). \quad (81)$$

Second, we proved above that $u'_2 \mapsto \Psi_2(u_1, u'_1, u_2, u'_2)$ is totally differentiable on \mathbb{R}_+^2 . Furthermore, by Theorem 4.1, it is invariant under the reflection orthogonal to the line containing u_1, u'_1 , and u_2 . We conclude that the partial derivative at u_2 orthogonal to this line is zero. It is therefore natural to decompose the partial derivative we analysed above into two: one along the line and one orthogonal to the line. The derivative along the line can be rewritten in terms of the function $F(\cdot, 0)$, using the rotation invariance of Theorem 4.1. By doing so, we get

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \Psi_2(u_1, u'_1, u_2, u_2 + (\varepsilon, 0)) = \left(\partial_{x'_1} |(|x'_1, y'_1|)| \right) (F(n \cdot |u_2|, 0) - F(|u_2|, 0)). \quad (82)$$

By sending n to infinity, the first term disappears by dominated convergence, so that

$$\lim_{n \rightarrow \infty} \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} \Psi_2(u_1, u'_1, u_2, u_2 + (\varepsilon, 0)) = - \left(\partial_{x'_1} |(|x'_1, y'_1|)| \right) F(|u_2|, 0) \quad (83)$$

$$= \partial_{x'_1} [I_F(|(|x'_1, y'_1|)|)], \quad (84)$$

where the second line follows from the definition of I_F .

Combining (81) and (84) yields $-F(x'_1, y'_1) = \partial_{x'_1} (I_F(|(|x'_1, y'_1|)|))$ as desired.

Step 3: Proof of Equation (78) on $\mathbb{R}_-^2 \times \mathbb{R}_-^2 \times \mathbb{R}_+^2 \times \mathbb{R}_+^2$ when $y_2 = 0$. Fix $(u_1, u'_1, u_2) \in (\mathbb{R}_-^2)^2 \times \mathbb{R}_+^2$. If $u'_2 = u_2$, then both sides of Equation (78) equal zero. To extend Equation (78) to every $u'_2 \in \mathbb{R}_+^2$ on the same horizontal line as u_2 , it suffices to prove that the horizontal partial derivatives of both sides with respect to u'_2 are equal. This follows from Equations (77) and (79).




Figure 5: Step 2 in the proof of Lemma 6.12

Step 4: Proof of Equation (78) to all of \mathcal{D}_2 . The left and right of Equation (78) are compatible with respect to the additivity property first introduced in Equation (10), and under applying a fixed isometry of the plane to all four points. One may therefore extend Equation (78) from $(\mathbb{R}_-^2)^2 \times (\mathbb{R}_+^2)^2$ to all of \mathcal{D}_2 . \square

6.4 Concentration on the sub-diagonal

For every convergence sequence, the function Ψ_2 is entirely determined by the function I_F introduced above. The problem therefore reduces to identifying I_F . If I_F happens to be a multiple of log, then Ψ_2 is a scaled version of Ψ_2^{GFF} , the two-point function of the GFF.

The remainder of this section is devoted to analysing I_F . The analysis consists of two steps. First, we show that for any convergence sequence, the corresponding measure μ is supported on $\{|b| \leq a\}$. In the second step, we use this information to derive further properties of I_F , which ultimately yield the dichotomy stated in Theorem 6.1. The first step is essential and constitutes a central component of the article. We establish it now.

Theorem 6.14. *For every convergence sub-sequence $(\delta_n)_n$, we have $\mu[\{|b| > a\}] = 0$.*

We start with two remarks motivating the proof. In the argument, we leverage these two perspectives and the definition of F in terms of μ to derive Theorem 6.14.

Remark 6.15. The fact that F is the horizontal partial derivative of a radially symmetric function imposes strong constraints on F .

Remark 6.16. Consider the definition of F and suppose for a second that $|b| = a$ almost everywhere. Then for any $(x, y) \in \mathbb{R}_+^2$, we get

$$F(x, y) = \int ae^{-ax-iby} d\mu(a, b) = \text{Re} \left(\int ae^{-ax-i|b|y} d\mu(a, b) \right) = \text{Re}(F(x + iy, 0)). \quad (85)$$

Although we will not really justify that $|b| = a$, the above perspective suggests that the expression $F(x + iy, 0)$ provides an interesting, alternative way to pass two-dimensional data to F .

Below, we shall use basic complex analysis tools such as holomorphic extensions and contour integrals. Issues such as singularities (poles) and branch cuts then become important. In this context, we shall write $\sqrt{\cdot}$ for the unique branch cut

$$\sqrt{\cdot} : (\mathbb{C} \setminus \mathbb{R}_{\leq 0}) \rightarrow \mathbb{C}_+ \quad (86)$$

which maps positive real numbers to positive real numbers. This means that we explicitly discard real nonpositive function values for the function $\sqrt{\cdot}$.

Lemma 6.17. *For any $(x, y) \in \mathbb{C}_+ \times \mathbb{R}$, we have*

$$F(x, y) = \frac{x}{\sqrt{x^2 + y^2}} F(\sqrt{x^2 + y^2}, 0). \quad (87)$$

Proof. If x is real, then this follows from Equation (79) in Lemma 6.12. For fixed y , both sides in Equation (87) are holomorphic functions in $x \in \mathbb{C}_+$, thanks to Lemma 6.11(iii). Since they coincide on the positive real axis, they must be equal for all $x \in \mathbb{C}_+$. \square

The following auxiliary lemma is a standard computation of a Fourier transform. We omit the proof.

Lemma 6.18. *For fixed $\sigma, \alpha \in \mathbb{R}_{>0}$, the Fourier transform of $\frac{\sigma^2}{t^2 + \sigma^2} \frac{\sin(\alpha t)}{t}$ satisfies*

$$\int_{\mathbb{R}} \frac{\sigma^2}{t^2 + \sigma^2} \frac{\sin(\alpha t)}{t} e^{-ibt} dt = \pi \int_{b-\alpha}^{b+\alpha} \frac{\sigma}{2} e^{-\sigma|y|} dy = \pi \cdot (\mathbb{1}[|\cdot| \leq \alpha] * \frac{\sigma}{2} e^{-\sigma|\cdot|})(b) \quad (88)$$

for any $b \in \mathbb{R}$.

Proof of Theorem 6.14. Since ae^{-a} is a strictly positive function on $\mathbb{R}_{>0}$, it suffices to prove that, for any $\lambda > 1$, we have

$$F(1, 0) = \int ae^{-a} d\mu(a, b) = \int \mathbb{1}[|b| \leq \lambda a] ae^{-a} d\mu(a, b). \quad (89)$$

Fix $\lambda > 1$. The idea is to mollify the indicator and apply a double Fourier transform, so that extra factors appear that are compatible with the general definition of F and Equation (87). First, for $\sigma > 0$, define J_σ as the integral on the right in Equation (89) but with a mollified indicator:

$$J_\sigma := \int \left((\mathbb{1}[|\cdot| \leq \lambda a] * \frac{\sigma}{2} e^{-\sigma|\cdot|})(b) \right) ae^{-a} d\mu(a, b). \quad (90)$$

By dominated convergence, J_σ converges to the integral on the right in Equation (89) as σ tends to infinity. Thus, to prove (89), it suffices to show that

$$J_\sigma \xrightarrow{\sigma \rightarrow \infty} F(1, 0). \quad (91)$$

Plugging in the formula of Lemma 6.18 above yields

$$J_\sigma = \int \left(\int_{\mathbb{R}} \frac{\sigma^2}{t^2 + \sigma^2} \frac{\sin(\lambda at)}{\pi t} e^{-ibt} dt \right) ae^{-a} d\mu(a, b). \quad (92)$$

Since the integrand is of order $O(ae^{-a}/(1+t^2))$, which is integrable in this product measure, we can apply Fubini's theorem to the product measure $dt \otimes \mu$. Rearranging yields

$$J_\sigma = \frac{1}{\pi} \int_{\mathbb{R}} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} \left(\int \sin(\lambda at) ae^{-a-ibt} d\mu(a, b) \right) dt. \quad (93)$$

Using that $\sin(\lambda at) = \frac{1}{2i}(e^{i\lambda at} - e^{-i\lambda at})$, the definition of F gives

$$2i \int \sin(\lambda at) ae^{-a-ibt} d\mu(a, b) = F(1 - i\lambda t, t) - F(1 + i\lambda t, t) =: F_-(t) - F_+(t) \quad (94)$$

Plugging this into our previous expression for J_σ enables to write

$$J_\sigma = \frac{1}{2\pi i} \int_{\mathbb{R}} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} (F_-(t) - F_+(t)) dt. \quad (95)$$

We are going to use complex analysis to study this integral, see the Figure 6 for the different poles and integration paths.

Let us make some remarks regarding Equation (95).




Figure 6: The paths γ_ε^+ and γ_ε^- run along \mathbb{R} and surround 0 using a small half-circle in the upper and lower half-plane, respectively. The function $F_- - F_+$ is analytic on the strip $|\text{Im}(z)| < \sigma$, hence J_σ may be written as its integral along γ_ε^- . When replacing the integral of F_- along γ_ε^- with that along γ_ε^+ , one should take into account the pole of F_- at 0.

- Equation (87) implies that

$$F_\pm(t) = \frac{1 \pm i\lambda t}{\sqrt{(1 \pm i\lambda t)^2 + t^2}} F(\sqrt{(1 \pm i\lambda t)^2 + t^2}, 0). \quad (96)$$

- The function F_- is *a priori* defined on \mathbb{R} . We may extend F_- and (96) as equal holomorphic functions on the open set $\mathcal{N}_- \supset \mathbb{R}$, where $\mathcal{N}_- \subset \mathbb{C}$ is the largest subset of \mathbb{C} such that the square root in the definition of F_- is well-defined. Indeed, as soon as the square root is well defined, $F(\sqrt{(1 \pm i\lambda t)^2 + t^2}, 0)$ is also defined, as the square root takes values in \mathbb{C}_+ . Similar considerations apply to F_+ , which extends to a holomorphic function on the similarly defined \mathcal{N}_+ .
- The integrand in Equation (95) is a holomorphic function on $(\mathcal{N}_- \cap \mathcal{N}_+) \setminus \{\pm\sigma i\}$. In particular, there is no simple pole at $t = 0$, because the factor $1/t$ is cancelled by the vanishing difference of holomorphic functions $F_-(t) - F_+(t)$.

Let γ_ε^- and γ_ε^+ denote the following paths:

- γ_ε^- first runs from $-\infty$ to $-\varepsilon$ along \mathbb{R} , then follows a half-circle of radius ε in the lower half-plane from $-\varepsilon$ to ε , and finally runs from ε to $+\infty$ along \mathbb{R} ,
- γ_ε^+ is defined similarly, except that the half-circle runs in the upper half-plane.

If $\varepsilon \in (0, \sigma)$ is so small that a ball of radius 2ε around 0 is contained in $\mathcal{N}_- \cap \mathcal{N}_+$, then the Cauchy integral theorem yields

$$J_\sigma = \frac{1}{2\pi i} \int_{\gamma_\varepsilon^-} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} (F_-(t) - F_+(t)) dt. \quad (97)$$

We now argue that the functions F_\pm are uniformly bounded on γ_ε^- . First, since they are continuous, it suffices to bound these functions on \mathbb{R} . But that is easy, since for real t , we get $|F_\pm(t)| = |F(1 \pm i\lambda t, t)| \leq F(1, 0) < \infty$ (see the bound in Lemma 6.11(iv)). We may now split the integral in Equation (97), which yields

$$J_\sigma = \frac{1}{2\pi i} \left(\int_{\gamma_\varepsilon^-} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_-(t) dt - \int_{\gamma_\varepsilon^-} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_+(t) dt \right). \quad (98)$$

For the integral on the left, we would like to replace the path γ_ε^- by γ_ε^+ . The integrand is holomorphic in the region enclosed by the two paths, except for a simple pole at $t = 0$. By the residue theorem, we get

$$\int_{\gamma_\varepsilon^-} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_-(t) dt = 2\pi i F_-(0) + \int_{\gamma_\varepsilon^+} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_-(t) dt. \quad (99)$$

Since $F_-(0) = F(1, 0)$,

$$J_\sigma = F(1, 0) + \frac{1}{2\pi i} \left(\int_{\gamma_\varepsilon^+} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_-(t) dt - \int_{\gamma_\varepsilon^-} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_+(t) dt \right) \quad (100)$$

$$= F(1, 0) + \frac{1}{\pi} \operatorname{Im} \left(\int_{\gamma_\varepsilon^+} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_-(t) dt \right), \quad (101)$$

since the two integrals in the first line are equal up to complex conjugation. We will now prove that

$$\lim_{\sigma \rightarrow \infty} \int_{\gamma_\varepsilon^+} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_-(t) dt = 0. \quad (102)$$

Recall that F_- is a holomorphic function on $\mathcal{N}_- \supset \mathbb{R}$, and that F_- is uniformly bounded on \mathbb{R} by $F(1, 0)$. The following claim asserts an even better control on the function F_- .

Claim. *All of the following hold true:*

- (i) *The set \mathcal{N}_- contains $\mathbb{H}^+ := \{z \in \mathbb{C} : \operatorname{Im}(z) > 0\}$,*
- (ii) *The function $|F_-|$ is uniformly bounded on \mathbb{H}^+ ,*
- (iii) *For any $s \in \mathbb{H}^+$, we have $\lim_{\sigma \rightarrow \infty} F_-(\sigma s) = 0$.*

We shall first see that the Claim implies Equation (102). By another application of the Cauchy integral theorem, and a change of variables $t = \sigma s$,

$$\int_{\gamma_\varepsilon^+} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_-(t) dt = \int_{\mathbb{R} + \frac{\sigma}{2}i} \frac{1}{t} \frac{\sigma^2}{t^2 + \sigma^2} F_-(t) dt = \int_{\mathbb{R} + \frac{1}{2}i} \frac{1}{s} \frac{1}{s^2 + 1} F_-(\sigma s) ds. \quad (103)$$

By the dominated convergence theorem, the right-hand side tends to zero as $\sigma \rightarrow \infty$, which implies Equation (102) and Equation (91) follows. Thus, the proof is complete, modulo the Claim which is proved below. \square

To conclude this section, we prove the claim invoked in the argument above. Notice that we have not used the property that $\lambda > 1$; we shall do so now (we use that $\lambda^2 - 1 > 0$).

Proof of the Claim. Introduce $\alpha : \mathbb{C} \rightarrow \mathbb{C}$, $t \mapsto (1 - i\lambda t)^2 + t^2$ so that

$$F_-(t) = \frac{1 - i\lambda t}{\sqrt{\alpha(t)}} F(\sqrt{\alpha(t)}, 0). \quad (104)$$

We now make some simple observations (see also Figure 7).

- For any $t \in \mathbb{C} \setminus \mathbb{R}_{<0}$, we have $\operatorname{Re}(\sqrt{t}) \geq 1$ if and only if t is on or to the right of the parabolic curve

$$\{(1 + it)^2 : t \in \mathbb{R}\} = \{1 - t^2 + 2it : t \in \mathbb{R}\} \subset \mathbb{C}. \quad (105)$$

- From the definition of $\alpha(t)$, it is immediate that $\alpha(t)$ lies on or to the right of this parabola for real t . Thus, we get $\operatorname{Re}(\sqrt{\alpha(t)}) \geq 1$ for any $t \in \mathbb{R}$.
- For $x \in \mathbb{R}$ and $y \in \mathbb{R}_{\geq 0}$, we may write out $\alpha(x + iy)$ explicitly:

$$\alpha(x + iy) = 1 + 2\lambda y + (\lambda^2 - 1)(y^2 - x^2) - 2ix(\lambda + (\lambda^2 - 1)y). \quad (106)$$

From this, it is immediate that:

$$\operatorname{Re}(\alpha(x + iy)) \geq \operatorname{Re}(\alpha(x)); \quad |\operatorname{Im}(\alpha(x + iy))| \geq |\operatorname{Im}(\alpha(x))|. \quad (107)$$

In particular, since $\alpha(x)$ lies on or to the right of the parabola, $\alpha(x + iy)$ also lies on or to the right of the parabola. We conclude that $\operatorname{Re}(\sqrt{\alpha(t)}) \geq 1$ for any $t \in \mathbb{H}^+$.




Figure 7

We have now proved $\mathcal{N}_- \supset \mathbb{H}^+$, i.e. Property (i) of the claim. Also, with the help of Lemma 6.11(iv) we get that, on \mathbb{H}^+ ,

$$|F(\sqrt{\alpha(t)}, 0)| \leq F(\operatorname{Re}(\sqrt{\alpha(t)}), 0) \leq F(1, 0). \quad (108)$$

This yields Property (ii) since the factor $(1 - i\lambda t)/\sqrt{\alpha(t)}$ is also bounded on \mathbb{H}^+ . It remains to establish Property (iii). Fix $s \in \mathbb{H}^+$ and set $t = \sigma s$. The factor $(1 - i\lambda t)/\sqrt{\alpha(t)}$ stays bounded like before; it suffices to show that $\operatorname{Re}(\alpha(\sigma s))$ converges to infinity as σ tends to infinity, so that Lemma 6.11(iv) implies that

$$|F(\sqrt{\alpha(\sigma s)}, 0)| \leq F(\operatorname{Re}(\sqrt{\alpha(\sigma s)}), 0) \rightarrow 0. \quad (109)$$

The explicit formula for $\alpha(\sigma s)$ above enables us to choose $\varepsilon = \varepsilon_s \in (0, \pi)$ such that, for any $R > 0$ and sufficiently large σ ,

$$\alpha(\sigma s) \in \{re^{i\theta} : \theta \in (-(\pi - \varepsilon), \pi - \varepsilon), r > R^2\}. \quad (110)$$

By the standard way of viewing the square root (half the angle, square root of the modulus), this implies that

$$\liminf_{\sigma \rightarrow \infty} \operatorname{Re}(\sqrt{\alpha(\sigma s)}) \geq \sin\left(\frac{\pi - \varepsilon}{2}\right) \cdot R. \quad (111)$$

Since R was arbitrary, the limit is infinite. \square

Remark 6.19 (The role of $\lambda > 1$). Since our main result implies that $\mu[\{a \neq b\}] = 0$, we may wonder how the argument fails when $\lambda < 1$. In that case, $\lambda^2 - 1 < 0$, and Equation (106) implies that $\alpha(t)$ has a zero in the upper half-plane at some point $t = iy$ with $y > 0$. The function F_- (Equation (104)) exhibits singular behaviour around this point, which breaks the above argument.

6.5 Derivation of the dichotomy

For any convergence sequence, we know that the associated analytic function I_F is decreasing, since $I'_F(x) = -F(x, 0) \leq 0$. If $I_F(s) \propto -\log s$, then $\Psi_2 \propto \Psi_2^{\text{GFF}}$ by Equation (78) of Lemma 6.12. To derive the dichotomy of Theorem 6.1, we must understand what happens when the analytic function I_F is not of this form. It turns out to be beneficial to encode the property that $I_F(s) \propto -\log s$ in terms of a new function Ξ , defined as

$$\Xi : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{\geq 0}, s \mapsto -sI'_F(s) = sF(s, 0). \quad (112)$$

Note that $I_F(s) \propto -\log s$ if and only if Ξ is constant. We now prove the following result on Ξ .

Theorem 6.20. For any convergence sequence $(\delta_n)_n$, the function Ξ is bounded and

$$-\Xi'(s) = s\Delta I_F(|u|)|_{u=se_1} = s \int (a^2 - b^2)e^{-as} d\mu(a, b) \geq 0. \quad (113)$$

In particular, Ξ is a non-increasing function.

Before diving into the proof, let us extend the applicability of Lemma 6.8.

Lemma 6.21 (Extension of Lemma 6.8). Let $(\delta_n)_n$ denote a convergence sequence. For every $\mathbf{u} = (u_1, u'_1, u_2, u'_2) \in (\mathbb{R}^2)^4$ whose first coordinates of u_1 and u'_1 are strictly smaller than those of u_2 and u'_2 ,

$$\Psi_2(\mathbf{u}) := \lim_{n \rightarrow \infty} \Phi_2^{(\delta_n)}(\mathbf{u}) = \int \chi_{\mathbf{u}}(a, b) d\mu(a, b). \quad (114)$$

Proof. Fix \mathbf{u} satisfying the assumptions of the lemma. Above, we proved already that $\Psi_2(\mathbf{u})$ exists, and that the function $\chi_{\mathbf{u}}$ is μ -integrable. Indeed, any term of the form $1 - e^{-ax-iby}$ is of order $O(a)$ thanks to Theorem 6.14. Thus, the left and right of Equation (114) are well-defined; it remains to prove equality.

Set $e_1 = (1, 0)$. We may use the additivity property to write the height difference $h(u'_2/\delta_n) - h(u_2/\delta_n)$ as the sum of height differences along the following path:

$$\frac{u_2}{\delta_n} \rightarrow \frac{u_2 + ne_1}{\delta_n} \rightarrow \frac{u'_2 + ne_1}{\delta_n} \rightarrow \frac{u'_2}{\delta_n}. \quad (115)$$

Set $\mathbf{s}^n := (u_1, u'_1, u_2, u_2 + ne_1)$ and $\mathbf{t}^n := (u_1, u'_1, u'_2 + ne_1, u'_2)$. We claim that

$$\Psi_2(\mathbf{u}) = \lim_{n \rightarrow \infty} (\Psi_2(\mathbf{s}^n) + \Psi_2(\mathbf{t}^n)). \quad (116)$$

For this to be true, we must justify that the correlation function of the middle step in the path does not appear. Yet, it tends to zero thanks to the regularity estimate (Theorem 4.5).

We may therefore apply (twice) the case where the two last points are on the same horizontal line to get

$$\Psi_2(\mathbf{u}) = \lim_{n \rightarrow \infty} \int (\chi_{\mathbf{s}^n} + \chi_{\mathbf{t}^n})(a, b) d\mu(a, b). \quad (117)$$

Finally, the dominated convergence theorem, where we use Theorem 6.14 to find the desired dominating function, enables us to insert the missing part of the function $\chi_{\mathbf{u}}$ to get

$$\Psi_2(\mathbf{u}) = \int \chi_{\mathbf{u}}(a, b) d\mu(a, b). \quad (118)$$

This concludes the proof. □

Proof of Theorem 6.20. To see that Ξ is bounded, recall from the definitions that

$$\Xi(s) = s \int ae^{-as} d\mu(a, b). \quad (119)$$

The bounds on $\mu \in \mathcal{M}$ thus imply that Ξ is bounded.

We now focus on proving (113). The equality on the right of (113) is Theorem 6.14. For the equality on the left, note that the radial symmetry of $I_F(|\cdot|)$ enables us to replace ∂_x by $s\partial_{yy}$, to get that

$$-\Xi'(s) = I'_F(s) + sI''_F(s) \quad (120)$$

$$= (\partial_x + s\partial_{xx})I_F(|(x, y)|)|_{(x, y)=se_1} \quad (121)$$

$$= s(\partial_{yy} + \partial_{xx})I_F(|(x, y)|)|_{(x, y)=se_1} \quad (122)$$

$$= s\Delta I_F(|u|)|_{u=se_1}. \quad (123)$$

To get (113), it only remains to justify that

$$\Delta I_F(|u|)|_{u=se_1} = \int (a^2 - b^2)e^{-as} d\mu(a, b). \quad (124)$$

Set $u_1 = -ne_1$, $u'_1 = (0, 0)$, $u_2 = se_1$, and $u'_2 = u_2 + (x, y)$ (with $|(x, y)|$ tiny). Using Lemma 6.21 and letting n tend to infinity gives

$$\lim_{n \rightarrow \infty} \Psi_2(u_1, u'_1, u_2, u'_2) = \int e^{-as}(e^{-ax-iby} - 1) d\mu(a, b). \quad (125)$$

Using Lemma 6.12 and letting n tend to infinity, we get another expression for the left-hand side of the previous equality:

$$\lim_{n \rightarrow \infty} \Psi_2(u_1, u'_1, u_2, u'_2) = I_F(|se_1 + (x, y)|) - I_F(|se_1|). \quad (126)$$

Equalling the two previous identities gives

$$I_F(|se_1 + (x, y)|) - I_F(|se_1|) = \int e^{-as}(e^{-ax-iby} - 1) d\mu(a, b). \quad (127)$$

Now, calculating the Laplacian in (x, y) at the point $(0, 0)$ yields

$$\Delta I_F(|u|)|_{u=se_1} = \int (a^2 - b^2)e^{-as} d\mu(a, b). \quad (128)$$

We used that we may differentiate under the integral thanks to the dominated convergence theorem, using that $a \wedge \frac{1}{a}$ is integrable, and that $|b| \leq a$ (Theorem 6.14). \square

Before concluding the proof of the dichotomy, we summarise what has been established so far. For any sequence $(\delta_n)_n$ tending to zero, we may extract a sub-sequence $(\delta'_n)_n$ and find some analytic function $I_F : \mathbb{R}_{>0} \rightarrow \mathbb{R}$ such that all of the following hold true:

- For any $\mathbf{u} \in \mathcal{D}_2$,

$$\Psi_2(\mathbf{u}) := \lim_{n \rightarrow \infty} \Phi_2(\mathbf{u}/\delta_n) = I_F(|u_2 - u_1|) + I_F(|u'_2 - u'_1|) - I_F(|u'_2 - u_1|) - I_F(|u_2 - u'_1|). \quad (129)$$

- The function $-sI'_F(s)$ is nonnegative, bounded, and non-increasing.

We are now ready to prove the dichotomy result.

Proof of Theorem 6.1. We start by observing that trivially at least one of the following three statements must hold:

- (i) Either there exists σ such that for every convergence sequence $(\delta_n)_n$, $I_F = -\frac{\sigma^2}{2\pi} \log$,
- (ii) Or there exist two different σ, σ' and two convergence sequences $(\delta_n)_n$ and $(\delta'_n)_n$ such that I_F is equal respectively to $-\frac{\sigma^2}{2\pi} \log$ and $-\frac{(\sigma')^2}{2\pi} \log$,

(iii) Or for some convergence sequence $(\delta_n)_n$, the function I_F is not proportional to $-\log$. The first two cases clearly enter into the framework of the dichotomy of Theorem 6.1. Indeed, Lemma 6.12(ii) then implies that $\Phi_2^{(\delta_n)}$ converges to its GFF counterpart. As a consequence, (i) corresponds to the first case of the dichotomy, and (ii) to the second. To conclude the proof, we only need prove that (iii) also implies the second case of the dichotomy. From now on, fix some convergence sequence $(\delta_n)_n$ such that $I_F \not\propto -\log$.

Using Theorem 6.20, define $\sigma, \sigma' \in \mathbb{R}_{\geq 0}$ such that

$$\frac{\sigma^2}{2\pi} = \lim_{s \rightarrow 0} -sI'_F(s) > \lim_{s \rightarrow \infty} -sI'_F(s) = \frac{(\sigma')^2}{2\pi}. \quad (130)$$

From the properties in the previous lemma we see that, for any $s \in \mathbb{R}$, we get

$$\lim_{\theta \rightarrow 0} I_F(\theta s) - I_F(\theta) = -\frac{\sigma^2}{2\pi} \log s; \quad \lim_{\Theta \rightarrow \infty} I_F(\Theta s) - I_F(\Theta) = -\frac{(\sigma')^2}{2\pi} \log s. \quad (131)$$

In particular, Lemma 6.12(ii) implies that for any $\mathbf{u} \in \mathcal{D}_2$,

$$\lim_{\theta \rightarrow 0} \lim_{n \rightarrow \infty} \Phi_2^{(\delta_n/\theta)}(\mathbf{u}) = \sigma^2 \Psi_2^{\text{GFF}}(\mathbf{u}); \quad \lim_{\Theta \rightarrow \infty} \lim_{n \rightarrow \infty} \Phi_2^{(\delta_n/\Theta)}(\mathbf{u}) = (\sigma')^2 \Psi_2^{\text{GFF}}(\mathbf{u}). \quad (132)$$

This means that we can take two “scaling limits of the sub-sequential scaling limit” : one obtained by “zooming in” and the other by “zooming out”. We present the zooming-in construction; the zooming-out case is analogous.

Fix a dense countable family $(\mathbf{u}_k)_k$ of quadruplets in \mathcal{D}_2 . Choose a sequence $(\theta_k)_k$ tending to zero such that, for every k ,

$$|\Psi_2(\theta_k \mathbf{u}_i) - \sigma^2 \Psi_2^{\text{GFF}}(\mathbf{u}_i)| \leq \frac{1}{k} \quad \forall i \leq k. \quad (133)$$

Next, choose an increasing sequence $(n_k)_k$ such that, for every k , $\delta_{n_k}/\theta_k \leq \frac{1}{k}$ and

$$|\Phi_2^{(\delta_{n_k})}(\theta_k \mathbf{u}_i) - \Psi_2(\theta_k \mathbf{u}_i)| \leq \frac{1}{k} \quad \forall i \leq k. \quad (134)$$

With this choice the sequence $(\delta_{n_k}/\theta_k)_k$ satisfies

$$|\Phi_2^{(\delta_{n_k}/\theta_k)}(\mathbf{u}_i) - \sigma^2 \Psi_2^{\text{GFF}}(\mathbf{u}_i)| = |\Phi_2^{(\delta_{n_k})}(\theta_k \mathbf{u}_i) - \sigma^2 \Psi_2^{\text{GFF}}(\mathbf{u}_i)| \leq \frac{2}{k} \quad \forall i \leq k. \quad (135)$$

Hence, for every i , the sequence $(\Phi_2^{(\delta_{n_k}/\theta_k)}(\mathbf{u}_i))$ converges pointwise as k tends to infinity to $\sigma^2 \Psi_2^{\text{GFF}}(\mathbf{u}_i)$. By Corollary 4.6, this means that $\Phi_2^{(\delta_{n_k}/\theta_k)}$ converges to $\sigma^2 \Psi_2^{\text{GFF}}$ uniformly on compact subsets of \mathcal{D}_2 . The same construction yields a sequence converging to $(\sigma')^2 \Psi_2^{\text{GFF}}$. \square

7 Sub-sequential GFF limits for multi-point functions

This section is dedicated to proving the following result.

Theorem 7.1. *Fix $\sigma \in \mathbb{R}_{\geq 0}$ and a sequence $(\delta_n)_n$ tending to zero. Consider the following statement for fixed k :*

$$\Phi_k^{(\delta_n)}|_{\mathcal{D}_k} \xrightarrow[n \rightarrow \infty]{} \sigma^k \Psi_k^{\text{GFF}} \quad \text{uniformly on compact subsets of } \mathcal{D}_k. \quad (136)$$

If this statement holds true for $k = 2$, then it holds true for all $k \in \mathbb{Z}_{\geq 1}$.

Recall Corollary 4.6, which guarantees precompactness. Hence, after passing to a sub-sequence, we may assume that the functions converge. Throughout this section, we therefore fix σ and $(\delta_n)_n$ such that:

- $\Phi_2^{(\delta_n)}|_{\mathcal{D}_2}$ converges to $\sigma^2\Psi_2^{\text{GFF}}$,
- $\Phi_k^{(\delta_n)}|_{\mathcal{D}_k}$ converges to some continuous function Ψ_k .

It suffices to prove that $\Psi_k = \sigma^k\Psi_k^{\text{GFF}}$ for every k .

The proof goes in two steps. First, we prove that the scaling limit Ψ_k of the $\Phi_k^{(\delta_n)}$ is harmonic in each coordinate, harvesting the concentration of the measure μ on $\{b = \pm a\}$. Once harmonicity has been established, we identify $\Psi_k = \sigma^k\Psi_k^{\text{GFF}}$ as the unique function that is harmonic and satisfies a few other properties (limit at infinity, behaviour when merging points) that Ψ_k possesses. The second step uses induction on k , with the base case $k = 2$ being our hypothesis.

7.1 Harmonicity

Proposition 7.2. *The function Ψ_k is harmonic in each coordinate.*

In order to prove this proposition, we introduce the set

$$\mathcal{H} := \{\mathbf{u} = (u_1, u'_1, \dots, u_k, u'_k) \in \mathcal{D}_k : \Psi_k \text{ is harmonic in } u_1 \text{ at } \mathbf{u}\}. \quad (137)$$

Since Ψ_k is naturally invariant or antisymmetric under various permutations of its arguments, it suffices to show that $\mathcal{H} = \mathcal{D}_k$. Observe also that, for harmonicity in u_1 , the position of u'_1 plays no role. We begin with the following lemma, which asserts that as long as u_1 does not lie in the convex hull of the points $u_2, u'_2, \dots, u_k, u'_k$, harmonicity basically follows from the arguments established so far.

Lemma 7.3. *We have*

$$\mathcal{Z}_0 := \{\mathbf{u} = (u_1, u'_1, \dots, u_k, u'_k) \in \mathcal{D}_k : u_1 \notin \text{ConvexHull}(\{u_2, u'_2, \dots\})\} \subset \mathcal{H}. \quad (138)$$

Proof. Fix $\mathbf{u} \in \mathcal{Z}_0$. Let $C_r \subset \mathbb{R}^2$ denote the circle of radius $r > 0$ centred at u_1 . It suffices to prove that if r is smaller than the distance from u_1 to the convex hull, then

$$\Upsilon := \int_{C_r} \Psi(u_1, s, u_2, u'_2, u_3, u'_3, \dots) ds = 0 \quad (139)$$

(the integral is just the uniform probability measure on C_r).

By applying an isometry of the plane, we may assume without loss of generality that C_r is strictly on the left of the vertical axis $\{0\} \times \mathbb{R}$, while the convex hull is strictly on the right of this axis.

To prove that $\Upsilon = 0$ we are going to use the spectral representation of general observables, the Cauchy–Schwarz inequality (Theorem 4.15), and harmonicity of $\sigma^2\Psi_2^{\text{GFF}}$. For fixed n , introduce the following observables belonging to \mathfrak{A}^- and \mathfrak{A}^+ (cf. Equation (42)) respectively:

$$\mathbf{X}_n(h) := \int_{C_r} (h(s/\delta_n) - h(u_1/\delta_n)) ds; \quad (140)$$

$$\mathbf{Y}_n(h) := \prod_{i=2}^k (h(u'_i/\delta_n) - h(u_i/\delta_n)). \quad (141)$$

By Theorem 4.15, we have

$$|\mathbb{E}_{\text{Cyl}_L}[\mathbf{X}_n \mathbf{Y}_n]| \leq \sqrt{\mathbb{E}_{\text{Cyl}_L}[\mathbf{X}_n \mathbf{X}_n^\dagger] \mathbb{E}_{\text{Cyl}_L}[\mathbf{Y}_n^\dagger \mathbf{Y}_n]}. \quad (142)$$

The left-hand side converges to $|\Upsilon|$ as we take L and then n to infinity. It suffices to prove that the right-hand side converges to 0 in the same double limit. Yet,

$$(\mathbf{X}_n \mathbf{X}_n^\dagger)(h) = \left(\int_{C_r} (h(s/\delta_n) - h(u_1/\delta_n)) \, ds \right) \left(\int_{C_r} (h^\dagger(s/\delta_n) - h^\dagger(u_1/\delta_n)) \, ds \right). \quad (143)$$

The $\mathbb{E}_{\text{Cyl}_L}$ -expectation of this random variable tends to zero in the double limit by uniform convergence and harmonicity of $\sigma^2 \Psi_2^{\text{GFF}}$.

It remains to demonstrate that the expectation of $\mathbf{Y}_n^\dagger \mathbf{Y}_n$ stays uniformly bounded in the double limit. But this is obvious since it tends to the bounded function

$$\Psi_{2k-2}(u_2, u'_2, \dots, u_k, u'_k, (u_2)^\dagger, (u'_2)^\dagger, \dots, (u_k)^\dagger, (u'_k)^\dagger) < \infty. \quad (144)$$

Here, we remind that $(u)^\dagger$ has been introduced in Subsection 4.5.3 and corresponds to the reflection of u over $\{\frac{1}{2}\} \times \mathbb{R}$. This concludes the proof. \square

We now extend harmonicity from \mathcal{Z}_0 to all of \mathcal{D}_k (see Figure 8 for an illustration). We start by a technical lemma.

Lemma 7.4. *Suppose that $\tilde{\mathbf{u}} \in \mathcal{D}_k$ may be obtained from $\mathbf{u} = (u_1, u'_1, \dots, u_k, u'_k) \in \mathcal{D}_k$ by choosing a line $L \subset \mathbb{R}^2$ such that u_i and u'_i lie strictly on the same side of L for each i , and then moving all points on one side of the line by a distance of $\lambda_0 \in \mathbb{R}_{\geq 0}$ perpendicularly away from L .*

Then, $\mathbf{u} \in \mathcal{H}$ whenever $\mathcal{N}_{\tilde{\mathbf{u}}} \subset \mathcal{H}$ for some neighbourhood $\mathcal{N}_{\tilde{\mathbf{u}}}$ of $\tilde{\mathbf{u}}$.

Proof. By applying an isometry of the plane and permuting the different $\{u_i, u'_i\}$, we can assume without loss of generality that $L = \{0\} \times \mathbb{R}$, that the pairs $\{u_i, u'_i\}$ lie on the left of L for $i \leq \ell$ and on the right of L for $i > \ell$, and that we move the points on the right of L to the right by $\lambda_0 \in \mathbb{R}_{\geq 0}$.

Set $e_1 = (1, 0)$. For $s \in \mathbb{R}^2$ and $\lambda \in \mathbb{R}_{\geq 0}$, define

$$\mathbf{u}^{-,s} := (u_1, s, u_2, u'_2, \dots, u'_\ell); \quad \mathbf{u}^{+,\lambda} := (u_{\ell+1} + \lambda e_1, u'_{\ell+1} + \lambda e_1, \dots, u'_k + \lambda e_1). \quad (145)$$

For any $r > 0$, we let C_r denote the circle of radius r around u_1 . We claim that the function

$$\Upsilon_r(\lambda) := \int_{C_r} \Psi_k(\mathbf{u}^{-,s}, \mathbf{u}^{+,\lambda}) \, ds \quad (146)$$

satisfies $\Upsilon_r(0) = 0$ for small enough r , which would establish the harmonicity and the lemma.

Fix r at least so small that the convex hull of C_r does not intersect L or $\cup_{i>1} \{u_i, u'_i\}$. Since $\Upsilon_r : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a continuous function and $\Upsilon_r(\lambda) = 0$ for $\lambda \approx \lambda_0$ and r small enough and $\mathcal{N}_{\tilde{\mathbf{u}}} \subset \mathcal{H}$, it suffices to demonstrate that $\Upsilon_r|_{(0,\infty)}$ is an analytic function.

For fixed n , introduce the following observables belonging to \mathfrak{A}^- and \mathfrak{A}^+ respectively:

$$\mathbf{X}_n(h) := \left(\int_{C_r} (h(s/\delta_n) - h(u_1/\delta_n)) \, ds \right) \prod_{i=2}^{\ell} (h(u'_i/\delta_n) - h(u_i/\delta_n)); \quad (147)$$

$$\mathbf{Y}_n(h) := \prod_{i=\ell+1}^k (h(u'_i/\delta_n) - h(u_i/\delta_n)). \quad (148)$$

Theorem 4.15 implies that for any $\lambda > 0$,

$$\Upsilon_r(\lambda) = \lim_{n \rightarrow \infty} \lim_{L \rightarrow \infty} \int (1-a)^{\lfloor \lambda/\delta_n \rfloor} \, d\mu_{\mathbf{X}_n, \mathbf{Y}_n, L}, \quad (149)$$




Figure 8

where $\mu_{\mathbf{X}_n, \mathbf{Y}_n, L}$ satisfies the following bound on its total variation

$$\|\mu_{\mathbf{X}_n, \mathbf{Y}_n, L}\| \leq \sqrt{\mathbb{E}_{\mathbb{C}Y_{1L}}[\mathbf{X}_n \mathbf{X}_n^\dagger] \mathbb{E}_{\mathbb{C}Y_{1L}}[\mathbf{Y}_n^\dagger \mathbf{Y}_n]}. \quad (150)$$

Now, Theorem 4.5 gives the following. Since there is some constant $C < \infty$ such that $\mathbb{E}_{\mathbb{Z}^2}[\mathbf{X}_n \mathbf{X}_n^\dagger] \mathbb{E}_{\mathbb{Z}^2}[\mathbf{Y}_n^\dagger \mathbf{Y}_n] \leq C^2$ for any n , we may find a complex measure $\mu_{\mathbf{X}, \mathbf{Y}}$ on $[0, \infty]$ such that $\|\mu_{\mathbf{X}, \mathbf{Y}}\| \leq C$ and $\mu_{\mathbf{X}, \mathbf{Y}}$ is a sub-sequential limit in the weak topology on complex measures on the compact set $[0, \infty]$ of the family $(\mu_{\mathbf{X}_n, \mathbf{Y}_n, L, \delta_n})_{n, L}$ as L and then n tend to infinity.

In particular, these statements imply that for any $\lambda > 0$, we have

$$\Upsilon_r(\lambda) = \int e^{-a\lambda} d\mu_{\mathbf{X}, \mathbf{Y}}(a). \quad (151)$$

Then, $\Upsilon_r|_{(0, \infty)}$ is the Laplace transform of a finite complex measure on $[0, \infty]$, or equivalently the linear combination (with complex coefficients) of four Laplace transforms of finite positive measure on $[0, \infty]$. It is automatically analytic. \square

We are now ready to prove full harmonicity.

Proof of Proposition 7.2. Consider the set (see Figure 8)

$$\mathcal{D}'_k := \{\mathbf{u} \in \mathcal{D}_k : \min_{i \neq j} |u_i - u_j| > 1000k^2 \max_i |u'_i - u_i|\}. \quad (152)$$

We proceed in two steps. We first show that $\mathcal{D}'_k \subset \mathcal{H}$, and then show that $\mathcal{D}_k \subset \mathcal{H}$, which is equivalent to the claim.

Step 1: proof that $\mathcal{D}'_k \subset \mathcal{H}$. Recall the definition of \mathcal{Z}_0 from Lemma 7.3. Define inductively the sets \mathcal{Z}_{t+1} (for $t \in \mathbb{Z}_{\geq 0}$) as the collection of elements $\mathbf{u} \in \mathcal{D}_k$ obtained from some $\tilde{\mathbf{u}} \in \mathcal{Z}_t$ via the operation described in the previous lemma. By construction, the previous lemma implies that $\cup_t \mathcal{Z}_t \subset \mathcal{H}$.

It suffices to prove that $\mathcal{D}'_k \subset \mathcal{Z}_2$. Fix $\mathbf{u} \in \mathcal{D}'_k$. We can choose two lines L and L' (see Figure 8) such that

- $\mathbb{R}^2 \setminus (L \cup L')$ is composed of four infinite “triangular” connected components T_1, \dots, T_4 indexed in clockwise order around the intersection point $L \cap L'$,
- u_1 and u'_1 lie in T_1 ,
- each pair $\{u_i, u'_i\}$ (for $i \geq 2$) is contained either in T_2 or in T_4 .

Translate all points in T_2 by a distance $\lambda \gg 1$ in the direction perpendicular to the line separating T_1 and T_2 . Next, translate all points in T_4 by a distance $\lambda' \gg 1$ in the direction perpendicular to the line separating T_1 and T_4 .

After these translations, it is straightforward to verify that u_1 and u'_1 no longer lie in the convex hull of the remaining points. The resulting $2k$ -tuple therefore belongs to \mathcal{Z}_0 . By definition of the sets \mathcal{Z}_t , the configuration before the second translation lies in \mathcal{Z}_1 , and the original configuration \mathbf{u} lies in \mathcal{Z}_2 .

Step 2: proof that $\mathcal{D}_k \subset \mathcal{H}$. Fix $\mathbf{u} = (u_1, u'_1, \dots, u_k, u'_k) \in \mathcal{D}_k$. We write $\Psi_k(\mathbf{u})$ as a sum of increments that are of the form $\Psi_k(\mathbf{v})$ for $\mathbf{v} \in \mathcal{D}'_k$.

Since harmonicity in u_1 does not depend on the position of u'_1 , we may assume without loss of generality that u'_1 is extremely close to u_1 . Let $(\gamma_i)_{i=2, \dots, k}$ denote a family of disjoint smooth curves $[0, 1] \rightarrow \mathbb{R}^2$ of finite length where γ_i starts at u_i and ends at u'_i . Then, for any $N \in \mathbb{Z}_{\geq 1}$, linearity of expectation implies that

$$\Psi_k(\mathbf{u}) = \sum_{n_2=1}^N \cdots \sum_{n_k=1}^N \Psi_k(u_1, u'_1, \gamma_2(\frac{n_2-1}{N}), \gamma_2(\frac{n_2}{N}), \dots, \gamma_k(\frac{n_k-1}{N}), \gamma_k(\frac{n_k}{N})). \quad (153)$$

By continuity, for N sufficiently large, all terms on the right correspond to $2k$ -tuples in \mathcal{D}'_k . Since each term is harmonic in u_1 , it follows that $\Psi(\mathbf{u})$ is harmonic in u_1 as well. \square

7.2 Characterisation of GFF correlation functions

We are now ready to prove Theorem 7.1. The idea is to proceed inductively.

Proof of Theorem 7.1. We wish to prove that $\Psi_k = \sigma^k \Psi_k^{\text{GFF}}$ by inducting on $k \in 2\mathbb{Z}_{\geq 1}$. The base case $k = 2$ holds true by hypothesis. Now, fix k and suppose that the statement holds true for all $\ell < k$.

Let us gather the properties of Ψ_k established so far. We view Ψ_k as a function in $u_1 \in \mathbb{R}^2 \setminus \cup_{i>1} \{u_i, u'_i\}$ while all other arguments are fixed. This function satisfies the following properties.

- **Harmonicity:** Ψ_k is harmonic in u_1 by Proposition 7.2.
- **Full-plane asymptotics:** Ψ_k converges to a constant in \mathbb{R} as $|u_1| \rightarrow \infty$. Indeed, the regularity estimate (Theorem 4.5) implies that $\Psi_k(\mathbf{u}) \rightarrow 0$ as $|u_1|, |u'_1| \rightarrow \infty$, which implies that $\Psi_k(\mathbf{u})$ tends to a finite constant as u'_1 is fixed and $|u_1| \rightarrow \infty$.
- **Fusion asymptotics:** for any $i > 1$ and $v \in \{u_i, u'_i\}$, Theorem 4.8 implies that as u_1 tends to v ,

$$\Psi_k(\mathbf{u}) = \Psi_2(u_1, u'_1, u_i, u'_i) \Psi_{k-2}((u_j, u'_j)_{j \notin \{1, i\}}) + O(1), \quad (154)$$

where $\Psi_2 = \sigma^2 \Psi_2^{\text{GFF}}$ and $\Psi_{k-2} = \sigma^{k-2} \Psi_{k-2}^{\text{GFF}}$ by the induction hypothesis.

- **Value at one specific point:** $\Psi_k(\mathbf{u}) = 0$ at $u_1 = u'_1$.

By extension theorems around singularities for harmonic functions [Ces14], these four properties (harmonicity together with the analysis of the singularities) fully determine the function Ψ_k , and yield $\Psi_k = \sigma^k \Psi_k^{\text{GFF}}$. This completes the proof. \square

8 Sub-sequential GFF limits

This section is dedicated to proving the following result.

Theorem 8.1 (Convergence criterion). *Fix $\sigma \in \mathbb{R}_{\geq 0}$ and let $(\delta_n)_n$ denote a sequence tending to zero. If for all $k \in \mathbb{Z}_{\geq 1}$,*

$$\Phi_k^{(\delta_n)}|_{\mathcal{D}_k} \xrightarrow[n \rightarrow \infty]{} \sigma^k \Psi_k^{\text{GFF}} \quad \text{uniformly on compacts of } \mathcal{D}_k, \quad (155)$$

then the height function converges to $\sigma \cdot \text{GFF}$ along the scaling sequence $(\delta_n)_n$ in the sense of Definition 2.7.

Proof. The proof relies on a single input: the regularity estimate of Theorem 4.5. The theorem is therefore valid for any model for which similar regularity estimates hold.

Fix $\sigma \in [0, \infty)$ and a sequence $(\delta_n)_n$ tending to zero. Recall Definition 2.7 which describes GFF convergence. Our objective is to prove convergence of finite-dimensional marginals, and convergence in law.

Convergence of finite-dimensional marginals. The generalised test functions with finite Dirichlet energy form a Hilbert space where the inner product is given by the Dirichlet form. Therefore, it suffices to consider a single generalised test function φ with finite Dirichlet energy, and prove that

$$\langle h^{(\delta_n)}, \varphi \rangle \xrightarrow{n \rightarrow \infty} \mathcal{N}(0, \sigma^2 \int G_{\mathbb{R}^2}(u, v) d\varphi(u) d\varphi(v)) \quad (156)$$

in law. Since a normal distribution is determined by its moments, it suffices to prove that all moments of the random variable $\langle h^{(\delta_n)}, \varphi \rangle$ converge to the desired limits [Bil95, Example 30.1 and Theorem 30.2], which is precisely what we will do.

Without loss of generality, we may assume that $|\varphi|(\mathbb{R}^2) = 2$ and $\varphi = \varphi_+ - \varphi_-$ where the two measures are probability measures. We first observe that

$$\langle h^{(\delta_n)}, \varphi \rangle = \int h^{(\delta_n)}(u') d\varphi_+(u') - \int h^{(\delta_n)}(u) d\varphi_-(u) \quad (157)$$

$$= \int \int (h^{(\delta_n)}(u') - h^{(\delta_n)}(u)) d\varphi_+(u') d\varphi_-(u). \quad (158)$$

Fix $k \in \mathbb{Z}_{\geq 1}$. Nothing needs to be proven for odd values of k since the corresponding moment is zero. We therefore restrict to k even. In the discrete setting, Fubini's theorem can be applied to get

$$\mathbb{E}_{\mathbb{Z}^2}[\langle h^{(\delta_n)}, \varphi \rangle^k] = \mathbb{E}_{\mathbb{Z}^2} \left[\left(\iint (h^{(\delta_n)}(u') - h^{(\delta_n)}(u)) d\varphi_+(u') d\varphi_-(u) \right)^k \right] \quad (159)$$

$$= \int \Phi_k^{(\delta_n)}(\mathbf{u}) d\tilde{\varphi}^k(\mathbf{u}) \quad (160)$$

where

$$d\tilde{\varphi}^k(\mathbf{u}) := d\varphi_-(u_1) d\varphi_+(u'_1) \cdots d\varphi_-(u_k) d\varphi_+(u'_k). \quad (161)$$

The integrand converges to $\sigma^k \cdot \Psi_k^{\text{GFF}}$ as $n \rightarrow \infty$ by our assumption (Equation (155)). Therefore, it suffices to justify an application of the dominated convergence theorem to ensure the convergence to the corresponding moments of the GFF.

By definition,

$$\Phi_k^{(\delta_n)}(\mathbf{u}) = \Phi_k\left(\left[\frac{u_1}{\delta_n}\right], \dots, \left[\frac{u'_k}{\delta_n}\right]\right) \quad (162)$$

where $[z] \in F(\mathbb{Z}^2)$ denotes the face whose lower-left corner is $(\lfloor z_1 \rfloor, \lfloor z_2 \rfloor)$ for any $z \in \mathbb{R}^2$. We shall find a dominating function using the regularity estimate (Theorem 4.5). We may do so separately for each pairing π of $\{1, \dots, k\}$ involved in (30). We focus on the pairing $\pi = \{(1, 2), \dots, (k-1, k)\}$. Since the integrals factorise over the pairs, it suffices to prove that the following integral converges:

$$\int d\tilde{\varphi}^2(\mathbf{u}) \sup_n \begin{cases} e^{-\alpha_k S_n(\mathbf{u})} & \text{if } S_n(\mathbf{u}) \geq 20k^2; \\ 1 \vee -S'_n(\mathbf{u}) & \text{if } S_n(\mathbf{u}) < 20k^2; \end{cases} \quad (163)$$

where $S_n^\#(\mathbf{u}) := S_{\mathbb{R}^2}^\#(\{[\frac{u_1}{\delta_n}], [\frac{u'_1}{\delta_n}]\}, \{[\frac{u_2}{\delta_n}], [\frac{u'_2}{\delta_n}]\})$, with $\#$ denoting either no superscript or the apostrophe.

Since $\tilde{\varphi}^2(\mathbf{u})$ is a probability measure, it suffices to prove that

$$\int d\tilde{\varphi}^2(\mathbf{u}) \sup_n \max\{0, -S'_n(\mathbf{u})\} < \infty. \quad (164)$$

By going back to the definitions of $S_{\mathbb{R}^2}$ and $S'_{\mathbb{R}^2}$, it is quite straightforward to find a constant η such that

$$\max\{0, -S'_n(\mathbf{u})\} \leq \max\{0, -S_{\mathbb{R}^2}(\mathbf{u})\} + \eta. \quad (165)$$

Thus, it suffices to prove that

$$\int \max\{0, -S_{\mathbb{R}^2}(\mathbf{u})\} d\tilde{\varphi}^2(\mathbf{u}) < \infty. \quad (166)$$

Now, recall the explicit formula for $S_{\mathbb{R}^2}$. Since $\tilde{\varphi}^2$ is compactly supported, Equation (166) follows from the fact that φ has finite Dirichlet energy.

Convergence in the negative regularity Hölder space. Fix $U \subset \mathbb{R}^2$ bounded and open and $\alpha \in (-1, 0)$. It suffices to prove that the family of random distributions $(h^{(\delta_n)})|_U$ is tight in $\mathcal{C}^\alpha(U)$, since any sub-sequential limit must then coincide with $\sigma \cdot \text{GFF}$ by the convergence of finite-dimensional marginals proved above.

Furlan and Mourrat established a general tightness criterion for random fields in negative regularity Hölder spaces [FM17, Theorem 1.1]. We apply this theorem (let us borrow notation from the paper) with $\beta = 0$ and with some integer $p > -\frac{1}{\alpha}$. Notice also that, since our random function $h^{(\delta_n)}$ is stationary, we do not need to take a supremum over x . Recall that $\mathcal{T}_1((-1, 1)^2)$ denotes the set of generalised test functions supported in $(-1, 1)^2$ and whose density with respect to the Lebesgue measure is Lipschitz. By [FM17, Theorem 1.1], it suffices to prove that for any $\varphi \in \mathcal{T}_1((-1, 1)^2)$, we have

$$\sup_{\delta \in (0, 1)} \mathbb{E}_{\mathbb{Z}^2} [|\langle h^{(\delta)}, \varphi \rangle|^p] < \infty. \quad (167)$$

Naturally, the supremum over n in [FM17, Theorem 1.1] is encoded in our scaling parameter δ , which is why it does not appear explicitly. But Equation (167) was already established above: the generalised test function φ has finite Dirichlet energy since it is Lipschitz, and the uniform bound in Equation (167) comes from the existence of the dominating function discussed above. To finish, [FM17, Theorem 1.1] implies tightness in the topological space induced by all semi-norms $(\|\cdot\|_{\mathcal{C}^\alpha(V)})_V$, which implies tightness in $\mathcal{C}^\alpha(U)$ in particular.

For tightness in the Besov spaces $\mathcal{B}_{p,q}^\alpha$ for $\alpha \in (-1, 0)$, $p \in [1, \infty)$, and $q \in [1, \infty]$, notice that the bounds in Equation (167) match the hypothesis in [FM17, Theorem 2.30], which implies the desired result. By [BL76, Theorem 6.2.4], this also implies the desired tightness in the Sobolev space $W^{\alpha,p}(U)$. \square

9 Proof of the main result (Theorem 2.8)

Proof of Theorem 2.8. Theorems 6.1, 7.1, and 8.1 jointly imply that, for $\mathbf{a} = \mathbf{b} = 1$ and for any fixed $\mathbf{c} \in [\sqrt{3}, 2]$, one of the following two statements holds true:

- There exists some $\sigma \in \mathbb{R}_{\geq 0}$ such that the height function converges to $\sigma \cdot \text{GFF}$,
- There exist two distinct standard deviations $\sigma, \sigma' \in \mathbb{R}_{\geq 0}$ and two sequences $(\delta_n)_n$ and $(\delta'_n)_n$ tending to zero, such that the height function converges to $\sigma \cdot \text{GFF}$ along the scaling sequence $(\delta_n)_n$, and to $\sigma' \cdot \text{GFF}$ along the scaling sequence $(\delta'_n)_n$.

Yet, we may rule out the second case, since σ and σ' cannot be distinct. Indeed, the combination of Theorems 4.3 and 4.4 yields that both σ^2 and $(\sigma')^2$ should be equal to $1/\arcsin \frac{c}{2}$. Thus, we are in the first case with $\sigma^2 = 1/\arcsin \frac{c}{2}$. \square

We record an additional result that may be useful in future works. Although our proof of the main theorem required a somewhat indirect route, we are now in a position to establish the convergence of the spectral measures. In particular, the sub-sequential limiting measure μ introduced earlier is in fact unique, and is simply the limit of the spectral measures.

Theorem 9.1. *Fix $\mathbf{a} = \mathbf{b} = 1$ and $\mathbf{c} \in [\sqrt{3}, 2]$. Then, if μ_L denotes the measure given by Theorem 4.12 (cf. the explicit construction in Equation (219) that follows), then*

$$\lim_{\delta \rightarrow 0} \lim_{L \rightarrow \infty} \mu_L^{(\delta)} = \frac{1}{2}(\delta_{b=a} + \delta_{b=-a}) \cdot \frac{\sigma^2}{2\pi a} da. \quad (168)$$

Proof. Since $(\mu_L)_{L \in 2\mathbb{Z}_{>0}} \subset \mathcal{M}$, the family is precompact. It therefore suffices to show that any subsequential limit μ coincides with $\frac{1}{2}(\delta_{b=a} + \delta_{b=-a}) \cdot \frac{\sigma^2}{2\pi a} da$.

We invoke the convergence of $\Phi_2^{(\delta)}$ to $\sigma^2 \cdot \Psi_2^{\text{GFF}}$. This yields

$$I_F(s) - I_F(1) = -\frac{\sigma^2}{2\pi} \log s, \quad (169)$$

which, together with (79), implies

$$F(s, 0) = \frac{\sigma^2}{2\pi s}. \quad (170)$$

Let us turn to the measure μ to see what the previous formula implies. We already know that $\mu[\{|b| > a\}] = 0$. To prove that $\mu[\{|b| < a\}] = 0$, observe that (113) gives

$$\int (a^2 - b^2) e^{-a} d\mu(a, b) = -\Xi'(1) = 0, \quad (171)$$

since $\Xi(s) = sF(s, 0) = \sigma^2/(2\pi)$ is constant. We therefore obtain that $\mu[\{|b| < a\}] = 0$.

It remains to identify the marginal in the a -variable and show that

$$d\mu(a) = \frac{\sigma^2}{2\pi a} da. \quad (172)$$

By the definition and the formula for F , we get that for every $x > 0$,

$$\int a e^{-ax} d\mu(a) = \frac{\sigma^2}{2\pi x}. \quad (173)$$

The right-hand side is precisely the Laplace transform of the measure $\frac{\sigma^2}{2\pi a} da$. Since a σ -finite positive measure is uniquely determined by its Laplace transform, this identifies μ and completes the proof. \square

10 Proof for the anisotropic case (Theorem 3.3)

The convergence of multi-point correlation functions and finite-dimensional marginals is extended to the *anisotropic* six-vertex models as follows: first express these quantities in terms of the macroscopic loops of a random-cluster model on an associated isoradial graph, and then invoke the universality result of [Dum+20], which ensures that the scaling limit of the loop ensemble is identical to that of the isotropic case. Consequently, the limiting correlations coincide with those obtained in the isotropic setting.

Proof of Theorem 3.3. Fix $\Delta \in [-1, -1/2]$. Consider the six-vertex model with weights $\mathbf{a}, \mathbf{b}, \mathbf{c}$ given by (20). Throughout this section, we restrict to the regime where $\Delta(\mathbf{a}, \mathbf{b}, \mathbf{c})$ is set to be equal to Δ . By definition of this parametrization, ζ is fixed and the only remaining degree of freedom is the angle θ . We write $\mathbb{E}_{\mathbf{L}_\theta(\mathbb{Z}^2)}^{6V}$ for the corresponding anisotropic six-vertex measure on $\mathbf{L}_\theta(\mathbb{Z}^2)$, and denote by h its associated height function.

The Baxter–Kelland–Wu correspondence [BKW76] maps this six-vertex model on $\mathbf{L}_\theta(\mathbb{Z}^2)$ to the random-cluster model on an isoradial graph $\mathbb{L}(\theta)$ (a rotated version of a rectangular lattice) with the corresponding isoradial weights; see [Dum+20] for details. The explicit construction of this random-cluster model will not be needed here. We will only use the fact that its loop representation is naturally supported on $\mathbf{L}_\theta(\mathbb{Z}^2)$. Following the notation of [Dum+20; Ave+a], we denote by $\phi_{\delta\mathbb{L}(\theta),q}$ the law of the loop ensemble of the random-cluster measure on $\delta\mathbb{L}(\theta)$ with the relevant isoradial weights.

Convergence of multi-point correlation functions. Fix $k \in 2\mathbb{Z}_{\geq 1}$ and $\mathbf{u} \in \mathcal{D}_k$. We assume without loss of generality that all $2k$ points in \mathbf{u} are distinct. In this section only, we shall write $\mathbf{u} = (u_1, u_2, \dots, u_{2k})$. Write \mathfrak{F} for the powerset of $\{1, \dots, 2k\}$ with subsets of cardinal 0, 1, and $2k$ removed.

For a given loop ensemble \mathcal{L} and $S \subset \{1, \dots, 2k\}$, let \mathcal{N}_S denote the number of loops in \mathcal{L} that surround all points in $\{u_i : i \in S\}$ and none of the points in $\{u_j : j \notin S\}$. For $S = \{v\}$, we write simply \mathcal{N}_v .

Now [Ave+a, Proposition 2.5] asserts that we may find real polynomials $P_R \in \mathbb{R}[(X_S)_{S \in \mathfrak{F}}]$ associated with each $R \subset \{1, \dots, k\}$ such that the following holds true: for every $\mathbf{u} \in \mathcal{D}_k$, and every mesh size δ for which none of the points lie on edges of $\delta\mathbf{L}_\theta(\mathbb{Z}^2)$ (so that one may unambiguously decide whether the points x_i are surrounded by the loops in \mathcal{L}),

$$\mathbb{E}_{\mathbf{L}_\theta(\mathbb{Z}^2)}^{6V} \left[\prod_{i=1}^k (h^{(\delta)}(u_{2i-1}) - h^{(\delta)}(u_{2i})) \right] = \sum_{R \subset \{1, \dots, k\}} \phi_{\delta\mathbb{L}(\theta),q} \left[\prod_{i \in R} (\mathcal{N}_{2i-1} - \mathcal{N}_{2i}) P_R((\mathcal{N}_S)_{S \in \mathfrak{F}}) \right]. \quad (174)$$

Moreover, [Ave+a, Lemma 4.5] asserts that we may restrict to “large” loops. More precisely, fix $\varepsilon > 0$, and let $\mathcal{N}_j^{\text{big}}$ be the number of loops contributing to \mathcal{N}_j that are *not* contained in the ball of radius ε centred at u_j . Then, there exists a universal constant $c > 0$ such that for every $\varepsilon > 0$, the terms on the right in (174) satisfy

$$\begin{aligned} & \phi_{\delta\mathbb{L}(\theta),q} \left[\prod_{i \in R} (\mathcal{N}_{2i-1} - \mathcal{N}_{2i}) P_R((\mathcal{N}_S)_{S \in \mathfrak{F}}) \right] \\ &= \phi_{\delta\mathbb{L}(\theta),q} \left[\prod_{i \in R} (\mathcal{N}_{2i-1}^{\text{big}} - \mathcal{N}_{2i}^{\text{big}}) P_R((\mathcal{N}_S)_{S \in \mathfrak{F}}) \right] + O(\varepsilon^c). \quad (175) \end{aligned}$$

Combining the previous two displayed equations, we obtain that the k -point correlations of the anisotropic six-vertex model can be approximated (with an error tending to zero as ε tends to zero) in terms of a functional depending only on the loops of radius at least ε in the loop ensemble obtained from the random-cluster model on $\delta\mathbb{L}(\theta)$. The universality of the random-cluster model on isoradial graphs, proved in [Dum+20], implies that the law of the loop ensembles for different values of θ are equal. In particular, the one with angle θ is the same as the one with angle $\pi/2$, which is related to the *isotropic* six-vertex model. We deduce that for every $\eta > 0$ and every compact set $K \subset \mathcal{D}_k$, there exists $\delta_0(k, K, \eta) > 0$

such that for every $0 < \delta < \delta_0$ and $\mathbf{u} \in K$,

$$\left| \mathbb{E}_{\mathbf{L}_\theta(\mathbb{Z}^2)}^{6V} \left[\prod_{i=1}^k (h^{(\delta)}(u_{2i}) - h^{(\delta)}(u_{2i-1})) \right] - \Phi_k^{(\delta)}(\mathbf{u}) \right| \leq \eta, \quad (176)$$

where we recall that $\Phi_k^{(\delta)}(\mathbf{u})$ denotes the correlation function of the isotropic case. In particular, uniform convergence of the k -point correlations on compact subsets of \mathcal{D}_k for the isotropic case $\theta = \pi/2$ (Theorem 2.8) immediately transfers to the anisotropic case with arbitrary θ .

Convergence of finite-dimensional marginals. Having established convergence of the multi-point correlation functions, it is natural to attempt the strategy used in the isotropic setting. The problem with this strategy is that it requires a dominating function (for applying the dominated convergence theorem), which we found in terms of uniform bounds on the correlation functions. At this point, a technical difficulty arises: Theorem 4.5 relies on crossing estimates for the spin representation, and these are not readily available in the anisotropic case. Implementing such an approach would therefore require revisiting the entire analysis without the symmetry given by the $\pi/2$ rotation.

To bypass this issue, we work directly with Equation (18) instead. Let $h_\theta^{(\delta)}$ be the height function of the anisotropic six-vertex model on $\mathbf{L}_\theta(\mathbb{Z}^2)$, and fix a finite Dirichlet energy test function φ . Recall that if the characteristic function of $\langle h_\theta^{(\delta)}, \varphi \rangle$ converges pointwise to that of a Gaussian, then this sequence converges to that Gaussian in law in the weak topology.

For a loop ensemble \mathcal{L} , define (whenever the infinite product is well-defined)

$$A_\varphi(\mathcal{L}) := \prod_{\ell \in \mathcal{L}} \cos_\mu \varphi(\text{int}(\ell)). \quad (177)$$

Let $\mathcal{L}^{(\delta)}$ and $\mathcal{L}_\theta^{(\delta)}$ denote the loop ensembles associated with the random-cluster model on $\delta\mathbb{L}(\frac{\pi}{2})$ and $\delta\mathbb{L}(\theta)$, respectively. Extracting a sub-sequence of $(\delta_n)_n$ if necessary, we may assume that $\mathcal{L}^{(\delta_n)}$ and $\mathcal{L}_\theta^{(\delta_n)}$ converge in law to loop ensembles $\mathcal{L}^{\text{cont}}$ and $\mathcal{L}_\theta^{\text{cont}}$. Let ϕ^{cont} and $\phi_\theta^{\text{cont}}$ denote their laws. The result of [Ave+a] ensures that for every $t \in \mathbb{R}$, the quantities $A_{t\varphi}(\mathcal{L}^{\text{cont}})$ and $A_{t\varphi}(\mathcal{L}_\theta^{\text{cont}})$ are almost surely well-defined and integrable, and that

$$\phi^{\text{cont}}[A_{t\varphi}(\mathcal{L}^{\text{cont}})] = \lim_{n \rightarrow \infty} \phi_{\delta_n \mathbb{L}(\frac{\pi}{2})}[A_{t\varphi}(\mathcal{L}^{(\delta_n)})] = \lim_{n \rightarrow \infty} \mathbb{E}_{\mathbb{Z}^2}^{6V} [e^{it\langle h^{(\delta_n)}, \varphi \rangle}], \quad (178)$$

$$\phi_\theta^{\text{cont}}[A_{t\varphi}(\mathcal{L}_\theta^{\text{cont}})] = \lim_{n \rightarrow \infty} \phi_{\delta_n \mathbb{L}(\theta)}[A_{t\varphi}(\mathcal{L}_\theta^{(\delta_n)})] = \lim_{n \rightarrow \infty} \mathbb{E}_{\mathbf{L}_\theta(\mathbb{Z}^2)}^{6V} [e^{it\langle h_\theta^{(\delta_n)}, \varphi \rangle}], \quad (179)$$

where the second equality in each line is simply (18). The universality result in [Dum+20] asserts that the two limiting loop ensembles have the same law, and therefore the quantities on the left are the same. Since we already identified the characteristic function in the isotropic case, this also implies the desired pointwise convergence of the characteristic function in the anisotropic case. \square

Remark 10.1. The fact that we do not get regularity for the anisotropic case explains why we refrain from stating convergence in Hölder spaces, as establishing it would require substantial additional technical work. Nevertheless, we expect that such a result could be obtained without encountering any fundamental obstacles.

Part C

Ingredient 4: Spectral representation of correlation functions

This part derives the results stated in Ingredient 4, namely Theorems 4.12 and 4.15. The value of $L \in 2\mathbb{Z}_{\geq 1}$ is fixed throughout this part. The two theorems follow from the basic symmetries on the cylinder (translation invariance, reflection invariance, and invariance under global arrow flip). We stress that while the transfer matrix of the six-vertex model has an extremely rich structure, the purpose of this part is underlining almost the *opposite* fact: demonstrating that the properties required for Part B do not require this structure, and are very general.

11 Transfer matrix formalism

Recall that $\text{Cyl}_L = (V(\text{Cyl}_L), E(\text{Cyl}_L))$ is the nearest-neighbour graph on the vertex set $V(\text{Cyl}_L) := \mathbb{Z} \times (\mathbb{Z}/L\mathbb{Z})$. Write E_i for the *horizontal edges* between $\{i\} \times (\mathbb{Z}/L\mathbb{Z})$ and $\{i+1\} \times (\mathbb{Z}/L\mathbb{Z})$. The set E_i is called the *i-th column of horizontal edges*. We also write $E_{ii'}$ for the union of $E_i \cup E_{i+1} \cup \dots \cup E_{i'}$ with the set of all *vertical edges* between those columns (i.e. with endpoints in $\{i+1, \dots, i'\} \times (\mathbb{Z}/L\mathbb{Z})$).

A *column configuration* is an element $\kappa \in \{\pm 1\}^{\mathbb{Z}/L\mathbb{Z}}$ encoding arrow orientations in a column E_i of *horizontal arrows*. There are L vertical positions; $\kappa_j = +1$ means that the arrow in position j is oriented to the right; $\kappa_j = -1$ means that the arrow is oriented to the left. Write \mathfrak{C} for the set of *balanced* column configurations, meaning that $\sum_j \kappa_j = 0$. Similarly, the *vertical column configuration* along the vertical line $\{k\} \times \mathbb{R}/L\mathbb{Z}$ is an element $\alpha^k \in \{\pm 1\}^{\mathbb{Z}/L\mathbb{Z}}$ encoding arrow orientations of *vertical arrows*. There are L vertical positions; $\alpha_j^k = +1$ means that the arrow between positions j and $j+1$ points up; $\alpha_j^k = -1$ means that the arrow points down. An arrow configuration on $E_{ii'}$ may be written as a family $((\kappa^k)_{i \leq k \leq i'}, (\alpha^k)_{i < k \leq i'})$, where the κ^k encode the horizontal arrows in column E_k , and the α^k the vertical arrows on $\{k\} \times \mathbb{R}/L\mathbb{Z}$.

For $i \leq i'$ and $\zeta, \zeta' \in \mathfrak{C}$, let $\mathbb{P}_{\zeta\zeta'}^{ii'}$ denote the probability measure on $\{\pm 1\}^{E_{ii'}}$ with boundary conditions ζ and ζ' , defined by

$$Z_{\zeta\zeta'}^{ii'} \mathbb{P}_{\zeta\zeta'}^{ii'} [((\kappa^k)_{i \leq k \leq i'}, (\alpha^k)_{i < k \leq i'})] := \mathbb{1}[(\kappa^i, \kappa^{i'}) = (\zeta, \zeta')] \cdot \mathbb{1}[\text{ice rule}] \cdot \mathbf{c}^{\#\{\mathbf{c}\text{-vertices}\}}, \quad (180)$$

where $Z_{\zeta\zeta'}^{ii'}$ is a normalisation factor given by

$$Z_{\zeta\zeta'}^{ii'} := \sum_{(\kappa^k)_{i \leq k \leq i'}, (\alpha^k)_{i < k \leq i'}} \mathbb{1}[(\kappa^i, \kappa^{i'}) = (\zeta, \zeta')] \cdot \mathbb{1}[\text{ice rule}] \cdot \mathbf{c}^{\#\{\mathbf{c}\text{-vertices}\}}. \quad (181)$$

We collect a few basic properties.

Lemma 11.1 (Basic properties of the probability measures).

(i) **One-step law.** For any $i \in \mathbb{Z}$ and $\kappa, \kappa' \in \{\pm 1\}^{\mathbb{Z}/L\mathbb{Z}}$,

$$Z_{\kappa\kappa'}^{i(i+1)} \mathbb{P}_{\kappa\kappa'}^{i(i+1)} [\alpha^{i+1}] = \mathbb{1}[\text{ice rule for } (\kappa, \alpha^{i+1}, \kappa')] \cdot \mathbf{c}^{\#\{\mathbf{c}\text{-vertices in } (\kappa, \alpha^{i+1}, \kappa')\}}. \quad (182)$$

(ii) **Composition rule.** For any $i \leq i' \leq i''$ and $\kappa, \kappa'' \in \{\pm 1\}^{\mathbb{Z}/L\mathbb{Z}}$,

$$Z_{\kappa\kappa''}^{ii''} \mathbb{P}_{\kappa\kappa''}^{ii''} := \sum_{\kappa' \in \mathfrak{C}} Z_{\kappa\kappa'}^{ii'} Z_{\kappa'\kappa''}^{i'i''} (\mathbb{P}_{\kappa\kappa'}^{ii'} \times \mathbb{P}_{\kappa'\kappa''}^{i'i''}). \quad (183)$$

(iii) **Torus measure.** We have

$$Z_{\mathbb{T}_{M,L}} \mathbb{P}_{\mathbb{T}_{M,L}}[\{\text{balanced}\}] \mathbb{P}_{\mathbb{T}_{M,L}}[\cdot | \{\text{balanced}\}] = \sum_{\kappa \in \mathfrak{C}} Z_{\kappa\kappa}^{0M} \mathbb{P}_{\kappa\kappa}^{0M}. \quad (184)$$

(iv) **Reflection symmetry.** Let $i \leq i'$ and let \mathbf{X} denote a random variable which is measurable in terms of the arrows in $E_{ii'}$. (Recall that the arrows encode the gradient of the height function.) Let \mathbf{X}^\dagger denote the random variable measurable in terms of the arrows in $E_{(-i')(-i)}$ obtained by composing \mathbf{X} with the reflection of the gradient over \mathcal{R} . Then,

$$\mathbb{E}_{\kappa\kappa'}^{ii'}[\mathbf{X}] = \mathbb{E}_{\kappa'\kappa}^{(-i')(-i)}[\mathbf{X}^\dagger]. \quad (185)$$

(v) **Horizontal shift invariance.** Let $i \leq i'$ and let \mathbf{X} denote a random variable which is measurable in terms of the arrows in $E_{ii'}$. For $n \in \mathbb{Z}$, let $\tau_{(n,0)}$ denote the shift by $(n, 0)$, and let $\tau_{(n,0)}(\mathbf{X})$ denote the random variable measurable in terms of $E_{(i+n)(i'+n)}$ obtained by composing \mathbf{X} with the shift. Then,

$$\mathbb{E}_{\kappa\kappa'}^{ii'}[\mathbf{X}] = \mathbb{E}_{\kappa\kappa'}^{(i+n)(i'+n)}[\tau_{(n,0)}(\mathbf{X})]. \quad (186)$$

Proof. The first three follow from bookkeeping manipulations of the sums. The final two are symmetries of the six-vertex model. \square

The previous lemma reveals a Hilbert space structure. Let $\Omega := \ell^2(\mathfrak{C})$ denote the complex Hilbert space with orthonormal basis $(e_\kappa)_{\kappa \in \mathfrak{C}}$. Use the symbol \dagger for the Hermitian conjugate of vectors and operators. Associate a Hilbert space $\Omega_i := \Omega$ to each column of horizontal edges E_i . For each *valid triple*, that is, a triple (\mathbf{X}, i, i') , where $i \leq i'$ and where \mathbf{X} is a complex-valued operator depending only on the edges in $E_{ii'}$, we define the operator $\mathfrak{o}_{\mathbf{X}}^{ii'} : \Omega_i \rightarrow \Omega_{i'}$ via

$$e_{\kappa'}^\dagger \mathfrak{o}_{\mathbf{X}}^{ii'} e_\kappa := Z_{\kappa\kappa'}^{ii'} \mathbb{E}_{\kappa\kappa'}^{ii'}[\mathbf{X}]. \quad (187)$$

Within this framework, the above lemma readily implies the following corollary.

Corollary 11.2 (Basic properties of operators).

(i) **One-step partition function.** The operator $t(\pi/2) := \mathfrak{o}_1^{i(i+1)}$, which corresponds to the one-step partition function, is independent of the choice of i , and satisfies

$$e_{\kappa'}^\dagger t(\pi/2) e_\kappa := \sum_{\alpha \in \{\pm 1\}^L} \mathbb{1}[\text{ice rule for } (\kappa, \alpha, \kappa')] \cdot \mathfrak{c}^{\#\{\mathfrak{c}\text{-vertices in } (\kappa, \alpha, \kappa')\}}. \quad (188)$$

(ii) **Composition rule.** For any two valid triples (\mathbf{X}, i, i') and (\mathbf{Y}, i', i'') , the composition $(\mathbf{X}\mathbf{Y}, i, i'')$ is a valid triple, and

$$\mathfrak{o}_{\mathbf{X}\mathbf{Y}}^{ii''} = \mathfrak{o}_{\mathbf{Y}}^{i'i''} \circ \mathfrak{o}_{\mathbf{X}}^{ii'}. \quad (189)$$

(iii) **Torus measure.** For any valid triple $(\mathbf{X}, 0, M)$,

$$Z_{\mathbb{T}_{M,L}} \mathbb{P}_{\mathbb{T}_{M,L}}[\{\text{balanced}\}] \mathbb{E}_{\mathbb{T}_{M,L}}[\mathbf{X} | \{\text{balanced}\}] = \text{Trace } \mathfrak{o}_{\mathbf{X}}^{0M}. \quad (190)$$

(iv) **Reflection symmetry.** For any valid triple (\mathbf{X}, i, i') where \mathbf{X} is real-valued,

$$(\mathfrak{o}_{\mathbf{X}}^{ii'})^\dagger = \mathfrak{o}_{\mathbf{X}^\dagger}^{(-i')(-i)}. \quad (191)$$

(v) **Horizontal shift-invariance.** For any valid triple (\mathbf{X}, i, i') and $n \in \mathbb{Z}_{\geq 0}$,

$$\mathfrak{o}_{\tau_{(n,0)}(\mathbf{X})}^{(i+n)(i'+n)} = \mathfrak{o}_{\mathbf{X}}^{ii'}. \quad (192)$$

We now turn to the transfer matrix $t(\pi/2)$ itself.

Lemma 11.3 (Basic properties of the transfer matrix). *The operator $t(\pi/2)$ is a Hermitian Perron-Frobenius matrix with a single block. Therefore, it is diagonal in some orthonormal basis $(v_k)_k$ with real eigenvalues, and the basis may be labelled such that the corresponding eigenvalues $(\lambda_k(\pi/2))_k$ satisfy*

$$\lambda_0(\pi/2) > |\lambda_1(\pi/2)| \geq |\lambda_2(\pi/2)| \geq \lambda_3(\pi/2) \geq \dots \quad (193)$$

Finally, v_0 may be chosen such that it has positive real entries in the basis $(e_\kappa)_\kappa$.

Proof. The operator $t(\pi/2)$ is a real symmetric matrix in the basis $(e_\kappa)_\kappa$. Therefore, it is Hermitian. It may be checked directly that it is a Perron-Frobenius matrix. It is easy to see that $t(\pi/2)^k$ has positive entries for sufficiently large k . As a consequence, it has a single block. The rest follows. \square

Define

$$T(\pi/2) := \frac{t(\pi/2)}{\lambda_0(\pi/2)}; \quad \Lambda_k(\pi/2) := \frac{\lambda_k(\pi/2)}{\lambda_0(\pi/2)}; \quad \mathfrak{D}_{\mathbf{X}}^{ii'} := \mathfrak{o}_{\mathbf{X}}^{ii'} / \lambda_0(\pi/2)^{i'-i}. \quad (194)$$

Remark 11.4 (Lemma 11.3). The eigenspace of $T(\pi/2)$ associated with the top eigenvalue 1 is one-dimensional. The convention that v_0 has positive real entries simply fixes the complex phase of the normalised eigenvector corresponding to $\lambda_0(\pi/2)$. The spectrum of $T(\pi/2)$ is supported on $(-1, 1]$ since -1 cannot be an eigenvalue of $T(\pi/2)$ due to the Perron-Frobenius property.

By construction, for any valid triple (\mathbf{X}, i, i') , we have

$$\mathbb{E}_{\text{Cyl}_L}[\mathbf{X}] = \lim_{M \rightarrow \infty} \mathbb{P}_{\mathbb{T}_{M,L}}[\mathbf{X} | \{\text{balanced}\}] = \lim_{M \rightarrow \infty} \frac{\text{Trace}(\mathfrak{o}_{\mathbf{X}}^{ii'} t(\pi/2)^{M-(i'-i)})}{\text{Trace} t(\pi/2)^M} = v_0^\dagger \mathfrak{D}_{\mathbf{X}}^{ii'} v_0. \quad (195)$$

Indeed, the rightmost equality is obtained by observing that only the largest eigenvalue of $t(\pi/2)$ contributes.

We conclude this subsection by observing that Lemma 4.11 is now straightforward.

Proof of Lemma 4.11. See Equation (195) above. \square

12 Spectral representation of general observables (Theorem 4.15)

Proof of Theorem 4.15. Consider the setting of the theorem, and fix L . Define the embeddings

$$\mathcal{E}^- : \mathfrak{A}^- \rightarrow \Omega, \quad \mathbf{X} \mapsto \mathfrak{D}_{\mathbf{X}}^{(-i)0} v_0; \quad (196)$$

$$\mathcal{E}^+ : \mathfrak{A}^+ \rightarrow \Omega^\dagger, \quad \mathbf{Y} \mapsto v_0^\dagger \mathfrak{D}_{\mathbf{Y}}^{0i}, \quad (197)$$

where i is chosen so large that \mathbf{X} and \mathbf{Y} are measurable with respect to the arrows in $E_{(-i)0}$ and E_{0i} respectively. The definitions do not depend on i since the composition rule implies that

$$v_0^\dagger \mathfrak{D}_{\mathbf{Y}}^{0(i+1)} = v_0^\dagger T(\pi/2) \mathfrak{D}_{\mathbf{Y}}^{0i} = v_0^\dagger \mathfrak{D}_{\mathbf{Y}}^{0i} \quad (198)$$

(the proof for \mathcal{E}^- is the same). Notice that \mathcal{E}^- and \mathcal{E}^+ are linear maps that map real vector spaces to complex vector spaces.

Step 1: Definition of $\mu_{\mathbf{X},\mathbf{Y},L}$. For \mathbf{X} and \mathbf{Y} , define the measure $\mu_{\mathbf{X},\mathbf{Y},L}$ on $[0, 2)$ via

$$\mu_{\mathbf{X},\mathbf{Y},L} := \sum_k (\mathcal{E}^+(\mathbf{Y})v_k)(v_k^\dagger \mathcal{E}^-(\mathbf{X})) \delta_{1-\Lambda_k(\pi/2)}. \quad (199)$$

(Above, δ_x denotes the Dirac measure at the point x .)

Step 2: Proof of (i). Using Equation (195) and the composition rule in the first equality, horizontal shift invariance in the second and the definition of the embeddings in the third, we get

$$\mathbb{E}_{\text{Cyl}_L}[\mathbf{X} \cdot \tau_{(n,0)}(\mathbf{Y})] = v_0^\dagger \mathfrak{D}_{\tau_{(n,0)}(\mathbf{Y})}^{n(i+n)} T(\pi/2)^n \mathfrak{D}_{\mathbf{X}}^{(-i)0} v_0 \quad (200)$$

$$= v_0^\dagger \mathfrak{D}_{\mathbf{Y}}^{0i} T(\pi/2)^n \mathfrak{D}_{\mathbf{X}}^{(-i)0} v_0 \quad (201)$$

$$= \mathcal{E}^+(\mathbf{Y})T(\pi/2)^n \mathcal{E}^-(\mathbf{X}). \quad (202)$$

Using (199), this implies the desired formula

$$\mathbb{E}_{\text{Cyl}_L}[\mathbf{X} \cdot \tau_{(n,0)}(\mathbf{Y})] = \int (1-a)^n d\mu_{\mathbf{X},\mathbf{Y},L}(a). \quad (203)$$

Step 3: Proof of (ii). For \mathbf{X} , reflection symmetry and the definition of the embeddings imply

$$\mathcal{E}^+(\mathbf{X}^\dagger) = v_0^\dagger \mathfrak{D}_{\mathbf{X}^\dagger}^{0i} = v_0^\dagger (\mathfrak{D}_{\mathbf{X}}^{-i0})^\dagger = (\mathfrak{D}_{\mathbf{X}}^{-i0} v_0)^\dagger = (\mathcal{E}^-(\mathbf{X}))^\dagger. \quad (204)$$

As a consequence, $\mu_{\mathbf{X},\mathbf{X}^\dagger,L}$ is a positive measure since for every k ,

$$(\mathcal{E}^+(\mathbf{X}^\dagger)v_k)(v_k^\dagger \mathcal{E}^-(\mathbf{X})) = (\mathcal{E}^-(\mathbf{X})^\dagger v_k)(v_k^\dagger \mathcal{E}^-(\mathbf{X})) = \|\mathcal{E}^-(\mathbf{X})^\dagger v_k\|^2 \geq 0. \quad (205)$$

It remains to prove the Cauchy-Schwarz inequality (Equation (46)). Using (203) (for $n = 0$) and the definition of the total variation, we obtain

$$|\mathbb{E}_{\text{Cyl}_L}[\mathbf{X}\mathbf{Y}]| = |\mu_{\mathbf{X},\mathbf{Y},L}([0, 2])| \leq \|\mu_{\mathbf{X},\mathbf{Y},L}\| \leq \sum_k \left| (\mathcal{E}^+(\mathbf{Y})v_k)(v_k^\dagger \mathcal{E}^-(\mathbf{X})) \right|. \quad (206)$$

Applying the Cauchy-Schwarz inequality to the sum on the right-hand side gives

$$\left(\sum_k |(\mathcal{E}^+(\mathbf{Y})v_k)(v_k^\dagger \mathcal{E}^-(\mathbf{X}))| \right)^2 \leq \left(\sum_k \|\mathcal{E}^+(\mathbf{Y})v_k\|^2 \right) \left(\sum_k \|v_k^\dagger \mathcal{E}^-(\mathbf{X})\|^2 \right). \quad (207)$$

Using (203) (for $n = 0$) in the other direction together with the positivity of $\mu_{\mathbf{X},\mathbf{X}^\dagger,L}$ gives

$$\sum_k \|v_k^\dagger \mathcal{E}^-(\mathbf{X})\|^2 = \|\mu_{\mathbf{X},\mathbf{X}^\dagger,L}\| = \int d\mu_{\mathbf{X},\mathbf{X}^\dagger,L}(a) = \mathbb{E}_{\text{Cyl}_L}[\mathbf{X}\mathbf{X}^\dagger]. \quad (208)$$

Applying the same argument, we similarly obtain

$$\sum_k \|\mathcal{E}^+(\mathbf{Y})v_k\|^2 = \mathbb{E}_{\text{Cyl}_L}[\mathbf{Y}^\dagger \mathbf{Y}]. \quad (209)$$

Combining (206)–(209) yields the Cauchy-Schwarz inequality, completing the proof. \square

13 Spectral representation for the two-point function (Theorem 4.12)

Before diving into the proof of Theorem 4.12, let us start by introducing two more operators – the vertical-arrow and up-shift operators – and discuss their basic properties. To measure arrow-arrow correlations between vertical arrows, we introduce some new operators encoding the direction of the vertical arrow between positions j and $j + 1$. For $j \in (\mathbb{Z}/L\mathbb{Z})$, set $s_j(\pi/2) := \mathfrak{o}_{\alpha_j}^{01}$, or equivalently

$$e_{\kappa'}^\dagger s_j(\pi/2) e_\kappa := Z_{\kappa\kappa'}^{01} \mathbb{E}_{\kappa\kappa'}^{01}[\alpha_j] \quad (210)$$

$$= \sum_{\alpha \in \{\pm 1\}^L} \alpha_j \cdot \mathbb{1}[\text{ice rule for } (\kappa, \alpha, \kappa')] \cdot \mathbf{c}^{\#\{\mathbf{c}\text{-vertices in } (\kappa, \alpha, \kappa')\}}. \quad (211)$$

Write $S(\pi/2) := \mathfrak{D}_{\alpha_0}^{01} = s_0(\pi/2)/\lambda_0(\pi/2)$.

Remark 13.1. For Step 3, it suffices to compute the arrow-arrow correlation between two vertical arrows. We shall see that the general case follows from the addition rule and a mixing property. We therefore work exclusively with the operator that measures vertical arrows and avoid introducing operators for horizontal arrows altogether.

Next, define the up-shift operator. For any $\kappa \in \mathfrak{C}$, we let $\kappa^\uparrow \in \mathfrak{C}$ denote the column configuration such that $(\kappa^\uparrow)_i = \kappa_{i-1}$. Similarly, we define $\kappa^\downarrow \in \mathfrak{C}$ such that $(\kappa^\downarrow)^\uparrow = \kappa$. They are called the *up* and *down* shift respectively. Define the *up-shift operator* $T(0) : \Omega \rightarrow \Omega$ by

$$e_{\kappa'}^\dagger T(0) e_\kappa := \mathbb{1}[\kappa' = \kappa^\uparrow]. \quad (212)$$

By going back to the definition of $s_j(\pi/2)$, it is easy to see that

$$s_j(\pi/2) = T(0)^j s_0(\pi/2) T(0)^{-j}; \quad (213)$$

$$\mathfrak{D}_{\alpha_j}^{01} = \frac{s_j(\pi/2)}{\lambda_0(\pi/2)} = T(0)^j S(\pi/2) T(0)^{-j}. \quad (214)$$

Let us now collect some basic properties of these operators.

Lemma 13.2 (The shift operator). *The basis $(v_k)_k$ introduced in Lemma 11.3 may be chosen such that it diagonalises $T(0)$ as well. If $\Lambda_k(0)$ denotes the eigenvalue associated with v_k , then the $\Lambda_k(0)$ are L -th roots of unity, and $\Lambda_0(0) = 1$.*

Proof. The operator $T(0)$ is normal since $T(0)$ commutes with $T(0)^\dagger = T(0)^{-1}$. Moreover, $T(0)$ and $T(\pi/2)$ commute thanks to shift-invariance of the six-vertex model. Therefore $T(0)$ and $T(\pi/2)$ are co-diagonalisable. We suppose without loss of generality that $T(0)$ was already diagonal in the basis $(v_k)_k$. This can be done without breaking the properties of $(v_k)_k$ from Lemma 11.3. Indeed, in Lemma 11.3 we simply chose an arbitrary orthonormal basis diagonalising $T(\pi/2)$, and then ordered it, but the ordering does not play a role here (in particular because the eigenspace of the $\Lambda_0(\pi/2)$ is one dimensional).

Since $T(0)^L = \text{Identity}$, its eigenvalues are L -th roots of unity. Finally, v_0 is the Perron-Frobenius eigenvector of $T(\pi/2)$, and we imposed that all its entries in the basis $(e_\kappa)_\kappa$ are strictly positive. Since $T(0)$ has non-negative coefficients (it is a permutation matrix) in the basis $(e_\kappa)_\kappa$, this means that the entries of $T(0)v_0$ are also positive. Since $T(0)v_0 = \Lambda_0(0)v_0$, it forces $\Lambda_0(0) = 1$. \square

Lemma 13.3 (The edge measurement operator). *The operator $S(\pi/2)$ is anti-Hermitian, and $v_0^\dagger S(\pi/2)v_0 = 0$.*

Proof. Focus on the first property. It is straightforward to see that $s_0(\pi/2)$ is anti-Hermitian, by going back to its definition in the basis $(e_\kappa)_\kappa$. Indeed, the definitions imply that

$$e_{\kappa'}^\dagger s_0(\pi/2) e_\kappa = \sum_{\alpha \in \{\pm 1\}^L} \alpha_j \cdot \mathbb{1}[\text{ice rule for } (\kappa, \alpha, \kappa')] \cdot \mathbf{c}^{\#\{\mathbf{c}\text{-vertices in } (\kappa, \alpha, \kappa')\}} \quad (215)$$

$$= \sum_{\alpha \in \{\pm 1\}^L} \alpha_j \cdot \mathbb{1}[\text{ice rule for } (\kappa', -\alpha, \kappa)] \cdot \mathbf{c}^{\#\{\mathbf{c}\text{-vertices in } (\kappa', -\alpha, \kappa)\}} \quad (216)$$

$$= -e_{\kappa'}^\dagger s_0(\pi/2) e_{\kappa'}. \quad (217)$$

For the second property, notice that Equation (195) yields

$$v_0^\dagger S(\pi/2) v_0 = v_0^\dagger \mathfrak{D}_{\alpha_0}^{01} v_0 = \mathbb{E}_{\text{Cyl}_L}[\text{the orientation of a fixed vertical edge}] = 0, \quad (218)$$

where the expectation is zero since the measure is invariant under a global arrow flip. \square

Proof of Theorem 4.12. For any even integer L , define the measure

$$\mu_L := \sum_{k>0} \frac{|v_k^\dagger S(\pi/2) v_0|^2}{(1 - \Lambda_k(\pi/2))^2} \delta_{(1 - \Lambda_k(\pi/2), -i \log \Lambda_k(0))}. \quad (219)$$

The support of μ_L is contained in $(0, 2) \times [-\pi, \pi)$, since $|\Lambda_k(\pi/2)| < 1$ for every $k > 0$. This inequality also ensures that the denominator $1 - \Lambda_k(\pi/2)$ is nonzero for every $k > 0$, thereby validating the definition.

The proof now proceeds in three steps. First, we prove Equation (41) under the constraint that $y_1 = y_2 = 0$. Then, we extend it to the generality of Theorem 4.12. Finally, we conclude the proof.

Step 1: Equation (41) under the constraint that $y_1 = y_2 = 0$. In this step, we consider \mathbf{u} where u_1 and u'_1 are on the same horizontal line, and that the same holds true for u_2 and u'_2 . In that case, $\Phi_{\text{Cyl}_{L,2}}(\mathbf{u})$ can be written as a sum of two-point functions for \mathbf{u} with (u_1, u'_1) and (u_2, u'_2) horizontally adjacent. The additivity property of Equation (10) therefore reduces the problem to proving Equation (41) in this special case.

We work out the left and right of Equation (41). Consider \mathbf{u} . When $u'_1 - u_1 = u'_2 - u_2 = (1, 0)$, we are just calculating an arrow-arrow correlation between vertical arrows. Recall the definition of $\mathbf{u}_2 - \mathbf{u}'_1 = (x'_1, y'_1)$, and that $y(u)$ denotes the second coordinate of $u \in \mathbb{R}^2$. For the left-hand side, we get

$$\Phi_{\text{Cyl}_{L,2}}(\mathbf{u}) = \mathbb{E}_{\text{Cyl}_L}[(h(u'_1) - h(u_1))(h(u'_2) - h(u_2))] \quad (220)$$

$$= v_0^\dagger \mathfrak{D}_{\alpha_{y(u_2)}}^{01} T(\pi/2)^{x'_1} \mathfrak{D}_{\alpha_{y(u_1)}}^{01} v_0 \quad (221)$$

$$= v_0^\dagger T(0)^{y(u_2)} S(\pi/2) T(0)^{-y(u_2)} T(\pi/2)^{x'_1} T(0)^{y(u_1)} S(\pi/2) T(0)^{-y(u_1)} v_0 \quad (222)$$

$$= v_0^\dagger S(\pi/2) T(\pi/2)^{x'_1} T(0)^{-y'_1} S(\pi/2) v_0. \quad (223)$$

In the last line we use that $T(\pi/2)$ and $T(0)$ commute. By expanding $S(\pi/2)$ in the basis $(v_k)_k$, and using the anti-Hermitian property, we get

$$\Phi_{\text{Cyl}_{L,2}}(\mathbf{u}) = - \sum_{k>0} |v_k^\dagger S(\pi/2) v_0|^2 \Lambda_k(\pi/2)^{x'_1} \Lambda_k(0)^{-y'_1}. \quad (224)$$

We omitted the $k = 0$ term since $v_0^\dagger S(\pi/2) v_0 = 0$ due to flip symmetry (218).

For the right hand side of Equation (41) we get, using that $x_1 = x_2 = 1$ and $y_1 = y_2 = 0$,

$$\begin{aligned} & \int -a^2(1-a)^{x'_1} e^{-iby'_1} d\mu_L(a, b) \\ &= - \sum_{k>0} \frac{|v_k^\dagger S(\pi/2)v_0|^2}{(1-\Lambda_k(\pi/2))^2} (1-\Lambda_k(\pi/2))^2 \Lambda_k(\pi/2)^{x'_1} \Lambda_k(0)^{-y'_1}. \end{aligned} \quad (225)$$

This matches the expression we found for $\Phi_{\text{Cyl}_L, 2}(\mathbf{u})$.

Step 2 : Equation (41) in the general case. We now relax the condition that each pair (u_i, u'_i) lies on a horizontal line.

We shall first relax the condition that $y_1 = 0$, and then the condition that $y_2 = 0$. In fact, the two proofs are the same, and we focus on the first step.

Fix \mathbf{u} such that $y_1 \neq 0$ and $y_2 = 0$. The idea is to write $\Phi_{\text{Cyl}_L, 2}(\mathbf{u})$ as the sum of three terms, where two terms fall under the umbrella of Step 4, and where the third term is an error term that tends to zero in a certain limit.

More precisely, we write the height difference $h(u'_1) - h(u_1)$ as a sum of increments along the path $(u_1, u_1 - ne_1, u'_1 - ne_1, u'_1)$, with $e_1 = (1, 0)$. The first and third steps are covered by Step 4, thus giving

$$\begin{aligned} & \Phi_{\text{Cyl}_L, 2}(\mathbf{u}) - \Phi_{\text{Cyl}_L, 2}(u_1 - ne_1, u'_1 - ne_1, u_2, u'_2) \\ &= \int ((1-a)^{x_2} - 1) (1-a)^{x'_1} e^{-iby'_1} \left(1 - (1-a)^{x_1} e^{-iby_1}\right) (1 - (1-a)^n) d\mu_L(a, b). \end{aligned} \quad (226)$$

Now let n tend to infinity. Since μ_L is a finite sum of Dirac measures with $0 < a < 2$ almost surely, the right-hand side tends to

$$\int ((1-a)^{x_2} - 1) (1-a)^{x'_1} e^{-iby'_1} \left(1 - (1-a)^{x_1} e^{-iby_1}\right) d\mu_L(a, b). \quad (227)$$

It suffices to prove that as n tends to infinity, the following error term vanishes:

$$\Phi_{\text{Cyl}_L, 2}(u_1 - ne_1, u'_1 - ne_1, u_2, u'_2) = \mathbb{E}_{\text{Cyl}_L} \left[(h(u'_1 - ne_1) - h(u_1 - ne_1))(h(u'_2) - h(u_2)) \right]. \quad (228)$$

But this expectation may be written as

$$v_0^\dagger \mathfrak{D}_2 T(\pi/2)^n \mathfrak{D}_1 v_0, \quad (229)$$

where each \mathfrak{D}_i measures one of the two height differences. As n tends to infinity, this tends to $v_0^\dagger \mathfrak{D}_2 v_0 v_0^\dagger \mathfrak{D}_1 v_0$, since $\Lambda_0(\pi/2) = 1$ and all other eigenvalues have a modulus strictly smaller than 1. But $v_0^\dagger \mathfrak{D}_1 v_0 = v_0^\dagger \mathfrak{D}_2 v_0 = 0$ as each factor encodes the expectation of a single height difference, which is zero by global flip symmetry. This proves the case that $y_1 \neq 0$ and $y_2 = 0$; the general case is proved similarly.

Step 3: Conclusion We constructed the measure μ_L above, we observed that it is supported on $(0, 2) \times [-\pi, \pi)$, and we established Equation (41) for all the desired points \mathbf{u} . The fact that $\mu_L[\{|b| \in (0, 2\pi/L)\}] = 0$ follows from the observation that the eigenvalues e^{-ib} of $T(0)$ are L -th roots of unity.

Finally, we want to prove that μ_L is invariant under the reflection $(a, b) \mapsto (a, -b)$. But this is immediate, as we may simply replace μ_L by its symmetrised version $(\mu_L + \bar{\mu}_L)/2$ owing to the real-valuedness of the correlation function, which is invariant under the reflection by construction and clearly still satisfies all the other properties of Theorem 4.12. \square

Remark 13.4. The above reflection $(a, b) \mapsto (a, -b)$ corresponds to applying a complex conjugation to the weighted eigenvalues of the two operators. Complex conjugation is an element of the Galois group of the field extension $\mathbb{R} \subset \mathbb{C}$. Since all operators are real-valued in the basis $(e_\kappa)_\kappa$, the Galois group preserves the weighted joint spectra of $T(\pi/2)$ and $T(0)$. In particular, the formula in (219) is itself invariant under the reflection $(a, b) \mapsto (a, -b)$.

Part D

Ingredient 3: Regularity estimates and qualitative behaviour

The purpose of this part is to prove the following intermediate results:

- Existence of the infinite-volume six-vertex measure (Theorem 2.2),
- The results stated in Ingredient 3 (Theorems 4.5, 4.8, 4.9, 4.10, and Corollary 4.7).

These results have already been applied in Part B, and we will apply them again in Part E (where we derive the remaining missing ingredients).

The proofs in this part are based on the Fortuin–Kasteleyn–Ginibre (FKG) inequality and the Russo–Seymour–Welsh (RSW) theory. We already mentioned that several representations of the six-vertex model satisfy this FKG inequality for $\mathbf{c} \geq 1$. We found the *spin representation* of the six-vertex model the most convenient for formalising the proof (see [Lis22; Lis21; GL25]). We expect that the proofs can also be written down in terms of the FKG inequality for the absolute value of the height function (see [Dum+24]), but we do not pursue this route here.

This part is organised as follows. Sections 14–16 introduce the spin representation and its properties. Although the results are more or less known (for the spin representation, the results may be found, for example, in [GL25]), we still state everything precisely because we need some subtle variations of the known results, as well as minor extensions. Everything is written in a self-contained fashion, except for the circuit estimate, which was proved in [Dum+24] and [GL25].

Once the language and standard results for the spin representation have been established, we prove the desired intermediate results one by one in Sections 17–23.

14 Spin representation and RSW theory

14.1 Motivation of the spin representation

Recall that $h : F(\mathbb{Z}^2) \rightarrow \mathbb{Z}$ denotes the height function representation of the six-vertex model, which is a function differing by ± 1 on neighbouring faces and which assigns an even number to the face whose south-west corner is $(0, 0)$. We may partition $F(\mathbb{Z}^2)$ into two so that h assigns even numbers to F_\circ (the *even* faces) and odd numbers to F_\bullet (the *odd* faces).

The gradient of h can be recovered from the values of h modulo 4, that is, from the function

$$h/4\mathbb{Z} : F(\mathbb{Z}^2) \rightarrow \mathbb{Z}/4\mathbb{Z}, x \mapsto h(x) + 4\mathbb{Z}. \quad (230)$$

Moreover, this height function modulo 4 can be encoded in terms of a family $(\sigma_\circ, \sigma_\bullet) \in \{\pm\}^{F_\circ(\mathbb{Z}^2)} \times \{\pm\}^{F_\bullet(\mathbb{Z}^2)}$ of *spins*. Indeed, we may simply define $(\sigma_\circ, \sigma_\bullet)$ as follows:

- If x is even, then

$$\sigma_\circ(x) := \begin{cases} + & \text{if } h(x) \in 0 + 4\mathbb{Z}, \\ - & \text{if } h(x) \in 2 + 4\mathbb{Z}; \end{cases} \quad (231)$$

- If x is odd, then

$$\sigma_{\bullet}(x) := \begin{cases} + & \text{if } h(x) \in 1 + 4\mathbb{Z}, \\ - & \text{if } h(x) \in 3 + 4\mathbb{Z}. \end{cases} \quad (232)$$

The spin representation of a height function (that is, the height function modulo 4) thus encodes the height function up to constant shifts by constant multiples of 4. Spins at even faces are called *even spins* and spins at odd faces are called *odd spins*.

The spin representation turns out to be useful because of two reasons:

- The even spins satisfy the Fortuin–Kasteleyn–Ginibre inequality,
- The even spins have a “smallest” and “largest” value, namely $-$ and $+$ respectively, setting it apart from the height function which is unbounded and has no smallest or largest value.

Although $F(\mathbb{Z}^2)$ is a set of faces, we often identify each face with its face centre. This way we may view $F(\mathbb{Z}^2)$ as a vertex set embedded in \mathbb{R}^2 , and we may identify edges between faces with line segments between face centres. Note that this coincides with the standard dual graph of \mathbb{Z}^2 , but we avoid referring to it explicitly since several other graphs will appear later on.

14.2 Formal definition of the spin representation

Definition 14.1 (Even and odd sublattices). We shall identify each face in $F(\mathbb{Z}^2)$ with its face centre. The set of even faces F_{\circ} is endowed with an edge set E_{\circ} such that the four neighbours of a face at (i, j) are given by $(i \pm 1, j \pm 1)$. We shall also simply write F_{\circ} for the graph (F_{\circ}, E_{\circ}) . Similar definitions apply to the odd sublattice F_{\bullet} . See Figure 9.

Definition 14.2 (Even domain). In the context of percolation theory, we shall identify each edge $\{x, y\} \subset \mathbb{R}^2$ with the straight closed line segment connecting the vertices x and y . Let $\partial\mathcal{D}$ denote a finite self-avoiding circuit through the even sublattice F_{\circ} . We let $F_{\circ}(\mathcal{D})$ denote the subgraph of F_{\circ} consisting of the vertices and edges which are entirely contained in the closure of the set of points surrounded by $\partial\mathcal{D}$ (in particular, $\partial\mathcal{D} \subset F_{\circ}(\mathcal{D})$). Its edge set is denoted $E_{\circ}(\mathcal{D})$ and its vertex set simply $F_{\circ}(\mathcal{D})$. The subgraph $F_{\bullet}(\mathcal{D}) \subset F_{\bullet}$ is defined similarly, except that obviously $\partial\mathcal{D}$ cannot be a subgraph of $F_{\bullet}(\mathcal{D})$. The triple $\mathcal{D} := (\partial\mathcal{D}, F_{\circ}(\mathcal{D}), F_{\bullet}(\mathcal{D}))$ is called an *even domain*, see Figure 9. We also write $F(\mathcal{D})$ for the subgraph of the graph $F(\mathbb{Z}^2)$ (where faces sharing an edge are neighbours) induced by the vertex set $F_{\circ}(\mathcal{D}) \cup F_{\bullet}(\mathcal{D})$. We shall not define or use odd domains.

Definition 14.3 (Consistent spin configurations). Let \mathcal{D} denote an even domain. A *spin configuration* on \mathcal{D} is a pair

$$(\sigma_{\circ}, \sigma_{\bullet}) \in \{\pm\}^{F_{\circ}(\mathcal{D})} \times \{\pm\}^{F_{\bullet}(\mathcal{D})}. \quad (233)$$

We call such a spin configuration *consistent* if for any edge $uv \in E_{\bullet}(\mathcal{D})$ with dual edge $xy = uv^* \in E_{\circ}(\mathcal{D})$ (here dual means intersecting it in its middle), we have

$$\sigma_{\bullet}(u) = \sigma_{\bullet}(v) \quad \text{or} \quad \sigma_{\circ}(x) = \sigma_{\circ}(y) \quad (\text{or both}). \quad (234)$$

In this case we write $\sigma_{\circ} \perp \sigma_{\bullet}$; see Figure 9 for a consistent spin configuration.

The ice rule and the consistency rule are two sides of the same coin: they are the necessary and sufficient conditions for the existence of the associated height function.

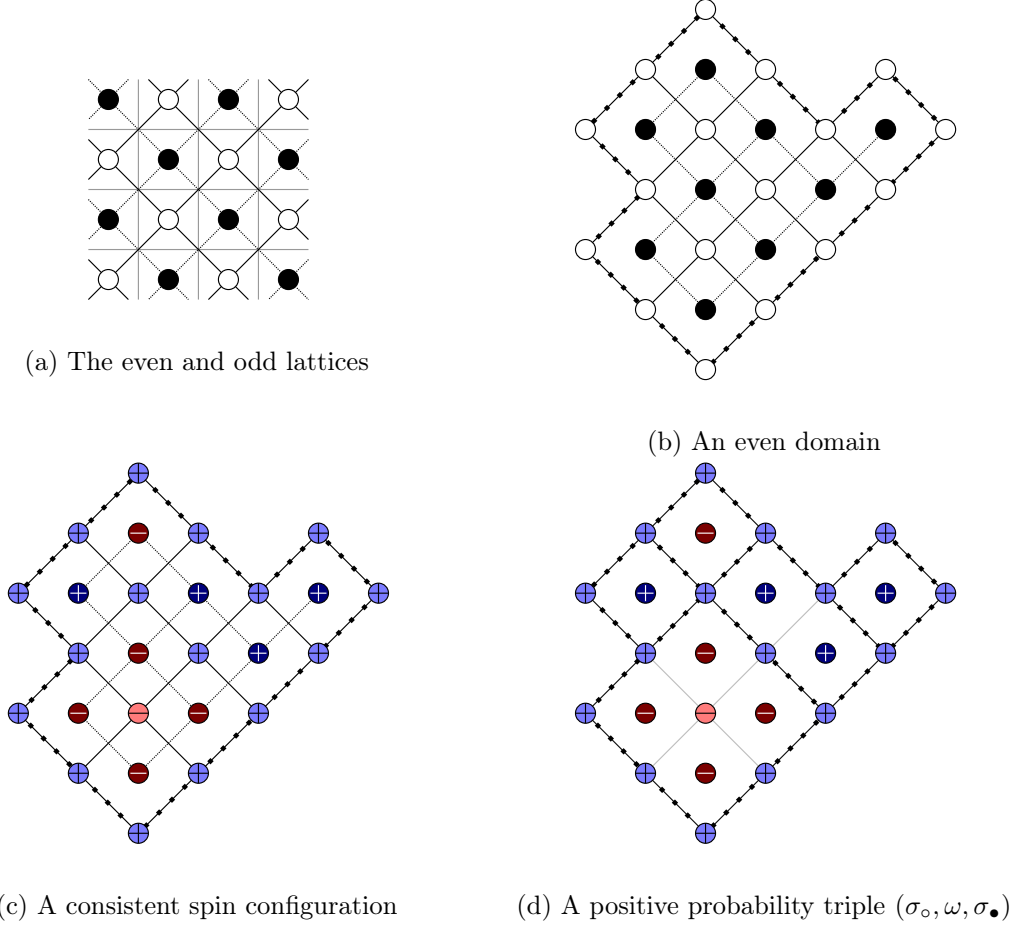


Figure 9

Definition 14.4 (Agreement edges). We also introduce an associated set of *agreement edges*

$$A(\sigma_o) := \{xy \in E_o(\mathcal{D}) : \sigma_o(x) = \sigma_o(y)\}. \quad (235)$$

We introduce the same notation for σ_\bullet . Thus, two spin configurations are consistent if and only if the *complements* of the agreement edges do not cross each other.

In the definitions below, we use the symbol \propto to denote “proportional to;” The reader should bear in mind that a normalization constant is implicitly present to ensure that the measures are probability measures.

We would like to define a probability measure $\mu_{\mathcal{D}}^+$ in which the random pair $(\sigma_o, \sigma_\bullet)$ follows the distribution

$$\mu_{\mathcal{D}}^+[(\sigma_o, \sigma_\bullet)] \propto \mathbf{1}[\sigma_o|_{\partial\mathcal{D}} \equiv +] \cdot \mathbf{1}[\sigma_o \perp \sigma_\bullet] \cdot \mathbf{c}^{\#A(\sigma_o)} \cdot \mathbf{c}^{\#A(\sigma_\bullet)}. \quad (236)$$

We shall prove that this is just a transformation of the six-vertex model with fixed boundary conditions at parameters $\mathbf{a} = \mathbf{b} = 1$ and $\mathbf{c} \in [1, 2]$.

We interpret the above formula as follows: σ_o and σ_\bullet are independent ferromagnetic Ising models (on dual graphs), conditioned on the event that their domain walls do not intersect. The Ising model σ_o comes with fixed boundary conditions, while σ_\bullet comes with free boundary conditions.

After conditioning on σ_o , the distribution of σ_\bullet may be interpreted as an Ising model, except that some of its coupling constants are set to infinity (due to the indicator function

$\mathbb{1}[\sigma_\circ \perp \sigma_\bullet]$). This conditional Ising model has an FK–Ising coupling. Rather than taking Equation (236) as a definition, we shall directly define the probability measure $\mu_{\mathcal{D}}^+$ on a larger space, which also incorporates this FK–Ising coupling. In what follows, the associated percolation configuration will be denoted by ω .

Definition 14.5 (Spin measure). Consider an even domain \mathcal{D} . Let $\Omega = \Omega_{\mathcal{D}}$ denote the sample space

$$\Omega := \{\pm\}^{F_\circ(\mathcal{D})} \times \{0, 1\}^{E_\circ(\mathcal{D})} \times \{\pm\}^{F_\bullet(\mathcal{D})}. \quad (237)$$

A typical element is denoted $(\sigma_\circ, \omega, \sigma_\bullet)$. We often identify ω with the set $\{xy \in E_\circ(\mathcal{D}) : \omega(xy) = 1\}$, and use the standard percolation terminology. Below, for an edge $xy \in E_\circ$, let xy^* be the unique edge in E_\bullet sharing the same middle. Define the probability measure $\mu_{\mathcal{D}}^+$ on $(\sigma_\circ, \omega, \sigma_\bullet) \in \Omega$ by

$$\mu_{\mathcal{D}}^+[(\sigma_\circ, \omega, \sigma_\bullet)] \propto \mathbb{1}[\sigma_\circ|_{\partial\mathcal{D}} \equiv +] \cdot \mathbb{1}[\partial\mathcal{D} \subset \omega] \cdot \mathbb{1}[\omega \subset A(\sigma_\circ)] \quad (238)$$

$$\cdot \mathbb{1}[xy \in \omega \text{ or } xy^* \in A(\sigma_\bullet) \text{ (or both) for any } xy \in E_\circ(\mathcal{D})] \quad (239)$$

$$\cdot \mathbf{c}^{\#A(\sigma_\circ)} \cdot \left(\frac{1}{\mathbf{c}}\right)^{\#\omega} \cdot \left(1 - \frac{1}{\mathbf{c}}\right)^{\#(A(\sigma_\circ) \setminus \omega)}. \quad (240)$$

See Figure 9 for a sample from this measure.

Lemma 14.6. *The definition of $\mu_{\mathcal{D}}^+$ is consistent with Equation (236).*

Proof. For fixed $(\sigma_\circ, \sigma_\bullet)$ one recovers Equation (236) by summing over ω . \square

14.3 Basic properties of the spin representation

Recall that an edge $xy \subset \mathbb{R}^2$ is identified with the straight line segment from x to y . For example, we view ω as a (closed) subset of \mathbb{R}^2 ; each edge is represented by a closed line segment between the two face centres.

Lemma 14.7 (Flip symmetry). *Consider the measure $\mu_{\mathcal{D}}^+$ conditional on (σ_\circ, ω) . Then the distribution of σ_\bullet is given by flipping a fair coin for each bounded connected component of $\mathbb{R}^2 \setminus \omega$.*

Proof. It is immediate from the definition of $\mu_{\mathcal{D}}^+$ that conditionally on (σ_\circ, ω) , the distribution of σ_\bullet is uniform in the set of configurations $\{\pm\}^{F_\bullet(\mathcal{D})}$ satisfying

$$xy^* \in A(\sigma_\bullet) \text{ for any } xy \in E_\circ(\mathcal{D}) \setminus \omega, \quad (241)$$

see also Figure 9. This leads to the distribution stated in the lemma. \square

We now consider the marginal law of the pair (σ_\circ, ω) . This marginal is particularly useful, as we often disregard the odd-spin configuration σ_\bullet . For a set $A \subset \mathbb{R}^2$, let $f(A)$ be the number of connected components of $\mathbb{R}^2 \setminus A$.

Lemma 14.8 $((\sigma_\circ, \omega)$ marginal). *For every $(\sigma_\circ, \omega) \in \{\pm\}^{F_\circ(\mathcal{D})} \times \{0, 1\}^{E_\circ(\mathcal{D})}$,*

$$\mu_{\mathcal{D}}^+[(\sigma_\circ, \omega)] \propto \mathbb{1}[\sigma_\circ|_{\partial\mathcal{D}} \equiv +] \mathbb{1}[\partial\mathcal{D} \subset \omega] \mathbb{1}[\omega \subset A(\sigma_\circ)] 2^{f(\omega)} \mathbf{c}^{\#A(\sigma_\circ) - \#\omega} \left(1 - \frac{1}{\mathbf{c}}\right)^{\#(A(\sigma_\circ) \setminus \omega)}. \quad (242)$$

Proof. The proof is related to the previous lemma. For fixed (σ_\circ, ω) , one needs to count the number of configurations σ_\bullet which satisfy the condition in Equation (241). This number is simply given by $2^{f(\omega)-1}$ (the -1 is due to the fact that the unique unbounded connected component is not involved). This leads to the expression in the lemma. \square

We now state a simple combinatorial observation. To understand the law of (σ_\circ, ω) , it is clearly important to evaluate $f(\omega)$, that is, to count the number of connected components of $\mathbb{R}^2 \setminus \omega$. Since $\omega \subset A(\sigma_\circ)$, each edge in ω connects two spins with the same sign. Thus, we may write $A(\sigma_\circ) = A^-(\sigma_\circ) \cup A^+(\sigma_\circ)$ and $\omega = \omega^- \cup \omega^+$, where:

- $A^-(\sigma_\circ)$ and ω^- connect faces with σ_\circ -spin $-$,
- $A^+(\sigma_\circ)$ and ω^+ connect faces with σ_\circ -spin $+$,
- No face is incident to an edge of $A^-(\sigma_\circ)$ and an edge of $A^+(\sigma_\circ)$,
- No face is incident to an edge of ω^- and an edge of ω^+ .

In particular, we observe that the following formula holds:

$$f(\omega) = -1 + \sum_{\# \in \{+, -\}} f(\omega^\#). \quad (243)$$

We shall also write $\bar{\omega}$ for the ordered pair of two percolation configurations $(\omega_\circ^+, \omega_\circ^-)$. We would like to rewrite the weight in Lemma 14.8 in terms of these new objects. The (long) expression in the following lemma shows that ω_\circ^+ and ω_\circ^- only interact via σ_\circ , that is, they are independent after conditioning on σ_\circ .

Lemma 14.9. *For any $(\sigma_\circ, \bar{\omega}) \in \{\pm\}^{F_\circ(\mathcal{D})} \times \{0, 1\}^{E_\circ(\mathcal{D})} \times \{0, 1\}^{E_\circ(\mathcal{D})}$, we have*

$$\mu_{\mathcal{D}}^+[(\sigma_\circ, \bar{\omega})] \propto \mathbf{c}^{\#A(\sigma_\circ)} \quad (244)$$

$$\times \mathbf{1}[\partial\mathcal{D} \subset \omega_\circ^+] \mathbf{1}[\omega_\circ^+ \subset A^+(\sigma_\circ)] 2^{f(\omega^+)} \left(\frac{1}{\mathbf{c}}\right)^{\#\omega^+} \cdot \left(1 - \frac{1}{\mathbf{c}}\right)^{\#(A^+(\sigma_\circ) \setminus \omega^+)} \quad (245)$$

$$\times \mathbf{1}[\omega_\circ^- \subset A^-(\sigma_\circ)] 2^{f(\omega^-)} \left(\frac{1}{\mathbf{c}}\right)^{\#\omega^-} \cdot \left(1 - \frac{1}{\mathbf{c}}\right)^{\#(A^-(\sigma_\circ) \setminus \omega^-)}. \quad (246)$$

Proof. This follows from Equation (243) and straightforward manipulations. \square

In Equation (243), the simple topology of the plane \mathbb{R}^2 plays an important role; the equation does not immediately generalise to, for example, the torus. When working with the torus, we must therefore slightly modify our setup (see Section 18 and Lemma 18.5).

Until now we defined all our measures with $+$ boundary conditions, but they can equally be defined with $-$ boundary conditions; we write $\mu_{\mathcal{D}}^-$ for this measure.

14.4 Six-vertex height function with fixed boundary conditions

The spin representation is directly related to the six-vertex model. Consider an even domain \mathcal{D} and $(\sigma_\circ, \omega, \sigma_\bullet) \sim \mu_{\mathcal{D}}^+$. Given $(\sigma_\circ, \sigma_\bullet)$, one recovers the gradient of the height function via the definition of the spins at the beginning of this section (Page 55). More precisely, for adjacent faces uv with $u \in F_\circ$ and $v \in F_\bullet$, then

$$h(v) - h(u) = \sigma_\circ(u)\sigma_\bullet(v). \quad (247)$$

This defines the gradient of h on $F(\mathcal{D})$ in the measure $\mu_{\mathcal{D}}^+$. The gradient is turned into a non-gradient function by imposing that it equals 0 on $\partial\mathcal{D}$.

The function h so defined is a *height function* on $F(\mathcal{D})$: an integer-valued function $h : F(\mathcal{D}) \rightarrow \mathbb{Z}$ which differs by exactly ± 1 on adjacent faces and which, as such, preserves the parity of each face. By Lemma 14.6 (Equation (236)), its law is given by:

$$\mu_{\mathcal{D}}^+[h] \propto \mathbf{1}[h|_{\partial\mathcal{D}} \equiv 0] \cdot \mathbf{c}^{\#A(h)}, \quad (248)$$

where $A(h) \subset E_\circ(\mathcal{D}) \cup E_\bullet(\mathcal{D})$ denotes the *agreement diagonals*. The factor $\mathbf{c}^{\#A(h)}$ is consistent with the definition of the six-vertex model (Equation (2)) since $\mathbf{a} = \mathbf{b} = 1$ and since \mathbf{c} -type vertices induce one more agreement diagonal than \mathbf{a} -type or \mathbf{b} -type vertices.

Definition 14.10 (Level lines). From now on, for any $k \in 2\mathbb{Z}$, write

$$A^k = A^k(h) := \{uv \in E_o(\mathcal{D}) : h(u) = h(v) = k\}; \quad (249)$$

$$\omega^k := \omega \cap A^k. \quad (250)$$

The percolation $\omega^k \subset E_o(\mathcal{D})$ is called the *level line* of height k .

Notice that

$$\omega^+ = \bigcup_{k \in 4\mathbb{Z}} \omega^k; \quad \omega^- = \bigcup_{k \in 4\mathbb{Z}+2} \omega^k. \quad (251)$$

Notice that $(\omega^k)_{k \in 2\mathbb{Z}}$ does not just partition ω : this partition also has the property that each connected component of ω is contained in one ω^k . If a connected component α of ω is a subset of ω^k , then $h(x) = k$ for any x incident to α , and we simply say that α has height k .

Before, we saw that the odd spins can only change sign if they are separated by ω -edges (see for example Lemma 14.7). This immediately implies an intermediate value theorem, which also motivates the terminology of *level lines*.

Lemma 14.11 (Intermediate value theorem). *Consider the measure $\mu_{\mathcal{D}}^{\pm}$ in some even domain \mathcal{D} . Fix a target height $k \in 2\mathbb{Z}$ as well as two faces $u, v \in F(\mathcal{D})$ and two heights $a, b \in \mathbb{Z}$ with $a < k < b$. Then almost surely, the following statement holds true: if $h(u) = a$ and $h(v) = b$ and if $\gamma : u \rightarrow v$ is any continuous path in \mathbb{R}^2 from u to v , then γ hits ω^k .*

14.5 Markov property of the spin representation

The spin representation satisfies a Markov property along even domains, a feature that will be fundamental to the analysis below. We present this property in the current section.

If \mathcal{D} and \mathcal{B} are two even domains, then we write $\mathcal{D} \subset \mathcal{B}$, and say that \mathcal{D} is *contained in* \mathcal{B} , whenever $E_o(\mathcal{D}) \subset E_o(\mathcal{B})$. Consider two even domains $\mathcal{D} \subset \mathcal{B}$. Recall the definition of the sample space $\Omega_{\mathcal{D}}$ (Definition 14.5), and write $\pi_{\mathcal{D}} : \Omega_{\mathcal{B}} \rightarrow \Omega_{\mathcal{D}}$ for the natural projection map (which simply erases the values of the spins and edges not relevant to $\Omega_{\mathcal{D}}$). Write $\pi_{\mathcal{D}}^c$ for the complementary projection map, so that $\pi_{\mathcal{D}} \times \pi_{\mathcal{D}}^c$ is the identity map on $\Omega_{\mathcal{D}}$.

Lemma 14.12 (Markov property). *Consider two even domains $\mathcal{D} \subset \mathcal{B}$. For $\#, \flat \in \{\pm\}$ such that $\mu_{\mathcal{B}}^{\#}[\{\partial\mathcal{D} \subset \omega^{\flat}\}] > 0$,*

- *The law $\mu_{\mathcal{B}}^{\#}[\cdot | \{\partial\mathcal{D} \subset \omega^{\flat}\}]$ of $\pi_{\mathcal{D}}(\sigma_{\circ}, \bar{\omega}, \sigma_{\bullet})$ is the same as the law $\mu_{\mathcal{D}}^{\flat}$ of $(\sigma_{\circ}, \bar{\omega}, \sigma_{\bullet})$,*
- *The random variables $\pi_{\mathcal{D}}$ and $\pi_{\mathcal{D}}^c$ are independent in $\mu_{\mathcal{B}}^{\#}[\cdot | \{\partial\mathcal{D} \subset \omega^{\flat}\}]$.*

Proof. Take the expression in Lemma 14.9 and insert an extra indicator for the conditional event. It is then straightforward to work out that the weight factorises over \mathcal{D} and its complement, as desired. \square

Since the distribution of σ_{\bullet} conditional on $(\sigma_{\circ}, \bar{\omega})$ is very simple (Lemma 14.7), it makes sense to focus the analysis entirely on $(\sigma_{\circ}, \bar{\omega})$. From now on, we shall write

$$\Omega_{\mathcal{D}}^{\circ} := \{\pm\}^{F_o(\mathcal{D})} \times \{0, 1\}^{E_o(\mathcal{D})}. \quad (252)$$

For any $\mathcal{D} \subset \mathcal{B}$, we write $\pi_{\mathcal{D}}^{\circ} : \Omega_{\mathcal{B}}^{\circ} \rightarrow \Omega_{\mathcal{D}}^{\circ}$ for the associated projection map, and $\pi_{\mathcal{D}}^{c,\circ}$ for the natural complementary projection map such that $\pi_{\mathcal{D}}^{\circ} \times \pi_{\mathcal{D}}^{c,\circ} = \pi_{\mathcal{B}}^{\circ}$. Finally, we shall often drop the subscript \circ from σ_{\circ} , ω , $\bar{\omega}$, and ω_{\circ}^{\pm} for brevity.

We want to prove one more Markov property, for so-called polar domains. A *polar domain* is an even domain \mathcal{D} together with a partition $\partial\mathcal{D}$ into two segments labelled $\partial^+\mathcal{D}$




Figure 10: A polar domain.

and $\partial^-\mathcal{D}$. More precisely, we impose that the vertices of these segments partition the vertices on $\partial\mathcal{D}$, and that the edges of the segments partition $E(\partial\mathcal{D})$ except that the two edges connecting the endpoints do not belong to any part (see Figure 10).

Lemma 14.13 (Markov property for polar domains). *Consider a polar domain $(\mathcal{D}, \partial^+\mathcal{D}, \partial^-\mathcal{D})$ as well as another even domain $\mathcal{B} \supset \mathcal{D}$. Consider the event*

$$\mathcal{E} := \{\partial^+\mathcal{D} \subset \omega^+ \text{ and } \partial^-\mathcal{D} \subset \omega^-\}. \quad (253)$$

Let $\# \in \{\pm\}$. If $\mu_{\mathcal{B}}^{\#}[\mathcal{E}] > 0$, then in the conditional measure $\nu := \mu_{\mathcal{B}}^{\#}[\cdot | \mathcal{E}]$, there is a Markov property over $\partial\mathcal{D}$ for (σ_{\circ}, ω) . More precisely, $\pi_{\mathcal{D}}^{\circ}$ and $\pi_{\mathcal{D}}^{c,\circ}$ are independent in ν .

Proof. On the event E , the values of σ_{\circ} on $\partial\mathcal{D}$ are known. Thus, the only factor in the weight of Lemma 14.9 that may make $\pi_{\mathcal{D}}^{\circ}$ and $\pi_{\mathcal{D}}^{c,\circ}$ interact, is the factor

$$\prod_{b \in \pm} 2^{f(\omega^b)}. \quad (254)$$

Notice that on the event E , we have

$$f(\omega^b) = 1 + \#\text{conn. comp. of } \mathbb{R}^2 \setminus \omega^b \text{ which are entirely surrounded by } \partial\mathcal{D} \quad (255)$$

$$+ \#\text{conn. comp. of } \mathbb{R}^2 \setminus \omega^b \text{ which are entirely outside } \partial\mathcal{D}. \quad (256)$$

The first term on the right is $\pi_{\mathcal{D}}^{\circ}$ -measurable, and the second term is $\pi_{\mathcal{D}}^{c,\circ}$ -measurable, proving the desired factorisation and independence. \square

15 FKG inequality of the spin representation

We now state and prove the Fortuin-Kasteleyn-Ginibre (FKG) inequality. We also discuss several of its immediate consequences.

Definition 15.1 (Increasing functions and the FKG property). A random variable is called \circ -increasing (or simply *increasing*) if it may be written as an increasing function of the triple $(\sigma_{\circ}, \omega^+, -\omega^-)$. Notice the minus sign in the last entry; this means that ω_{\circ}^- -open edges are lower in the partial order on such triples. A random variable X is called \circ -decreasing (or simply *decreasing*) whenever $-X$ is \circ -increasing. An event is called \circ -increasing or \circ -decreasing if its indicator function is \circ -increasing or \circ -decreasing respectively.

A probability measure μ is said to have \circ -FKG (or simply *FKG*) whenever

$$\text{Cov}_{\mu}[X, Y] := \mu[XY] - \mu[X]\mu[Y] \geq 0 \quad (257)$$

for any bounded \circ -increasing random variables X and Y which are measurable in terms of finitely many spins and edges.

The following statement captures the essence of this subsection.

Lemma 15.2. *The measures $\mu_{\mathcal{D}}^+$ and $\mu_{\mathcal{D}}^-$ have \circ -FKG for any even domain \mathcal{D} .*

This lemma was proved in [GL25], but we shall state and prove a more general version of it. More precisely, we shall define *lattice events*, which are events \mathcal{E} with the property that conditioning on \mathcal{E} preserves FKG. Many events of interest can be partitioned into lattice events, which is very useful when applying the FKG inequality.

Definition 15.3 (Lattice event). A \circ -lattice event (or *lattice event*) is an event of the form

$$\{\omega|_Q = \zeta\} \cap \{\sigma_{\circ}|_B = \xi\}, \quad (258)$$

where B and Q contain vertices and edges of F_{\circ} respectively, and such that B contains all endpoints of edges in Q . Naturally, we require that $\zeta \in \{0, 1\}^Q$ and $\xi \in \{\pm\}^B$. Notice that if $Q = B = \emptyset$, then \mathcal{E} is just the entire sample space.

Proposition 15.4 (FKG inequality). *Consider an even domain \mathcal{D} and a \circ -lattice event \mathcal{E} . Let $\# \in \{\pm\}$. If $\mu_{\mathcal{D}}^{\#}[\mathcal{E}] > 0$, then the conditional probability measure $\mu_{\mathcal{D}}^{\#}[\cdot|\mathcal{E}]$ has \circ -FKG.*

We closely follow [LO24, Theorem 2.8] and [GL25, Proposition 4.9]. The reader may choose to skip this technical proof on a first read. We focus on the case $\# = +$; the proof of the other case is identical. The proof of \circ -FKG follows from the following lemma which is proved below.

Lemma 15.5. *Fix a domain \mathcal{D} and a lattice event \mathcal{E} such that $\mu_{\mathcal{D}}^+[\mathcal{E}] > 0$. Then, all of the following hold true in the conditional measure $\nu := \mu_{\mathcal{D}}^+[\cdot|\mathcal{E}]$.*

- (i) *The weights of σ_{\circ} satisfy the FKG lattice condition, which implies the FKG inequality for σ_{\circ} [FKG71].*
- (ii) *Conditional on σ_{\circ} , the percolations ω^+ and ω^- are independent.*
- (iii) *Conditional on σ_{\circ} , the law of ω^+ satisfies the FKG inequality.*
- (iv) *The conditional law of ω^+ is stochastically increasing in σ_{\circ} .*
- (v) *Conditional on σ_{\circ} , the law of ω^- satisfies the FKG inequality.*
- (vi) *The conditional law of ω^- is stochastically decreasing in σ_{\circ} .*

Before proving the lemma, let us derive the FKG inequality.

Proof of Proposition 15.4. As mentioned, we focus on the $+$ case. Fix \mathcal{D} and \mathcal{E} such that $\mu_{\mathcal{D}}^+[\mathcal{E}] > 0$ and set $\nu = \mu_{\mathcal{D}}^+[\cdot|\mathcal{E}]$. We apply the so-called *tower property* for the FKG inequality.

Lemma 15.5(ii)–(vi) implies that conditionally on σ_{\circ} , the pair $(\omega^+, -\omega^-)$ satisfies the FKG inequality, and that the conditional law of the pair $(\omega^+, -\omega^-)$ is stochastically increasing in σ_{\circ} .

Now, let X and Y denote two bounded \circ -increasing functions. Assert that

$$\nu[XY] = \nu[\nu[XY|\sigma_{\circ}]] \geq \nu[\nu[X|\sigma_{\circ}]\nu[Y|\sigma_{\circ}]] \geq \nu[\nu[X|\sigma_{\circ}]]\nu[\nu[Y|\sigma_{\circ}]] = \nu[X]\nu[Y]. \quad (259)$$

This standard trick (see [LO24]) is proved as follows. The tower property implies the two equalities. The first inequality is the conditional FKG of the pair $(\omega^+, -\omega^-)$. Since the law of $(\omega^+, -\omega^-)$ is stochastically increasing in σ_{\circ} , we see that $\nu[X|\sigma_{\circ}]$ and $\nu[Y|\sigma_{\circ}]$ are increasing functions of σ_{\circ} . The second inequality then follows from the FKG inequality for σ_{\circ} (Lemma 15.5(i)). \square

Proof of Lemma 15.5. We omit the subscript and write $\sigma = \sigma_{\circ}$. To prove the lemma, we first find an appropriate decomposition of ν and then derive the items one by one.

Step 1: Decomposition of ν . Since $\mathcal{E} = \{\omega|_Q = \zeta\} \cap \{\sigma|_B = \tau\}$ has positive probability, the edge set Q may be written as the disjoint union

$$Q^+ \cup Q^- \cup Q^\emptyset, \quad (260)$$

where $Q^\emptyset = \{\zeta = 0\}$ and where τ is equal to $+1$ on the endpoints of Q^+ and to -1 on the endpoints of Q^- . Without loss of generality, $\partial\mathcal{D} \subset Q^+$. Inserting the indicator functions for the conditioning event in the expression of Lemma 14.9 yields

$$\nu[(\sigma, \bar{\omega})] \propto \mathbb{1}[\sigma|_B = \tau] \cdot \mathbf{c}^{\#A(\sigma)} \cdot 2^{f(\omega^+)} Y_+(\sigma, \omega^+) \cdot 2^{f(\omega^-)} Y_-(\sigma, \omega^-); \quad (261)$$

where

$$Y_\pm(\sigma, \omega^\pm) := \mathbb{1}[Q^\pm \subset \omega^\pm] \mathbb{1}[\omega^\pm \cap Q^\emptyset = \emptyset] \mathbb{1}[\omega^\pm \subset A^\pm(\sigma)] \cdot \left(\frac{1}{\mathbf{c}}\right)^{\#(\omega^\pm \cap (A^\pm(\sigma) \setminus Q))} \left(1 - \frac{1}{\mathbf{c}}\right)^{\#((A^\pm(\sigma) \setminus Q) \setminus \omega^\pm)}. \quad (262)$$

Step 2: Proof of (ii). For fixed σ , the above weight may be written as a product of one factor depending only on ω^+ , and one depending only on ω^- . This implies the desired independence.

Step 3: Proof of (i). Proving (i) is the most delicate part of the argument; once it is established, the remaining claims follow easily. To compute the weight of a spin configuration σ , we must sum over ω^+ and ω^- in Equation (261). Since in this expression only $2^{f(\omega^+)} Y_+(\sigma, \omega^+)$ depends on ω^+ , and only $2^{f(\omega^-)} Y_-(\sigma, \omega^-)$ on ω^- , the sums may be carried out separately. This is almost the same for ω^+ and ω^- . We first focus on ω^+ , and then explain how to adapt the argument for ω^- .

The key idea is the following. Conditional on σ , the percolation configuration ω^+ behaves like the dual of a random-cluster model on $F_\bullet(\mathcal{D})$ with cluster weight $q = 2$, corresponding to an Ising model on the odd faces via the Edwards–Sokal coupling. The law and partition function of both the random-cluster model and the Ising model are well understood, and this correspondence allows us to verify the claims below. For completeness, we give full detail.

Introduce the couplings

$$a_{uv}^\pm(\sigma) = \begin{cases} 1 & uv^* \in Q^\pm, \\ 1/\mathbf{c} & uv^* \in A^\pm(\sigma) \setminus Q, \\ 0 & \text{otherwise.} \end{cases} \quad (263)$$

and the Ising model (on $F_\bullet(\mathcal{D})$) partition function $\mathcal{Z}_{\text{Ising}}(a)$ with couplings a defined by

$$\mathcal{Z}_{\text{Ising}}(a) = \sum_{\tilde{\sigma} \in \{\pm 1\}^{F_\bullet(\mathcal{D})}} \prod_{uv \in E_\bullet(\mathcal{D})} (a_{uv})^{\mathbb{1}[\tilde{\sigma}(u) \neq \tilde{\sigma}(v)]}. \quad (264)$$

Our objective is to compute $\sum_{\omega^+} 2^{f(\omega^+)} Y_+(\sigma, \omega^+)$. First, rewrite

$$2^{f(\omega^+)} = 2 \sum_{\tilde{\sigma} \in \{\pm 1\}^{F_\bullet(\mathcal{D})}} \mathbb{1}[\tilde{\sigma} \perp \omega^+], \quad (265)$$

where $\tilde{\sigma} \perp \omega^+$ means that $\tilde{\sigma}$ is constant on each connected component of $\mathbb{R}^2 \setminus \omega^+$. The prefactor two on the right compensates for the unbounded face counted in $f(\omega^+)$. Inserting the previous formula and exchanging the sums yields

$$\sum_{\omega^+} 2^{f(\omega^+)} Y_+(\sigma, \omega^+) = 2 \sum_{\tilde{\sigma} \in \{\pm 1\}^{F_\bullet(\mathcal{D})}} \sum_{\omega^+, \tilde{\sigma} \perp \omega^+} Y_+(\sigma, \omega^+). \quad (266)$$

Although the expression for Y_+ is lengthy, the inner sum (over $\omega^+ \perp \tilde{\sigma}$) is easy to compute. In fact, the indicators and the requirement $\omega^+ \perp \tilde{\sigma}$ simply tell us that some edges must be open or closed. The sum over the remaining edges can be performed independently. By carrying out this computation, carefully bookkeeping the conditions on ω^+ , one obtains

$$\sum_{\omega^+} 2^{f(\omega^+)} Y_+(\sigma, \omega^+) = 2\mathcal{Z}_{\text{Ising}}(a^+(\sigma)), \quad (267)$$

Similarly, one obtains

$$\sum_{\omega^-} 2^{f(\omega^-)} Y_-(\sigma, \omega^-) = \mathcal{Z}_{\text{Ising}}(a^-(\sigma)) \quad (268)$$

(contrary to Equation (267), we do not need the prefactor two, as the unbounded face of $\mathbb{R}^2 \setminus \omega^-$ intersects $F_\bullet(\mathcal{D})$ and therefore its sign is already accounted for). Putting the two expressions together yields

$$\nu[\{\sigma\}] \propto \mathbb{1}[\sigma|_{\partial\mathcal{D}} \equiv +] \cdot \mathbb{1}[\sigma|_B = \tau] \cdot \mathbf{c}^{\#A(\sigma)} \cdot \mathcal{Z}_{\text{Ising}}(a^+(\sigma)) \cdot \mathcal{Z}_{\text{Ising}}(a^-(\sigma)). \quad (269)$$

We now verify that each factor satisfies the FKG lattice condition:

$$g(\sigma \vee \sigma')g(\sigma \wedge \sigma') \geq g(\sigma)g(\sigma'), \quad (270)$$

where $\sigma \vee \sigma'$ and $\sigma \wedge \sigma'$ denote the pointwise maximum and minimum of σ and σ' .

The first two factors trivially satisfy the condition as they are constant (recall that we are interested in configurations with positive $\nu = \nu_{\mathcal{D}}^+[\cdot|\mathcal{E}]$). The third one is classical: the map $\sigma \mapsto \mathbf{c}^{\#A(\sigma)}$ corresponds to ferromagnetic Ising interactions and it is straightforward to check the desired inequality. We now handle $\mathcal{Z}_{\text{Ising}}(a^\pm(\sigma))$. We focus on $+$, the case $-$ being similar. The definition of a^+ (Equation (263)) and the inequality $\mathbf{c} \geq 1$ imply that

$$a_{uv}^+(\sigma \vee \sigma') \geq a_{uv}^+(\sigma) \vee a_{uv}^+(\sigma'); \quad a_{uv}^+(\sigma \wedge \sigma') = a_{uv}^+(\sigma) \wedge a_{uv}^+(\sigma'). \quad (271)$$

Suppose now for a second that the map

$$a \mapsto \mathcal{Z}_{\text{Ising}}(a) \quad (272)$$

satisfies the FKG lattice condition over $a \in [0, 1]^{E_\bullet(\mathcal{D})}$. Since $\mathcal{Z}_{\text{Ising}}(a)$ is increasing in a , Equation (271) and the FKG lattice condition for (272) imply the desired Equation (270) for $g(\sigma) = \mathcal{Z}_{\text{Ising}}(a^+(\sigma))$. The FKG lattice condition for Equation (272) is classical (after renormalising the partition function in a way that does not affect the FKG lattice condition), and may be found in [Frö82, Equations (35), (36), (38)], [Cha98, Proposition A.1], [LO24, Lemma 6.1 and Equation (7)], or [GL25].

We have established that all factors satisfy the FKG lattice condition, and therefore the same holds true for their product $\nu[\{\sigma\}]$. By [FKG71], this implies that σ satisfies the FKG property under ν . This completes the proof of (i).

Proof of (iii)–(vi). Let us start with (iii) and (v). Conditionally on σ , the percolation ω^+ is the dual of a random-cluster model with cluster weight $q = 2$, which is well-known to satisfy the FKG inequality. The same holds true for ω^- .

We now prove (iv); (vi) being derived similarly. Conditionally on σ , the coupling constants of the random-cluster model are encoded in $a^+(\sigma)$. Notice that $a^+(\sigma)$ (the inverse of the coupling strengths) is *increasing* in σ . This means that law of ω^+ , which is the dual of the random-cluster model, is stochastically *increasing* in σ . This proves (iv) and concludes the proof. \square

We now state a classical consequence of the monotonicity properties established above, enabling to “push” domains “away”, and to compare boundary conditions. It is used to circumvent the lack of independence in the model. Similar statements can be found in the theory of Ising and random-cluster models. This result will play a central role in the probabilistic analysis of the spin configuration that follows.

Corollary 15.6 (Monotonicity in domains and boundary conditions). *Consider two even domains $\mathcal{D} \subset \mathcal{B}$ as well as a bounded \circ -increasing random variable X . Then*

$$\mu_{\mathcal{D}}^{-}[X] \leq \mu_{\mathcal{B}}^{-}[X] \leq \mu_{\mathcal{B}}^{+}[X] \leq \mu_{\mathcal{D}}^{+}[X]. \quad (273)$$

Proof. Let \mathcal{M} be an even domain containing \mathcal{B} which is so big that $\partial\mathcal{B}$ does not intersect $\partial\mathcal{M}$. Then the two events $\{\partial\mathcal{B} \subset \omega^{\pm}\}$ have a positive probability in $\mu_{\mathcal{M}}^{+}$. We claim that

$$\mu_{\mathcal{B}}^{+}[X] = \mu_{\mathcal{M}}^{+}[X | \partial\mathcal{B} \subset \omega^{+}] \geq \mu_{\mathcal{M}}^{+}[X]. \quad (274)$$

Indeed, the equality is the Markov property, and the inequality is the FKG inequality. This proves the claim. Similarly, we get

$$\mu_{\mathcal{B}}^{-}[X] = \mu_{\mathcal{M}}^{+}[X | \partial\mathcal{B} \subset \omega^{-}] \leq \mu_{\mathcal{M}}^{+}[X], \quad (275)$$

leading to

$$\mu_{\mathcal{B}}^{-}[X] \leq \mu_{\mathcal{B}}^{+}[X]. \quad (276)$$

The other inequalities in the statement of the lemma are similar; for example, on the right side, we obtain

$$\mu_{\mathcal{B}}^{+}[X] \leq \mu_{\mathcal{B}}^{+}[X | \partial\mathcal{D} \subset \omega^{+}] = \mu_{\mathcal{D}}^{+}[X]. \quad (277)$$

This proves the lemma. \square

16 Percolation estimates for the spin configuration

From now on, we shall consistently write

$$\llbracket n \rrbracket := [-n, n]; \quad \llbracket n \times m \rrbracket := [-n, n] \times [-m, m]; \quad (278)$$

$$\text{Ball}_n := \llbracket n \times n \rrbracket; \quad \text{Ball}_n(x) := \text{Ball}_n + x; \quad (279)$$

$$\text{Annulus}_{R,r} := \text{Ball}_R \setminus \text{Ball}_r; \quad \text{Annulus}_{R,r}(x) := \text{Annulus}_{R,r} + x. \quad (280)$$

A *rectangle* is a subset of \mathbb{R}^2 of the form $R = \llbracket a \times b \rrbracket + x$. Its four sides (closed line segments which are subsets of ∂R) are denoted Right_R , Top_R , Left_R , and Bottom_R in the obvious way.

16.1 Percolation events associated with the spin representation

Recall that we view ω^{+} and ω^{-} as random subsets of \mathbb{R}^2 (each edge is viewed as the closed line segment between its two endpoints). Our objective is to understand the random geometry of $\bar{\omega} = (\omega^{+}, \omega^{-})$. In this section we define some useful percolation events: first “simple” events (defined in terms of either ω^{+} or ω^{-}), then “alternating” events (defined in terms of both ω^{+} and ω^{-}).




Figure 11: An element of the generic percolation event $\{A \stackrel{D}{\leftrightarrow} B\}$




Figure 12: Elements of $\text{Hor}(R)$, $\text{Ver}(R)$, $\text{Circuit}(D \setminus H)$, and $\text{Arm}(D \setminus H)$

Definition 16.1 (Simple percolation events). Define

$$\{A \stackrel{D}{\leftrightarrow} B\} := \{S \subset \mathbb{R}^2 : S \cap D \text{ contains a path from } A \text{ to } B\}; \quad (281)$$

see Figure 11. Define the *horizontal* and *vertical crossings* of rectangles to be the sets

$$\text{Hor}(R) := \{\text{Left}_R \stackrel{R}{\leftrightarrow} \text{Right}_R\}; \quad (282)$$

$$\text{Ver}(R) := \{\text{Top}_R \stackrel{R}{\leftrightarrow} \text{Bottom}_R\}. \quad (283)$$

If $A \subset \mathbb{R}^2$ is a topological annulus, define the *circuit* and *arm* events

$$\text{Circuit}(A) := \{S \subset \mathbb{R}^2 : S \cap A \text{ contains a non-contractible circuit in } \mathbb{R}^2 \setminus A\}; \quad (284)$$

$$\text{Arm}(A) := \{S \subset \mathbb{R}^2 : S \cap A \text{ has a path connecting the boundary components of } A\}. \quad (285)$$

See Figure 12 for an illustration of all four events we just defined.

Remark 16.2 (Combining circuit events). Circuit events are versatile as they can be combined to create (horizontal or vertical) crossings (see Figure 13); for example, for $k \geq 1$, we have

$$\bigcap_{i=0}^{k-1} \text{Circuit}(\text{Annulus}_{2,1}((2i, 0))) \subset \text{Hor}([-1, 2k-1] \times [1, 2]). \quad (286)$$

They can also be combined to create more complicated percolation events. For example, for $r \in \mathbb{Z}_{\geq 1}$, we may find a subset $Z \subset \mathbb{R}^2$ of cardinal $\#Z = 4r$ such that

$$\bigcap_{z \in Z} \text{Circuit}(\text{Annulus}_{2,1}(z)) \subset \text{Circuit}(\text{Annulus}_{r+1,r}). \quad (287)$$




Figure 14: Elements of $\text{AltHor}_4(R)$, $\text{AltVer}_4(R)$, $\text{AltCircuit}_4(D \setminus H)$, and $\text{AltArm}_4(D \setminus H)$. Full lines depict ω^+ ; dotted lines ω^- .

references therein) that at criticality, rectangles with a fixed aspect ratio have crossing probabilities – that is, the probability that the percolation configuration contains a connected component crossing the rectangle – that remain uniformly bounded away from both 0 and 1 as the size of the rectangle tends to infinity. This reflects the fact that connected components in the critical regime qualitatively exhibit scale-free behaviour and possess fractal-like geometric properties.

Originally developed for Bernoulli percolation, the Russo–Seymour–Welsh (RSW) theory has become indispensable in the analysis of critical phenomena. Over the past fifteen years, the theory has been significantly extended to encompass many dependent percolation models [DHN11; DST17; DT15; DTT18; DMT21; DT19; Lam23; KT23]. In this broader context, particular attention must be paid to the dependence between the configuration inside a given rectangle and its exterior. To be applicable, crossing estimates must therefore be uniform with respect to boundary conditions.

When studying the emergence of large connected components, an especially convenient geometric setting is that of an annulus rather than a rectangle. In this setting, one asks whether the percolation configuration contains a path that remains within the annulus and encircles its inner boundary. We adopt this annular framework throughout the present section.

The subsection states the key RSW-type input for our percolation-type arguments. The theorem below yields that ω^- satisfies RSW-type estimates, even when boundary conditions are least favourable (that is, + boundary conditions).

For any rectangle R , we shall write μ_R^+ for $\mu_{\mathcal{D}}^+$, where \mathcal{D} is the *largest even domain* whose face centres are all contained in R .

Theorem 16.4 (Circuit estimate). *There exists a constant $c_{\text{circuit}} > 0$ (independent of $\mathbf{c} \in [1, 2]$) such that for any $r \in [4, \infty)$ and $x \in \mathbb{R}^2$,*

$$\mu_{\text{Ball}_{2r}(x)}^+[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x))] \geq c_{\text{circuit}}. \quad (292)$$

The same inequality remains true with + and – interchanged.

The result tells us that ω^- -circuits have a uniformly positive probability of appearing, even with the *worst possible* boundary conditions (in the sense of increasing functions). The uniformity in boundary conditions renders the result very flexible.

The theorem was first proved in [Dum+24, Theorem 1.4]. The article uses an input coming from the Bethe Ansatz to derive certain crossing estimates, then uses the FKG inequality (for the absolute value of the height function) to turn these crossing estimates into circuit estimates.

A second proof not relying on the Bethe Ansatz was given in [GL25], via a *renormalisation inequality*. This technique was first used in [DST17] in the context of the random-cluster model. The renormalisation inequality essentially asserts that if the theorem is *false*, then the circuit probability decays exponentially fast in the radius r , which in turn implies that the model is localised (contradicting the known delocalisation).

We now turn to the proof. While Theorem 16.4 is essentially identical to [Dum+24, Theorem 1.4]; the two theorems are stated in slightly different settings. We will now “translate” [Dum+24, Theorem 1.4] into Theorem 16.4. The reader may choose to skip this technical proof on a first read. We start with a lemma.

Recall that the event $\text{AltCircuit}_{20}^{\bar{\omega}}(\text{Annulus}_{2n,n})$ means that the annulus contains 20 circuits which (from outside to inside) alternately belong to ω^+ , ω^- , ω^+ , etc.

Lemma 16.5. *There exist constants $c > 0$ and $N \in \mathbb{Z}_{\geq 0}$ (both independent of $\mathbf{c} \in [1, 2]$) such that for every $n \geq N$,*

$$\mu_{\text{Ball}_{(2n+2)}}^+[\text{AltCircuit}_{20}^{\bar{\omega}}(\text{Annulus}_{2n,n})] \geq c. \quad (293)$$

Proof. The first translation problem stems from the fact that boundary conditions are enforced differently in [Dum+24]. In [Dum+24], boundary conditions are imposed on the heights of an $F(\mathbb{Z}^2)$ -path of *adjacent* faces, as opposed to an F_{\circ} -path of *diagonally adjacent even faces* which is what we do here.

Fix $n \in \mathbb{Z}_{\geq 0}$ and let \mathcal{D}^n denote the largest even domain whose face centres are contained in $\text{Ball}_{(2n+2)}$. Let γ^n denote the self-avoiding $F(\mathcal{D}^n)$ -circuit which surrounds the largest area (this path alternately visits even and odd faces). It is straightforward to see that, as a subset of \mathbb{R}^2 , Ball_{2n} is surrounded by γ^n .

Since each face visited by γ^n lies on $\partial\mathcal{D}^n$ or is adjacent to it, we have

$$\mu_{\text{Ball}_{(2n+2)}}^+[\{|h|_{\gamma^n} \leq 1\}] = 1. \quad (294)$$

This is good news, because the conditional measure $\mu_{\text{Ball}_{(2n+2)}}^+[\cdot | h|_{\gamma^n}]$ has boundary conditions like in [Dum+24]. We may now directly apply [Dum+24, Theorem 1.4] with $\ell = 1$ and $k = 100$, to find a constant $c > 0$ such that for sufficiently large n ,

$$\mu_{\text{Ball}_{(2n+2)}}^+[\text{Circuit}^{\{h \geq 100\}}(\text{Annulus}_{2n,n}) | h|_{\gamma^n}] \geq c. \quad (295)$$

In particular,

$$\mu_{\text{Ball}_{(2n+2)}}^+[\text{Circuit}^{\{h \geq 100\}}(\text{Annulus}_{2n,n})] \geq c. \quad (296)$$

Since $\text{Annulus}_{2n,n}$ contains a circuit of $\{h \leq 10\}$ almost surely, the intermediate value theorem asserts that we may find ω^k -circuits at heights $k = 12, 14, \dots, 98$ in the annulus. This implies the claim since the circuits alternately belong to ω^+ and ω^- . \square

Proof of Theorem 16.4. Let N and c be provided by Lemma 16.5. We distinguish two cases: either $r \leq N$ or $r > N$.

Step 1: the case $r \leq N$. In any measure $\mu_{\mathcal{D}}^+$, every edge not incident to $\partial\mathcal{D}$ is ω^- -open with some uniformly positive probability $\eta > 0$ independent of \mathcal{D} and the chosen edge. Thus, using the FKG inequality for ω^- , a uniform lower bound is given by η^M where M is the supremum over the minimal circuit lengths within annuli with $r \leq N$.

Step 2: the case $r > N$. Without loss of generality, $x \in [-2, 2]^2$. Let $n = \lceil r \rceil$. Then

$$\text{AltCircuit}_{20}^{\bar{\omega}}(\text{Annulus}_{2n,n}) \subset \text{AltCircuit}_2^{\bar{\omega}}(\text{Annulus}_{2r,r}(x)). \quad (297)$$

Thus, the claim implies that

$$\mu_{\text{Ball}(2n+2)}^+[\text{AltCircuit}_2^{\bar{\omega}}(\text{Annulus}_{2r,r}(x))] \geq c. \quad (298)$$

Let γ denote the largest self-avoiding ω^+ -circuit in the annulus above (or $\gamma := \emptyset$ if such a circuit does not exist), and let $\mathcal{D}_\gamma \subset \text{Ball}_{2r}(x)$ denote the corresponding even domain (so that $\partial\mathcal{D}_\gamma = \gamma$). By the tower property and the Markov property, we have

$$\begin{aligned} & \mu_{\text{Ball}(2n+2)}^+[\text{AltCircuit}_2^{\bar{\omega}}(\text{Annulus}_{2r,r}(x))] \\ &= \int_{\{\gamma \neq \emptyset\}} \mu_{\mathcal{D}_\gamma}^+[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x))] d\mu_{\text{Ball}(2n+2)}^+[\gamma] \geq c. \end{aligned} \quad (299)$$

Since $\mu_{\mathcal{D}}^+[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x))]$ is increasing in \mathcal{D} and $\mathcal{D}_\gamma \subset \text{Ball}_{2r}(x)$ almost surely, we get from Corollary 15.6 that

$$\mu_{\text{Ball}_{2r}(x)}^+[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x))] \geq c. \quad (300)$$

This concludes the proof of the second step. Taken together, the two steps establish the theorem with $c_{\text{circuit}} := \eta^M \wedge c$. \square

We record an easy consequence of the previous theorem together with the FKG inequality. This statement will prove convenient later on.

Corollary 16.6 (Circuit estimate). *For any $r \in [4, \infty)$ and $x \in \mathbb{R}^2$, and for any even domain \mathcal{D} such that $\partial\mathcal{D}$ surrounds $\text{Ball}_{2r}(x)$, we have*

$$\mu_{\mathcal{D}}^+[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x))|\mathcal{E}] \geq c_{\text{circuit}} \quad (301)$$

for any event \mathcal{E} that is measurable in terms of the even spins and edges which do not intersect $\text{Ball}_{2r}(x)$. The same inequality remains true with ω^- replaced by ω^+ .

Proof. The event \mathcal{E} may be written as a partition of finitely many \circ -lattice events which are measurable in terms of the even spins and edges which do not intersect $\text{Ball}_{2r}(x)$. Without loss of generality, we may assume that \mathcal{E} itself is of this type. Let $\gamma \subset E_\circ$ denote the boundary of the largest even domain in $\text{Ball}_{2r}(x)$. By the FKG inequality (Proposition 15.4) for $\mu_{\mathcal{D}}^+[\cdot|\mathcal{E}]$, we get

$$\mu_{\mathcal{D}}^+[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x))|\mathcal{E}] \geq \mu_{\mathcal{D}}^+[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x))|\mathcal{E} \cap \{\gamma \subset \omega^+\}] \quad (302)$$

$$= \mu_{\text{Ball}_{2r}(x)}^+[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x))] \quad (303)$$

$$> c_{\text{circuit}}. \quad (304)$$

The equality is just the Markov property; the inequality on the right is Theorem 16.4. \square

17 Full-plane spin representation (incl. Theorem 4.10)

The previous circuit estimate enables us to define a full plane analogue of the spin representation.

17.1 Definition in the full plane

We define an infinite volume version of our spin representation. We also show that the properties obtained in finite volume extend to the infinite volume setting. Introduce the set

$$\Omega_{\mathbb{Z}^2} := \{\pm\}^{F_\circ} \times \{0, 1\}^{E_\circ} \times \{\pm\}^{F_\bullet} \quad (305)$$

Theorem 17.1 (Full-plane limit). *There exists a unique probability measure $\mu_{\mathbb{Z}^2}$ on $\Omega_{\mathbb{Z}^2}$ such that for any random variable X measurable in terms of finitely many spins and edges,*

$$\lim_{\mathcal{D} \nearrow \mathbb{Z}^2} \mu_{\mathcal{D}}^+[X] = \lim_{\mathcal{D} \nearrow \mathbb{Z}^2} \mu_{\mathcal{D}}^-[X] = \mu_{\mathbb{Z}^2}[X]. \quad (306)$$

Moreover, $\mu_{\mathbb{Z}^2}$ satisfies the following properties:

- (i) $\mathbb{R}^2 \setminus \omega$ contains $\mu_{\mathbb{Z}^2}$ -almost surely no unbounded connected components.
- (ii) Conditional spin flip property for σ^\bullet (analogous to Lemma 14.7),
- (iii) Markov property (analogous to Lemma 14.12),
- (iv) Markov property for polar domains (analogous to Lemma 14.13),
- (v) FKG inequality (analogous to Proposition 15.4),
- (vi) Circuit estimate (analogous to Corollary 16.6).

Remark 17.2. We have not yet proved well-definedness of $\mathbb{P}_{\mathbb{Z}^2}$ (Theorem 2.2). This is done later, in Section 18, where we also prove that the law of the gradient of h in $\mu_{\mathbb{Z}^2}$ is precisely $\mathbb{P}_{\mathbb{Z}^2}$. This motivates our interest in the measure $\mu_{\mathbb{Z}^2}$.

Before proving Theorem 17.1, we introduce a lemma.

Lemma 17.3. *The measures $\mu_{\mathcal{D}}^+$ and $\mu_{\mathcal{D}}^-$ restricted to even spins and edges converge to the same limit, which we denote $\mu_{\mathbb{Z}^2, \circ}$.*

Proof. Corollary 15.6 implies that for any \circ -increasing positive random variable X which is measurable in terms of finitely many spins and edges, the limits

$$\ell^- := \lim_{\mathcal{D} \nearrow \mathbb{Z}^2} \mu_{\mathcal{D}}^-[X]; \quad \ell^+ := \lim_{\mathcal{D} \nearrow \mathbb{Z}^2} \mu_{\mathcal{D}}^+[X] \quad (307)$$

are well-defined (respectively as increasing and decreasing sequences) and satisfy $\ell^- \leq \ell^+$. It therefore suffices to prove the other inequality $\ell^+ \leq \ell^-$. This will be done by showing that even under $\mu_{\mathcal{D}}^+[X]$, there is a probability tending to 1 as \mathcal{D} tends to \mathbb{Z}^2 (this uses Corollary 16.6) that there is a circuit in ω^- surrounding the vertices that serve to measure X .

More formally, fix $r \geq 4$ so large that X is measurable in terms of the spins and edges in Ball_{r-2} . Define the event $\mathcal{A}_{n,m} := \text{Circuit}^{\omega^-}(\text{Ball}_{2^m r} \setminus \text{Ball}_{2^n r})$. Fix $m \in \mathbb{Z}_{\geq 0}$. If \mathcal{D} is an even domain whose perimeter surrounds $\text{Ball}_{2^m r}$, then Corollary 16.6 implies that

$$\mu_{\mathcal{D}}^+[\mathcal{A}_{n,n+1} | (\mathcal{A}_{n+1,m})^c] \geq c_{\text{circuit}} \quad (308)$$

for any $n = 0, 1, \dots, m-1$. By induction, this yields

$$\mu_{\mathcal{D}}^+[\mathcal{A}_{0,m}] \geq 1 - (1 - c_{\text{circuit}})^m. \quad (309)$$

The right-hand side tends to 1 as $m \rightarrow \infty$. Let γ^m denotes the largest self-avoiding ω^- -circuit contributing to $\mathcal{A}_{0,m}$ (if the event occurs), and set $\gamma^m := \emptyset$ otherwise. Then

$$\mu_{\mathcal{D}}^+[\mathbb{1}_{\mathcal{A}_{0,m}} \cdot X] = \int_{\{\gamma^m \neq \emptyset\}} \mu_{\mathcal{D}, \gamma^m}^-[X] d\mu_{\mathcal{D}}^+[\gamma^m] \leq \ell^- \quad (310)$$

where \mathcal{D}_{γ^m} is the even domain whose boundary is γ^m . The identity is the Markov property, and the inequality is again Corollary 15.6. If we first take $\mathcal{D} \nearrow \mathbb{Z}^2$ and then $m \rightarrow \infty$, the left-hand side converges to ℓ^+ , proving the desired inequality. This finishes the proof of the lemma. \square

Proof of Theorem 17.1. By reasoning as for the proof of the previous lemma, we observe that Corollary 16.6 implies that $\mu_{\mathbb{Z}^2, \circ}$ -almost surely all connected components of $\mathbb{R}^2 \setminus \omega$ are bounded. Lemma 14.7 therefore extends to the full-plane limit, and $\mu_{\mathbb{Z}^2}$ may simply be obtained from $\mu_{\mathbb{Z}^2, \circ}$ by flipping coins for the odd spins σ^\bullet in each connected component of $\mathbb{R}^2 \setminus \omega$. The properties stated in the result now immediately follow (by passing to the limit) from the properties of $\mu_{\mathcal{D}}^+$ and $\mu_{\mathcal{D}}^-$. \square

17.2 Flip domination (Theorem 4.10)

Now that the infinite-volume measure is defined, we derive the flip domination property *en passant*. (In fact, we only formally prove that the law of h in $\mu_{\mathbb{Z}^2}$ and $\mathbb{P}_{\mathbb{Z}^2}$ is the same in Lemma 18.7.)

Proof of Theorem 4.10. The proof is illustrated by Figure 15. Fix the $F(\mathbb{Z}^2)$ -circuit γ , and define F_γ as in the statement of the theorem. Fix $n \in 2\mathbb{Z}$. It suffices to prove that, for any even domain \mathcal{D} containing γ , and for any positive probability event $\mathcal{E} \subset \{h|_\gamma \leq n\}$ that is measurable in terms of $h|_{F(\mathcal{D}) \setminus F_\gamma}$, the height function $2n - h_{F_\gamma}$ stochastically dominates $h|_{F_\gamma}$ under $\mu_{\mathcal{D}}^+[\cdot | \mathcal{D}]$.

The idea is to simply explore the outermost ω_n -loops within γ . Such loops may intersect γ ; this is not a problem. Conditional on this exploration:

- The unrevealed heights (surrounded by ω_n -loop) are flip-symmetric around n ,
- The revealed heights are at most n .

This implies the desired stochastic domination. \square

18 Torus and cylinder spin representation (Theorem 2.2)

This section adapts the analysis developed above to the settings of the torus and the cylinder. While no essential difficulties arise, special care is required in defining the corresponding measures so as to preserve the FKG property. For completeness, we present the relevant results below.




Figure 15: Flip domination: perfect flip symmetry around the height n in the unexplored region; heights at most n in the explored region.

18.1 Definitions on the torus

We want to replace fixed boundary conditions outside some even domain \mathcal{D} by periodic boundary conditions on the cylinder $\mathbb{T}_{M,L}$ defined in the introduction. The graphs $F(\mathbb{T}_{M,L})$, $F_\circ(\mathbb{T}_{M,L})$, and $F_\bullet(\mathbb{T}_{M,L})$ are defined in the obvious way, in analogy with their definitions for \mathcal{D} in Section 14.

Definition 18.1 (Torus spin measure). For any torus $\mathbb{T}_{M,L}$, let $\Omega = \Omega_{\mathbb{T}_{M,L}}$ denote the sample space

$$\Omega := \{\pm\}^{F_\circ(\mathbb{T}_{M,L})} \times \{0,1\}^{E_\circ(\mathbb{T}_{M,L})} \times \{\pm\}^{F_\bullet(\mathbb{T}_{M,L})}. \quad (311)$$

A typical element is denoted by $(\sigma_\circ, \omega, \sigma_\bullet)$. We often identify ω with the set $\{\omega = 1\} \subset E_\circ(\mathbb{T}_{M,L})$, and use the standard percolation terminology. Introduce a new, special event

$$\mathcal{E}_8 := \{\text{the height gain of } h \text{ along any closed } F(\mathbb{T}_{M,L})\text{-circuit, belongs to } 8\mathbb{Z}\}. \quad (312)$$

Define the probability measure $\mu_{\mathbb{T}_{M,L}}$ on $(\sigma_\circ, \omega, \sigma_\bullet) \in \Omega$ by

$$\mu_{\mathbb{T}_{M,L}}[(\sigma_\circ, \omega, \sigma_\bullet)] \propto \mathbb{1}[(\sigma_\circ, \sigma_\bullet) \in \mathcal{E}_8] \cdot \mathbb{1}[\omega \subset A(\sigma_\circ)] \quad (313)$$

$$\cdot \mathbb{1}[xy \in \omega \text{ or } xy^* \in A(\sigma_\bullet) \text{ (or both) for any } xy \in E_\circ(\mathbb{T}_{M,L})] \quad (314)$$

$$\cdot \mathbf{c}^{\#A(\sigma_\circ)} \cdot \left(\frac{1}{\mathbf{c}}\right)^{\#\omega} \cdot \left(1 - \frac{1}{\mathbf{c}}\right)^{\#(A(\sigma_\circ) \setminus \omega)}. \quad (315)$$

Remark 18.2. Compare Definition 18.1 with Definition 14.5. The only difference is that we dropped the first two indicators in Equation (238) (related to the boundary conditions on $\partial\mathcal{D}$) and inserted an indicator for the event \mathcal{E}_8 (which constrains the height gain on loops which wind nontrivially around the torus).

Remark 18.3. Samples from $\mu_{\mathbb{T}_{M,L}}$ are naturally interpreted as six-vertex configurations on $\mathbb{T}_{M,L}$, but the induced law is *not* given by the measure $\mathbb{P}_{\mathbb{T}_{M,L}}$ defined in the introduction. The reasons are as follows. In $\mathbb{P}_{\mathbb{T}_{M,L}}$, with positive probability, the arrow configuration induces a height gain in $2\mathbb{Z} \setminus 4\mathbb{Z}$ along some nontrivial loops. For the spin representation to be well-defined, the height gain along nontrivial loops must lie in $4\mathbb{Z}$, but in the previous definition we further constrain the height gain to belong to $8\mathbb{Z}$ (via the event \mathcal{E}_8). This is done to make the FKG inequality work (this is discussed below in further detail).

18.2 Markov property and the FKG inequality on the torus

Let us first state a Markov property. The maps $\pi_{\mathcal{D}}$ and $\pi_{\mathcal{D}}^c$ are defined the same as before. The proof is identical to the original proof (Lemma 14.12).

Lemma 18.4 (Torus Markov property). *Consider a torus $\mathbb{T}_{M,L}$ and an even domain \mathcal{D} such that (as a subset of $\mathbb{T}_{M,L}$) $\partial\mathcal{D}$ is a contractible self-avoiding circuit in $F_\circ(\mathbb{T}_{M,L})$. For $\# \in \{\pm\}$,*

- *The laws $\mu_{\mathbb{T}_{M,L}}[\cdot | \{\partial\mathcal{D} \subset \omega^\#\}]$ of $\pi_{\mathcal{D}}(\sigma_\circ, \bar{\omega}, \sigma_\bullet)$ and $\mu_{\mathcal{D}}^\#$ of $(\sigma_\circ, \bar{\omega}, \sigma_\bullet)$ are the same,*
- *The random variables $\pi_{\mathcal{D}}$ and $\pi_{\mathcal{D}}^c$ are independent in $\mu_{\mathbb{T}_{M,L}}[\cdot | \{\partial\mathcal{D} \subset \omega^\#\}]$.*

We also mention an FKG inequality.

Lemma 18.5 (Torus FKG inequality). *Consider a torus $\mathbb{T}_{M,L}$ and an \circ -lattice event \mathcal{E} . If $\mu_{\mathbb{T}_{M,L}}[\mathcal{E}] > 0$, then the conditional probability measure $\mu_{\mathbb{T}_{M,L}}[\cdot | \mathcal{E}]$ has \circ -FKG.*

Proof. The proof is entirely the same as the original proof (Proposition 15.4). There is, however, one critical point where we should pay attention. The decomposition into two independent Ising models is a bit more complicated due to the nontrivial topology. One may work out, however, that the decomposition in Equation (261) remains valid on the torus for the measure $\mu_{\mathbb{T}_{M,L}}$ defined above (one should only replace the ambient space \mathbb{R}^2 by $(\mathbb{R}/L\mathbb{Z}) \times (\mathbb{R}/M\mathbb{Z})$ when counting the number of connected components). Crucially, Equation (261) remains valid thanks to the event \mathcal{E}_8 in Definition 18.1, and would fail otherwise. We refer to [GL25, Section 5] for details. \square

18.3 Circuit estimates on the torus

With the Markov property and the FKG inequality, we can essentially bring back our measures to the finite domain setup discussed before. Let us now describe how to do this.

Lemma 18.6 (Torus circuit estimate). *Fix M, L . Then, for any $4 \leq r < (M \wedge L)/4$ and $x \in \mathbb{R}^2$,*

$$\mu_{\mathbb{T}_{M,L}}[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x)) | \mathcal{E}] \geq c_{\text{circuit}} \quad (316)$$

for any event \mathcal{E} that is measurable in terms of the even spins and edges which do not intersect $\text{Ball}_{2r}(x)$. The same inequality remains true with ω^- replaced by ω^+ .

Proof. The proof is identical to the proof of Corollary 16.6. \square

18.4 Definition and main properties on the cylinder

Recall the definitions of Cyl_L and $\mathbb{P}_{\text{Cyl}_L}$. The graphs $F(\text{Cyl}_L)$, $F_\circ(\text{Cyl}_L)$, and $F_\bullet(\text{Cyl}_L)$ are defined in the obvious way, in analogy with their definitions for \mathcal{D} and $\mathbb{T}_{M,L}$.

Lemma 18.7 (Torus limits). *All of the following statements hold true.*

(i) **Full-plane limit.** *For any random variable X which is measurable in terms of finitely many spins and edges,*

$$\lim_{M,L \rightarrow \infty} \mu_{\mathbb{T}_{M,L}}[X] = \mu_{\mathbb{Z}^2}[X], \quad (317)$$

where $\mu_{\mathbb{Z}^2}$ is the measure defined in Section 14 (Theorem 17.1).

(ii) **Cylinder limit.** *For any $L \in 2\mathbb{Z}_{\geq 1}$, there exists a probability measure μ_{Cyl_L} on the sample space*

$$\Omega_{\text{Cyl}_L} := \{\pm\}^{F_\circ(\text{Cyl}_L)} \times \{0,1\}^{E_\circ(\text{Cyl}_L)} \times \{\pm\}^{F_\bullet(\text{Cyl}_L)} \quad (318)$$

such that for any random variable X which is measurable in terms of finitely many spins and edges,

$$\lim_{M \rightarrow \infty} \mu_{\mathbb{T}_{M,L}}[X] = \mu_{\text{Cyl}_L}[X]. \quad (319)$$

Moreover, all of the following are true:

- The Markov property (Lemma 18.4) passes to the limit,
- The FKG inequality (Lemma 18.5) passes to the limit,
- The circuit estimate (Lemma 18.6) passes to the limit,
- The conditional spin flip property for σ^\bullet (Lemma 14.7) holds true,
- As a measure on six-vertex configurations, μ_{Cyl_L} equals $\mathbb{P}_{\text{Cyl}_L}$.

(iii) **Full-plane as a limit of cylinder measures.** *We have*

$$\lim_{L \rightarrow \infty} \mu_{\text{Cyl}_L}[X] = \mu_{\mathbb{Z}^2}[X] \quad (320)$$

for any random variable X which is measurable in terms of finitely many spins and edges, due to the first item. If X is also measurable in terms of h , then we furthermore have

$$\lim_{L \rightarrow \infty} \mu_{\text{Cyl}_L}[X] = \mu_{\mathbb{Z}^2}[X] = \mathbb{E}_{\mathbb{Z}^2}[X]. \quad (321)$$

In particular, Theorem 2.2 holds true.

Proof. The proof of the first item is identical to the proof of Theorem 17.1: one simply uses the circuit estimate (Lemma 18.6) and the Markov property (Lemma 18.4) to compare the measure $\mu_{\mathbb{T}_{M,L}}$ with the measure $\mu_{\mathcal{D}}^+$ for a large enough even domain \mathcal{D} .

The third item clearly follows from the first two since the limit in the first item may be taken in any order. We therefore focus on the second item.

Recall from the introduction that $\{\text{balanced}\}$ denotes the event that a configuration is *balanced*, meaning that in each column of horizontal arrows, half of the arrows point to the left and the other half to the right. This event constrains the global topology, and equals the event that the height gain along any closed $F(\text{Cyl}_L)$ -circuit that winds in the vertical direction, but not the horizontal direction, equals 0. The limit

$$\lim_{M \rightarrow \infty} \mu_{\mathbb{T}_{M,L}}[\cdot | \{\text{balanced}\}] \quad (322)$$

is well-defined and equal to $\mathbb{P}_{\text{Cyl}_L}$ since we are essentially dealing with a recurrent Markov chain. We must therefore prove that $\{\text{balanced}\}$ has high probability in the $M \rightarrow \infty$ limit.

We now make one observation. If for some $n \in \mathbb{Z}$, the slice $[n, n+1] \times (\mathbb{R}/L\mathbb{Z})$ contains a nontrivial ω -path, then the gradient of the height function is zero along that path and therefore the event $\{\text{balanced}\}$ occurs. Write S_n for this event.

As discussed before, it is easy to find a uniform lower bound $\eta > 0$ on the probability that an edge is ω^+ -open, even if we condition on the states of all edges not incident to that edge. This implies that for any $k = 0, \dots, M/2 - 1$, we have

$$\mu_{\mathbb{T}_{M,L}}[S_{2k} | (S_0 \cup S_2 \cup S_4 \cup \dots \cup S_{2k-2})^c] \geq \eta^L, \quad (323)$$

and therefore

$$\mu_{\mathbb{T}_{M,L}}[\{\text{balanced}\}] \geq 1 - (1 - \eta^L)^{M/2}. \quad (324)$$

The lower bounds clearly tends to 1 as M tends to infinity.

The first three properties obviously pass to the limit, and the last property was already proved above. The conditional spin flip property for σ^\bullet is not obvious (it is not true on the torus due to the event \mathcal{E}_8), but it is recovered in the limit thanks to the Markov chain structure of the six-vertex model on the cylinder. \square

19 Bounds on crossing counts (Theorem 4.9)

The purpose of this section is to establish a log-linear bound on the number of alternating rectangle crossings (see the statement below). We also prove the log-quadratic bound on arm events (Theorem 4.9 in the introduction).

19.1 The rectangle crossings case

We start with a (far from optimal) bound on crossing counts.

Recall that $\#\text{AltHor}^{\omega}(R)$ counts the number of alternating horizontal crossings of the rectangle R .




Figure 16: Lemma 19.4. Suppose that we want to attach an ω^- -percolation to an ω^- -open boundary segment, but we have conditioned on some “messy” event inside the unions of the domains \mathcal{D}_i (the shaded area). The lemma tells us that we can still open part of the circuit which connects to the boundary. The proof uses the Markov property for polar domains to remove the conditioning on the adverse event.

Lemma 19.1 (Log-linear bound on rectangle crossing counts). *There exists a constant $c_{\text{lin}} \in (0, 1)$ (independent of $\mathbf{c} \in [1, 2]$) with the following property. For any $n \in \mathbb{Z}_{\geq 0}$ and any rectangle $R := \llbracket w \times \rho w \rrbracket + u$, where $w \in \mathbb{R}_{\geq 1}$ and $\rho \in \mathbb{R}_{\geq 1}$, and $u \in \mathbb{R}^2$,*

$$\mu_{\mathbb{Z}^2}[\{\#\text{AltHor}^{\bar{\omega}}(R) \geq 2n\}|\mathcal{E}] \leq (1 - c_{\text{lin}}^\rho)^n, \quad (325)$$

where \mathcal{E} is any event that is measurable in terms of the even spins and edges at a ℓ^∞ -distance at least $w/5$ of R .

Moreover, the same bounds hold with $\mu_{\mathbb{Z}^2}$ replaced by μ_{Cyl_L} , provided that $2(\rho + \frac{1}{5})w \leq L$, and with the understanding that R is now regarded as a rectangle on the cylinder. Analogous bounds also hold for vertical crossings of the rectangle, both in the full plane and on the cylinder.

Remark 19.2. Requiring a macroscopic buffer zone between the rectangle and the spins or edges used to determine \mathcal{E} is standard in this type of result. In fact, the conclusion may fail without such a buffer if \mathcal{E} is allowed to depend on all edges and spins outside R .

Remark 19.3. In the context of the lemma, summing over n gives that

$$\mu_{\mathbb{Z}^2}[\#\text{AltHor}^{\bar{\omega}}(R)|\mathcal{E}] \leq 2c_{\text{lin}}^{-\rho}. \quad (326)$$

For the proof of this lemma on rectangle crossings, we first require a strengthened version of the circuit estimate in which the annulus is allowed to be partially “scarred”, provided that the boundary conditions induced on the remaining region are favorable. This variation is illustrated by Figure 16.

Lemma 19.4 (Partial circuit estimate). *Fix $r \in [4, \infty)$ and $x \in \mathbb{R}^2$, and consider the following setup.*

- Let $(\mathcal{D}_i)_i$ denote a finite family of even domains.
- Let \mathcal{B} denote the set of edges in $\cup_i \partial \mathcal{D}_i$ that are contained in $\text{Ball}_{2r}(x)$.
- Let $\mathcal{U} := \cup_i E_o(\mathcal{D}_i)$.

Then, for any event \mathcal{E} measurable in terms of the even spins and edges contained in \mathcal{U} , and such that $\mu_{\mathbb{Z}^2}[\mathcal{E} \cap \{\mathcal{B} \subset \omega^-\}] > 0$, we have

$$\mu_{\mathbb{Z}^2}[\text{Circuit}^{\omega^- \cup \mathcal{U}}(\text{Annulus}_{2r,r}(x))|\mathcal{E} \cap \{\mathcal{B} \subset \omega^-\}] \geq c_{\text{circuit}}. \quad (327)$$

Proof. Without loss of generality, \mathcal{E} is an \circ -lattice event. Let γ denote the smallest self-avoiding even circuit surrounding $\text{Ball}_{2r}(x)$. By using the FKG inequality, we may




Figure 17: When exploring the paths from top to bottom, each time we explore an ω^+ -crossing, there is a universally positive probability that this crossing connects to the bottom, blocking any further horizontal ω^+ -crossings from appearing.

furthermore assume that \mathcal{E} satisfies the following “worst-case” description: $\gamma \subset \mathcal{U}$, and $\mathcal{E} \subset \{\gamma \subset \omega^+\}$. We now claim that

$$\mu_{\mathbb{Z}^2}[\text{Circuit}^{\omega^- \cup \mathcal{U}}(\text{Annulus}_{2r,r}(x)) | \mathcal{E} \cap \{\mathcal{B} \subset \omega^-\}] \quad (328)$$

$$= \mu_{\mathbb{Z}^2}[\text{Circuit}^{\omega^- \cup \mathcal{U}}(\text{Annulus}_{2r,r}(x)) | \{\gamma \subset \omega^+\} \cap \{\mathcal{B} \subset \omega^-\}] \quad (329)$$

$$\geq \mu_{\mathbb{Z}^2}[\text{Circuit}^{\omega^-}(\text{Annulus}_{2r,r}(x)) | \{\gamma \subset \omega^+\}] \quad (330)$$

$$\geq c_{\text{circuit}}. \quad (331)$$

The equality is the Markov property (Lemma 14.12) and the Markov property for polar domains (Lemma 14.13). The first inequality is inclusion of events and FKG. The second inequality is the circuit estimate (Corollary 16.6). \square

Proof of Lemma 19.1. Define two sequences of paths, as follows:

- α_1^+ is the highest horizontal ω^+ -crossing of R ,
- α_i^- is the highest horizontal ω^- -crossing of R below α_i^+ , for $i = 1, 2, \dots$,
- α_i^+ is the highest horizontal ω^+ -crossing of R below α_{i-1}^- , for $i = 2, 3, \dots$,
- If such a crossing does not exist, we set it (and all subsequent ones) equal to \emptyset .

Notice that $\{\#\text{AltHor}^{\omega}(R) \geq 2n\} = \{\alpha_n^- \neq \emptyset\}$. For Lemma 19.1, it suffices to prove that for any i , we have

$$\mu_{\mathbb{Z}^2}[\{\alpha_i^- \neq \emptyset\} | \mathcal{E} \cap \{\alpha_i^+ \neq \emptyset\}] \leq 1 - c_{\text{lin}}^\rho \quad (332)$$

for some fixed $c_{\text{lin}} \in (0, 1)$.

We split into two cases, depending on the value of w . First suppose that $w \leq 1000$. Let $\eta > 0$ denote a uniform constant such that the probability that an even spin is $+$ is at least η , even after conditioning on \mathcal{E} and all the other even spins. Then

$$\mu_{\mathbb{Z}^2}[\{\text{all even spins in } R \text{ below } \alpha_i^+ \text{ are valued } +\} | \mathcal{E} \cap \{\alpha_i^+ \neq \emptyset\}] \geq (\eta^{1000000})^\rho. \quad (333)$$

Since this event is disjoint from $\{\alpha_i^- \neq \emptyset\}$, any value $c_{\text{lin}} \leq \eta^{1000000}$ will work.

We are left with the (more interesting) case $w > 1000$. By applying Lemma 19.4 at least $\lceil 100\rho \rceil$ times and applying the FKG inequality, we observe that with a $\mu_{\mathbb{Z}^2}[\cdot | \mathcal{E} \cap \{\alpha_i^+ \neq \emptyset\}]$ -probability of at least $(c_{\text{circuit}}^{100})^\rho$, the path α_i^+ is connected to Bottom_R within the rectangle (see Remark 16.2 and Figure 17). In this case, it is impossible that the event $\{\alpha_i^- \neq \emptyset\}$ occurs.

This proves that the value $c_{\text{lin}} = \eta^{1000000} \wedge c_{\text{circuit}}^{100}$ works. \square




Figure 18: Illustration of the reduction from annulus arms to rectangle crossings. Each arm also crosses one of the four rectangles R_1 , R_2 , R_3 , and R_4 in the “easy” direction.

19.2 The annulus arms case

We now adapt the previous results to the context of arms crossing an annulus. For $r \in \mathbb{R}_{\geq 4}$, $R \in \mathbb{R}_{\geq 2r}$, and $x \in \mathbb{R}^2$, define the random variable

$$K_{R,r,x} := (1 - \mathbf{1}_{\text{AltCircuit}_2^{\bar{\omega}}(\text{Annulus}_{R,r}(x))}) \vee \#\text{AltArm}^{\bar{\omega}}(\text{Annulus}_{R-1,r+1}(x)), \quad (334)$$

which, roughly speaking, counts alternating crossings from inside to outside in the annulus. Notice that this random variable is measurable with respect to the even edges which are entirely contained in $\text{Annulus}_{R,r}(x)$.

Lemma 19.5 (Linear bound on arm exponents). *There exists a constant $c'_{\text{arm}} > 0$ (independent of $\mathbf{c} \in [1, 2]$) with the following property. For $r \in \mathbb{R}_{\geq 4}$, $R \in \mathbb{R}_{\geq 2r}$, $x \in \mathbb{R}^2$, and $n \in \mathbb{Z}_{\geq 0}$,*

$$\mu_{\mathbb{Z}^2}[\{K_{R,r,x} \geq n\} | \mathcal{E}] \leq (r/R)^{c'_{\text{arm}} n} \quad (335)$$

for any \mathcal{E} with $\mu_{\mathbb{Z}^2}[\mathcal{E}] > 0$ that is measurable with respect to the even spins and edges which do not intersect $\text{Ball}_R(x)$.

The same bounds hold true if $\mu_{\mathbb{Z}^2}$ is replaced by μ_{Cyl_L} , provided that $L > 2R$.

Remark 19.6. In the context of the lemma, summing over n gives that

$$\mu_{\mathbb{Z}^2}[K_{R,r,x} | \mathcal{E}] \leq \frac{(r/R)^{c'_{\text{arm}}}}{1 - 2^{-c'_{\text{arm}}}}. \quad (336)$$

Proof. It suffices to consider the case that $R = 2^m r$. We first prove this for $m = 1$.

Let us start by proving that with uniformly positive probability, $K_{R,r,x} = 0$. Since the event $\text{AltCircuit}_2^{\bar{\omega}}(\text{Annulus}_{R,r}(x))$ bars any arms from appearing, it suffices to show that this event occurs with uniformly positive probability. If $r \leq 100$ then we may simply choose two disjoint even circuits in $\text{Annulus}_{R,r}(x)$, and lower bound the probability that they are open for ω^+ and ω^- (see the part of the proof of Lemma 19.1 where $w \leq 1000$). For $r > 100$, we may combine 1000 circuit estimates (Corollary 16.6) and the FKG inequality to show that $\text{AltCircuit}_2^{\bar{\omega}}(\text{Annulus}_{R,r}(x))$ occurs with probability at least $c_{\text{circuit}}^{1000}$ (see also the part of the proof of Lemma 19.1 where $w > 1000$).

To finish the proof for the case that $m = 1$, it suffices to find constants $c > 0$ and $N \in 20\mathbb{Z}_{\geq 1000}$ such that for any $n \in \mathbb{Z}_{\geq 1}$, we have

$$\mu_{\mathbb{Z}^2}[\{K_{R,r,x} \geq Nn\} | \mathcal{E}] \leq e^{-cn}. \quad (337)$$

Observe that any continuous path connecting $\text{Ball}_r(x)$ to $\partial\text{Ball}_R(x)$ must necessarily traverse at least one of the following four rectangles in the “easy” direction:

$$R_1 := ([5r/4, 6r/4] \times [-7r/4, 7r/4]) + x \quad (338)$$

and its images R_2, R_3, R_4 , defined by the rotations by an angle $\frac{\pi}{2}, \pi$, and $\frac{3\pi}{2}$ around x . See Figure 18 for an illustration. In particular, we get

$$\{K_{R,r,x} \geq Nn\} \subseteq \text{AltHor}_{\frac{\bar{\omega}_{Nn}}{10}}(R_1) \cup \text{AltVer}_{\frac{\bar{\omega}_{Nn}}{10}}(R_2) \cup \text{AltHor}_{\frac{\bar{\omega}_{Nn}}{10}}(R_3) \cup \text{AltVer}_{\frac{\bar{\omega}_{Nn}}{10}}(R_4) \quad (339)$$

(the division by 10 royally suffices; the safety margin compensates a few crossings that we may lose by how we set up the definitions). Lemma 19.1 then implies

$$\mu_{\mathbb{Z}^2}[\{K_{R,r,x} \geq Nn\}|\mathcal{E}] \leq 4(1 - c_{\text{lin}}^{14})^{Nn/20}. \quad (340)$$

It is then straightforward to find good values for c and N . This concludes the case $m = 1$.

For the general case $m \geq 1$, we notice that

$$K_{2^m r, r, x} \leq \min_{i=0, \dots, m-1} K_{2^{i+1} r, 2^i r, x}, \quad (341)$$

and then use the $m = 1$ case m times at m disjoint concentric annuli. \square

Theorem 4.9 in the introduction is a direct corollary of the following stronger result.

The random variable $K_{R,r,x}$ was defined in equation (334). The following lemma improves on Lemma 19.5.

Theorem 19.7 (Quadratic bound on arm exponents). *For any $\mathbf{c} \in [1, 2]$, there exists a constant $c_{\text{arm}} \in (0, 10^{-9})$ with the following property. For $r \in \mathbb{R}_{\geq 4}$, $R \in \mathbb{R}_{\geq 2r}$, $x \in \mathbb{R}^2$, and $n \in \mathbb{Z}_{\geq 0}$,*

$$\mu_{\mathbb{Z}^2}[\{K_{R,r,x} \geq n\}|\mathcal{E}] \leq (r/R)^{c_{\text{arm}} n^2} \quad (342)$$

for any \mathcal{E} that is measurable with respect to the even spins and edges which do not intersect $\text{Ball}_R(x)$.

The same bounds hold true if $\mu_{\mathbb{Z}^2}$ is replaced by μ_{Cyl_L} , provided that $L > 2R$.




Figure 19: Each arm of the annulus $\text{Annulus}_{3r/2, r}(x)$ must touch a large number of small boxes, creating arm events around those smaller boxes.

Proof. Just like in the proof of Lemma 19.5, it suffices to consider the case that $R = 2r$. We shall also suppose that $x = (0, 0)$; this makes no difference to the proof.

Recall the constant of c'_{arm} from Lemma 19.5 on the linear bound on arm exponents. Since Lemma 19.5 already handles Equation (342) for small values of n , it suffices to restrict ourselves to the case that $n = 100 \cdot K \cdot 2^k$ for integers $k \geq 1$, where K is a constant to be fixed later.

Let $\bar{r} := 2^{-k-2} \cdot r$. The annulus $\text{Annulus}_{3r/2,r}(x)$ can be tiled by exactly $5 \cdot 2^{2k+2}$ squares of side length $2\bar{r}$. Write S for the set of centres of these squares. Now for any percolation $\omega \subset \mathbb{R}^2$ which belongs to $\text{Arm}(\text{Annulus}_{3r/2,r})$, we may find at least 2^k elements u_1, \dots, u_{2^k} in S such that $\omega \in \text{Arm}(\text{Annulus}_{2\bar{r},\bar{r}}(u_i))$ for each $i = 1, \dots, 2^k$; see Figure 19. By a pigeonhole argument, we get

$$K_{2r,r,x} \geq n \implies \sum_{u \in S} K_{2\bar{r},\bar{r},u} \geq 25 \cdot K \cdot 2^{2k}. \quad (343)$$

We must prove that the event on the right has a $\mu_{\mathbb{Z}^2}[\cdot | \mathcal{E}]$ -probability of at most $2^{-c_{\text{arm}} n^2}$.

We call a set $S' \subset S$ *separating* if the boxes $(\text{Ball}_{2\bar{r}}(u))_{u \in S'}$ are disjoint. We shall use the following input from Lemma 19.5: there exists some fixed constant $p > 0$ such that, for any separating $S' \subset S$, the random variables $(K_{2\bar{r},\bar{r},u})_{u \in S'}$ in $\mu_{\mathbb{Z}^2}[\cdot | \mathcal{E}]$ are stochastically dominated by an i.i.d. family of random variables having a geometric random variable of parameter p .

Write \mathcal{S} for some partition of S into 25 separating subsets. The pigeonhole principle implies that

$$\mu \left[\left\{ \sum_{u \in S} K_{2\bar{r},\bar{r},u} \geq 25 \cdot K \cdot 2^{2k} \right\} \middle| \mathcal{E} \right] \leq 25 \max_{S' \in \mathcal{S}} \mu \left[\left\{ \sum_{u \in S'} K_{2\bar{r},\bar{r},u} \geq K \cdot 2^{2k} \right\} \middle| \mathcal{E} \right]. \quad (344)$$

The probability on the right may be bounded by the probability that the sum of $5 \cdot 2^{2k+2}$ i.i.d. geometric random variables of parameter $p > 0$ is at least $K \cdot 2^{2k}$. By standard large deviation estimates, this probability decays like $e^{-c2^{2k}}$ for some constant $c > 0$, provided that K is some sufficiently large fixed constant (depending only on p). Since n is of order 2^k , this gives the desired quadratic bound. \square

20 Level line tree and branching function

20.1 Heuristic of the level line tree

Before formally describing the level line tree, we first give an analogy with the Gaussian free field, and then proceed with some preliminary remarks on the combinatorial structure important for the definition.

Discussion of the coupling of the GFF with CLE(4). The Gaussian free field in a disk with zero boundary conditions has a natural coupling with CLE(4) [BP25]. Conditional on the CLE(4) loops, we orient them counterclockwise or clockwise (independently and with equal probability). The GFF is then morally equal to some constant λ times the net winding of the loops around each point. In this picture, the conditional variance between two points equals λ^2 times the number of loops surrounding both points.

We can also think of the CLE(4) loops as a rooted tree: the root is the whole disk, the other tree vertices are simply connected subsets of the disk whose boundary is a CLE(4) loop, and each tree vertex points to its parent, the smallest tree vertex strictly containing it. Each point in the disk may be associated with the set of tree vertices containing it, which may be interpreted as a tree path starting at the root and going downwards (it may be finite or infinite, typically it is infinite). In this formalism, the number of loops (and therefore the conditional covariance) may be expressed in terms of the depth of the lowest common ancestor of the tree paths of the two points.

In practice, we will be interested in the expectation of products of height differences. The height difference between two points of the disk is naturally expressed in terms of a

tree path between them. This tree path may be finite or infinite on both ends (for typical points, it is a bi-infinite paths). The natural path $x \rightarrow y$ is the concatenation of the tree path from x to the lowest common ancestor of x and y , and the tree path from that ancestor to y . The conditional expectation of the product of k height differences is then equal to some function of the way that the k tree paths intersect each other.

Notes on the combinatorial structure of the level line tree. We now put ourselves in the context of the measure $\mu_{\mathcal{D}}^+$, which means that the boundary height is equal to 0. Suppose for a second that we explore the outermost ω_{\circ}^- -circuits. Then all of the following hold true:

- On each ω_{\circ}^- -circuit, the height is equal to $+2$ or -2 , depending on the value of the odd spins outside the circuit,
- The odd spins on each connected component of $\mathbb{R}^2 \setminus \omega$ are fair coin flips (Lemma 14.7),
- The Markov property (Lemma 14.12) implies a renewal property: within each circuit $\partial\mathcal{B}$, the conditional law of the spins and edges within \mathcal{B} is given by $\mu_{\mathcal{B}}^-$ (which is nothing more than $\mu_{\mathcal{B}}^+$ with all even spins flipped).

This enables the definition of a tree structure bearing some resemblance to the tree defined above for the GFF. At the same time, there are some differences.

- In the discrete the tree is finite, which makes things easier.
- The ± 2 -labels of the outermost ω_{\circ}^- -circuits are *not* independent, because it may happen that two such circuits are surrounded by the same connected component of odd spins in Lemma 14.7, in which case they have the same label.
- There is a notion of odd and even circuits: starting from $\mu_{\mathcal{D}}^+$, one explores the outermost ω_{\circ}^- -circuits, then, within each circuit, the outermost ω_{\circ}^+ -circuits, et cetera. This parity issue is unrelated to the ± 2 -labels (more precisely, the parity issue is related to the parity of the even spins, while the ± 2 -labels relate to odd spins).
- The precise height at a face is not determined by the tree structure alone; sometimes we require small corrections, for example when the face itself is odd (and we need a correction of the form ± 1).

These issues slightly complicate the definition of the tree structure (see the comments following the definition). Nevertheless, we can still apply the same logic as for the tree in the GFF case to bound correlations, as we will see for instance in the next section.

20.2 Formal definition of the level line tree

Definition 20.1 (Level line tree). Consider a sample $(\sigma_{\circ}, \bar{\omega}, \sigma_{\bullet})$ from the measure $\mu_{\mathcal{D}}^+$ for some even domain \mathcal{D} . The associated *level line tree* $\mathcal{X} = (V_{\circ}(\mathcal{X}), V_{\bullet}(\mathcal{X}), E(\mathcal{X}))$ is a rooted directed tree, defined as follows.

- The *even* vertices $V_{\circ}(\mathcal{X})$ are defined as follows.
 - First, it has a bipartition $V_{\circ}(\mathcal{X}) = V_{\circ}^+(\mathcal{X}) \cup V_{\circ}^-(\mathcal{X})$.
 - $V_{\circ}^+(\mathcal{X})$ is the partition of $\{\sigma_{\circ} = +\}$ such that two faces are in the same member of the partition if and only if they belong to the same connected component of $\mathbb{R}^2 \setminus \omega_{\circ}^-$.
 - $V_{\circ}^-(\mathcal{X})$ is the partition of $\{\sigma_{\circ} = -\}$ such that two faces are in the same member if and only if they belong to the same connected component of $\mathbb{R}^2 \setminus \omega_{\circ}^+$.
 - The root of the tree is the member of $V_{\circ}^+(\mathcal{X})$ containing $\partial\mathcal{D}$.
- The *odd* vertices $V_{\bullet}(\mathcal{X})$ of the tree are formed by the bounded connected components of $\mathbb{R}^2 \setminus \omega$.
- Each vertex (other than the root) has exactly one outgoing edge:

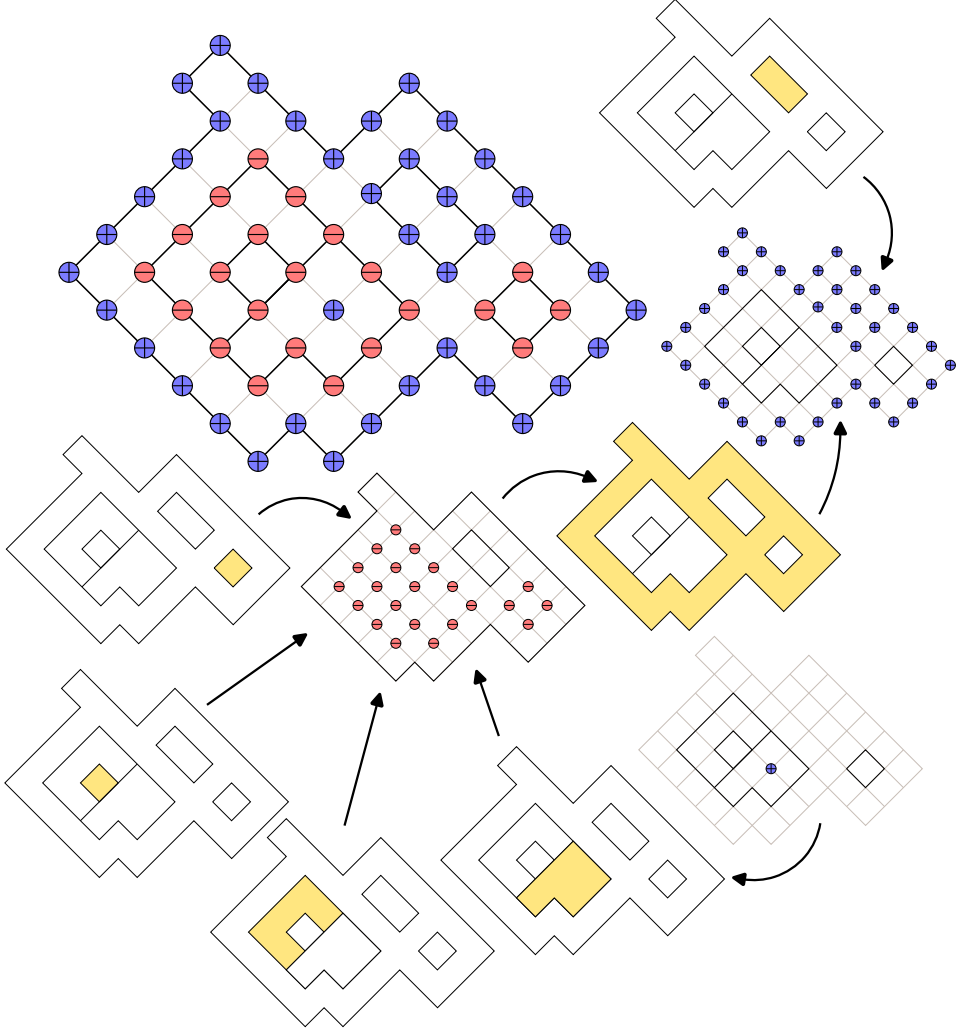


Figure 20: The directed level line tree \mathcal{X} associated with a configuration $(\sigma_\circ, \bar{\omega})$.

- Each *odd* vertex v_\bullet points towards the *even* vertex v_\circ containing the outer boundary of the connected component $v_\bullet \subset \mathbb{R}^2 \setminus \omega$,
- Each *even* vertex v_\circ labelled \pm points toward the odd vertex v_\bullet such that:
 - * v_\circ and v_\bullet lie in the same connected component C of $\mathbb{R}^2 \setminus \omega^\mp$,
 - * v_\bullet is the unique such odd vertex pointing towards $V_\circ^\mp(\mathcal{X})$ (in fact, the other odd vertices in the same connected component C point back to v_\circ).

For each vertex v other than the root, we let $p(v)$ denote the parent vertex. The *depth* of a vertex is defined as its distance to the root. Notice that vertices in $V_\circ^+(\mathcal{X})$, $V_\circ^-(\mathcal{X})$, and $V_\bullet(\mathcal{X})$ are at depth $4\mathbb{Z}$, $4\mathbb{Z} + 2$, and $2\mathbb{Z} + 1$ respectively. Notice also that each *odd* vertex has *at most* one child pointing towards it.

Remark 20.2. The odd spins are also integrated into the tree as odd vertices. This facilitates the definition of the height function at every face of the graph.

Remark 20.3. One surprising aspect lies in the fact that two ω^+ -circuits can surround each other, and still belong to the same even vertex of the tree. This is natural because no height difference is realised when no ω^- -circuits separates the two.

Ultimately, the definition is set up in such a way that the following lemma works.

The lemma below says that conditionally on \mathcal{X} , the law of $h(v)$ is given by the sum of k independent ± 2 -valued coin flips if the depth of v is $2k$, and by the sum of k independent ± 2 -valued coins and one ± 1 -valued coin if the depth of v is $2k + 1$.

Lemma 20.4 (Basic properties of the level line tree). *Let \mathcal{D} denote an even domain, and consider $\mu_{\mathcal{D}}^+$. Then, all of the following hold true.*

- (i) \mathcal{X} is measurable in terms of $(\sigma_{\circ}, \bar{\omega})$,
- (ii) Conditional on \mathcal{X} , an independent fair coin flip $\sigma_{\bullet}(X)$ is attached to each element $X \in V_{\bullet}(\mathcal{X})$, so that $\sigma_{\bullet}(x) = \sigma_{\bullet}(X)$ for every $x \in X$,
- (iii) h is almost surely constant on each vertex of \mathcal{X} ,
- (iv) Conditional on \mathcal{X} , the height function h has the following law:
 - $h(r) = 0$ on the root vertex r ,
 - $(h(v_{\bullet}) - h(p(v_{\bullet})))_{v_{\bullet} \in V_{\bullet}(\mathcal{X})}$ has the law of independent fair ± 1 -valued coin flips,
 - $h(v_{\circ}) - h(p(v_{\circ})) = h(p(v_{\circ})) - h(p(p(v_{\circ})))$ for any even non-root vertex v_{\circ} .

Here, we recall that $p(v)$ denotes the parent vertex, as introduced in Definition 20.1.

Proof. While this lemma is important, its proof follows straightforwardly from the definitions. Property (i) is immediate as the definition only involves $(\sigma_{\circ}, \bar{\omega})$. For Property (ii), observe that the odd vertices are precisely the connected components of $\mathbb{R}^2 \setminus \omega$, so that Lemma 14.7 applies. Properties (iii) and (iv) follow from the relation between spins and the height function, detailed in Equation (247). \square

Definition 20.5 (Tree path). Let \mathcal{D} denote an even domain, and consider the level line tree \mathcal{X} in $\mu_{\mathcal{D}}^+$. For any $u, v \in F(\mathcal{D})$, we define the *tree path* p^{uv} as the unique path in \mathcal{X} starting at the vertex containing u and ending at the vertex containing v . It is viewed as a set of vertices of \mathcal{X} .

20.3 Branching function

We now introduce a convenient tool to analyse the covariance structure of the height function.

Definition 20.6 (Branching function). Let \mathcal{D} denote an even domain. Consider the level line tree \mathcal{X} in $\mu_{\mathcal{D}}^+$. Let $p^{u\partial\mathcal{D}}$ be the unique \mathcal{X} -path from u to $\partial\mathcal{D}$. The *branching function* is the random \mathcal{X} -measurable function

$$\psi = \psi_{\mathcal{D}} : F(\mathbb{Z}^2) \times F(\mathbb{Z}^2) \rightarrow \mathbb{Z}, (u, v) \mapsto \begin{cases} |p^{u\partial\mathcal{D}} \cap p^{v\partial\mathcal{D}}| - 1 & \text{if } u, v \in F(\mathcal{D}), \\ 0 & \text{otherwise.} \end{cases} \quad (345)$$

Its diagonal is denoted

$$\psi^* = \psi_{\mathcal{D}}^* : F(\mathbb{Z}^2) \rightarrow \mathbb{Z}, u \mapsto \psi_{\mathcal{D}}(u, u) = \begin{cases} \text{the } \mathcal{X}\text{-depth of } u & \text{if } u \in F(\mathcal{D}), \\ 0 & \text{otherwise.} \end{cases} \quad (346)$$

Lemma 20.7 (Basic properties of the branching function). *The branching function ψ and its diagonal ψ^* satisfy the following properties.*

- (i) ψ^* is a graph homomorphism from $F(\mathcal{D})$ to \mathbb{Z} that equals 0 on $\partial\mathcal{D}$.
- (ii) ψ^* preserves the parity of the faces.
- (iii) ψ^* is equal to the maximum of h over all possible realisations of h given \mathcal{X} .

(iv) ψ encodes the conditional covariance of each pair of faces as follows:

$$\mu_{\mathcal{D}}^+[h(u)h(v)|\mathcal{X}] = \begin{cases} 4k & \text{if } \psi(u, v) = 2k, \\ 4k + 1 & \text{if } \psi(u, v) = 2k + 1 \text{ and } \psi^*(u) = \psi^*(v) = 2k + 1, \\ 4k + 2 & \text{if } \psi(u, v) = 2k + 1 \text{ and } \psi^*(u) \vee \psi^*(v) > 2k + 1. \end{cases} \quad (347)$$

Proof. The first two follow from the definition. Conditionally on \mathcal{X} , the height function is obtained by sampling the labels $\sigma_{\bullet}(X)$ for every $X \in V_{\bullet}(\mathcal{X})$ and attributing the spin to each odd faces (see the comment following the proof of Lemma 20.4). It is easy to see that if the labels $\sigma_{\bullet}(X)$ are chosen so that they maximise the height function, then we obtain $h = \psi^*$.

For the fourth property, we make the link with the GFF picture sketched in the beginning of this section. Each odd vertex is labelled \pm by flipping independent fair coins. Each ω^+/ω^- -circuit is associated with a height gain of ± 2 , and so the conditional covariance of $h(u)$ and $h(v)$ is (roughly speaking) equal to 4 times the number of alternating ω^+/ω^- -circuits surrounding both u and v . We must be a bit more careful at odd heights, because the total height contains an additional ± 1 -contribution. To make the distinction between the second and third case in Equation (347), observe that since no odd \mathcal{X} -vertex has more than one child, $\psi(u, v) = 2k + 1$ implies that either u or v is at depth $2k + 1$ in the tree. \square

Definition 20.8 (Maximal domains). Let $D \subset \mathbb{R}^2$ denote any simply connected set, and consider an even spin configuration $(\sigma_{\circ}, \bar{\omega})$. Let $\mathcal{M}^+(D)$ denote the set of maximal even domains \mathcal{D} subject to the following two conditions:

$$\partial\mathcal{D} \subset D \quad \text{and} \quad \partial\mathcal{D} \subset \omega^+. \quad (348)$$

The set $\mathcal{M}^-(\mathcal{D})$ is defined similarly by replacing ω^+ with ω^- .

We also let $\partial\mathcal{M}^{\pm}(D) \subset \mathbb{R}^2$ denote the union of $\partial\mathcal{D} \subset \mathbb{R}^2$ over all $\mathcal{D} \in \mathcal{M}^{\pm}(D)$.

Lemma 20.9 (Recursion relation for the branching function). *Consider the independent coupling of all measures $(\mu_{\mathcal{B}}^+)_{\mathcal{B}}$ over all even domains. More precisely, define the probability measure $\mathbb{P} := \prod_{\mathcal{B}} \mu_{\mathcal{B}}^+$, where the product is over all even domains. Let \mathcal{D} denote a fixed even domain. Then, in the probability measure \mathbb{P} , $\psi_{\mathcal{D}}$ and $(\psi_{\mathcal{D}} \wedge 2) + \sum_{\mathcal{B} \in \mathcal{M}^-(\mathcal{D})} \psi_{\mathcal{B}}$ have the same distribution.*

Proof. Suppose that we simply explore the domains $\mathcal{M}^-(\mathcal{D})$ and everything that happens outside the union $\cup \mathcal{M}^-(\mathcal{D})$ of those domains. Conditionally on $\mathcal{M}^-(\mathcal{D})$, the configuration inside $\cup \mathcal{M}^-(\mathcal{D})$ is given by the independent product

$$\prod_{\mathcal{B} \in \mathcal{M}^-(\mathcal{D})} \mu_{\mathcal{B}}^- \quad (349)$$

(this follows from the spatial Markov property). Notice that by definition of $\mathcal{M}^-(\mathcal{D})$ and by the independent flip symmetry in each smaller domain, we get

$$\psi_{\mathcal{D}}(u, v) \leq 2 \quad (350)$$

for any u, v unless u and v belong to the same domain $\mathcal{B} \in \mathcal{M}^-(\mathcal{D})$. Moreover, if u and v belong to the same domain $\mathcal{B} \in \mathcal{M}^-(\mathcal{D})$, then

$$\psi_{\mathcal{D}}(u, v) = 2 + \psi_{\mathcal{B}}(u, v). \quad (351)$$

This yields the lemma via Equation (349). \square

We conclude by gathering two useful properties related to the branching function.

Lemma 20.10 (Monotonicity for the branching function). *The following holds true.*

- (i) **FKG inequality.** *The law of ψ satisfies the FKG inequality in any measure $\mu_{\mathcal{D}}^+$,*
- (ii) **Monotonicity in domains.** *The law of ψ in $\mu_{\mathcal{D}}^+$ is stochastically increasing in the domain \mathcal{D} .*

Proof. We prove that $\psi \wedge 2k$ satisfies both properties for all $k \geq 1$ by inducting on k . The base case $k = 1$ is easy since $\psi \wedge 2$ is a decreasing function of the triple $(\sigma_{\circ}, \omega^+, -\omega^-)$, for which both properties are known by \circ -FKG (see Lemmas 15.2 and 15.6).

By similar reasoning, the law of $\cup \mathcal{M}^-(\mathcal{D})$ in $\mu_{\mathcal{D}}^+$ is also stochastically increasing in \mathcal{D} .

We now start the induction step: suppose that $\psi \wedge 2k$ satisfies both properties. Then the law of $\sum_{\mathcal{B} \in \mathcal{M}^-(\mathcal{D})} (\psi_{\mathcal{B}} \wedge 2k)$ is increasing in $\cup \mathcal{M}^-(\mathcal{D})$, and satisfies the FKG inequality. By the tower property (see for instance in the proof of Proposition 15.4) for the FKG inequality with respect to conditioning on $\cup \mathcal{M}^-(\mathcal{D})$, this implies that

$$(\psi_{\mathcal{D}} \wedge 2) + \sum_{\mathcal{B} \in \mathcal{M}^-(\mathcal{D})} (\psi_{\mathcal{B}} \wedge 2k) = \psi_{\mathcal{D}} \wedge (2k + 2). \quad (352)$$

satisfies both properties as well. This concludes the induction step and the proof. \square

21 Regularity estimate in full plane (Theorem 4.5)

We are now in a position to establish the regularity estimates stated in Theorem 4.5. It is proved in three steps: first, we express the correlation function Φ_k in terms of the level line tree (Subsection 21.1), second we introduce a way to relate the geometry of the level line tree to arm events (Subsections 21.2 and 21.3), and third we use this relation to bound the correlation function (Subsection 21.4).

21.1 Bounding the k -point correlations in terms of the level line tree

Recall Definition 20.5. We first prove the following lemma.

Lemma 21.1. *Fix $k \in 2\mathbb{Z}_{\geq 1}$ and $\mathbf{u} \in (F(\mathbb{Z}^2))^{2k}$, and let \mathcal{D} denote an even domain containing all faces in \mathbf{u} . Then*

$$\left| \mu_{\mathcal{D}}^+ \left[\prod_{i=1}^k (h(u'_i) - h(u_i)) \right] \right| \leq 2^k \sum_{\pi} \mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(u_i u'_i, u_j u'_j) \right], \quad (353)$$

where the sum runs over pairings of $\{1, \dots, k\}$, and where $I(vv', zz') := |p^{vv'} \cap p^{zz'} \cap V_{\bullet}(\mathcal{X})|$. Here, $p^{zz'}$ refers to the unique \mathcal{X} -path from z to z' .

Proof. The idea is to bound the conditional expectation of the product of the height differences, given \mathcal{X} . This conditional expectation may be bounded as follows. The height difference $h(u'_i) - h(u_i)$ may be written

$$h(u'_i) - h(u_i) = \sum_{n=0}^{|p^{u_i u'_i}|-1} h(p_{n+1}^{u_i u'_i}) - h(p_n^{u_i u'_i}) \quad (354)$$

where $|\cdot|$ denotes the length of the path.

We decompose the left-hand side of (353) by conditioning on \mathcal{X} and writing each height difference as in the previous displayed equation. We then expand the sum with respect to the product. At the end, we obtain a sum of products of (integer) powers of terms of the form $h(p_{n+1}^{u_i u'_i}) - h(p_n^{u_i u'_i})$. Each term in the sum contributes $-1, 0,$ or $1,$ so that it suffices to upper bound the number of terms with a nonzero contribution.

Since increments of h on edges of \mathcal{X} only interact on edges incident to the same odd vertex, we get, if $I(vv', zz') := |p^{vv'} \cap p^{zz'} \cap V_\bullet(\mathcal{X})|,$

$$\left| \mu_{\mathcal{D}}^+ \left[\prod_{i=1}^k (h(u'_i) - h(u_i)) \middle| \mathcal{X} \right] \right| \leq \sum_{\pi} \prod_{ij \in \pi} 4I(u_i u'_i, u_j u'_j), \quad (355)$$

where π runs over pairings of $\{1, \dots, k\}$ and paths are viewed as subsets of $V(\mathcal{X})$.

Clearly $|\pi| = k/2$. The tower property yields

$$\left| \mu_{\mathcal{D}}^+ \left[\prod_{i=1}^k (h(u'_i) - h(u_i)) \right] \right| \leq 2^k \sum_{\pi} \mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(u_i u'_i, u_j u'_j) \right], \quad (356)$$

where the $I(u_i u'_i, u_j u'_j)$ are viewed as random variables. \square

21.2 Key input for the regularity estimate

We now want to bound the right hand side of Equation (353). While the appropriate bound requires a bit of geometrical analysis, the kea idea is fairly simple and explained in the following lemma.

Lemma 21.2. *Let $u, u', v, v' \in \mathbb{Z}^2$ denote four points contained in some finite domain \mathcal{D} . Then all of the following hold true.*

- (i) *Suppose that $\{u, u'\}$ is contained in one connected component of $\mathbb{R}^2 \setminus \text{Annulus}_{R,r}(x)$ and $\{v, v'\}$ in the other connected component, where $r \geq 4, R \geq 2r,$ and $x \in \mathbb{R}^2$. Then*

$$I(uu', vv') \leq 4K_{R,r,x}, \quad (357)$$

where $K_{R,r,x}$ counts alternating arms in the annulus as defined in Equation (334).

- (ii) *If u and u' are neighbours (or v and v'), then*

$$I(uu', vv') \leq 4. \quad (358)$$

Proof. The second part is easy: if $u \sim u'$, then $I(\cdot) \leq |p^{uu'} \cap V_\bullet(\mathcal{X})| \leq 4$. For the first part, notice that as we walk along the path $p^{uu'} \cap p^{vv'}$, we discover alternating ω -circuits, which each have the property that they separate u from u' and v from v' . In particular, each such circuit creates an arm in the annulus. The factor 4 royally suffices. This implies the desired result. \square

To bound the right-hand side of Equation (353), we would like to proceed as follows: if we can find an annulus for each pair $ij \in \pi$ such that the $|\pi| = k/2$ annuli are disjoint, then we can simply apply the previous lemma and Remark 19.6 to bound the expectation of the product of the arm counts.

For the general case (when the scale separation between some of the pairs is small), we must do more work. To reduce to the case with good scale separation, we shall further decompose each variable $I(\cdot)$. This random variable clearly satisfies the triangular inequality $I(vv'', zz') \leq I(vv', zz') + I(v'v'', zz')$. In the next subsection, we decompose the right hand side of Equation (353), to the point that we can use Lemma 21.2 in combination with Lemma 19.5.

21.3 Organizing the points in a suitable fashion

The purpose of the following lemma is roughly as follows: if the points a, a', b, b' do not have good scale separation, then we may find a path from a to a' such that we can apply one of the two scenarios of Lemma 21.2 to each step of the path (against $\{b, b'\}$). The shortest appropriate path is called an *optimal path*.

Lemma 21.3. *Recall the discrete scale separation function $S'_{\mathbb{R}^2}$ in Equation (29). Then for each $k \in 2\mathbb{Z}_{\geq 1}$, there exists a constant $N_k \in \mathbb{R}_{>0}$ with the following properties. For any two pairs of points $\{a, a'\}$ and $\{b, b'\}$ in \mathbb{Z}^2 with $|a' - a| \leq |b' - b|$, there exists a path $p = (p(n))_{0 \leq n \leq \ell} \subset \mathbb{Z}^2$ of length at most $N_k \cdot S'_{\mathbb{R}^2}(\{a, a'\}, \{b, b'\})$ from a to a' such that for each $0 \leq n \leq \ell - 1$, one of the following two holds true:*

- $\gamma(n)$ and $\gamma(n + 1)$ are neighbours,
- $S_{\mathbb{R}^2}(\{\gamma(n), \gamma(n + 1)\}, \{b, b'\}) \geq 20k^2$.

The shortest such path is called optimal.

The lower bound $20k^2$ on the scale separation may appear arbitrary at this point. Its choice is motivated by the following problem. Broadly speaking, we want that the annuli do not overlap in order to apply Lemma 19.5. If the scale separation is large enough, then we may extract appropriate *disjoint* annuli. This is proved in the following “shrinking annuli” lemma which one can skip in a first reading.

Lemma 21.4 (Shrinking annuli lemma). *Fix $N \in \mathbb{Z}_{\geq 2}$ and a family $(R_i, r_i, x_i)_{i=1, \dots, N} \subset (0, \infty) \times (0, \infty) \times \mathbb{R}^2$ satisfying $R_i/r_i \geq e^{40N^2}$. Then, we may find some $(R'_i, r'_i)_{i=1, \dots, N} \subset (0, \infty) \times (0, \infty)$ such that:*

- (i) *The radii satisfy $r_i \leq r'_i \leq R'_i \leq R_i$,*
- (ii) *The radii satisfy $R'_i/r'_i = \sqrt[4N^2]{R_i/r_i} \geq e^{10}$ for each i ,*
- (iii) *The annuli $(\text{Annulus}_{R'_i, r'_i}(x_i))_{i=1, \dots, N}$ are pairwise disjoint.*

Proof. Set $Q = 2N^2$ and $\rho_i := \sqrt[2Q]{R_i/r_i} = \sqrt[4N^2]{R_i/r_i} \geq e^{10}$. Let \mathbb{P} denote the uniform probability measure on the random element $s \in \{0, \dots, Q - 1\}^N$. Define the random annuli

$$A_i := \text{Annulus}_{R'_i, r'_i}(x_i); \quad R'_i := R'_i(s_i) := r_i \cdot \rho_i^{2s_i+2}; \quad r'_i := r'_i(s_i) := r_i \cdot \rho_i^{2s_i+1} \quad (359)$$

in this probability space. These radii clearly satisfy the first two properties; it suffices to prove that the annuli are disjoint with positive \mathbb{P} -probability, a fact which would follow from $\mathbb{P}[\{A_i \cap A_j \neq \emptyset\}] \leq 2/N^2$ for any distinct i and j .

Let us turn to the proof of this fact. Fix $i \neq j$ and let $O \subset \{0, \dots, Q - 1\}^2$ denote the set of pairs (s_i, s_j) which lead to overlapping annuli. Fix $(a, b) \in O$; we then claim that

$$O \cap \{(a', b') : a' > a, b' \leq b\} = \emptyset \quad \text{or} \quad O \cap \{(a', b') : a' \leq a, b' > b\} = \emptyset. \quad (360)$$

Indeed, without loss of generality (by swapping i and j if necessary), we may assume that $R'_i(a) \geq R'_j(b)$. One can then see that for any $a' > a$ and $b' \leq b$, the annulus $\text{Annulus}_{R'_i(a'), r'_i(a')}(x_i)$ is disjoint from $\text{Ball}_{R'_j(b)}(x_j)$, hence from $\text{Annulus}_{R'_j(b'), r'_j(b')}(x_j)$.

With (360) in hand, we get that $|O| \leq 4Q$. This implies the desired bound $\mathbb{P}[\{A_i \cap A_j \neq \emptyset\}] \leq 4Q/Q^2 = 2/N^2$, and concludes the proof of the lemma. \square

21.4 Proof of the regularity estimate

We are now ready to dive into the proof of Theorem 4.5.




Figure 21: Left: An optimal path q_{ij} from u_i to u'_i for a pair $ij \in \pi_b$. Right: If $ij \in \pi_g$, then no optimal path is needed.

Proof of Theorem 4.5. Recall that the law of h is the same in $\mu_{\mathbb{Z}^2}$ and $\mathbb{E}_{\mathbb{Z}^2}$ (Lemma 18.7). We therefore only need to bound the correlations in an even domain with $+$ boundary conditions, and then to let the domain go to \mathbb{Z}^2 to obtain our result. From now on, we fix a large enough even domain \mathcal{D} containing all the points of \mathbf{u} .

By Lemma 21.1, it suffices to bound for every π ,

$$\mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(u_i u'_i, u_j u'_j) \right]. \quad (361)$$

Below, we shall view π in (353) as a set of *ordered* pairs $ij = (i, j)$, where we order each pair such that $|u'_i - u_i| \leq |u'_j - u_j|$. For each $ij \in \pi$, let q_{ij} denote the optimal path from u_i to u'_i (relative to the pair $\{u_j, u'_j\}$). Write $|q_{ij}|$ for the length of this path (this quantity is bounded via Lemma 21.3). Partition the pairs in π into two sets, as follows.

- The set π_g is defined as the set “good” pairs, that is, the pairs ij which satisfy $S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\}) \geq 20k^2$. This means that automatically, $|q_{ij}| = 1$ for any $ij \in \pi_g$.
- The remaining pairs are “bad”; the set of bad pairs is denoted π_b . For $ij \in \pi_b$, it is still possible that $|q_{ij}| = 1$ for $ij \in \pi_b$, namely when u_i and u'_i are neighbours.

Define $\mathcal{I}(\pi) := \prod_{ij \in \pi} \{0, \dots, |q_{ij}| - 1\}$; the path decomposition leads to

$$\mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(u_i u'_i, u_j u'_j) \right] \leq \sum_{\mathbf{n} \in \mathcal{I}(\pi)} \mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(q_{ij}(\mathbf{n}_{ij}) q_{ij}(\mathbf{n}_{ij} + 1), u_j u'_j) \right]. \quad (362)$$

Next, we are going to bound the random variables $I(\cdot)$ using Lemma 21.2. We distinguish three cases (see Figure 21).

- If $ij \in \pi_b$ and $q_{ij}(\mathbf{n}_{ij})$ and $q_{ij}(\mathbf{n}_{ij} + 1)$ are neighbours, then $I(\cdot) \leq 4$ deterministically.
- If $ij \in \pi_b$ and the two faces are not neighbours, then there exists an annulus $A_{ijn} := \text{Annulus}_{R_{ijn}, r_{ijn}}(x_{ijn})$ such that $\{q_{ij}(\mathbf{n}_{ij}), q_{ij}(\mathbf{n}_{ij} + 1)\}$ is contained in one connected component of $\mathbb{R}^2 \setminus A_{ijn}$, and $\{u_j, u'_j\}$ in the other connected component, and such that $R_{ijn}/r_{ijn} \geq e^{10k^2}$. In that case, Lemma 21.2 gives

$$I(\cdot) \leq 4K_{R_{ijn}, r_{ijn}, x_{ijn}}. \quad (363)$$

- If $ij \in \pi_g$, then we may argue as for the previous case, except that we may choose the annulus such that $R_{ijn}/r_{ijn} \geq e^{S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\})/2} \geq e^{10k^2}$.

Writing $\pi_{b, \mathbf{n}} \subset \pi_b$ for the pairs in Case (ii), we get

$$\mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(u_i u'_i, u_j u'_j) \right] \leq 4^k \sum_{\mathbf{n} \in \mathcal{I}(\pi)} \mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi_g \cup \pi_{b, \mathbf{n}}} K_{R_{ijn}, r_{ijn}, x_{ijn}} \right]. \quad (364)$$

The right-hand side is easy to upper bound if the annuli do not overlap (using Remark 19.6), but the disjointness is not guaranteed by our construction. We therefore apply the shrinking annuli lemma (Lemma 21.4) with $N = k/2$ to find a family of non-overlapping annuli $A'_{ijn} := \text{Annulus}_{R'_{ijn}, r'_{ijn}}(x_{ijn}) \subset A_{ijn}$ (with $ij \in \pi_g \cup \pi_{b,n}$) with radii $r_{ij} \geq 4$ and $R'_{ijn}/r'_{ijn} = \sqrt[k^2]{R_{ijn}/r_{ijn}} \geq e^{10}$.

Thanks to the inclusion of the annuli, the previous estimate turns into

$$\mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(u_i u'_i, u_j u'_j) \right] \leq 4^k \sum_{\mathbf{n} \in \mathcal{I}(\pi)} \mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi_g \cup \pi_{b,n}} K_{R'_{ijn}, r'_{ijn}, x_{ijn}} \right]. \quad (365)$$

The expectation on the right may now be bounded as follows: first we bound the expectation of $K_{(\cdot)}$ for the annulus with the largest outer radius, then the *conditional* expectation of the $K_{(\cdot)}$ with the next-largest outer radius, et cetera. Lemma 19.5 (more precisely Remark 19.6) asserts that the conditional expectation of each $K_{R'_{ijn}, r'_{ijn}, x_{ijn}}$ is bounded by $(R'_{ijn}/r'_{ijn})^{-c_{\text{arm}}}/(1 - 2^{-c_{\text{arm}}})$ in this procedure. We therefore get

$$\mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(u_i u'_i, u_j u'_j) \right] \leq \frac{4^k}{(1 - 2^{-c_{\text{arm}}})^{k/2}} \sum_{\mathbf{n} \in \mathcal{I}(\pi)} \prod_{ij \in \pi_g \cup \pi_{b,n}} (R'_{ijn}/r'_{ijn})^{-c_{\text{arm}}}. \quad (366)$$

Using the lower bounds on R'_{ijn}/r'_{ijn} provided by Lemma 21.4 and Lemma 21.3 (which bounds $|q_{ij}|$ and therefore $|\mathcal{I}(\pi)|$), we get the bound

$$\mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} I(u_i u'_i, u_j u'_j) \right] \leq C_k \prod_{ij} \begin{cases} e^{-(c_{\text{arm}}/2k^2)S_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\})} & \text{if } ij \in \pi_g, \\ 1 \vee -S'_{\mathbb{R}^2}(\{u_i, u'_i\}, \{u_j, u'_j\}) & \text{if } ij \in \pi_b. \end{cases} \quad (367)$$

As mentioned in the preamble of the proof, this implies the claim, as one can let \mathcal{D} to \mathbb{Z}^2 and sum over every π to get a bound on $|\Phi_k(\mathbf{u})|$. \square

22 Regularity estimate (cylinder, Corollary 4.7)

The proof is quite straightforward relative to the proof of Theorem 4.5. We can essentially work as for the full-plane case, except when the cylinder is very thin relative to the distance between the points (see the picture on the right in Figure 22). In that case, we must slightly modify our proof. This modification is very natural: rather than counting alternating arms in an annulus, we count alternating arms going through a thin subcylinder. This leads to an even better bound than the one we need.

Proof of Equation (34) in Corollary 4.7. We first prove that there is some constant C such that

$$|\Phi_{\text{Cyl}_L, 2}((0, 0), (k, 0), (2k, 0), (3k, 0))| \leq C \quad (368)$$


for any k and L . Fix $N = 10^9$. For simplicity we split in three cases.

Small values for k ($k \leq N$). If $k \leq N$, then $\Phi_{\text{Cyl}_L, 2}((0, 0), (k, 0), (2k, 0), (3k, 0)) \leq N^2$.

Values for k that are large, but smaller than the cylinder ($N < k \leq L/40$). This is the “full-plane case”. Set $u := (k/2, 0)$, $r = 5k/8$, and $R = 2r = 5k/4$. By arguing as in the proof of Theorem 4.5 (Section 21), we see that

$$|\Phi_{\text{Cyl}_L, 2}((0, 0), (k, 0), (2k, 0), (3k, 0))| \leq 8\mu_{\text{Cyl}_L} [K_{R, r, u}], \quad (369)$$

where $K_{R, r, u}$ is defined in (334); see Figure 22. This expectation is uniformly bounded by Remark 19.6.



Case $L \geq 40k$

Case $L < 40k$

Figure 22: Regularity estimates for the cylinder

Values where k is on the scale of the cylinder ($k > N \vee (L/40)$). This is the “cylinder case”. To illustrate why this case is different, suppose for a moment that L is in fact much smaller than k (for example, $L \approx k/100$). The problem that arises is that the annulus considered in the previous case wraps around the cylinder many times, and therefore the estimate that we used before does not make sense. To circumvent this problem, we are going to replace the annulus by a different shape, which leverages the topology of the cylinder. This leads to an *even better* upper bound, even though a constant bound suffices for our purposes.

Define the subcylinder $A = [\frac{5}{4}k, \frac{7}{4}k] \times (\mathbb{R}/L\mathbb{Z}) \subset \text{Cyl}_L$, and let K denote the $2\mathbb{Z}_{\geq 0}$ -valued random variable defined to be maximal subject to A having K alternating $\bar{\omega}$ crossings from left to right (see Figure 22). Notice that, just like in the definition of $\text{AltArm}_K^{\bar{\omega}}(A)$, there is not really a highest or leftmost crossing, due to the topology of A .

By arguing as before, we see that

$$|\Phi_{\text{Cyl}_L, 2}((0, 0), (k, 0), (2k, 0), (3k, 0))| \leq 8(1 + \mu_{\text{Cyl}_L}[K]). \quad (370)$$

Importantly, the aspect ratio $L/\frac{1}{2}k$ of A is upper bounded thanks to our assumption. We may therefore argue as before to see that K has exponentially decaying tails, uniformly in the choice of k and L (subject to the bound on the aspect ratio). This yields the desired uniform bound on $\mu_{\text{Cyl}_L}[K]$, which is also uniformly bounded by Remark 19.6. \square

Remark 22.1. Although not necessary, in this last case it is straightforward to see that the upper bound on the correlation function decays *exponentially fast* in k/L . This is consistent with an intuition coming from the transfer matrix perspective: the eigenvalues of $T(0)$ are L -th roots of unity, and therefore they cannot be too close to one (without being equal to one). This is obvious because $T(0)^L$ is the identity operator. Such a one-line proof does not exist for the eigenvalues of the Hermitian matrix $T(\pi/2)$, but it can be proved that there are no eigenvalues in the range $(1 - \varepsilon/L, 1)$ by RSW-type arguments.

Proof of Equation (35) in Corollary 4.7. We aim to prove that for any $L \in 2\mathbb{Z}_{\geq 1}$ and for any integers $0 < \ell \leq L/2$ and $k \geq \ell/8$,

$$|\Phi_{\text{Cyl}_L, 2}((0, 0), (0, \ell), (k, 0), (k, \ell))| \leq C(\ell/k)^c. \quad (371)$$

Again, we divide into three cases.

Values of $k \leq 10^6 \ell$. It suffices to bound the correlation function by a uniform constant. If $\ell \leq 1000$ then this is trivial. Suppose that $\ell > 1000$. Let

$$K := \#\text{AltArm}^{\bar{\omega}}(A); \quad (372)$$

$$A := (\llbracket \ell/10 \rrbracket \times [-\frac{\ell}{3}, \frac{4\ell}{3}]) \setminus (\llbracket \ell/20 \rrbracket \times [-\frac{\ell}{6}, \frac{7\ell}{6}]). \quad (373)$$

Then, the correlation function is bounded by $8(1 + \mu_{\text{Cyl}_L}[K])$. It is easy to prove that $\mu_{\text{Cyl}_L}[K]$ is uniformly bounded, by arguing as in the proof of Lemma 19.5.

Values of $k > 10^6 \ell$ satisfying $\sqrt{\ell k} \leq L$. This is the “full-plane case”. Fix $r := 4\ell \geq 4$ and $R := \frac{1}{10}\sqrt{\ell k}$. Recall the definition of $K_{r,R,(0,\ell/2)}$ from the statement of Lemma 19.5 (roughly speaking, it counts the number of alternating crossings from inside to outside in the annulus). The level line tree picture can be developed for the cylinder: without going into the details, it is not difficult to see that the two-point correlation function is bounded by the number of alternating circuits separating the two points in each pair (Lemma 21.2), which in turn is bounded by (see the picture on the left in Figure 22)

$$8\mu_{\text{Cyl}_L}[K_{r,R,(0,\ell/2)}] \leq \frac{(r/R)^{c_{\text{arm}}}}{1 - 2^{-c_{\text{arm}}}} \propto (\ell/k)^{c_{\text{arm}}/2}, \quad (374)$$

which is a bound of the desired form.

Values of $k > 10^6 \ell$ satisfying $\sqrt{\ell k} > L$. This is the “cylinder case”. First observe that $k/L \geq \sqrt{k/\ell} \geq 1000$. Define the subcylinders

$$U_0 := [L, 2L] \times (\mathbb{R}/L\mathbb{Z}); \quad (375)$$

$$U_1 := [3L, k-L] \times (\mathbb{R}/L\mathbb{Z}). \quad (376)$$

Write K for the $2\mathbb{Z}_{\geq 0}$ -valued random variable defined to be maximal subject to $\text{AltArm}^{\bar{\omega}}_K(U_0)$ occurring. Write $\mathcal{E} := \text{Circuit}^{\omega}(U_1)$. Then, the correlation function is bounded by

$$8\mu_{\text{Cyl}_L}[\mathbb{1}_{\mathcal{E}^c}(1 + K)]. \quad (377)$$

Yet, the conditional expectation $\mu_{\text{Cyl}_L}[K|\mathcal{E}^c]$ is uniformly bounded, and the probability $\mu_{\text{Cyl}_L}[\mathcal{E}^c]$ tends to zero exponentially fast in k/L . Thus, we get a bound of the form

$$|\Phi_{\text{Cyl}_L,2}((0,0), (0,\ell), (k,0), (k,\ell))| \leq Ce^{-ck/L} \leq Ce^{-c\sqrt{k/\ell}}. \quad (378)$$

The stretch-exponential decay is even stronger than the desired polynomial decay. \square

23 Mixing estimate (Theorem 4.8)

This section is split into two subsections. The first subsection contains the main proof of the mixing estimate. The second subsection analyses in further detail the covariance structure of the height function, which is used in the last step of the main proof. The ideas on the covariance structure are also used in the last section (Section 26).

23.1 Main part of the proof

Proof of Theorem 4.8. Fix $\mathbf{c} \in [1, 2]$, $k \in 2\mathbb{Z}_{\geq 3}$, and $\varepsilon := 1/10^9$ throughout this section. Consider $\mathbf{u} \in \mathcal{D}_k$.




Figure 23: The points are positioned such that the pairs (\bar{u}_i, \bar{u}'_i) for $i \geq 3$ have a relatively small diameter and are far from each other and from the point $(0, 0)$. For the proof of the mixing estimate, it is convenient to write the height difference $h(\bar{u}_i) - h(\bar{u}'_i)$ as $H_i + \tilde{H}_i$ for $i = 1, 2$, illustrated by the figure.

Step 1: Conveniently positioning the points \mathbf{u} . We first argue that without loss of generality, \mathbf{u} satisfies the following conditions (cf. Figure 23):

- $u_2 = (0, 0)$, $u'_2 = (-1, 0)$, $u'_1 = (1, 0)$,
- $|u_i - u_j| \geq 1/\varepsilon^2$ for any distinct $i, j \geq 2$,
- $|u'_i - u_i| \leq \varepsilon^2$ for any $i \geq 3$.

Indeed, by translating the system, we may assume that $u_2 = (0, 0)$. By decomposing each pair (u_i, u'_i) into paths of tiny steps (like in the proof of Proposition 7.2) and using additivity of the correlation functions, we may assume that $\inf_{i \neq j} |u_j - u_i| / \sup_i |u'_i - u_i| \geq \varepsilon^5$. We may then rescale the system such that the second and third properties are satisfied. Finally, the statement we are trying to prove does not depend on the position of u'_1 and u'_2 by the additivity property and the regularity bound, and therefore we may choose them as described in the claim.

Recall that u_1 is the variable of interest, and we would like to let u_1 tend to $(0, 0) = u_2$.

Step 2: Splitting the height increments. Consider $|u_1| < \varepsilon$ and $\delta < \varepsilon^2$, and write $\mathbf{w} := (u_3, u'_3, \dots) \in \mathcal{D}_{k-2}$. Write $\bar{u}_i \in F(\mathbb{Z}^2)$ for the face corresponding to u_i/δ , and define \bar{u}'_i , and $\bar{\mathbf{w}}$ similarly. It now suffices to find a universal constant $C < \infty$ (depending only on k) such that $|(\star)| \leq C$, where

$$(\star) := |\Phi_k(\bar{u}_1, \bar{u}'_1, \bar{u}_2, \bar{u}'_2, \bar{\mathbf{w}}) - \Phi_2(\bar{u}_1, \bar{u}'_1, \bar{u}_2, \bar{u}'_2) \Phi_{k-2}(\bar{\mathbf{w}})|. \quad (379)$$

We rephrase our correlation functions in terms of the spin measure $\mu_{\mathbb{Z}^2}$. The key step in the proof is to consider the random domain \mathcal{A} defined as the unique maximal even domain such that $0 = \bar{u}_2 \in F(\mathcal{A})$ and $\partial\mathcal{A} \subset \omega^+ \cap \text{Ball}_{(1/\delta)}$. We set $\mathcal{A} = \emptyset$ and $\partial\mathcal{A} = \{\bar{u}_2\}$ when such an even domain does not exist. Notice that almost surely h is constant on $\partial\mathcal{A}$; write $h(\partial\mathcal{A})$ for this constant value. Define the following random variables (cf. Figure 23):

$$\begin{aligned} P &:= \prod_{i=3}^k (h(\bar{u}'_i) - h(\bar{u}_i)); & a &:= \begin{cases} h(\partial\mathcal{A}) & \text{if } F(\mathcal{A}) \ni \bar{u}_1, \\ h(\bar{u}_1) & \text{otherwise;} \end{cases} \\ H_1 &:= a - h(\bar{u}'_1); & \tilde{H}_1 &:= h(\bar{u}_1) - a; \\ H_2 &:= h(\partial\mathcal{A}) - h(\bar{u}'_2); & \tilde{H}_2 &:= h(\bar{u}_2) - h(\partial\mathcal{A}). \end{aligned}$$

Then,

$$\Phi_k(\bar{u}_1, \bar{u}'_1, \bar{u}_2, \bar{u}'_2, \bar{\mathbf{w}}) = \mu_{\mathbb{Z}^2}[P(H_1 + \tilde{H}_1)(H_2 + \tilde{H}_2)]; \quad (380)$$

$$\Phi_2(\bar{u}_1, \bar{u}'_1, \bar{u}_2, \bar{u}'_2) = \mu_{\mathbb{Z}^2}[(H_1 + \tilde{H}_1)(H_2 + \tilde{H}_2)]; \quad (381)$$

$$\Phi_{k-2}(\bar{\mathbf{w}}) = \mu_{\mathbb{Z}^2}[P]; \quad (382)$$

$$(\star) = \text{Cov}_{\mu_{\mathbb{Z}^2}}[P, (H_1 + \tilde{H}_1)(H_2 + \tilde{H}_2)]. \quad (383)$$

Step 3: Rewriting (\star) as a finite sum of bounded terms. Conditionally on $\partial\mathcal{A}$, the Markov property over \mathcal{A} applies. This means that the tuples (P, H_1, H_2) and $(\tilde{H}_1, \tilde{H}_2)$ are independent. Moreover, we know that the second tuple has zero mean due to flip symmetry (Lemma 14.7). Thus, we get

$$(\star) = \text{Cov}_{\mu_{\mathbb{Z}^2}}[P, H_1 H_2 + \tilde{H}_1 \tilde{H}_2] = \text{Cov}_{\mu_{\mathbb{Z}^2}}[P, H_1 H_2] + \text{Cov}_{\mu_{\mathbb{Z}^2}}[P, \tilde{H}_1 \tilde{H}_2]. \quad (384)$$

Since P and $\tilde{H}_1 \tilde{H}_2$ are independent conditionally on \mathcal{A} , the second term on the right-hand side is bounded by

$$|\text{Cov}_{\mu_{\mathbb{Z}^2}}[\mu_{\mathbb{Z}^2}[P|\mathcal{A}], \mu_{\mathbb{Z}^2}[\tilde{H}_1 \tilde{H}_2|\mathcal{A}]]| \leq \sqrt{\mu_{\mathbb{Z}^2}[\mu_{\mathbb{Z}^2}[P|\mathcal{A}]^2] \text{Var}_{\mu_{\mathbb{Z}^2}}[\mu_{\mathbb{Z}^2}[\tilde{H}_1 \tilde{H}_2|\mathcal{A}]]}. \quad (385)$$

Combining yields

$$|(\star)| \leq |\text{Cov}_{\mu_{\mathbb{Z}^2}}[P, H_1 H_2]| + \sqrt{\mu_{\mathbb{Z}^2}[\mu_{\mathbb{Z}^2}[P|\mathcal{A}]^2] \text{Var}_{\mu_{\mathbb{Z}^2}}[\mu_{\mathbb{Z}^2}[\tilde{H}_1 \tilde{H}_2|\mathcal{A}]]}. \quad (386)$$

To get Equation (379), it suffices to bound each of the three probabilistic terms appearing on the right in (386) by a universal constant (depending only on k). This is neither short nor straightforward. So the three terms are handled in Lemmata 23.1, 23.2, and 23.4 below. \square

23.2 Proofs of the lemmata

Lemma 23.1. *In the context of the proof of Theorem 4.8 (in particular Equation (386)), there exists a constant $C < \infty$ (depending only on k) such that*

$$\mu_{\mathbb{Z}^2}[\mu_{\mathbb{Z}^2}[P|\mathcal{A}]^2] \leq C. \quad (387)$$

Proof. Since \mathcal{A} is ω^+ -measurable, we get

$$\mu_{\mathbb{Z}^2}[\mu_{\mathbb{Z}^2}[P|\mathcal{A}]^2] \leq \mu_{\mathbb{Z}^2}[\mu_{\mathbb{Z}^2}[P|\sigma_{\circ}, \bar{\omega}]^2]. \quad (388)$$

Recall the definition of the level line tree from Section 20. We shall argue as in Section 21 on the regularity estimate. Let \mathcal{D} be an extremely large domain. In Section 21, we argued in Equation (355) that

$$|\mu_{\mathcal{D}}^+[P|\mathcal{X}]| \leq \sum_{\pi} \prod_{ij \in \pi} 4I(\bar{u}_i \bar{u}'_i, \bar{u}_j \bar{u}'_j), \quad (389)$$

where π runs over the pairings of $\{3, 4, \dots, k\}$. In Section 21 our ultimate goal was to bound the first moment of the right-hand side, but now we want to bound its second moment.

This follows straightforwardly from Lemma 19.5. Indeed, later in Section 21 (Equation (363)), we bounded each intersection count by an arm count around \bar{u}_i :

$$I(\bar{u}_i \bar{u}'_i, \bar{u}_j \bar{u}'_j) \leq 4K_{(1/\delta), (\varepsilon/\delta), \bar{u}_i}. \quad (390)$$

Letting \mathcal{D} converge to \mathbb{Z}^2 yields

$$|\mu_{\mathbb{Z}^2}[P|\sigma_\circ, \bar{\omega}]| \leq \sum_{\pi} \prod_{ij \in \pi} 16K_{(1/\delta), (\varepsilon/\delta), \bar{u}_i}. \quad (391)$$

But the annuli corresponding to the \bar{u}_i do not overlap, and therefore we may use our (exponential) bounds on the tail of $K_{(\cdot)}$ from Lemma 19.5 (which bounds the second moment of each $K_{(1/\delta), (\varepsilon/\delta), \bar{u}_i}$) to see that the second moment of the random variable $\mu_{\mathbb{Z}^2}[P|\sigma_\circ, \bar{\omega}]$ is universally bounded (with a bound depending on k and \mathbf{c} only). \square

Let us turn to Lemma 23.2, which is slightly more involved but relies on similar ideas.

Lemma 23.2. *In the same context (of Equation (386)), there exists a constant $C < \infty$ (depending only on k) such that*

$$|\text{Cov}_{\mu_{\mathbb{Z}^2}}[P, H_1 H_2]| \leq C. \quad (392)$$

Before diving into its proof, let us state a convenient intermediary result.

Lemma 23.3. *In the same context (of Equation (386)), there exists a constant $C < \infty$ such that*

$$\mu_{\mathbb{Z}^2}[\mu_{\mathbb{Z}^2}[H_1 H_2|\sigma_\circ, \bar{\omega}]^2] \leq C. \quad (393)$$

Proof. Since $\partial\mathcal{A} \subset \omega^+$, the set $\partial\mathcal{A}$ is contained in a single \mathcal{X} -vertex. Identify $\partial\mathcal{A}$ with this \mathcal{X} -vertex. It therefore makes sense to consider the unique \mathcal{X} -path from \bar{u}'_2 to $\partial\mathcal{A}$, which we denote by $p^{\bar{u}'_2 \partial\mathcal{A}}$.

Let \mathcal{D} denote an extremely large domain. By arguing as in Section 21 (Equation (355)), we observe that

$$|\mu_{\mathcal{D}}^+[H_1 H_2|\mathcal{X}]| \leq 4|p^{\bar{u}'_2 \partial\mathcal{A}} \cap p^{\bar{u}'_1 \bar{u}_1} \cap V_\bullet(\mathcal{X})|. \quad (394)$$

Notice that $p^{\bar{u}'_2 \partial\mathcal{A}}$ is a truncated version of the \mathcal{X} -path $p^{\bar{u}'_2 \bar{u}_2}$, which will work in our favour.

The right-hand side of the previous display is a random variable, and our objective is to bound its second moment with a universal constant.

Define the following geometric objects (illustrated in Figure 24):

$$S := \text{Ball}_{\frac{1}{\delta}}; \quad A := \text{Annulus}_{\frac{4}{\delta}, \frac{1}{\delta}}; \quad R^1 := [\frac{1}{2\delta}, \frac{1}{\delta}] \times [-\frac{4}{\delta}, \frac{4}{\delta}], \quad (395)$$

and R^2, R^3, R^4 the rotations of R^1 by the angles $\pi/2, \pi,$ and $3\pi/2$ around $(0, 0)$.

It can be verified (by analysing Figure 24) that any self-avoiding loop $\gamma \subset \mathbb{R}^2$ surrounding exactly one point of the pair $\{\bar{u}_1, \bar{u}'_1\}$ and exactly one of $\{\bar{u}_2, \bar{u}'_2\}$ satisfies one of the following three properties:

1. $\gamma \subset S$,
2. γ crosses R^1 or R^3 horizontally, or R^2 or R^4 vertically,
3. γ realises an arm event for A .

To upper bound Equation (394), we observe that each ω -connected component in $p^{\bar{u}'_2 \partial\mathcal{A}} \cap p^{\bar{u}'_1 \bar{u}_1} \cap V_\bullet(\mathcal{X})$ must contribute to at least one of the three cases above. In fact, by definition of \mathcal{A} , at most two such components can be contained in S . Thus, we get

$$|p^{\bar{u}'_2 \partial\mathcal{A}} \cap p^{\bar{u}'_1 \bar{u}_1} \cap V_\bullet(\mathcal{X})| \leq Q, \quad (396)$$

where

$$Q := 24 + \left(\sum_{i=1,3} \#\text{AltHor}^{\bar{\omega}}(R^i) \right) + \left(\sum_{i=2,4} \#\text{AltVer}^{\bar{\omega}}(R^i) \right) + \#\text{AltArm}^{\bar{\omega}}(A). \quad (397)$$




Figure 24: For Lemma 23.3, each loop which surrounds exactly one point of each pair, must either be contained in the inner boundary of the annulus (the small square), or cross one of the four rectangles in the easy direction, or realise an arm event for the annulus. The figure contains an example for each of these three cases.

The term 24 compensates for the potential components in S (at most two), and a few other components that we may lose because of boundary effects when counting crossings (for example, in our definitions we always imposed that we start counting at the highest ω^+ -crossing, thus missing out on a potential higher ω^- -crossing).

Putting our bounds together and letting \mathcal{D} tend to \mathbb{Z}^2 yields

$$|\mu_{\mathbb{Z}^2}[H_1 H_2 | \sigma_\circ, \omega]| \leq 4Q. \quad (398)$$

It suffices to prove that $4Q$ has a uniformly bounded second moment. This follows from Lemma 19.1 (for the tails of the number of rectangle crossings) and Lemma 19.5 (for the tails of the number of arms). \square

We are now in a position to prove Lemma 23.2.

Proof of Lemma 23.2. Lemmata 23.1 and 23.3 imply that the expectations of the conditional expectations P and $H_1 H_2$ are uniformly bounded (since the second moments are). The product of the expectations is therefore bounded and it suffices to bound the expectation of the product of the two conditional expectations.

By arguing as before, we get

$$|\mu_{\mathcal{D}}^+[P H_1 H_2 | \mathcal{X}]| \leq \sum_{\pi} \prod_{ij \in \pi} 4I(\bar{u}_i \bar{u}'_i, \bar{u}_j \bar{u}'_j), \quad (399)$$

where we abusively write $\bar{u}_2 := \partial \mathcal{A}$. The sum runs over pairings of $\{1, \dots, k\}$. It suffices to bound the term corresponding to each pairing π separately.

Suppose first that $\{1, 2\} \notin \pi$. Then, each pair may be ordered such that $i \neq 1, 2$ for any $ij \in \pi$. This assumption enables us to bound each factor as in Lemma 23.1, yielding

$$\lim_{\mathcal{D} \nearrow \mathbb{Z}^2} \mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} 4I(\bar{u}_i \bar{u}'_i, \bar{u}_j \bar{u}'_j) \right] \leq \mu_{\mathbb{Z}^2} \left[\prod_{ij \in \pi} 16K_{(1/\delta), (\varepsilon/\delta), \bar{u}_i} \right]. \quad (400)$$

This leads to the desired uniform upper bound via Lemma 19.5.

Suppose now that $\{1, 2\} \in \pi$. Then, by a reasoning similar to the previous proof, we get

$$\lim_{\mathcal{D} \nearrow \mathbb{Z}^2} \mu_{\mathcal{D}}^+ \left[\prod_{ij \in \pi} 4I(\bar{u}_i \bar{u}'_i, \bar{u}_j \bar{u}'_j) \right] \leq \mu_{\mathbb{Z}^2} [4Q \prod_{ij \in \pi \setminus \{1, 2\}} 16K_{(1/\delta), (\varepsilon/\delta), \bar{u}_i}], \quad (401)$$

which is uniformly bounded by applying again our bounds for arms and crossings. \square

It remains to bound the last term in Equation (386). This proof is quite different.

Lemma 23.4. *In the same context (of Equation (386)), there exists some universal constant $C > 0$ (depending only on \mathbf{c} and k) such that*

$$|\mathrm{Var}_{\mu_{\mathbb{Z}^2}}[\mu_{\mathbb{Z}^2}[\tilde{H}_1\tilde{H}_2|\mathcal{A}]]| \leq C. \quad (402)$$

Proof. It follows immediately from the definitions and from the Markov property over \mathcal{A} that

$$\mu_{\mathbb{Z}^2}[\tilde{H}_1\tilde{H}_2|\mathcal{A}] = X_{\tilde{u}_1, \tilde{u}_2}(\mathcal{A}); \quad (403)$$

where

$$X_{v_1, v_2}(\mathcal{A}) := \begin{cases} \mu_{\mathcal{A}}^+[h(v_1)h(v_2)] & \text{if } F(\mathcal{A}) \ni v_1, v_2, \\ 0 & \text{otherwise.} \end{cases} \quad (404)$$

Write $X := X_{v_1, v_2}$. Note that X is an increasing function of the domain. Indeed, $X(\mathcal{D})$ can be expressed as the expectation of an increasing function of the branching function ψ (see Lemma 20.7), and the branching function is stochastically increasing in the choice of the domain (Lemma 20.10).

To prove the lemma, we are going to prove the following stronger statement. There exists a constant $C < \infty$ (depending only on \mathbf{c}) with the following properties. Let $B \subset \mathbb{R}^2$ denote any bounded simply connected set, and fix any two faces $v_1, v_2 \in F(\mathbb{Z}^2)$ (say v_2 is the face at $(0, 0)$ without loss of generality). Let \mathcal{B} denote the largest even domain such that $\partial\mathcal{B} \subset B$ and such that $v_2 \in F(\mathcal{B})$ (if v_2 exists), and set $\mathcal{B} = \emptyset$ otherwise. Then, we claim that

$$\mathrm{Var}_{\mu_{\mathbb{Z}^2}}[X_{v_1, v_2}(\mathcal{B})] \leq C. \quad (405)$$

The monotonicity in the domain gives that (405) implies (402). We are now going to prove Equation (405).

Let \mathcal{D}_n be the largest even domain contained in Ball_{e^n} . We claim that the increasing function

$$q : \mathbb{Z}_{\geq 1} \rightarrow \mathbb{R}, \quad n \mapsto X(\mathcal{D}_n) \quad (406)$$

is Lipschitz with a uniform Lipschitz constant. Since $q_n \leq 8^n$, it suffices to prove that the increments $q_{n+1} - q_n$ are bounded (uniformly in all input data) for all $n \geq 100$.

Consider the measure $\mu_{\mathcal{D}_{n+1}}^+$. Let \mathcal{B} denote the largest even domain containing v_2 and such that $\partial\mathcal{B} \subset \omega^-$. By the circuit estimate, there exists some uniform probability $p > 0$ such that, with probability at least p , the even domain \mathcal{D}_n contains an ω^+ -circuit which surrounds $(0, 0)$ and is connected to $\partial\mathcal{D}_{n+1}$. By inclusion, the event $\{\mathcal{B} \subset \mathcal{D}_n\}$ has probability at least p as well. By recursion (Lemma 20.9) and monotonicity (Lemma 20.10), we get

$$q_{n+1} \leq 4 + \mu_{\mathcal{D}_{n+1}}^+[X(\mathcal{B})] \leq 4 + pX(\mathcal{D}_n) + (1-p)X(\mathcal{D}_{n+1}) = 4 + pq_n + (1-p)q_{n+1}. \quad (407)$$

This implies the desired uniform Lipschitz bound $q_{n+1} - q_n \leq 4/p$.

Finally, define

$$\mathrm{Scale}(\mathcal{B}) := 0 \vee \sup\{n \in \mathbb{Z}_{\geq 1} : \mathcal{D}_n \subset \mathcal{B}\}. \quad (408)$$

To finish the proof, it suffices to prove that:

- The variance of $\mathrm{Scale}(\mathcal{B})$ is uniformly bounded, and
- $|X(\mathcal{B}) - q_{\mathrm{Scale}(\mathcal{B})}|$ is uniformly bounded.

Lemma 19.5 implies that the second moment of $\text{Scale}(B) - \text{Scale}(\mathcal{B})$ is uniformly bounded (by exploring from the outside towards the inside, we find the desired ω^+ -circuit with a uniformly positive probability at each scale, independently of the past). This implies that the variance of $\text{Scale}(\mathcal{B})$ is also uniformly bounded.

It suffices to prove the second statement. Since X is an increasing function, it suffices to uniformly upper bound $X(\mathcal{B}) - X(\mathcal{D}_{\text{Scale}(\mathcal{B})})$. By arguing as for the Lipschitz property above, we may find some uniform $p' > 0$ such that

$$\mu_{\mathcal{B}}^+ \left[\left\{ \begin{array}{l} \text{there is some } \omega^- \cap \mathcal{D}_{\text{Scale}(\mathcal{B})}\text{-circuit around } v_2 \\ \text{that is also connected to } \partial\mathcal{B} \end{array} \right\} \right] \geq p'. \quad (409)$$

Analogously to what we did above, we deduce that

$$X(\mathcal{B}) \leq 4 + p'X(\mathcal{D}_{\text{Scale}(\mathcal{B})}) + (1 - p')X(\mathcal{B}) \quad \text{and} \quad X(\mathcal{B}) - X(\mathcal{D}_{\text{Scale}(\mathcal{B})}) \leq 4/p'. \quad (410)$$

This concludes the proof. \square

Part E

Ingredient 2: Glimpse of scale invariance

Theorem 4.4 is proved in Section 24; Theorem 4.3 in Sections 25–26.

24 Free energy second derivative (Theorem 4.4)

It was proved in [Dum+22, Theorem 2] that the free energy f is twice differentiable at 0 whenever $\mathbf{c} \in [0, 2]$, with the value of the derivative depending on the parameter \mathbf{c} . Moreover, in that article, the second derivative $f''(0)$ is characterised, but not explicitly computed. We give the explicit computation below. Theorem 4.4 is an immediate consequence of Proposition 24.1 and Lemma 24.3 below.

The statements in this section use another parametrisation of the six-vertex model, namely the one given in Equation (20). Recall that this amounts to $\Delta = -\cos \zeta$.

In this section, we use the convention that the Fourier transform of any $F \in L^1(\mathbb{R})$, denoted \widehat{F} , where $\widehat{F} : \mathbb{R} \rightarrow \mathbb{C}$ is defined by $t \mapsto \int_{\mathbb{R}} e^{-itx} F(x) dx$.

24.1 Reduction to a Wiener–Hopf equation

The proposition below is essentially in [Dum+22], except that it is not written there as such. Below, we shall describe how to derive the proposition from the (intermediate) results stated in [Dum+22].

Proposition 24.1 ([Dum+22]). *For any fixed $\zeta \in [0, 2\pi/3]$, let $T : \mathbb{R} \rightarrow \mathbb{R}$ be the unique $L^1(\mathbb{R})$ -solution of*

$$T(x) - \int_0^\infty R(x-y)T(y)dy = \mathbf{e}(x), \quad (411)$$

where \mathbf{e} and R are given by Table 3 (in particular, R is defined via its Fourier transform). Then

$$f''(0) = -2 \frac{\int_0^\infty \frac{\mathbf{e}(x)}{\mathbf{e}(0)} T(x) dx}{\left(\frac{\pi}{\pi - \zeta} \int_0^\infty T(x) dx \right)^2} \quad \text{where} \quad \frac{\mathbf{e}(x)}{\mathbf{e}(0)} = \begin{cases} e^{-\frac{\pi}{\zeta}x} & \text{if } \zeta > 0, \\ e^{-\pi x} & \text{if } \zeta = 0. \end{cases} \quad (412)$$

	$\zeta = 0$	$\zeta \in (0, \frac{2}{3}\pi]$
$\widehat{R}(t)$	$\frac{e^{-\frac{ t }{2}}}{2 \cosh(\frac{t}{2})}$	$\frac{\sinh((\pi-2\zeta)\frac{t}{2})}{2 \sinh((\pi-\zeta)\frac{t}{2}) \cosh(\zeta\frac{t}{2})}$
$1 - \widehat{R}(t)$	$\frac{e^{\frac{ t }{2}}}{2 \cosh(\frac{t}{2})}$	$\frac{\sinh(\pi\frac{t}{2})}{2 \sinh((\pi-\zeta)\frac{t}{2}) \cosh(\zeta\frac{t}{2})}$
$\mathbf{e}(x)$	$e^{-\pi x} \mathbf{1}_{\mathbb{R}_{\geq 0}}(x)$	$\frac{1}{\zeta} e^{-\frac{\pi}{\zeta} x} \mathbf{1}_{\mathbb{R}_{\geq 0}}(x)$
$\widehat{\mathbf{e}}(t)$	$\frac{1}{\pi+it}$	$\frac{1}{\pi+i\zeta t}$
$\alpha(t), t \in \mathbb{H}^+$	$\left(\frac{1-it}{e}\right)^{\frac{-it}{2\pi}} \frac{\sqrt{2\pi}}{\Gamma(\frac{1}{2}-\frac{it}{2\pi})}$	$(1-\frac{\zeta}{\pi})^{-i(1-\frac{\zeta}{\pi})\frac{t}{2}} \cdot (\frac{\zeta}{\pi})^{-i\frac{\zeta}{\pi}\frac{t}{2}} \frac{\Gamma(1-i\frac{\zeta}{2})}{\Gamma(1-i(1-\frac{\zeta}{\pi})\frac{t}{2})} \frac{\sqrt{2(\pi-\zeta)}}{\Gamma(\frac{1}{2}-i\frac{\zeta}{\pi}\frac{t}{2})}$
$\alpha_+(0)^2$	$2 = 2\frac{\pi-\zeta}{\pi}$	$2\frac{\pi-\zeta}{\pi}$

Table 3: The summary of key quantities for the cases $\zeta = 0$ and $\zeta > 0$. The letter Γ refers to the Euler gamma function.

Remark 24.2. Basic properties of the integral equation were discussed in [Dum+22, Propositions 25 and 27]. Note that the integral kernel R has constant sign and integrates to $\widehat{R}(0) \in [-1/2, 1/2]$. By Young's convolution inequality, the integral equation (411) is contractive and therefore has a unique $L^1(\mathbb{R})$ -solution which may be written $T = \sum_{k=0}^{\infty} \mathfrak{R}^k(\mathbf{e})$ where $\mathfrak{R}(\mathbf{a}) := R * (\mathbf{1}_{\mathbb{R}_{\geq 0}} \mathbf{a})$.

Moreover, if $\zeta > 0$, then the integral kernel R has exponentially decaying tails, and so does the solution T . The case $\zeta = 0$ is similar; the integral kernel R and T have tails of order $O(x^{-2})$.

We stress that to solve (411), one first restricts it to $\mathbb{R}_{\geq 0}$, solves it on $L^1(\mathbb{R}_{\geq 0})$ and then uses the relation $T(x) = \int_0^{\infty} R(x-y)T(y)dy$ to extend T to $\mathbb{R}_{< 0}$.

Proof. First, [Dum+22, eqn. (60)] defines a function δf which it relates to $f''(0)$ by

$$f''(0) = - \lim_{\alpha \searrow 0} \frac{2}{\alpha^2} \cdot \delta f(Q(\frac{1-\alpha}{2})). \quad (413)$$

The function Q is defined in [Dum+22, eqn. (16)], but is not important to us. The asymptotics of $f(Q(\frac{1-\alpha}{2}))$ as $\alpha \searrow 0$ are then computed in [Dum+22, Sec. 7.2] for $0 \leq c < 2$ and [Dum+22, Sec. 7.3] for $c = 2$. Combining them with (413) we obtain

$$f''(0) = - \frac{2}{C_{\Delta}^2} \int_0^{\infty} \frac{\mathbf{e}(x)}{\mathbf{e}(0)} T(x) dx, \quad (414)$$

with C_{Δ} defined at the start of [Dum+22, Sec. 7]. The value of C_{Δ} is computed in [Dum+22, Prop. 25] as

$$C_{\Delta} = \frac{\pi}{\pi - \zeta} \int_0^{+\infty} T(\lambda) d\lambda. \quad (415)$$

Inserting this into (414), we obtain (412). \square

The purpose of the next sections is to compute the ratio of the integrals in Equation (412), or equivalently to derive the following result. Together with the previous proposition, it implies Theorem 4.4.

Lemma 24.3. *For any $\zeta \in [0, 2\pi/3]$ (corresponding to $\mathbf{c} \in [1, 2]$), we have*

$$\frac{\int_0^\infty \frac{\mathbf{e}(x)}{\mathbf{e}(0)} T(x) dx}{\left(\int_0^\infty T(x) dx \right)^2} = \frac{\pi^2}{4(\pi - \zeta)}. \quad (416)$$

We proceed in several steps. First, we rephrase the Wiener–Hopf equation in Fourier space. Then, we express T in terms of the solutions of a well-chosen Riemann–Hilbert problem. This new expression enables an explicit calculation of the ratio in Equation (412).

24.2 Fourier transform of the Wiener–Hopf equation

Equation (411) is an integral equation that can be solved explicitly via the Wiener–Hopf method [WH31], see [AF03, p. 7.4.1] for a modern exposition. Let $T = T_\uparrow - T_\downarrow$ where $T_\uparrow := \mathbb{1}_{\mathbb{R}_{\geq 0}} T$ and $T_\downarrow := -\mathbb{1}_{\mathbb{R}_{< 0}} T$. The Fourier transforms of T_\uparrow and T_\downarrow then extend to holomorphic functions on the lower and upper half plane, respectively. Equation (411) may then be written in Fourier space, leading to a Riemann–Hilbert problem for the corresponding holomorphic extensions. This Riemann–Hilbert problem can be solved uniquely and explicitly, and the solution for T is given by the inverse Fourier transform. Once this is done, it is straightforward to calculate the ratio of the two integrals in (412).

We now implement this strategy. Let $\mathbb{H}^\pm \subset \mathbb{C}$ denote the open upper and lower half plane. Let $\mathcal{O}(\mathbb{C} \setminus \mathbb{R})$ denote the set of holomorphic functions φ on $\mathbb{C} \setminus \mathbb{R}$, such that $\varphi|_{\mathbb{H}^+}$ has a continuous extension to $\overline{\mathbb{H}^+}$, and such that $\varphi|_{\mathbb{H}^-}$ has a continuous extension to $\overline{\mathbb{H}^-}$. For any such $\varphi \in \mathcal{O}(\mathbb{C} \setminus \mathbb{R})$, we define the continuous functions φ_+ and φ_- on \mathbb{R} via

$$\varphi_\pm(t) := \lim_{\lambda \rightarrow 0^+} \varphi(t \pm i\lambda). \quad (417)$$

In the context of our Wiener–Hopf equation, define

$$\tilde{T} : \mathbb{C} \setminus \mathbb{R}, t \mapsto \begin{cases} \tilde{T}_\uparrow(t) := \int_{\mathbb{R}} e^{-itx} T_\uparrow(x) dx & \text{if } t \in \mathbb{H}^-, \\ \tilde{T}_\downarrow(t) := \int_{\mathbb{R}} e^{-itx} T_\downarrow(x) dx & \text{if } t \in \mathbb{H}^+. \end{cases} \quad (418)$$

Notice that $\tilde{T} \in \mathcal{O}(\mathbb{C} \setminus \mathbb{R})$ by the dominated convergence theorem, $\lim_{|t| \rightarrow \infty} \tilde{T}(t) = 0$ by integration by parts, and $\tilde{T}_- = \tilde{T}_\uparrow$ and $\tilde{T}_+ = \tilde{T}_\downarrow$ by the definition of the Fourier transform.

In Fourier space, Equation (411) is written as

$$(1 - \widehat{R})\widehat{T}_\uparrow - \widehat{T}_\downarrow = (1 - \widehat{R})\tilde{T}_- - \tilde{T}_+ = \widehat{\mathbf{e}}. \quad (419)$$

We shall see below that this equation admits a unique solution $\tilde{T} \in \mathcal{O}(\mathbb{C} \setminus \mathbb{R})$ that tends to 0 as $|t| \rightarrow \infty$. We explicitly solve it via two Riemann–Hilbert problems: we first solve a variant of the equation without the driving term \mathbf{e} , and then we use our first solution to solve a second Riemann–Hilbert problem which incorporates the driving term.

24.3 Two Riemann–Hilbert problems

Branch cuts will start to play a role. We view the Gamma function Γ involved in the definition of α as a holomorphic function on the right half plane $\{t \in \mathbb{C} : \operatorname{Re}(t) > 0\}$, and we view the function $t \mapsto t^t$ as a holomorphic function on the set $\mathbb{C} \setminus \mathbb{R}_{\leq 0}$. Recall that $t \mapsto t^t$ extends continuously to the point $t = 0$, where it takes the value 1.

Define $\alpha \in \mathcal{O}(\mathbb{C} \setminus \mathbb{R})$ as follows:

$$\alpha(t) := \begin{cases} \text{the value given by Table 3} & \text{if } t \in \mathbb{H}^+, \\ 1/\alpha(-t) & \text{if } t \in \mathbb{H}^-. \end{cases} \quad (420)$$

Lemma 24.4 (First Riemann-Hilbert problem). *The function α satisfies the following properties:*

- (i) *The functions α and α_{\pm} do not vanish,*
- (ii) *The function α satisfies $\lim_{|t| \rightarrow \infty} \alpha(t) = 1$,*
- (iii) *We have $\alpha_-/\alpha_+ = 1 - \widehat{R}$.*

Proof. We view Γ as a holomorphic function that is non-vanishing on the right half plane, and that Stirling's approximation gives an estimate up to a factor of order $1 + o(1)$ as $|t| \rightarrow \infty$. The first two properties follow by basic manipulations.

Finally, we must show that for any $t \in \mathbb{R}$, we have

$$\frac{\alpha_-(t)}{\alpha_+(t)} = \frac{1}{\alpha_+(-t)\alpha_+(t)} = (1 - \widehat{R})(t). \quad (421)$$

This is an elementary consequence of the well-known identities

$$\Gamma(1 + i\lambda)\Gamma(1 - i\lambda) = \frac{\pi\lambda}{\sinh(\pi\lambda)}; \quad \Gamma\left(\frac{1}{2} + i\lambda\right)\Gamma\left(\frac{1}{2} - i\lambda\right) = \frac{\pi}{\cosh(\pi\lambda)}, \quad (422)$$

which are valid for any $\lambda \in \mathbb{C} \setminus i\mathbb{Z}$. □

Lemma 24.4 enables to rewrite (419) as

$$\alpha_- \widetilde{T}_- - \alpha_+ \widetilde{T}_+ = \alpha_+ \widehat{\mathbf{e}}. \quad (423)$$

Lemma 24.5. *There exists a unique solution $G \in \mathcal{O}(\mathbb{C} \setminus \mathbb{R})$ to the equation*

$$G_- - G_+ = \alpha_+ \widehat{\mathbf{e}} \quad (424)$$

with $\lim_{|t| \rightarrow \infty} G(t) = 0$. In particular, one has $\widetilde{T} = G/\alpha$ and $\widehat{T}_{\uparrow}(t) = \alpha_+(-t)G_-(t)$.

Proof. Equation (423) implies that $G = \alpha \widetilde{T}$ is a solution to Equation (424). This solution also satisfies $\widehat{T}_{\uparrow}(t) = \widetilde{T}_-(t) = G_-(t)/\alpha_-(t) = \alpha_+(-t)G_-(t)$.

It suffices to prove uniqueness of the solution G . Let $G' \in \mathcal{O}(\mathbb{C} \setminus \mathbb{R})$ denote another solution. Then $D := G - G' \in \mathcal{O}(\mathbb{C} \setminus \mathbb{R})$ satisfies

$$D_- - D_+ = 0 \quad \text{and} \quad \lim_{|t| \rightarrow \infty} D(t) = 0. \quad (425)$$

The first relation implies that D extends continuously to \mathbb{C} and by Morera's theorem this extension is an entire function. The limit at infinity being 0, we get $D = 0$ and $G = G'$ by the second relation. □

Lemma 24.6. *We have $\widehat{T}_{\uparrow} = \widetilde{T}_-$. The functions $\widehat{T}_{\uparrow} = \widetilde{T}_-$, and $\widetilde{T}|_{\mathbb{H}^-}$ take the explicit form*

$$\widehat{T}_{\uparrow}(t) = \alpha_+(-t)\alpha(t_{\zeta})\widehat{\mathbf{e}}(t); \quad \widetilde{T}|_{\mathbb{H}^-}(t) = \alpha(-t)\alpha(t_{\zeta})\widehat{\mathbf{e}}(t). \quad (426)$$

Proof. The unique solution $G \in \mathcal{O}(\mathbb{C} \setminus \mathbb{R})$ of Equation (424) is given by

$$G(t) := (\alpha(t_{\zeta}) - \alpha(t)\mathbf{1}_{\mathbb{H}^+}(t))\widehat{\mathbf{e}}(t) \quad \text{where} \quad t_{\zeta} := \begin{cases} i\pi/\zeta & \text{if } \zeta > 0, \\ i\pi & \text{if } \zeta = 0, \end{cases} \quad (427)$$

is the unique pole of $\widehat{\mathbf{e}}$ in \mathbb{C} . Indeed, the function G defined in (427) is clearly holomorphic and has the desired limit at infinity. The jump condition is straightforwardly verified.

The expressions for \widehat{T}_{\uparrow} and $\widetilde{T}|_{\mathbb{H}^-}$ now follow from the previous lemma. □

24.4 Proof of Lemma 24.3

We are now in a position to prove Lemma 24.3.

Proof of Lemma 24.3. Recall that $\tilde{T}|_{\mathbb{H}^-}$ is a holomorphic function that extends continuously to the real line, where it equals \widehat{T}_\uparrow . As a consequence, using Lemma 24.6 in the third equality gives

$$\int_0^\infty T(x)dx = \int_{\mathbb{R}} T_\uparrow(x)dx = \widehat{T}_\uparrow(0) = \alpha_+(0)\alpha(t_\zeta)\widehat{\mathbf{e}}(0) = \frac{1}{\pi}\alpha_+(0)\alpha(t_\zeta). \quad (428)$$

The lemma therefore follows from the value of $\alpha_+(0)$ given in Table 3 and the identity

$$\int_0^\infty \mathbf{e}(x)T(x)dx = \frac{|t_\zeta|}{2\pi^2}\alpha(t_\zeta)^2. \quad (429)$$

To prove the latter, observe first that T_\uparrow and \mathbf{e} belong to $L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. Indeed, we already know that $T \in L^1(\mathbb{R})$. Since T_\uparrow is also bounded, we get $T_\uparrow \in L^2(\mathbb{R})$. For \mathbf{e} the claim is obvious.

Plancherel's theorem can therefore be used to obtain

$$\int_0^\infty \mathbf{e}(x)T(x)dx = \int_{\mathbb{R}} \mathbf{e}(x)T_\uparrow(x)dx = \frac{1}{2\pi} \int_{\mathbb{R}} \widehat{\mathbf{e}}(-t)\widehat{T}_\uparrow(t)dt. \quad (430)$$

Now, $\tilde{T}|_{\mathbb{H}^-}$ is holomorphic and extends continuously to \mathbb{R} , where it equals \widehat{T}_\uparrow . Cauchy's theorem implies that for some small $\varepsilon > 0$,

$$\int_0^\infty \mathbf{e}(x)T(x)dx = \frac{1}{2\pi} \int_{\mathbb{R}-\varepsilon i} \widehat{\mathbf{e}}(-t)\tilde{T}(t)dt. \quad (431)$$

(We use that the integrand is $O(1/t^2)$ as $|t| \rightarrow \infty$.) Cauchy's integral formula (again using the $O(1/t^2)$ decay) gives

$$\int_0^\infty \mathbf{e}(x)T(x)dx = \frac{|t_\zeta|}{\pi}\tilde{T}(-t_\zeta) = \frac{|t_\zeta|}{\pi}\alpha(t_\zeta)^2\widehat{\mathbf{e}}(-t_\zeta) = \frac{|t_\zeta|}{2\pi^2}\alpha(t_\zeta)^2. \quad (432)$$

This is exactly (429), so that the proof is completed. \square

25 Proof of Theorem 4.3

This section establishes Theorem 4.3. The proof relies on Ingredient 3 (Part D). In particular, we use our bound on arm exponents (Theorem 4.9) and flip domination (Theorem 4.10). Finally, Proposition 25.3 plays an important role – it is stated in the current section, but its proof is deferred to Section 26.

The proof is structured as follows. We first state a preparatory identity, linking the free energy to a large deviations principle. Then, we present an intuitive version of our strategy, which has some clear flaws. This guides our rigorous proof, which consists of three steps. The first step entails Proposition 25.3, which is thus deferred to Section 26.

25.1 Preparatory identity

Let us first introduce the notion of *constant even height circuit* of h . It is a circuit γ of edges in E_\circ with the property that the height difference between the endpoints is equal to zero. By definition, all faces corresponding to endpoints in γ have the same height, which we denote by $h(\gamma)$. We view each edge of the circuit as a line segment in \mathbb{R}^2 connecting the centres of the two corresponding faces; so that the circuit itself can be seen as a random subset of \mathbb{R}^2 in this way.

Consider a *topological annulus*, that is, a subset $A \subset \mathbb{R}^2$ that is bounded and topologically equivalent to an annulus. For any topological annulus A and any $k > 0$, define

$$\text{Circuit}^{+k}(A) := \left\{ \begin{array}{l} A \text{ contains two even height circuits } \gamma_0, \gamma_+ \text{ of } h \text{ with nontrivial} \\ \text{winding in } A, \text{ such that } \gamma_0 \text{ is the outermost even height circuit} \\ \text{included in } A, \text{ and such that } h(\gamma_+) - h(\gamma_0) \geq k \end{array} \right\}. \quad (433)$$

Notice that this event is measurable in terms of the gradient of h , that is, it has a well-defined $\mathbb{P}_{\mathbb{Z}^2}$ -probability. We shall start by giving an alternative way of characterising $f''(0)$, where f is the free energy.

Proposition 25.1 (Variational principle). *For any $\alpha \in (0, 1)$, we have*

$$\lim_{\rho \rightarrow \infty} \lim_{L \rightarrow \infty} \frac{1}{4\rho L^2} \log \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+\alpha L}(A_{\rho,L})] = f(\alpha) - f(0), \quad (434)$$

where $A_{\rho,L} := [(\rho + 1)L \times L] \setminus [\rho L \times 0]$, while $[a \times b]$ is as defined in Equation (278).

Proposition 25.1 can be understood as a *variational principle*. It tells us that the likelihood of height function deviations of order L in a domain of size L can be expressed in terms of the free energy functional (which provides the entropy of the system at different slopes). For height functions, such a principle was first established in the work of Cohn, Kenyon, and Propp in the context of the domino tiling model [CKP01] (the domino tiling model is integrable, but the proof of the variational principle does not use the integrable structure). All the essential ingredients for the proof of Proposition 25.1 are already present in [CKP01]. The variational principle was later proved in a general finite-range setting (including the six-vertex model) in a work of Sheffield [She05]. Finally, we mention [LT24], which also establishes the variational principle for height functions which are potentially infinite-range (also including the six-vertex model). A proof of Proposition 25.1 may thus be found in any of the references [CKP01; She05; LT24]; we do not reproduce it here.

25.2 Naive strategy and its problems

To explain our strategy for establishing Theorem 4.3, we first “prove” the result by using two assumptions that are sensible, but which are not exactly true or, at least, which we cannot rigorously establish at this stage. Below, \approx means that the events have the same probability up to the relevant precision.

Assume first that

$$\mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+k}(A_{\rho,L})] \approx \mathbb{P}_{\mathbb{Z}^2} \left[\left\{ \begin{array}{l} \text{the average of } h \text{ on the line segment } [\rho L \times 0] \text{ minus} \\ \text{the average of } h \text{ on } \partial[(\rho + 1)L \times L] \text{ is at least } k \end{array} \right\} \right]. \quad (435)$$

If this assumption were true, then the GFF convergence (a hypothesis in Theorem 4.3) and a GFF calculation that can be made precise quite easily (we shall do this later in full

detail) would imply that for fixed k , we would have

$$\begin{aligned} & \lim_{L \rightarrow \infty} \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+k}(A_{\rho,L})] \\ & \approx \lim_{L \rightarrow \infty} \mathbb{P}_{\mathbb{Z}^2} \left[\left\{ \begin{array}{l} \text{the average of } h \text{ on the line segment } \llbracket \rho L \times 0 \rrbracket \text{ minus} \\ \text{the average of } h \text{ on } \partial \llbracket (\rho + 1)L \times L \rrbracket \text{ is at least } k \end{array} \right\} \right] \approx e^{-\frac{(4\rho)k^2}{2\sigma^2}}. \end{aligned} \quad (436)$$

Let us now make a second assumption, which is even more optimistic: suppose that the previous formula is true even when we set $k = \alpha L$ rather than keeping k fixed. Such an assumption is a big stretch, since GFF convergence covers events with a probability of order 1, while setting $k = \alpha L$ means passing to events in the large deviation regime.

Such an assumption would allow us to combine Equations (434) and (436) with $k = \alpha L$ to get

$$f(\alpha) - f(0) \approx -\frac{\alpha^2}{2\sigma^2}. \quad (437)$$

Since f is twice differentiable at $\alpha = 0$, this is equivalent to Theorem 4.3.

In the remainder of this section, we describe a rigorous proof of Theorem 4.3. It does not really establish the previous two assumptions, but it is inspired by the same ideas. Remark that the equality $\sigma^2 = -1/f''(0)$ may be viewed as the combination of two inequalities; unfortunately, we will have to treat the two inequalities separately in each step.

25.3 Step 1. Splitting the large deviation event into “independent” events

This step circumvents the problems with the second assumption. We essentially split up the large deviation event which has a probability of order e^{-cL^2} into $O(L^2)$ “independent” events which have a probability of order e^{-c} .

Definition 25.2 (Straightened annulus). An η -straightened path is a path in \mathbb{R}^2 which is a union of line segments of the square lattice graph $\eta\mathbb{Z}^2$. A (ρ, η) -straightened annulus is a topological annulus $A \subset \llbracket \rho + 1 \rrbracket \times \mathbb{R}$ such that its inner and outer boundaries are η -straightened paths, and such that the inner boundary crosses the vertical strip $[-\rho, \rho] \times \mathbb{R}$.

Proposition 25.3 (Splitting). Fix $\mathbf{c} \in [1, 2]$.

(i) **Straight bound.** For fixed $\rho, k, N \in 1000\mathbb{Z}_{\geq 1}$ with $N \geq 2k$,

$$\log \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+k}(A_{\rho,N})] \leq 4\rho N^2 \left(f\left(\frac{k-24}{N}\right) - f(0) \right). \quad (438)$$

In particular,

$$\limsup_{k \rightarrow \infty} \limsup_{N \rightarrow \infty} \frac{1}{4\rho k^2} \log \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+k}(A_{\rho,N})] \leq \frac{1}{2} f''(0). \quad (439)$$

(ii) **Curly bound.** For fixed $\rho \in 2\mathbb{Z}_{\geq 1}$ and $\varepsilon \in (0, 1/100)$, there exist constants $\eta > 0$ and $C, r = r(\rho, \varepsilon) < \infty$ such that, for sufficiently large $k \in 2\mathbb{Z}_{\geq 1}$ and for sufficiently large $N \in 2\mathbb{Z}_{\geq 1}$ (depending on k), we have

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{n} \log \left(\max_{(A_i) \in \mathfrak{A}(r, \rho, \eta, \varepsilon, n)} \prod_i \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+2\lceil (1-\varepsilon)k \rceil}(NA_i)] \right) \\ & \geq 4(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right) + 200k^2 \log c_{\text{circuit}} - C, \end{aligned} \quad (440)$$

where $\mathfrak{A}(r, \rho, \eta, \varepsilon, n)$ is the set of all families of $\lceil (1 - \varepsilon)n \rceil$ disjoint (ρ, η) -straightened annuli A_i for $1 \leq i \leq \lceil (1 - \varepsilon)n \rceil$ of diameter at most $2r = 2r(\rho, \varepsilon)$ that are contained in $\llbracket (\rho + 1) \times n \rrbracket$.

In particular,

$$\begin{aligned} \liminf_{k \rightarrow \infty} \liminf_{N \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{4(\rho+1)nk^2} \log \left(\max_{(A_i) \in \mathfrak{A}(r, \rho, \eta, \varepsilon, n)} \prod_i \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+2\lceil(1-\varepsilon)k/2\rceil}(NA_i)] \right) \\ \geq \frac{1}{2} f''(0) + \frac{100}{\rho+1} \log c_{\text{circuit}}. \end{aligned} \quad (441)$$

The proof of Proposition 25.3 is deferred to Section 26. We now turn to the second part of the proof, that corresponds to a variation of the first assumption we made.

25.4 Step 2. GFF counterparts of circuit events

First, notice that our first assumption is not yet unambiguously defined, since the notion of “average” requires the introduction of a probability measure, which we did not explicitly do. It turns out that, from the perspective of the GFF, there is a unique natural choice of probability measures. Let us describe this choice first. In this subsection, we shall work in the generality of an arbitrary topological annulus $A \subset \mathbb{R}^2$.

For any topological annulus A , let $\partial_{\text{ext}}A$ denote its exterior boundary, and let $\partial_{\text{int}}A$ denote its interior boundary. Introduce the function

$$H_A : \mathbb{R}^2 \rightarrow [0, 1], \quad x \mapsto \mathbb{P}[\{\text{a Brownian motion started at } x \text{ hits } \partial_{\text{int}}A \text{ before } \partial_{\text{ext}}A\}]. \quad (442)$$

This is the unique bounded harmonic extension of

$$\partial A \rightarrow \mathbb{R}, \quad x \mapsto \begin{cases} 1 & \text{if } x \in \partial_{\text{int}}A, \\ 0 & \text{if } x \in \partial_{\text{ext}}A, \end{cases} \quad (443)$$

to \mathbb{R}^2 . Let ν_A denote the “Laplacian measure” associated to A : it is defined as $\nu_A := -\Delta H_A$, and may also be characterised as the unique finite signed measure supported on ∂A such that for any smooth compactly supported function $F : \mathbb{R}^2 \rightarrow \mathbb{R}$, we have

$$\int F(x) d\nu_A(x) = \int \nabla F(x) \cdot \nabla H_A(x) dx. \quad (444)$$

Decompose $\nu_A := \nu_A^+ - \nu_A^-$ where both terms on the right are (positive) finite measures. By definition, ν_A^+ is supported on $\partial_{\text{int}}A$, and ν_A^- is supported on $\partial_{\text{ext}}A$. Moreover, Equation (444) with $F = H_A$ shows that

$$\nu_A^+(\partial_{\text{int}}A) = \nu_A^-(\partial_{\text{ext}}A) = \int |\nabla H_A(x)|^2 dx =: \text{Dirichlet}(H_A). \quad (445)$$

Define

$$\varphi_A := \frac{1}{\text{Dirichlet}(H_A)} \nu_A. \quad (446)$$

This measure is normalised in the sense that it decomposes as the difference $\varphi_A^+ - \varphi_A^-$ of two probability measures with disjoint support. We are now ready to state the main result of this subsection. Recall that a *convergence sequence* is a sequence $(\delta_n)_n$ tending to zero such that $h^{(\delta_n)}$ converges to a multiple of the GFF.

Proposition 25.4 (Formal version of the first assumption). *For $(\delta_n)_n$ a convergence sequence and A a topological annulus,*

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{k^2} \log \frac{\mathbb{P}_{\mathbb{Z}^2}[\{\langle h, \varphi_{A/\delta_n} \rangle \geq k\}]}{\mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+k}(A/\delta_n)]} = 0. \quad (447)$$

Here and below, Γ denotes a normalised GFF on \mathbb{R}^2 (in the sense of Section 2.3). To prove the previous proposition, we first need some very basic information on the behaviour of the random variable $\langle \Gamma, \varphi_A \rangle$.

Lemma 25.5 (GFF analysis of $\langle \Gamma, \varphi_A \rangle$). *For every topological annulus A ,*

(i) *The random variable $\langle \Gamma, \varphi_A \rangle$ has the distribution $\mathcal{N}(0, 1/\text{Dirichlet}(H_A))$,*

(ii) *We have*

$$\lim_{k \rightarrow \infty} \frac{1}{k^2} \log \mathbb{P}[\{\langle \Gamma, \varphi_A \rangle \geq k\}] = -\frac{1}{2} \text{Dirichlet}(H_A), \quad (448)$$

(iii) *The Gaussian process Γ may be decomposed as the sum*

$$\Gamma = \Gamma^A + \langle \Gamma, \varphi_A \rangle H_A, \quad (449)$$

where Γ^A is a Gaussian process that is independent of $\langle \Gamma, \varphi_A \rangle$.

Proof. Fix any finite Dirichlet energy generalised test function μ . Since the Green function $G_{\mathbb{R}^2}$ is the inverse of the negative-Laplacian, we get

$$\text{Dirichlet}(H_A) \text{Cov}[\langle \Gamma, \varphi_A \rangle, \langle \Gamma, \mu \rangle] = \text{Cov}[\langle \Gamma, \nu_A \rangle, \langle \Gamma, \mu \rangle] \quad (450)$$

$$= \int G_{\mathbb{R}^2}(x, y) d\nu_A(x) d\mu(y) \quad (451)$$

$$= \int d\mu(y) ((-\Delta)^{-1}(-\Delta)H_A)(y) \quad (452)$$

$$= \int H_A(y) d\mu(y). \quad (453)$$

In particular, we get

$$\text{Cov}[\langle \Gamma, \varphi_A \rangle, \langle \Gamma, \varphi_A \rangle] = \frac{1}{\text{Dirichlet}(H_A)^2} \int H_A(y) d\nu_A(y) \quad (454)$$

$$\stackrel{(444)}{=} \frac{1}{\text{Dirichlet}(H_A)^2} \int |\nabla H_A(y)|^2 dy \quad (455)$$

$$= \frac{1}{\text{Dirichlet}(H_A)}; \quad (456)$$

and similarly

$$\text{Cov}[\langle \Gamma, \varphi_A \rangle, \langle \Gamma, \delta_x - \varphi_A^- \rangle] = \frac{1}{\text{Dirichlet}(H_A)} \int H_A(y) d(\delta_x - \varphi_A^-)(y) \quad (457)$$

$$= \frac{H_A(x)}{\text{Dirichlet}(H_A)}. \quad (458)$$

Item (i) follows from (456). Item (ii) follows from item (i). Let us now focus on (iii). Since Γ is a centred Gaussian process, it must clearly have a decomposition of the form $\Gamma = \Gamma^A + \langle \Gamma, \varphi_A \rangle F$. It suffices to show that $F = H_A$, which is done via (456) and (458). \square

The next lemma is a direct corollary of the previous lemma.

Lemma 25.6. *For a convergent sequence $(\delta_n)_n$ and a topological annulus A ,*

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{k^2} \log \mathbb{P}_{\mathbb{Z}^2}[\{\langle h, \varphi_{A/\delta_n} \rangle \geq k\}] = \lim_{k \rightarrow \infty} \frac{1}{k^2} \log \mathbb{P}[\{\langle \sigma \Gamma, \varphi_A \rangle \geq k\}] \quad (459)$$

$$= -\frac{1}{2\sigma^2} \text{Dirichlet}(H_A), \quad (460)$$

where σ^2 is the variance of the limiting GFF along $(\delta_n)_n$.

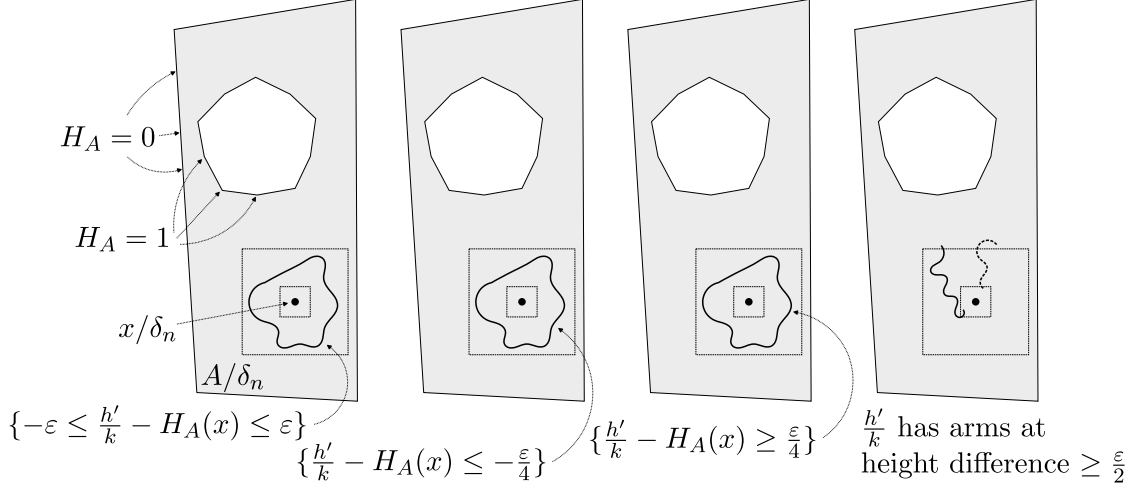


Figure 25: By topological considerations, one of the following three events must occur for the scaled height function $(h')^{(\delta_n)}/k$: the event in the statement of Lemma 25.7 (left picture), the event $E_- \cup E_+$ in the proof (middle pictures), or the event E_* in the proof (right picture).

The heavy lifting is now done in the following lemma. Below, a subset of $F(\mathbb{Z}^2)$ is seen as a subset of \mathbb{R}^2 by considering the union of the line segments joining the middles of adjacent faces.

Lemma 25.7. *Fix $\mathbf{c} \in [1, 2]$. For any convergence sequence $(\delta_n)_n$, any topological annulus A and any $\varepsilon > 0$, we may find some constant $\alpha > 0$ such that the following holds true. Let h' denote the representative of h such that $\varphi_{A/\delta_n}^-(h') \in [0, 2)$ (this depends implicitly on n). Suppose that $x \in A$ and $r > 0$ are chosen such that $\text{Ball}_r(x) \subset A$. Then,*

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P}_{\mathbb{Z}^2} \left[\left\{ \begin{array}{l} \{H_A(x) - \varepsilon \leq \frac{h'}{k} \leq H_A(x) + \varepsilon\} \text{ contains} \\ \text{a nontrivial circuit in } \text{Annulus}_{r, \alpha r}(x)/\delta_n \end{array} \right\} \middle| \{ \langle h, \varphi_{A/\delta_n} \rangle \geq k \} \right] = 1. \quad (461)$$

We start with an informal explanation of the statement of this lemma. The height function $h^{(\delta_n)}$ converges to a multiple of the GFF in the sense of finite-dimensional distributions. In particular, the distribution of $(h')^{(\delta_n)}/k$ in the conditional measure $\mathbb{P}_{\mathbb{Z}^2}[\cdot | \{ \langle h, \varphi_{A/\delta_n} \rangle \geq k \}]$ converges to $H_A + \sigma \Gamma^A/k$ as n tends to infinity (Lemma 25.5). This implies that in the double limit, $(h')^{(\delta_n)}/k$ converges to the deterministic function H_A , but only in the sense of finite-dimensional distributions.

The lemma provides a link between this distributional convergence, and the appearance of circuits in the level sets of $(h')^{(\delta_n)}/k$ (which are defined on the microscopic level). More precisely, it says that with high probability, we may find a microscopic circuit around any point $x \in A$ along which the height of $(h')^{(\delta_n)}/k$ is close to $H_A(x)$ (Figure 25).

Proof. Define the following three events:

$$E_- := \left\{ \left\{ \frac{h'}{k} \leq H_A(x) - \frac{\varepsilon}{4} \right\} \text{ contains a nontrivial} \right\}; \quad (462)$$

$$E_+ := \left\{ \left\{ \frac{h'}{k} \geq H_A(x) + \frac{\varepsilon}{4} \right\} \text{ contains a nontrivial} \right\}; \quad (463)$$

$$E_* := \left\{ \begin{array}{l} \text{we may find some } a \in 2\mathbb{Z} \text{ such that } \{h+a \leq 0\} \subset F(\mathbb{Z}^2) \\ \text{and } \{h+a \geq \frac{\varepsilon}{2}k\} \text{ both contain a path connecting the} \\ \text{two boundary components of } \text{Annulus}_{r,\alpha r}(x)/\delta_n \end{array} \right\}. \quad (464)$$

We claim that if the event in the statement of the lemma does *not* occur, then the event $E_- \cup E_+ \cup E_*$ must occur. The claim is true for topological reasons (Figure 25). Indeed, if the event in the statement does not occur, then one of

$$\{h'/k - H_A(x) \geq 7\varepsilon/8\} \quad \text{or} \quad \{h'/k - H_A(x) \leq -7\varepsilon/8\} \quad (465)$$

must contain an arm of the annulus. If the event $E_- \cup E_+$ does not occur, then both

$$\{h'/k - H_A(x) \geq -3\varepsilon/8\} \quad \text{and} \quad \{h'/k - H_A(x) \leq 3\varepsilon/8\} \quad (466)$$

must contain an arm of the annulus. But if one of (465) and both of (466) contain an arm, then E_* occurs. This proves the claim.

By the claim, it suffices to prove that the events E_- , E_+ , and E_* each occur with a low conditional probability in the limit. We first handle E_* by fixing a suitable value for α . Note that the bound on arm exponents (Theorem 4.9) implies that

$$\mathbb{P}_{\mathbb{Z}^2}[E_*] \leq \alpha^{c\varepsilon^2 k^2/4}. \quad (467)$$

If α is sufficiently small, then Lemma 25.6 implies that this probability is negligible compared to the probability of the conditioning event. Fix such an $\alpha > 0$ from now on.

It suffices to prove that E_- occurs with a conditional probability tending to zero (the case of E_+ is similar). Let π_n denote the uniform probability measure on a circle of radius α/δ_n centred at x/δ_n . By flip domination (Theorem 4.10), we get

$$\mathbb{P}_{\mathbb{Z}^2} \left[\left\{ \frac{\langle h, \pi_n - \varphi_{A/\delta_n}^- \rangle}{k} \leq H_A(x) - \frac{\varepsilon}{8} \right\} \middle| E_- \cap \{ \langle h, \varphi_{A/\delta_n} \rangle \geq k \} \right] \geq \frac{1}{2}. \quad (468)$$

But without E_- as conditioning event, Lemma 25.5 tells us that

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P}_{\mathbb{Z}^2} \left[\left\{ \frac{\langle h, \pi_n - \varphi_{A/\delta_n}^- \rangle}{k} \leq H_A(x) - \frac{\varepsilon}{8} \right\} \middle| \{ \langle h, \varphi_{A/\delta_n} \rangle \geq k \} \right] = 0. \quad (469)$$

This implies that the conditional probability of E_- tends to zero. \square

Proof of Proposition 25.4. Consider a convergence sequence (δ_n) . Let σ^2 be the variance of the limiting GFF. We view Equation (447) as the combination of two inequalities, which are treated somewhat differently. More precisely, it suffices to prove that, for fixed $\varepsilon > 0$,

$$\liminf_{k \rightarrow \infty} \liminf_{n \rightarrow \infty} \frac{\mathbb{P}_{\mathbb{Z}^2}[\{ \langle h, \varphi_{A/\delta_n} \rangle \geq k \}]}{\mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+k+4}(A/\delta_n)]} \geq \frac{1}{4}; \quad (470)$$

$$\liminf_{k \rightarrow \infty} \liminf_{n \rightarrow \infty} \frac{\mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+k}(A/\delta_n)]}{\mathbb{P}_{\mathbb{Z}^2}[\{ \langle h, \varphi_{A/\delta_n} \rangle \geq (1+4\varepsilon)k \}]} \geq \frac{1}{4}. \quad (471)$$

Proof of Equation (470). In fact, we shall prove this inequality for fixed k and n using flip domination. Define γ_0 to be the outermost even height circuit in A/δ_n , and γ_ℓ to be the outermost even height circuit in A/δ_n with $h(\gamma_\ell) - h(\gamma_0) = \ell$. These circuits may or may not be well-defined (if they are not, we set them equal to \emptyset). Introduce the event $C_\ell := \{\gamma_\ell \neq \emptyset\}$. For $\ell \in 2\mathbb{Z} \setminus \{0\}$, define the random variables

$$X_0 := h(\gamma_0) - \varphi_{A/\delta_n}^-(h); \quad (472)$$

$$X_\ell := \varphi_{A/\delta_n}^+(h) - h(\gamma_\ell). \quad (473)$$

On C_ℓ , we have that

$$\langle h, \varphi_{A/\delta_n} \rangle = X_0 + \ell + X_\ell. \quad (474)$$

Flip domination (Theorem 4.10) implies that for any even $\ell \geq 4$,

$$\mathbb{P}[\{X_0 \leq 0\} \cap C_\ell] \leq \mathbb{P}[\{X_0 \leq 0\} \cap C_{-\ell+4}] = \mathbb{P}[\{X_0 \geq 0\} \cap C_{\ell-4}], \quad (475)$$

where the equality is just total flip symmetry. Since also $C_{\ell-4} \subset C_\ell$, we get

$$\mathbb{P}[\{X_0 \geq 0\} \cap C_{\ell-4}] \geq \frac{1}{2}\mathbb{P}[C_\ell]. \quad (476)$$

Using flip domination (Theorem 4.10) again, we obtain

$$\mathbb{P}[\{X_{\ell-4} \geq 0\} | \{X_0 \geq 0\} \cap C_{\ell-4}] \geq \frac{1}{2}. \quad (477)$$

Putting things together, we arrive at

$$\mathbb{P}[\{\langle h, \varphi_{A/\delta_n} \rangle \geq \ell - 4\}] \geq \mathbb{P}[\{X_0 \geq 0\} \cap C_{\ell-4} \cap \{X_{\ell-4} \geq 0\}] \geq \frac{1}{4}\mathbb{P}[C_\ell]. \quad (478)$$

Setting $\ell = k + 4$ yields Equation (470).

Proof of Equation (471). Define h' as in Lemma 25.7. Using Lemma 25.7 and a union bound (see also Remark 16.2), we obtain

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P}_{\mathbb{Z}^2} \left[\left\{ \left\{ \frac{h'}{k} \leq \varepsilon \right\} \text{ and } \left\{ \frac{h'}{k} \geq 1 + 3\varepsilon \right\} \text{ contain} \right\} \left| \left\{ \langle h, \varphi_{A/\delta_n} \rangle \geq (1 + 4\varepsilon)k \right\} \right. \right] = 1. \quad (479)$$

By the intermediate value theorem, we deduce from this that

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbb{P}_{\mathbb{Z}^2} \left[\left\{ \left(\begin{array}{l} \text{there are two even height circuits} \\ \gamma_0, \gamma_+ \text{ winding nontrivially in } A/\delta_n \\ \text{with } h(\gamma_+) - h(\gamma_0) \geq (1 + \varepsilon)k \end{array} \right) \right\} \left| \left\{ \langle h, \varphi_{A/\delta_n} \rangle \geq (1 + 4\varepsilon)k \right\} \right. \right] = 1. \quad (480)$$

Equation (471) follows by combining this with (470), Lemma 25.6, and the inequality

$$\mathbb{P}[C_{\ell+2\ell'-4}] \geq \frac{1}{2}\mathbb{P}[C_\ell \cap C_{-\ell'}] \quad \text{for any even } \ell \geq \ell' \geq 2, \quad (481)$$

which can be obtained in the same way as in the proof of Equation (470) above. This concludes the proof. \square

25.5 Proof of Theorem 4.3

With everything in hand, the proof of the theorem will basically boil down to an analysis of the Dirichlet energy.

Proof of Theorem 4.3. Consider a convergence sequence $(\delta_n)_n$ and let σ^2 be the variance of the associated limiting GFF.

Proof of $f''(0) \geq -1/\sigma^2$. First, the straight bound in Proposition 25.3 combined with Proposition 25.4 and Lemma 25.6 yields

$$\frac{1}{2}f''(0) \geq -\frac{\text{Dirichlet}(H_{A_{\rho,1}})}{4\rho} \frac{1}{2\sigma^2}. \quad (482)$$

Introduce a new function

$$H'_\rho : \mathbb{R}^2 \rightarrow [0, 1], x \mapsto 0 \vee (1 - \text{dist}(x, \llbracket \rho \times 0 \rrbracket)) \quad (483)$$

which coincides with $H_{A_{\rho,1}}$ on ∂A_ρ . As harmonic functions minimise the Dirichlet energy,

$$\text{Dirichlet}(H_{A_{\rho,1}}) \leq \text{Dirichlet}(H'_\rho) = 4\rho + C, \quad (484)$$

where C is a constant independent of ρ . Letting ρ goes to infinity yields the result.

Proof of $f''(0) \leq -1/\sigma^2$. Consider the curly bound in Proposition 25.3(ii). Importantly, for fixed ρ and ε , there are only finitely many *shapes* of (ρ, η) -straightened annuli A_i of diameter at most $2r$ (where we say that two topological annuli have the same shape if they differ by a translation). Since there are only finitely many shapes, we may apply Proposition 25.4 and Lemma 25.6 to each of these shapes to get

$$-\lim_{n \rightarrow \infty} \min_{(A_i) \in \mathfrak{A}(r, \rho, \eta, \varepsilon, n)} \frac{1}{4(\rho+1)n} \sum_i \frac{\text{Dirichlet}(H_{A_i})}{2\sigma^2} \geq \frac{1}{2}f''(0) + \frac{100}{\rho+1} \log c_{\text{circuit}}. \quad (485)$$

Recall that $\mathfrak{A}(r, \rho, \eta, \varepsilon, n)$ is the set of all families of $K_n := \lceil (1-\varepsilon)n \rceil$ disjoint (ρ, η) -straightened annuli A_i of diameter at most $2r$ that are contained in $\llbracket (\rho+1) \times n \rrbracket$.

To deduce $f''(0) \leq -1/\sigma^2$ from (485), it suffices to demonstrate that

$$\sum_{i=1}^{K_n} \text{Dirichlet}(H_{A_i}) \geq 4\rho n(1-\varepsilon)^2, \quad (486)$$

since ρ and n can be taken large while ε can be taken small. Equation (486) follows from a straightforward calculation which we now describe.

Fix $(A_i) \in \mathfrak{A}(r, \rho, \eta, \varepsilon, n)$; we are now going to establish (486). Since the annuli are disjoint, we get

$$\sum_i \text{Dirichlet}(H_{A_i}) = \text{Dirichlet}(\sum_i H_{A_i}) \geq \int_{\llbracket \rho \times n \rrbracket} |\nabla_2(\sum_i H_{A_i})(z)|^2 dz. \quad (487)$$

For the avoidance of doubt: on the right, we only integrate the square of the *vertical* derivative over a *smaller* set than \mathbb{R}^2 , hence the inequality. By writing $z = (x, y)$, we may write this out explicitly as

$$\sum_i \text{Dirichlet}(H_{A_i}) \geq \int_{-\rho}^{\rho} \int_{-n}^n |\partial_y(\sum_i H_{A_i})(a, b)|^2 db da. \quad (488)$$

We now claim that for fixed $a \in [-\rho, \rho]$, we have

$$\int_{-n}^n |\partial_y(\sum_i H_{A_i})(a, b)|^2 db \geq 2n(1-\varepsilon)^2, \quad (489)$$

which clearly suffices for Equation (486). Notice that the total variation of $\sum_i H_{A_i}$ along the vertical line $\{a\} \times [-n, n]$ is at least $2K_n \geq 2(1-\varepsilon)n$, since each annulus contributes at least 2 to the total variation along this line. Equation (489) now follows from the lemma below. \square

Lemma 25.8. *Let $g : [-n, n] \rightarrow \mathbb{R}$ denote any continuous piecewise smooth function whose total variation is at least $2(1 - \varepsilon)n$. Then $\int g'(x)^2 dx \geq 2n(1 - \varepsilon)^2$.*

Proof. We may assume that $g(-n) = 0$ and that g is non-decreasing, by replacing it by $\tilde{g}(x) := \int_{-n}^x |g'(t)| dt$ if necessary (this function has the same total variation and Dirichlet energy). This also implies that $g(n) \geq 2(1 - \varepsilon)n$. But given the values $g(-n)$ and $g(n)$, we know which function \hat{g} minimises the one-dimensional Dirichlet energy: it is a one-dimensional harmonic function, that is, a linear interpolation. We therefore obtain

$$\int g'(x)^2 dx \geq \int \hat{g}'(x)^2 dx = \frac{g(n)^2}{2n} \geq 2n(1 - \varepsilon)^2. \quad (490)$$

This finishes the proof. \square

26 Proof of the splitting (Proposition 25.3)

We conclude this article by proving the two bounds in Proposition 25.3. This final section consists of two subsections; each one is dedicated to one of the two bounds. The first bound is significantly easier to prove than the second.

This section builds on the following ideas: the general spin representation in Section 14, the level line tree and the covariance structure in Section 20, and the combinatorial argument used in Section 19 (proof of Theorem 19.7).

26.1 Proof of the straight bound (Proposition 25.3(i))

The proof of Proposition 25.3(i) consists of four steps. First, we rewrite the quantity $\chi := \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+k}(A_{\rho, N})]$ in terms of the branching function of the spin representation. Second, we use the FKG inequality to prove (roughly) that the probability of a circuit in an extremely wide topological annulus $A_{2n\rho, N}$ is at least χ^{2n} . Third, we stack the annuli vertically to show that we can get many horizontal crossings in an extremely large rectangle with a probability lower bounded by some power of χ . Finally, we relate the final quantity to the free energy $f(\alpha)$.

Proof of Proposition 25.3(i). Fix ρ , k , and N .

Step 1: Rewriting in terms of the branching function. Recall that the law of h is the same in $\mathbb{P}_{\mathbb{Z}^2}$ and $\mu_{\mathbb{Z}^2}$. We shall work solely in the measure $\mu_{\mathbb{Z}^2}$. By inclusion of events, we get

$$\chi := \mu_{\mathbb{Z}^2}[\text{Circuit}^{+k}(A_{\rho, N})] \leq \mu_{\mathbb{Z}^2}[\text{AltCircuit}_{k/2-4}^{\bar{\omega}}(A_{\rho, N})]. \quad (491)$$

Recall the definition of maximal domains (Definition 20.8). Let $\mathcal{D} \in \mathcal{M}^+(A_{\rho, N})$ denote the largest domain such that $\partial\mathcal{D}$ surrounds the interior boundary of $A_{\rho, N}$ (and set $\mathcal{D} := \emptyset$ if such a domain does not exist). By the tower property and the Markov property, we obtain

$$\chi \leq \mu_{\mathbb{Z}^2}[\text{AltCircuit}_{k/2-4}^{\bar{\omega}}(A_{\rho, N})] = \mu_{\mathbb{Z}^2}^+[\mu_{\mathcal{D}}^+[\text{AltCircuit}_{k/2-4}^{\bar{\omega}}(A_{\rho, N})]]\mathbf{1}[\mathcal{D} \neq \emptyset]. \quad (492)$$

Now, recall the definition of the branching function and its diagonal ψ^* (Definition 20.6). Roughly speaking, $\psi^*/2$ counts the number of alternating disjoint circuits around each point. Since $\text{AltCircuit}_{k/2-4}^{\bar{\omega}}(A_{\rho, N}) \subset \text{Circuit}^{\{\psi^* \geq k-12\}}(A_{\rho, N})$, we may use the previous bound and monotonicity in domains (Lemma 20.10) to push \mathcal{D} away and get

$$\chi \leq \mu_{[(\rho+1)N \times N]}^+[\text{Circuit}^{\{\psi^* \geq k-12\}}(A_{\rho, N})]. \quad (493)$$

Step 2: Estimating the probability of a circuit in an extremely wide annulus.

We claim that for any $n \in \mathbb{Z}_{\geq 1}$,

$$\chi^{2n} \leq \mu_{\llbracket (2n\rho+1)N \times N \rrbracket}^+ [\text{Circuit}^{\{\psi^* \geq k-12\}}(A_{2n\rho, N})]. \quad (494)$$

We call $2n$ the *horizontal multiplicity*. To see that the claim is true, observe first that monotonicity in domains and the FKG inequality (Lemma 20.10) imply

$$\chi^{2n} \leq \mu_{\llbracket (2n\rho+1)N \times N \rrbracket}^+ \left[\bigcap_{i=1}^{2n} \text{Circuit}^{\{\psi^* \geq k-12\}}(A_{\rho, N} + ((2i - 2n - 1)\rho N, 0)) \right] \quad (495)$$

$$\leq \mu_{\llbracket (2n\rho+1)N \times N \rrbracket}^+ [\text{Circuit}^{\{\psi^* \geq k-12\}}(A_{2n\rho, N})]. \quad (496)$$

The second inequality is just inclusion of events (see Remark 16.2).

If we let $\mathcal{A} \in \mathcal{M}^-(\llbracket (2n\rho+1)N \times N \rrbracket)$ denote the largest domain containing the hole of the annulus, then boundary pushing yields

$$\chi^{2n} \leq \mu_{\mathbb{Z}^2} [\text{Circuit}^{\{\psi_{\mathcal{A}}^* \geq k-14\}}(A_{2n\rho, N}) | \mathcal{E}] \quad (497)$$

for any event \mathcal{E} measurable with respect to the even edges outside $\llbracket (2n\rho+1)N \times N \rrbracket$. By inclusion of events, we then conclude that

$$\chi^{2n} \leq \mu_{\mathbb{Z}^2} [\text{AltHor}_{k-20}^{\bar{\omega}}(\llbracket 2n\rho N \times N \rrbracket) | \mathcal{E}]. \quad (498)$$

Step 3: Stacking extremely wide annuli vertically. Fix $\rho', n \in \mathbb{Z}_{\geq 1}$. Apply the previous step $2n$ times in $2n$ vertically stacked rectangles, which have horizontal multiplicity $2\rho'n$. This yields

$$\chi^{4\rho'n^2} \leq \mu_{\mathbb{Z}^2} [E_+ \cap E_-], \quad (499)$$

where

$$E_{\pm} := \text{AltHor}_{n(k-22)}^{\bar{\omega}}(\llbracket \rho\rho'2nN \rrbracket \times \pm[0, 2nN]). \quad (500)$$

Step 4: Relating our quantity to the free energy. We first claim that

$$\mu_{\mathbb{Z}^2} [\text{AltCircuit}_{n(k-23)}^{\bar{\omega}}(A_{\rho\rho', 2nN}) | E_+ \cap E_-] \geq c^{n^2 N^2} \quad (501)$$

for some small fixed constant $c > 0$. We prove the claim via a straightforward *finite energy* or *surgery* argument. Indeed, condition on $E_+ \cap E_-$, and reveal the necessary edges to verify that this event occurs. Then, any *unrevealed* even edge still has a uniformly positive probability of being open (unless it connects two vertices with known opposite spins). By conditioning on the states of at most $100n^2 N^2$ such edges, we may wire up the horizontal crossings in such a way that the event $\text{AltCircuit}_{n(k-23)}^{\bar{\omega}}(A_{\rho\rho', 2nN})$ occurs. This proves the display above.

As a consequence, we get

$$\mu_{\mathbb{Z}^2} [\text{AltCircuit}_{n(k-23)}^{\bar{\omega}}(A_{\rho\rho', 2nN})] \geq c^{n^2 N^2} \cdot \chi^{4\rho'n^2}. \quad (502)$$

Finally, by flipping the coins (for each element of $V_{\bullet}(\mathcal{X})$, see Lemma 20.4) to determine the odd spins, we find

$$\mu_{\mathbb{Z}^2} [\text{Circuit}^{+2n(k-24)}(A_{\rho\rho', 2nN})] \geq 2^{-nk} \cdot c^{n^2 N^2} \cdot \chi^{4\rho'n^2}. \quad (503)$$

By the variational principle (Proposition 25.1), sending first n and then ρ' tend to infinity, we get

$$f\left(\frac{k-24}{N}\right) - f(0) \geq \frac{1}{4\rho N^2} \log \chi. \quad (504)$$

This is the desired inequality. \square

26.2 Proof of the curly bound (Proposition 25.3(ii))

We start the proof with the key step, which is called *ridge splitting* (Lemma 26.4 below). After the key step, we connect one side of the inequality in Lemma 26.4 to the variational principle (Proposition 25.1), see Lemma 26.6 below. Finally, we perform three simplification steps to connect the other side to the quantity in Proposition 25.3. The simplification steps are: cutoff of the domain diameter (Lemma 26.12), decoupling (Lemma 26.13), and straightening of the annuli (Lemma 26.14). The curly bound in Proposition 25.3 is an immediate corollary of Lemma 26.14.

The next lemma follows immediately from twice-differentiability of the free energy obtained via the Bethe Ansatz (a result that requires exact integrability). We want this section to be independent of such integrability, and therefore we provide an alternative proof of this lemma that only relies on the circuit estimate.

Lemma 26.1 (Lower bound on the free energy via RSW theory). *There exists a constant $c_{\text{fe}} < \infty$ such that $f(\alpha) - f(0) \geq -c_{\text{fe}}\alpha^2$ for any $\alpha < 1$.*

Proof. Since $f(\alpha)$ is a convex function, it suffices to bound $f(\alpha)$ for $\alpha \approx 0$. Recall Proposition 25.1. It is straightforward to see that the event $\text{AltCircuit}_{2\lceil \alpha L/4 \rceil}^{\bar{\omega}}(A_{\rho,L})$ is contained in the intersection of $O(\rho L^2 \alpha^2)$ circuit events, thus having a probability of at least $(c_{\text{circuit}})^{O(\rho L^2 \alpha^2)}$. The circuits are combined in such a way that they form alternating ω^+ -loops and ω^- -loops in the topological annulus (cf. Remark 16.2). Conditional on this event, the event $\text{Circuit}^{+\alpha L}(A_{\rho,L})$ has a probability of at least $e^{-O(\alpha L)}$ since every circuit is oriented upwards with probability $1/2$. This lower bound leads to the desired bound in the lemma. \square

26.2.1 Ridge splitting

In order to define more complicated events, it will be useful to introduce the following notion.

Definition 26.2 (Local level line forest). Recall Definition 20.1 for the level line tree and Definition 20.8 for maximal domains. Consider a fixed continuum domain $D \subset \mathbb{R}^2$. The *local level line forest* is the family of level line trees $\mathcal{L}^D := (\mathcal{X}_{\mathcal{D}})_{\mathcal{D} \in \mathcal{M}^+(D)}$. This means that we first find the maximal domains \mathcal{D} in D with $\partial \mathcal{D} \subset \omega^+$, and then construct the level line tree $\mathcal{X}_{\mathcal{D}}$ for each such domain \mathcal{D} .

We are interested in ω -circuits at a certain *forest depth*, which means that they are contained in a certain vertex of some level line tree $\mathcal{X}_{\mathcal{D}}$ at that distance from the root of $\mathcal{X}_{\mathcal{D}}$.

We can use this notion to define *ridge events*. The ridge event basically means that there is some outermost ω^+ -circuit in which we may find alternating ω^+/ω^- -circuits such that the k -th circuit still surrounds some path crossing $[a, b] \times \mathbb{R}$.

Definition 26.3 (Ridge events). Let $D \subset \mathbb{R}^2$ denote a domain, fix $k \in 2\mathbb{Z}_{\geq 1}$, and let $a \leq b$. Then, define

$$\text{Ridge}_k^{\bar{\omega}}(D, a, b) := \bigcup_p \text{AltCircuit}_k^{\bar{\omega}}(D \setminus p) = \bigcup_p \left\{ \text{some } \omega^- \cap D\text{-circuit at forest depth } 2k - 2 \text{ surrounds } p \right\}, \quad (505)$$

where the union runs over all paths p from the left of $[a, b] \times \mathbb{R}$ to the right (see Figure 26).




Figure 26: Left: the event $\text{Ridge}_4^{\bar{\omega}}(D, a, b)$. Right: the event $\{\#\text{Ridge}_2^{\bar{\omega}}(D, a, b) = 2\}$.




Figure 27: Left: the left event in Lemma 26.4. Right: the right event in Lemma 26.4.

We also define

$$\#\text{Ridge}_k^{\bar{\omega}}(D, a, b) := \#\left\{ \mathcal{D} \in \mathcal{M}^+(D) : \begin{array}{l} \text{some } \omega^- \cap D\text{-circuit at depth } 2k - 2 \text{ in } \mathcal{X}_{\mathcal{D}} \\ \text{surrounds some path } p \text{ that crosses } [a, b] \times \mathbb{R} \end{array} \right\}; \quad (506)$$

which simply counts the number of maximal domains in which the ridge event occurs.

The key step in the proof of the curly bound is *ridge splitting* (Lemma 26.4 and Figure 27). This lemma involves *ridge events*. On one side in the comparison, we consider the event that the large rectangle contains many maximal domains containing several very wide nested ω^+/ω^- -circuits (right side in Figure 27). On the other side, we chop the large rectangle into smaller ones, and consider the event that the smaller rectangles contain many maximal domains containing several nested ω^+/ω^- -circuits (left side in Figure 27).

For $\rho, k, N, \rho', n \in 2\mathbb{Z}_{\geq 1}$ with $k, N/k \geq 100$, define the following rectangles:

- The rectangle with side lengths $4n\rho'(\rho + 1)N$ and $4nN$ given by

$$R_{\text{Large}} := \llbracket 2n\rho'(\rho + 1)N \times 2nN \rrbracket, \quad (507)$$

- This rectangle may be sliced vertically into $2n\rho'$ rectangles of width $2(\rho + 1)N$ and height $4nN$, which we shall call R_i for $i = 1, \dots, 2n\rho'$ (indexed from left to right),
- Write $R'_i \subset R_i$ for the rectangle of width $2\rho N$ and height $4nN$ centred within R_i ,
- Write a_i and b_i for the x -coordinates of the left and right sides of R'_i respectively.

Define the event

$$\mathcal{E}_{\rho, k, N, \rho', n} := \left\{ \sum_{i=1}^{2n\rho'} \#\text{Ridge}_{k-4}^{\bar{\omega}}(R_i, a_i, b_i) \geq 4n^2\rho' \right\}. \quad (508)$$

Lemma 26.4 (Ridge splitting, cf. Figure 27). *For $\rho, k, N, \rho', n \in 2\mathbb{Z}_{\geq 1}$ with $k, N/k \geq 100$,*

$$\mu_{\mathbb{Z}^2}[\mathcal{E}_{\rho, k, N, \rho', n}] \geq (c_{\text{circuit}})^{100k^2n^2\rho'} \cdot \mu_{\mathbb{Z}^2}[\{\#\text{Ridge}_k^{\bar{\omega}}(R_{\text{Large}}, a_1, b_{2n\rho'}) \geq 2n\}]. \quad (509)$$

The idea of the proof is that if the event on the right occurs, then we can somehow locally rearrange our percolations $\bar{\omega}$ to create the event on the left. We do so by essentially rewiring the configurations in the vertical slits between the slightly slimmer rectangles R'_i . The rewiring is done by first exploring the correct event (conditional on the event on the right, see Figure 29), then using the circuit estimate to obtain the appropriate wiring (see Figure 30).

Before diving into the proof, let us first define an exploration process. This exploration process may succeed or fail, with \mathcal{S} denoting the event that it succeeds. This event contains the event on the right in Equation (509). Once the exploration process is done, we will prove that conditional on \mathcal{S} , the event $\mathcal{E}_{\rho,k,N,\rho',n}$ has a probability of at least $(c_{\text{circuit}})^{100k^2n^2\rho'}$. This will imply the lemma.

To facilitate the construction, define, for any $A, \omega \subset \mathbb{R}^2$ and any rectangle $R \subset \mathbb{R}^2$, the sets

$$\mathfrak{B}(A, \omega, R) := \left\{ u \in A \cap R : \begin{array}{l} \{u\} \text{ and } \partial A \text{ do not intersect the same} \\ \text{connected component of } R \setminus \omega \end{array} \right\}; \quad (510)$$

$$\mathfrak{C}(A, \omega, R) := \left\{ D : \begin{array}{l} D \text{ is a connected component of } \mathfrak{B} \text{ intersecting} \\ \text{both the left and right of } R \end{array} \right\}. \quad (511)$$

We think of \mathfrak{B} as the set of points *blocked* (or *shielded away*) from ∂A by ω within R , and of \mathfrak{C} as the set of *crossings* of such blocked points in R . See Figure 28.

In the construction below, $R'_{i,j} \subset R_i$ is defined to be the rectangle of width $2(\rho + j/k)N$ and height $4nN$ centred within R_i .

Remark 26.5. We think of $R'_{i,j}$ as “interpolating” between R'_i and R_i , in the sense that

$$R'_i = R'_{i,0} \subset R'_{i,1} \subset \cdots \subset R'_{i,k-1} \subset R'_{i,k} = R_i. \quad (512)$$

In the setting of the lemma, the exploration process is represented by a tree $\mathcal{T} = (V(\mathcal{T}), E(\mathcal{T}))$ with root R_{Large} . For each $x \in V(\mathcal{T})$ at distance j from the root, let $\mathcal{N}(x)$ denote the neighbours of x at distance $j+1$ from the root (the letter \mathcal{N} refers to *nesting* for reasons becoming clear shortly). The tree is constructed iteratively by explicitly constructing the function \mathcal{N} and the vertices at a distance j from the root (let $V_j(\mathcal{T})$ denote this set) as follows:

- $\mathcal{N}(R_{\text{Large}}) := \{\text{bounded connected components of } \mathbb{R}^2 \setminus \partial\mathcal{M}^+(R_{\text{Large}})\}$.
- For any $A \in V_1(\mathcal{T})$, set $\mathcal{N}(A) := \{\text{bounded connected components of } \mathbb{R}^2 \setminus \partial\mathcal{M}^-(A)\}$.
- For any $A \in V_2(\mathcal{T})$, set $\mathcal{N}(A) := \cup_{i=1}^{2n\rho'} \mathfrak{C}(A, \omega^+, R'_{i,k-2})$,
- For any $A \in V_3(\mathcal{T})$, set $\mathcal{N}(A) := \cup_{i=1}^{2n\rho'} \mathfrak{C}(A, \omega^-, R'_{i,k-3})$,

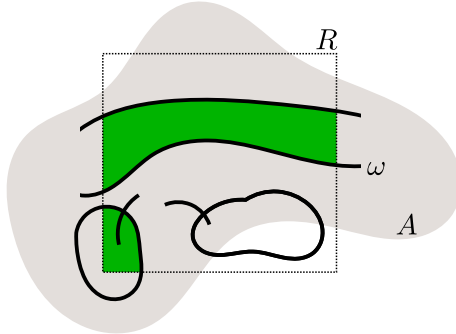


Figure 28: The green set depicts \mathfrak{B} . Only the wide green component belongs to \mathfrak{C} .




Figure 29: The nested sets in the exploration tree \mathcal{T} (assuming \mathcal{S} occurs).

- Then, repeat the last two steps, alternating $+$ and $-$, until $V_k(\mathcal{T})$ has been defined. With this definition, we are now in a position to prove the lemma.

Proof of Lemma 26.4. Let us first make some remarks. Notice that the tree is *nesting* in the sense that $A' \subset A$ for any $A' \in \mathcal{N}(A)$. Notice also that any $A \in \cup_{j>2} V_j(\mathcal{T})$ is contained in some R_i . Let us write $(V_{j,i})_i$ for the partition of $V_j(\mathcal{T})$ such that any $A \in V_{j,i}$ is contained in R_i .

Let \mathcal{S} denote the event

$$\mathcal{S} := \left\{ \text{there exist } 2n \text{ distinct elements } (O_\ell)_{\ell=1, \dots, 2n} \subset V_1(\mathcal{T}) \text{ such that for every } \right. \\ \left. 1 \leq \ell \leq 2n \text{ and } 1 \leq i \leq 2n\rho', \text{ the descendants of } O_\ell \text{ contain an element of } V_{k,i} \right\}. \quad (513)$$

Figure 29 contains a detailed illustration of this event, zooming in on a small rectangle and a single element O_ℓ between two rectangles R_i and R_{i+1} . It is tedious but straightforward to check that \mathcal{S} is included in the event on the right-hand side of Equation (509):

$$\mathcal{S} \subset \{ \# \text{Ridge}_{k-4}^{\bar{\omega}}(R_{\text{Large}}, a_1, b_{2n\rho'}) \geq 2n \}. \quad (514)$$

To finish the proof, introduce the event

$$\mathcal{C} := \left\{ \text{any } A \in \cup_{2 < j \leq k-2} V_j(\mathcal{T}) \text{ contains an } \omega^+ \text{-circuit (if } j \text{ is odd) or } \right. \\ \left. \text{an } \omega^- \text{-circuit (if } j \text{ is even) surrounding all elements of } \mathcal{N}(A) \right\}. \quad (515)$$

An impression of the event \mathcal{C} is given in Figure 30. In that figure, we already explored the tree \mathcal{T} (it is the same realisation as used for Figure 29). The dotted area represents the explored set, where we had to reveal states of the spins and edges to construct the tree. The event \mathcal{C} is the event that in the *unexplored* set, the percolations ω^+ and ω^- link up the already existing segments to form the desired circuits. In the figure, the circuits are formed of already explored segments (marked with blue) and unexplored segments (marked with red).

It is again tedious but straightforward to check that

$$\mathcal{C} \cap \mathcal{S} \subset \bigcap_{i=1}^{2n\rho'} \{ \# \text{Ridge}_{k-4}^{\bar{\omega}}(R_i, a_i, b_i) \geq 2n \} \subset \mathcal{E}_{\rho, k, N, \rho', n}. \quad (516)$$




Figure 30: The explored set (dotted area) and the event \mathcal{C} .

Thus, to finish the proof of the claim, it suffices to prove that

$$\mu_{\mathbb{Z}^2}[\mathcal{C}|\mathcal{S}] \geq (c_{\text{circuit}})^{100k^2n^2\rho'}. \quad (517)$$

Let η denote the set of even edges which have been revealed in the exploration process. For any $i = 1, \dots, 2n\rho'$ and $j = 0, \dots, k-1$, the set $R'_{i,j+1} \setminus R'_{i,j}$ consists of two vertical strips of width N/k and height $4nN$. Write $S_{i,j}^-$ and $S_{i,j}^+$ for the left and right strips respectively (see Figure 30). Then

$$\mathcal{C} \supset \mathcal{C}' := \bigcap_{i=1}^{2n\rho'} \bigcap_{j=3}^{k-2} \bigcap_{\#=\pm} \begin{cases} \text{Ver}^{\omega^+ \cup \eta}(S_{i,k-j}^\#) & \text{if } j \text{ is odd;} \\ \text{Ver}^{\omega^- \cup \eta}(S_{i,k-j}^\#) & \text{if } j \text{ is even.} \end{cases} \quad (518)$$

In Figure 30, this means, for example, that $S_{i,k-3}^+$ is crossed vertically by the union of the explored (dotted) set with ω^+ (red segment). It boils down to proving that

$$\mu_{\mathbb{Z}^2}[\mathcal{C}'|\mathcal{S}] \geq (c_{\text{circuit}})^{100k^2n^2\rho'}. \quad (519)$$

Notice that within each strip $S_{i,j}^\#$, the boundary conditions on the top and bottom of each unexplored domain are favourable for ω^+ (if j is odd) and ω^- (if j is even). We may therefore apply the partial circuit estimate (Lemma 19.4). The aspect ratio (the height divided by the width) of each strip is equal to $4nk$. The vertical crossing probability in each strip may thus be lower bounded by $(c_{\text{circuit}})^{25nk}$ (where the factor 25 royally suffices). Since there are $4n\rho'(k-4)$ such strips, we get the desired bound

$$\mu_{\mathbb{Z}^2}[\mathcal{C}|\mathcal{S}] \geq \mu_{\mathbb{Z}^2}[\mathcal{C}'|\mathcal{S}] \geq (c_{\text{circuit}})^{4n\rho'(k-4) \cdot 25nk} \geq (c_{\text{circuit}})^{100k^2n^2\rho'}. \quad (520)$$

This finishes the proof. \square

26.2.2 Connecting to the free energy

To finish the proof of the curly bound, we must connect the quantities on either side in Lemma 26.4 to the quantities on either side in Proposition 25.3. We start by relating the probability of $\mathcal{E}_{\rho,k,N,\rho',n}$ to the free energy, by analysing the right hand side in Lemma 26.4.

Lemma 26.6 (Analysis of the large deviation event). *For $\rho, k, N \in 2\mathbb{Z}_{\geq 1}$ with $k, N/k \geq 100$,*

$$\lim_{\rho' \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{\rho' n^2} \log \mu_{\mathbb{Z}^2}[\mathcal{E}_{\rho, k, N, \rho', n}] \geq 16(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right) + 100k^2 \log c_{\text{circuit}}. \quad (521)$$

Proof. We claim that

$$\mu_{\mathbb{Z}^2}[\{\# \text{Ridge}_{\tilde{c}_k}^{\tilde{\omega}}(R_{\text{Large}}, a_1, b_{2n\rho'}) \geq 2n\}] \geq \tilde{c}^{n^2 N^2} \mu_{\mathbb{Z}^2}[\text{AltCircuit}_{2n(k+4)}^{\tilde{\omega}}(A_{\rho'(\rho+1), 2nN})] \quad (522)$$

$$\geq \tilde{c}^{n^2 N^2} \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+4n(k+4)+8}(A_{\rho'(\rho+1), 2nN})]. \quad (523)$$

The second inequality is inclusion of events. The first inequality follows by a standard surgery argument (as in Step 4 in the proof for the straight bound). More precisely, conditional on the alternating crossing event, we may explore the horizontal rectangle crossings that it induces. Conditional on the exploration, we may now open $O(n^2 N^2)$ ω^+/ω^- -edges in order to realise the ridge event. Since each edge is open with a uniformly positive probability, it is easy to find the desired universal constant \tilde{c} .

By the variational principle (Proposition 25.1),

$$\begin{aligned} \lim_{\rho' \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{\rho' n^2} \log \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+4n(k+4)+8}(A_{\rho'(\rho+1), 2nN})] \\ = 16(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right). \end{aligned} \quad (524)$$

The previous lemma now yields the desired inequality. \square

Remark 26.7. By Lemma 26.1, we deduce that

$$\lim_{\rho' \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{\rho' n^2} \log \mu_{\mathbb{Z}^2}[\mathcal{E}_{\rho, k, N, \rho', n}] \geq 100k^2(-c_{\text{fe}}\rho + \log c_{\text{circuit}}). \quad (525)$$

26.2.3 Simplification steps

To close the gap between the event $\mathcal{E}_{\rho, k, N, \rho', n}$ and the curly bound in Proposition 25.3, we perform three simplification steps.

- **Diameter cutoff (Lemma 26.12).** First, we control the diameter of each domain: we impose that the outermost circuit of each ridge event has a diameter proportional to N .
- **Decoupling (Lemma 26.13).** Next, we split up the occurrence of many ridges in a single probability measure, into the occurrence of a single ridge in many independent probability measures.
- **Straightening (Lemma 26.14).** Finally, we impose that the ridge events occur in straightened domains (Definition 25.2).

These steps are slightly technical. Their order is not important, and we do not exclude the possibility that some steps may be combined.

Recall Lemma 26.4. Split R_{Large} into $4n^2 \rho'(\rho + 1)$ squares of side length $2N$, and let Q denote the set of centres of such squares. Recall the definition of the random variable $K_{2N, N, x}$ from Lemma 19.5 (cf. Theorem 19.7). We first prove the following lemma; its proof is similar to that of Theorem 19.7.

Lemma 26.8 (Bound on arm events). *Let $\rho, k, N, \rho', n \in 2\mathbb{Z}_{\geq 1}$ with $k, N/k \geq 100$. Then for any $\varepsilon, t > 0$, we get*

$$\frac{1}{\rho'n^2} \log \mu_{\mathbb{Z}^2} \left[\left\{ \sum_{x \in Q} \left\lfloor \frac{K_{2N, N, x}}{\varepsilon k} \right\rfloor \geq t|Q| \right\} \right] \leq 4(\rho + 1)(1 + t(1 - c_{\text{arm}}\varepsilon^2 k^2/9)) \log 2. \quad (526)$$

Remark 26.9. In particular, if $k \geq 9/(\varepsilon\sqrt{c_{\text{arm}}})$, and if

$$t := t_\varepsilon := \frac{1000}{\varepsilon^2 c_{\text{arm}}} (c_{\text{fe}} - \log c_{\text{circuit}}), \quad (527)$$

(recall that c_{fe} comes from Lemma 26.1 and c_{circuit} from Theorem 16.4) then the event $\{\sum_{x \in Q} \lfloor \frac{K_{2N, N, x}}{\varepsilon k} \rfloor \geq t|Q|\}$ has a much smaller probability than $\mathcal{E}_{\rho, k, N, \rho', n}$.

Proof. The union bound gives that

$$\mu_{\mathbb{Z}^2} \left[\left\{ \sum_{x \in Q} \left\lfloor \frac{K_{2N, N, x}}{\varepsilon k} \right\rfloor \geq t|Q| \right\} \right] \leq \sum_{\substack{T: Q \rightarrow \mathbb{Z}_{\geq 0}, \\ \sum_x T_x = \lceil t|Q| \rceil}} \mu_{\mathbb{Z}^2} \left[\left\{ \left(\left\lfloor \frac{K_{2N, N, x}}{\varepsilon k} \right\rfloor \right)_x \geq T \right\} \right] \quad (528)$$

Using a greedy algorithm to find a subset $Q' \subset Q$ such that the corresponding annuli do not overlap and such that $\sum_{x \in Q'} T(x) \geq T/9$, we obtain via Theorem 19.7 that for any T ,

$$\mu_{\mathbb{Z}^2} \left[\left\{ \left(\left\lfloor \frac{K_{2N, N, x}}{\varepsilon k} \right\rfloor \right)_x \geq T \right\} \right] \leq 2^{-c_{\text{arm}}\varepsilon^2 k^2 t|Q|/9}. \quad (529)$$

Bounding the number of functions T by $2^{(1+t)|Q|}$ implies that

$$\mu_{\mathbb{Z}^2} \left[\left\{ \sum_{x \in Q} \left\lfloor \frac{K_{2N, N, x}}{\varepsilon k} \right\rfloor \geq t|Q| \right\} \right] \leq 2^{(1+t)|Q| - c_{\text{arm}}\varepsilon^2 k^2 t|Q|/9} = 2^{4n^2 \rho'(\rho+1)(1+t(1-c_{\text{arm}}\varepsilon^2 k^2/9))}. \quad (530)$$

The result follows. \square

We now use the quantitative bound in the previous lemma to analyse our event $\mathcal{E}_{\rho, k, N, \rho', n}$. We first need some more definitions. Recall the definition of the level line forest from Definition 26.2.

Definition 26.10 (Local level lines forest with diameter cut-off). Consider a fixed continuum domain $D \subset \mathbb{R}^2$ and a configuration $(\sigma_\circ, \bar{\omega}, \sigma_\bullet)$. Recall that $\omega^\pm \cap D$ denotes the set of edges of ω^\pm which (as line segments embedded in \mathbb{R}^2) are entirely contained in D . Let $T_r(\omega^\pm \cap D)$ denote the set of edges in $\omega^\pm \cap D$ which belong to a connected components whose diameter is at most $2r$. Let $\mathcal{M}_r^+(D)$ denote the set of even domains \mathcal{D} which satisfy $\partial\mathcal{D} \subset T_r(\omega^+ \cap D)$ and which are maximal subject to this condition. The *local level line forest with diameter cut-off* is the family of level line trees $\mathcal{L}^{D, r} := (\mathcal{X}_{\mathcal{D}})_{\mathcal{D} \in \mathcal{M}_r^+(D)}$.

Definition 26.11 (Ridge events with diameter cut-off). The event $\text{Ridge}_{k, r}^{\bar{\omega}}(D, a, b)$ and the random variable $\#\text{Ridge}_{k, r}^{\bar{\omega}}(D, a, b)$ are defined exactly as before, except that they are defined with respect to the forest $\mathcal{L}^{D, r}$ instead of \mathcal{L}^D .

Let $\rho, k, N, \rho', n \in 2\mathbb{Z}_{\geq 1}$ with $N \geq 100k$. Recall the definition of t_ε from Remark 26.9. For $r := r(\rho, \varepsilon) := \frac{1000\rho}{\varepsilon} t_\varepsilon/8\rho$, define the event

$$\mathcal{E}_{\rho, \varepsilon, k, N, \rho', n}^* := \left\{ \sum_{i=1}^{2n\rho'} \#\text{Ridge}_{2^{\lfloor (1-\varepsilon)k/2 \rfloor}, rN}^{\bar{\omega}}(R_i, a_i, b_i) \geq (1-\varepsilon)4n^2\rho' \right\}. \quad (531)$$




Figure 31: Diameter cutoff. If alternating circuits cover a lot of vertical distance, then this creates many arm events (see the squares on the left), violating the bound obtained in Lemma 26.8.

Lemma 26.12 (Diameter cutoff). *Let $\rho, k, N, \rho', n \in 2\mathbb{Z}_{\geq 1}$ with $N \geq 100k$. Fix $\varepsilon \in (0, 1/100)$ and suppose that $k \geq 72\rho/(\varepsilon\sqrt{c_{\text{arm}}})$. Then*

$$\lim_{\rho' \rightarrow \infty} \lim_{n \rightarrow \infty} \frac{1}{\rho' n^2} \log \mu_{\mathbb{Z}^2}[\mathcal{E}_{\rho, \varepsilon, k, N, \rho', n}^*] \geq 16(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right) + 100k^2 \log c_{\text{circuit}}. \quad (532)$$

Proof. This proof is illustrated by Figure 31. By Lemma 26.8, it suffices to prove that

$$\mathcal{E}_{\rho, k, N, \rho', n} \setminus \left\{ \sum_{x \in Q} \lfloor \frac{K_{2N, N, x}}{(\varepsilon/8\rho)k} \rfloor \geq t_{\varepsilon/8\rho}|Q| \right\} \subset \mathcal{E}_{\rho, \varepsilon, k, N, \rho', n}^*. \quad (533)$$

Define the auxiliary event

$$\mathcal{A} := \left\{ \begin{array}{l} \text{the forests } (\mathcal{L}^{R_i})_i \text{ contain at most } 2\varepsilon n^2 \rho' \text{ trees} \\ \text{that contain a vertex of depth } 2\lfloor \varepsilon k/4 \rfloor \text{ that} \\ \text{has a diameter larger than } 2rN \end{array} \right\}. \quad (534)$$

Suppose that \mathcal{A} does not occur. Then, every “big” tree (that is, a tree that contains a vertex of depth $2\lfloor \varepsilon k/4 \rfloor$ that has a diameter larger than $2rN$) contributes at least $r/2$ to the sum in the arm event (of Lemma 26.8), so that the arm event certainly occurs. Therefore, it suffices to prove that

$$\mathcal{E}_{\rho, k, N, \rho', n} \cap \mathcal{A} \subset \mathcal{E}_{\rho, \varepsilon, k, N, \rho', n}^*. \quad (535)$$

Yet, this is straightforward to see: if a tree in $(\mathcal{L}^{R_i})_i$ contributes to the trees counted in $\mathcal{E}_{\rho, k, N, \rho', n}$ and does not belong to the trees counted in \mathcal{A} , then it must necessarily contain a tree contributing to those counted in $\mathcal{E}_{\rho, \varepsilon, k, N, \rho', n}^*$. \square

Lemma 26.13 (Decoupling step). *Fix ρ, ε, k , and N as in the previous lemma. Then*

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \log \max_{(\mathcal{D}_i)_{1 \leq i \leq \lfloor (1-\varepsilon)n \rfloor}} \prod_i \mu_{\mathbb{Z}^2}[\text{Ridge}_{2\lfloor (1-2\varepsilon)k/2 \rfloor}^{\bar{\omega}}(\mathcal{D}_i, -\rho N, \rho N)] \\ \geq 4(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right) + 100k^2 \log c_{\text{circuit}}. \end{aligned} \quad (536)$$

where the maximum runs over all families of disjoint even domains \mathcal{D}_i which have a diameter of at most $2r_{\rho, \varepsilon}N$, and which are contained in $\llbracket (\rho + 1)N \times nN \rrbracket$.

Proof. Consider the measure $\mu_{\mathbb{Z}^2}$. Let $(\mathcal{D}_{ij})_{j=1,\dots,|\mathcal{M}_{rN}^+(R_i)|}$ denote the family of even domains in $\mathcal{M}_{rN}^+(R_i)$ (where the index j is consistently chosen in a suitable way, for example via a dictionary order on \mathbb{R}^2). Recall that $(\mathcal{D}_{ij})_j$ is a family of disjoint even domains of diameter at most $2rN$ contained in R_i . The idea of the proof is to show that:

- We do not lose too much probability by further conditioning on the exact number of even domains \mathcal{D}_{ij} where a ridge occurs,
- We do not lose too much probability by selecting on beforehand the domains \mathcal{D}_{ij} where the ridges occur.

The proof is mostly technical, but we shall provide full detail.

For any even domain \mathcal{D} , let $\eta_{\mathcal{D}}$ denote the set of even edges that are not contained in \mathcal{D} . For each ij , define the events

$$A_{ij} := \text{Ridge}_{2^{\lceil(1-\varepsilon)k/2\rceil}, rN}(\mathcal{D}_{ij}, a_i, b_i); \quad p_{ij} := \mu_{\mathcal{D}_{ij}}^+[A_{ij}]; \quad (537)$$

$$B_{ij} := \{R'_i \text{ is crossed vertically by } \omega^+ \cup \eta_{\mathcal{D}_{ij}}\}; \quad q_{ij} := \mu_{\mathcal{D}_{ij}}^+[B_{ij}]. \quad (538)$$

Observe that A_{ij} and B_{ij} are disjoint and that

$$\mathcal{E}_{\rho, \varepsilon, k, N, \rho', n}^* = \{\sum_{ij} \mathbb{1}_{A_{ij}} \geq (1 - \varepsilon)4n^2\rho'\}. \quad (539)$$

Set $T := \lceil(1 - \varepsilon)4n^2\rho'\rceil$ and define the events

$$\mathcal{A} := \{\sum_{ij} \mathbb{1}_{A_{ij}} = T\}; \quad (540)$$

$$\mathcal{B} := \{\sum_{ij} \mathbb{1}_{B_{ij}^c} = T\}; \quad (541)$$

$$\mathcal{G} := \{\text{the event } A_{ij} \cup B_{ij} \text{ occurs for every } ij\}. \quad (542)$$

Our first aim is to prove that

$$\mu_{\mathbb{Z}^2}[\mathcal{A} \cap \mathcal{G} | \mathcal{E}_{\rho, \varepsilon, k, N, \rho', n}^*] \geq (c_{\text{circuit}})^{100n^2\rho'}. \quad (543)$$

We first claim that $\mu_{\mathbb{Z}^2}$ -almost surely

$$\prod_{ij} q_{ij} = \mu_{\mathbb{Z}^2} \left[\bigcap_{ij} B_{ij} \middle| (\mathcal{D}_{ij})_{ij} \right] \geq (c_{\text{circuit}})^{100n^2\rho'}. \quad (544)$$

To prove the claim, we first use the tower property and the Markov property, and then the FKG inequality to obtain

$$\prod_{ij} q_{ij} = \left(\bigotimes_{ij} \mu_{\mathcal{D}_{ij}}^+ \right) \left[\bigcap_{ij} B_{ij} \right] \quad (545)$$

$$= \mu_{\mathbb{Z}^2} \left[\bigcap_{ij} B_{ij} \middle| (\mathcal{D}_{ij})_{ij} \right] \quad (546)$$

$$= \mu_{\mathbb{Z}^2} \left[\{\text{each } R_i \text{ is crossed vertically by } \omega^+\} \middle| \left\{ 2R_{\text{Large}} \cap \left(\bigcap_{ij} \eta_{\mathcal{D}_{ij}} \right) \subset \omega^+ \right\} \right] \quad (547)$$

$$\geq \mu_{\mathbb{Z}^2} [\{\text{each } R_i \text{ is crossed vertically by } \omega^+\}] \quad (548)$$

$$\geq (c_{\text{circuit}})^{100n^2\rho'}. \quad (549)$$

This proves the claim. It is now straightforward to prove Equation (543) by running the following exploration process. We inspect \mathcal{D}_{ij} one by one, counting how often the event A_{ij}

occurs. Let Q denote this count. If $Q < T$, then we first ask if A_{ij} occurs. If it does not occur, then we condition on the event that B_{ij} occurs (which happens with a probability of at least q_{ij} since A_{ij} and B_{ij} are disjoint, and we know that A_{ij} does not occur). If $Q = T$, then we simply condition on the event that B_{ij} occurs. The product of the probabilities of the additional conditioning events in this algorithm, is clearly lower bounded by $\prod_{ij} q_{ij}$. The desired lower bound then follows by Equation (544). Combining with the previous lemma, this yields the asymptotic bound

$$\mu_{\mathbb{Z}^2}[\mathcal{A} \cap \mathcal{G}] \geq \exp \left[(n^2 \rho') \left(16(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right) + 400k^2 \log c_{\text{circuit}} + o(1) \right) \right]. \quad (550)$$

Now, notice that $\mathcal{A} \cap \mathcal{G} = \mathcal{A} \cap \mathcal{B}$. Conditional on \mathcal{B} , let $(\mathcal{D}'_\ell)_{\ell=1, \dots, T} \subset (\mathcal{D}_{ij})_{ij}$ denote the set of even domains where B_{ij} does not occur. By the tower property and the Markov property, we get

$$\mu_{\mathbb{Z}^2}[\mathcal{A} \cap \mathcal{B}] = \int_{\mathcal{B}} \prod_{\ell=1}^T \mu_{\mathcal{D}'_\ell}^+[A_\ell | B_\ell^c] d\mu_{\mathbb{Z}^2}[(\mathcal{D}'_\ell)_\ell] \leq \max_{(\mathcal{D}_\ell)_\ell} \prod_{\ell=1}^T \mu_{\mathcal{D}_\ell}^+[A_\ell | B_\ell^c], \quad (551)$$

where the maximum runs over all families of T disjoint even domains \mathcal{D}_ℓ of diameter at most $2rN$ such that each \mathcal{D}_ℓ is contained in some R_i .

Set

$$B_{\mathcal{D}} := \{[\rho N] \times \mathbb{R} \text{ is crossed vertically by } \omega^+ \cup \eta_{\mathcal{D}}\}. \quad (552)$$

By stacking the rectangles $(R_i)_i$ vertically, we get the following bound:

$$\max_{(\mathcal{D}_\ell)_{1 \leq \ell \leq \lceil (1-\varepsilon)4n^2\rho' \rceil}} \prod_{\ell} \mu_{\mathcal{D}_\ell}^+[\text{Ridge}_{2\lceil(1-\varepsilon)k/2\rceil, rN}^{\bar{\omega}}(\mathcal{D}_\ell, -\rho N, \rho N) | B_{\mathcal{D}}^c] \geq \mu_{\mathbb{Z}^2}[\mathcal{A} \cap \mathcal{G}], \quad (553)$$

where the maximum runs over tuples of disjoint even domains \mathcal{D}_ℓ which are of diameter at most $2rN$ and which are contained in $R_{\text{Stacked}} := [(\rho + 1)N \times 4n^2\rho'N]$.

Following the same arguments as those employed to go from Equation (496) to Equation (497) in the proof of the straight bound, we obtain that for every \mathcal{D} contained in R_{Stacked} ,

$$\mu_{\mathcal{D}}^+[\text{Ridge}_{k+2, rN}^{\bar{\omega}}(\mathcal{D}, -\rho N, \rho N) | B_{\mathcal{D}}^c] \leq \mu_{\mathbb{Z}^2}[\text{Ridge}_k^{\bar{\omega}}(\mathcal{D}, -\rho N, \rho N)]. \quad (554)$$

Together with Equations (550) and (553), this inequality finishes the proof. \square

Recall the notions of *straightened paths* and *straightened annuli* from Definition 25.2.

Lemma 26.14 (Straightening step). *Fix $\rho \in 2\mathbb{Z}_{\geq 1}$ and $\varepsilon \in (0, 1/100)$. Then, we may find constants $\eta \in (0, 1/100)$ and $r, C < \infty$ such that the following holds true. Fix $k, N \in 2\mathbb{Z}_{\geq 1}$ with $k \geq 72\rho/(\varepsilon\sqrt{c_{\text{arm}}})$ and $N \geq 100k/\eta$. Then,*

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \log \max_{(A_i)_{1 \leq i \leq \lceil (1-4\varepsilon)n \rceil}} \prod_i \mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+4\lceil(1-8\varepsilon)k/2\rceil}(NA_i)] \\ \geq 4(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right) + 200k^2 \log c_{\text{circuit}} - C, \end{aligned} \quad (555)$$

where the maximum runs over all families of disjoint (ρ, η) -straightened annuli A_i of diameter at most $2r$ and contained in $[(\rho + 1) \times n]$.

Proof. Fix $\rho, \varepsilon, r = r_{\rho, \varepsilon}$, and k as in the previous step. We start with a series of simple bounds.

Fix $\eta > 0$ very small, and suppose that $N \in 2\mathbb{Z}_{\geq 1}$ satisfies $N > 100/\eta$. Recall the definition of the random variable $K_{R, r', x}$ from Lemma 19.5 (see also Theorem 19.7). Then, for any $x \in \mathbb{R}^2$,

$$\mu_{\mathbb{Z}^2}[\{K_{N/2, 2\eta N, x} \geq \varepsilon k/10\}] \leq (4\eta)^{c_{\text{arm}} \varepsilon^2 k^2 / 100}. \quad (556)$$

By a union bound over the points $x \in S := (\eta N \mathbb{Z})^2 \cap \llbracket 2(\rho + 1)N \times 4r_{\rho, \varepsilon} N \rrbracket$, we deduce

$$\mu_{\mathbb{Z}^2}[\{\max_{x \in S} K_{N/2, 2\eta N, x} \geq \varepsilon k/10\}] \leq 1000 \rho r_{\rho, \varepsilon} \cdot (4\eta)^{c_{\text{arm}} \varepsilon^2 k^2 / 100 - 2} =: p_{\rho, \varepsilon, \eta, k}. \quad (557)$$

By inclusion of events, it is then easy to deduce that

$$\mu_{\mathbb{Z}^2}[\{\max_{x \in \llbracket (\rho+1)N \times 2r_{\rho, \varepsilon} N \rrbracket} K_{N, \eta N, x} \geq \varepsilon k/10\}] \leq p_{\rho, \varepsilon, \eta, k}. \quad (558)$$

The probability that this event occurs at least εn times in n independent samples, is at most $(2(p_{\rho, \varepsilon, \eta, k})^\varepsilon)^n$. For sufficiently small $\eta = \eta_{\rho, \varepsilon}$ (depending only on ρ and ε), this probability is much smaller than the probability of the event in the previous lemma (the decoupling step, Lemma 26.13).

Now, consider the previous lemma. We view the product over the probabilities as a single probability of a cylinder event \mathcal{R} in the product measure. In this product measure, the probability

$$\mathcal{R} \cap \{\text{in at least } (1 - 3\varepsilon)n \text{ of the domains, too many arms contribute to } K_{N, \eta N, x}\} \quad (559)$$

has the same asymptotics as the event \mathcal{R} itself. By choosing the domains in which the arms occur, we get the bound

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \max_{(\mathcal{D}_i)_{1 \leq i \leq \lceil (1-3\varepsilon)n \rceil}} \prod_i \mu_{\mathbb{Z}^2}[\mathcal{E}_{\mathcal{D}_i}] \quad (560)$$

$$\geq 4(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right) + 100k^2 \log c_{\text{circuit}} - \log 2; \quad (561)$$

where

$$\mathcal{E}_{\mathcal{D}} := \text{Ridge}_{2\lceil (1-2\varepsilon)k/2 \rceil}^{\bar{\omega}}(\mathcal{D}_i, -\rho N, \rho N) \cap \{\max_{x \in \mathcal{D}_i} K_{N, \eta N, x} \leq \varepsilon k/10\}. \quad (562)$$

(the extra $\log 2$ comes from the combinatorial choice of domains).

Let us now study event $\mathcal{E}_{\mathcal{D}}$. It is easy to see that

$$\mathcal{E}_{\mathcal{D}} \subset \bigcup_A \text{AltCircuit}_{2\lceil (1-4\varepsilon)k/2 \rceil}^{\bar{\omega}}(NA), \quad (563)$$

where the union is over all (ρ, η) -straightened annuli A such that NA is completely surrounded by $\partial \mathcal{D}$. Notice that the number of such annuli is at most $2^{99\rho r_{\rho, \varepsilon} / \eta_{\rho, \varepsilon}^2}$ (by simply upper bounding the number of $\eta \mathbb{Z}^2$ -edges that may or may not belong to the boundary of the annulus). Thus, we get

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \max_{(A_i)_{1 \leq i \leq \lceil (1-3\varepsilon)n \rceil}} \prod_i \mu_{\mathbb{Z}^2}[\text{AltCircuit}_{2\lceil (1-4\varepsilon)k/2 \rceil}^{\bar{\omega}}(NA_i)] \quad (564)$$

$$\geq 4(\rho + 1)N^2 \left(f\left(\frac{2(k+4)}{N}\right) - f(0) \right) + 100k^2 \log c_{\text{circuit}} - \frac{100\rho r_{\rho, \varepsilon}}{\eta_{\rho, \varepsilon}} \log 2, \quad (565)$$

where the maximum is over all families of disjoint (ρ, η) -straightened annuli A_i such that each annulus has a diameter of at most $2r_{\rho, \varepsilon}$ and remains contained in $\llbracket(\rho + 1) \times 2n\rrbracket$.

To finish the proof, notice simply that

$$\mathbb{P}_{\mathbb{Z}^2}[\text{Circuit}^{+4\lceil(1-8\varepsilon)k/2\rceil}(NA_i)] \geq 2^{-2\lceil(1-4\varepsilon)k/2\rceil} \mu_{\mathbb{Z}^2}^{\bar{\omega}}[\text{AltCircuit}_{2\lceil(1-4\varepsilon)k/2\rceil}^{\bar{\omega}}(NA_i)], \quad (566)$$

by flipping the coins corresponding to each level line. This implies the desired bound. \square

Acknowledgements. This project has received funding from the Swiss National Science Foundation and the NCCR SwissMAP. HDC acknowledges the support from the Simons collaboration on localization of waves. The work of KKK is supported by the ERC Project LDRAM: ERC-2019-ADG Project 884584. KKK acknowledges the support from CNRS and from the joint AND-DFG TSF24 project ANR-24-CE92-0033. PL acknowledges the support from the French National Research Agency (ANR), project number ANR-23-CPJ1-0150-01.

References

- [ADH] E. Averous, H. Duminil-Copin, and T. He. “Two-arm exponent for the critical planar random-cluster model with cluster weight $q \in [1, 4]$ ”. Manuscript in preparation.
- [AF03] Mark J. Ablowitz and A. S. Fokas. *Complex Variables: Introduction and Applications*. Cambridge University Press, 2003.
- [Aff86] Ian Affleck. “Universal Term in the Free Energy at a Critical Point and the Conformal Anomaly”. In: *Physical Review Letters* 56.7 (1986), pp. 746–748.
- [Ang+21] Morris Ang, Gefei Cai, Xin Sun, and Baojun Wu. *Integrability of Conformal Loop Ensemble: Imaginary DOZZ Formula and Beyond*. 2021. arXiv: [2107.01788](https://arxiv.org/abs/2107.01788) [math-ph].
- [Ave+a] E. Averous, H.-B. Chen, H. Duminil-Copin, T. He, D. Krachun, I. Manolescu, and J. Xia. “Rotational invariance of any six-vertex model scaling limit with $-1 \leq \Delta \leq -\frac{1}{2}$ ”. Manuscript in preparation.
- [Ave+b] E. Averous, H. Duminil-Copin, T. He, and D. Krachun. “Critical exponents for the planar random-cluster model with cluster-weight close to 4”. Manuscript in preparation.
- [Bak61] George A. Baker. “Application of the Padé Approximant Method to the Investigation of Some Magnetic Properties of the Ising Model”. In: *Physical Review* 124.3 (1961), pp. 768–774.
- [Bax72] Rodney J Baxter. “Partition Function of the Eight-Vertex Lattice Model”. In: *Annals of Physics* 70.1 (1972), pp. 193–228.
- [Bax82] Rodney J. Baxter. *Exactly Solved Models in Statistical Mechanics*. London; New York: Academic Press, 1982.
- [BCN86] H. W. J. Blöte, John L. Cardy, and M. P. Nightingale. “Conformal Invariance, the Central Charge, and Universal Finite-Size Amplitudes at Criticality”. In: *Physical Review Letters* 56.7 (1986), pp. 742–745.
- [BFM14a] G. Benfatto, P. Falco, and V. Mastropietro. “Universality of One-Dimensional Fermi Systems, I. Response Functions and Critical Exponents”. In: *Communications in Mathematical Physics* 330.1 (2014), pp. 153–215.

- [BFM14b] G. Benfatto, P. Falco, and V. Mastropietro. “Universality of One-Dimensional Fermi Systems, II. The Luttinger Liquid Structure”. In: *Communications in Mathematical Physics* 330.1 (2014), pp. 217–282.
- [Bil95] Patrick Billingsley. *Probability and Measure*. 3rd ed. Wiley Series in Probability and Mathematical Statistics. New York: Wiley, 1995. 593 pp.
- [Bis09] Marek Biskup. “Reflection Positivity and Phase Transitions in Lattice Spin Models”. In: *Methods of Contemporary Mathematical Statistical Physics*. Ed. by Marek Biskup, Anton Bovier, Frank den Hollander, Dima Ioffe, Fabio Martinelli, Karel Netocný, Fabio Toninelli, and Roman Kotecký. Berlin, Heidelberg: Springer, 2009, pp. 1–86.
- [BKW76] R. J. Baxter, S. B. Kelland, and F. Y. Wu. “Equivalence of the Potts Model or Whitney Polynomial with an Ice-Type Model”. In: *Journal of Physics A: Mathematical and General* 9.3 (1976), pp. 397–406.
- [BL76] Jöran Bergh and Jörgen Löfström. *Interpolation Spaces: An Introduction*. Ed. by S. S. Chern, J. L. Doob, J. Douglas, A. Grothendieck, E. Heinz, F. Hirzebruch, E. Hopf, S. Mac Lane, W. Magnus, M. M. Postnikov, F. K. Schmidt, W. Schmidt, D. S. Scott, K. Stein, J. Tits, B. L. Van Der Waerden, B. Eckmann, and J. K. Moser. Vol. 223. Grundlehren Der Mathematischen Wissenschaften. Berlin, Heidelberg: Springer, 1976.
- [BP25] Nathanaël Berestycki and Ellen Powell. *Gaussian Free Field and Liouville Quantum Gravity*. Cambridge Studies in Advanced Mathematics. Cambridge: Cambridge University Press, 2025.
- [BPZ84] A. A. Belavin, A. M. Polyakov, and A. B. Zamolodchikov. “Infinite Conformal Symmetry in Two-Dimensional Quantum Field Theory”. In: *Nuclear Physics B* 241.2 (1984), pp. 333–380.
- [Car08] John Cardy. “Conformal Field Theory and Statistical Mechanics”. In: *Les Houches, Summer School on Exact methods in low-dimensional statistical physics and quantum computing* (2008).
- [Car86] John L. Cardy. “Operator Content of Two-Dimensional Conformally Invariant Theories”. In: *Nuclear Physics B* 270 (1986), pp. 186–204.
- [Car92] J. L. Cardy. “Critical Percolation in Finite Geometries”. In: *Journal of Physics A: Mathematical and General* 25.4 (1992), pp. L201–L206.
- [Ces14] Annalisa Cesaroni. “Removable Singularities of Harmonic Functions”. 2014.
- [CF24] Federico Camia and Yu Feng. “Logarithmic Correlation Functions in 2D Critical Percolation”. In: *Journal of High Energy Physics* 2024.8 (2024), p. 103.
- [Cha+21] Nishant Chandgotia, Ron Peled, Scott Sheffield, and Martin Tassy. “Delocalization of Uniform Graph Homomorphisms from \mathbb{Z}^2 to \mathbb{Z} ”. In: *Communications in Mathematical Physics* 387.2 (2021), pp. 621–647.
- [Cha98] L. Chayes. “Discontinuity of the Spin-Wave Stiffness in the Two-Dimensional XY Model”. In: *Communications in Mathematical Physics* 197.3 (1998), pp. 623–640.
- [Che+a] H.-B. Chen, H. Duminil-Copin, T. He, F. Jacopin, D. Krachun, I. Manolescu, and J. Xia. “Critical exponents for the planar random-cluster model with cluster-weight $q = 4$ ”. Manuscript in preparation.

- [Che+b] H.-B. Chen, H. Duminil-Copin, T. He, D. Krachun, I. Manolescu, and J. Xia. “One-arm exponent for the critical planar random-cluster model with cluster weight $q \in [1, 4]$ ”. Manuscript in preparation.
- [Che+14] Dmitry Chelkak, Hugo Duminil-Copin, Clément Hongler, Antti Kemppainen, and Stanislav Smirnov. “Convergence of Ising Interfaces to Schramm’s SLE Curves”. In: *Comptes Rendus Mathématique* 352.2 (2014), pp. 157–161.
- [CHI15] Dmitry Chelkak, Clément Hongler, and Konstantin Izyurov. “Conformal Invariance of Spin Correlations in the Planar Ising Model”. In: *Annals of Mathematics* 181 (2015), pp. 1087–1138.
- [CHI22] Dmitry Chelkak, Clément Hongler, and Konstantin Izyurov. *Correlations of Primary Fields in the Critical Ising Model*. 2022. arXiv: [2103.10263](https://arxiv.org/abs/2103.10263) [math-ph].
- [CKP01] Henry Cohn, Richard Kenyon, and James Propp. “A Variational Principle for Domino Tilings”. In: *Journal of the American Mathematical Society* 14.2 (2001), pp. 297–346.
- [CN06] Federico Camia and Charles M. Newman. “Two-Dimensional Critical Percolation: The Full Scaling Limit”. In: *Communications in Mathematical Physics* 268.1 (2006), pp. 1–38.
- [CN07] Federico Camia and Charles M. Newman. “Critical Percolation Exploration Path and SLE6: A Proof of Convergence”. In: *Probability Theory and Related Fields* 139.3 (2007), pp. 473–519.
- [CN26] Federico Camia and Rongvoram Nivesvivat. *Boundary Operators in the Brownian Loop Soup*. 2026. arXiv: [2601.02755](https://arxiv.org/abs/2601.02755) [math-ph].
- [CRV23] Baptiste Cerclé, Rémi Rhodes, and Vincent Vargas. “Probabilistic Construction of Toda Conformal Field Theories”. In: *Annales Henri Lebesgue* 6 (2023), pp. 31–64.
- [CS12] Dmitry Chelkak and Stanislav Smirnov. “Universality in the 2D Ising Model and Conformal Invariance of Fermionic Observables”. In: *Inventiones mathematicae* 189.3 (2012), pp. 515–580.
- [Dav+16] François David, Antti Kupiainen, Rémi Rhodes, and Vincent Vargas. “Liouville Quantum Gravity on the Riemann Sphere”. In: *Communications in Mathematical Physics* 342.3 (2016), pp. 869–907.
- [DGT24] Chiara Damiolini, Angela Gibney, and Nicola Tarasca. “On Factorization and Vector Bundles of Conformal Blocks from Vertex Algebras”. In: *Annales Scientifiques de l’École Normale Supérieure* (2024).
- [DHN11] Hugo Duminil-Copin, Clément Hongler, and Pierre Nolin. “Connection Probabilities and RSW-type Bounds for the Two-Dimensional FK Ising Model”. In: *Communications on Pure and Applied Mathematics* 64.9 (2011), pp. 1165–1198.
- [DLM18] Hugo Duminil-Copin, Jhih-Huang Li, and Ioan Manolescu. “Universality for the Random-Cluster Model on Isoradial Graphs”. In: *Electronic Journal of Probability* 23 (2018), pp. 1–70.
- [DM22] Hugo Duminil-Copin and Ioan Manolescu. “Planar Random-Cluster Model: Scaling Relations”. In: *Forum of Mathematics, Pi* 10 (2022), e23.
- [DMS97] Philippe Di Francesco, Pierre Mathieu, and David Sénéchal. *Conformal Field Theory*. Graduate Texts in Contemporary Physics. New York, NY: Springer, 1997.

- [DMT21] Hugo Duminil-Copin, Ioan Manolescu, and Vincent Tassion. “Planar Random-Cluster Model: Fractal Properties of the Critical Phase”. In: *Probability Theory and Related Fields* 181.1 (2021), pp. 401–449.
- [DS09] T. C. Dorlas and M. Samsonov. *On the Thermodynamic Limit of the 6-Vertex Model*. 2009. arXiv: [0903.2657 \[cond-mat\]](#).
- [DS11] Bertrand Duplantier and Scott Sheffield. “Liouville Quantum Gravity and KPZ”. In: *Inventiones mathematicae* 185.2 (2011), pp. 333–393.
- [DS57] Cyril Domb and M. F. Sykes. “On the Susceptibility of a Ferromagnetic above the Curie Point”. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 240.1221 (1957), pp. 214–228.
- [DST17] Hugo Duminil-Copin, Vladas Sidoravicius, and Vincent Tassion. “Continuity of the Phase Transition for Planar Random-Cluster and Potts Models with $1 \leq q \leq 4$ ”. In: *Communications in Mathematical Physics* 349.1 (2017), pp. 47–107.
- [DT15] H. Duminil-Copin and V. Tassion. “RSW and Box-Crossing Property for Planar Percolation”. In: *IAMP proceedings* (2015).
- [DT19] Hugo Duminil-Copin and Vincent Tassion. *Renormalization of Crossing Probabilities in the Planar Random-Cluster Model*. 2019. arXiv: [1901.08294 \[math\]](#).
- [DTT18] H. Duminil-Copin, V. Tassion, and A. Teixeira. “The Box-Crossing Property for Critical Two-Dimensional Oriented Percolation”. In: *Probability Theory and Related Fields* 171.3 (2018), pp. 685–708.
- [Dub11] Julien Dubédat. “Topics on Abelian Spin Models and Related Problems”. In: *Probability Surveys* 8 (2011), pp. 374–402.
- [Dum+20] Hugo Duminil-Copin, Karol Kajetan Kozłowski, Dmitry Krachun, Ioan Manolescu, and Mendes Oulamara. *Rotational Invariance in Critical Planar Lattice Models*. 2020. arXiv: [2012.11672 \[math\]](#).
- [Dum+21] Hugo Duminil-Copin, Maxime Gagnebin, Matan Harel, Ioan Manolescu, and Vincent Tassion. “Discontinuity of the Phase Transition for the Planar Random-Cluster and Potts Models with $q > 4$ ”. In: *Annales scientifiques de l’École Normale Supérieure* 54.6 (2021), pp. 1363–1413.
- [Dum+22] Hugo Duminil-Copin, Karol Kajetan Kozłowski, Dmitry Krachun, Ioan Manolescu, and Tatiana Tikhonovskaia. “On the Six-Vertex Model’s Free Energy”. In: *Communications in Mathematical Physics* 395.3 (2022), pp. 1383–1430.
- [Dum+24] Hugo Duminil-Copin, Alex M. Karrila, Ioan Manolescu, and Mendes Oulamara. “Delocalization of the Height Function of the Six-Vertex Model”. In: *Journal of the European Mathematical Society* 26.11 (2024), pp. 4131–4190.
- [Dum23] Hugo Duminil-Copin. “100 Years of the (Critical) Ising Model on the Hypercubic Lattice”. In: *Proceedings of the International Congress Mathematicians 2022* 1 (2023), pp. 164–210.
- [DV95] C. Destri and H. J. de Vega. “Unified Approach to Thermodynamic Bethe Ansatz and Finite Size Corrections for Lattice Models and Field Theories”. In: *Nuclear Physics B* 438.3 (1995), pp. 413–454.
- [DW25] Chiara Damiolini and Lukas Woike. *Modular Functors from Conformal Blocks of Rational Vertex Operator Algebras*. 2025. arXiv: [2507.05845 \[math\]](#).

- [dW85] H. J. de Vega and F. Woynarovich. “Method for Calculating Finite Size Corrections in Bethe Ansatz Systems: Heisenberg Chain and Six-Vertex Model”. In: *Nuclear Physics B* 251 (1985), pp. 439–456.
- [EF63] John W. Essam and Michael E. Fisher. “Padé Approximant Studies of the Lattice Gas and Ising Ferromagnet below the Critical Point”. In: *The Journal of Chemical Physics* 38.4 (1963), pp. 802–812.
- [El+12] Sheer El-Showk, Miguel F. Paulos, David Poland, Slava Rychkov, David Simmons-Duffin, and Alessandro Vichi. “Solving the 3D Ising Model with the Conformal Bootstrap”. In: *Physical Review D* 86.2 (2012), p. 025022.
- [FGK25] Saskia Faulmann, Frank Göhmann, and Karol K. Kozłowski. “Low-Temperature Spectrum of the Quantum Transfer Matrix of the XXZ Chain in the Massless Regime”. In: *Probability and Mathematical Physics* 6.4 (2025), pp. 1507–1658.
- [Fis64] Michael E. Fisher. “Correlation Functions and the Critical Region of Simple Fluids”. In: *Journal of Mathematical Physics* 5.7 (1964), pp. 944–962.
- [Fis66] Michael E. Fisher. “Quantum Corrections to Critical-Point Behavior”. In: *Physical Review Letters* 16.1 (1966), pp. 11–14.
- [Fis67] Michael E. Fisher. “The Theory of Condensation and the Critical Point”. In: *Physique Physique Fizika* 3.5 (1967), pp. 255–283.
- [Fje+06] Jens Fjelstad, J Urgen Fuchs, Ingo Runkel, and Christoph Schweigert. “TFT Construction of RCFT Correlators V: Proof of Modular Invariance and Factorisation”. In: *Theory and Applications of Categories* 16.16 (2006), pp. 342–433.
- [FK72] C. M. Fortuin and P. W. Kasteleyn. “On the Random-Cluster Model: I. Introduction and Relation to Other Models”. In: *Physica* 57.4 (1972), pp. 536–564.
- [FKG71] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. “Correlation Inequalities on Some Partially Ordered Sets”. In: *Communications in Mathematical Physics* 22.2 (1971), pp. 89–103.
- [FM17] Marco Furlan and Jean-Christophe Mourrat. “A Tightness Criterion for Random Fields, with Application to the Ising Model”. In: *Electronic Journal of Probability* 22 (2017), pp. 1–29.
- [Frö82] Jürg Fröhlich. “On the Triviality of $\lambda\phi_d^4$ Theories and the Approach to the Critical Point in $d \geq 4$ Dimensions”. In: *Nuclear Physics B* 200.2 (1982), pp. 281–296.
- [FRS02] Jürgen Fuchs, Ingo Runkel, and Christoph Schweigert. “TFT Construction of RCFT Correlators I: Partition Functions”. In: *Nuclear Physics B* 646.3 (2002), pp. 353–497.
- [FRS04a] Jürgen Fuchs, Ingo Runkel, and Christoph Schweigert. “TFT Construction of RCFT Correlators II: Unoriented World Sheets”. In: *Nuclear Physics B* 678.3 (2004), pp. 511–637.
- [FRS04b] Jürgen Fuchs, Ingo Runkel, and Christoph Schweigert. “TFT Construction of RCFT Correlators: III: Simple Currents”. In: *Nuclear Physics B* 694.3 (2004), pp. 277–353.

- [FRS05] Jürgen Fuchs, Ingo Runkel, and Christoph Schweigert. “TFT Construction of RCFT Correlators IV: Structure Constants and Correlation Functions”. In: *Nuclear Physics B* 715.3 (2005), pp. 539–638.
- [Gaw] K. Gawedzki. “Conformal field theory”. To appear.
- [GGM12] Alessandro Giuliani, Rafael L. Greenblatt, and Vieri Mastropietro. “The Scaling Limit of the Energy Correlations in Non-Integrable Ising Models”. In: *Journal of Mathematical Physics* 53.9 (2012), p. 095214.
- [GL25] Alexander Glazman and Piet Lammers. “Delocalisation and Continuity in 2D: Loop $O(2)$, Six-Vertex, and Random-Cluster Models”. In: *Communications in Mathematical Physics* 406.5 (2025), p. 108.
- [GMT17] Alessandro Giuliani, Vieri Mastropietro, and Fabio Lucio Toninelli. “Height Fluctuations in Interacting Dimers”. In: *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* 53.1 (2017), pp. 98–168.
- [GP23] Alexander Glazman and Ron Peled. “On the Transition between the Disordered and Antiferroelectric Phases of the 6-Vertex Model”. In: *Electronic Journal of Probability* 28 (2023), pp. 1–53.
- [GRV16] Christophe Garban, Rémi Rhodes, and Vincent Vargas. “Liouville Brownian Motion”. In: *The Annals of Probability* 44.4 (2016). arXiv: [1301.2876 \[math\]](https://arxiv.org/abs/1301.2876).
- [GRV19] Colin Guillarmou, Rémi Rhodes, and Vincent Vargas. “Polyakov’s Formulation of 2d Bosonic String Theory”. In: *Publications Mathématiques de l’IHÉS* 130 (2019), pp. 111–185.
- [GT19] Alessandro Giuliani and Fabio Lucio Toninelli. “Non-Integrable Dimer Models: Universality and Scaling Relations”. In: *Journal of Mathematical Physics* 60.10 (2019), p. 103301.
- [Gui+24] Colin Guillarmou, Antti Kupiainen, Rémi Rhodes, and Vincent Vargas. “Conformal Bootstrap in Liouville Theory”. In: *Acta Mathematica* 233.1 (2024), pp. 33–194.
- [Gus80] E. Gusev. “Weak Convergence of Wave Values in Quantum Heisenberg Model”. In: *Reports on Mathematical Physics* 18.3 (1980), pp. 399–410.
- [Gus85] E. V. Gusev. “Thermodynamics and Excited States of the Heisenberg Model”. In: *Theoretical and Mathematical Physics* 63.2 (1985), pp. 527–532.
- [GZ26] Bin Gui and Hao Zhang. “Analytic Conformal Blocks of C_2 -Cofinite Vertex Operator Algebras III: The Sewing-Factorization Theorems”. In: *Proceedings of the London Mathematical Society* 132.2 (2026).
- [Hon] Clément Hongler. “Conformal Invariance of Ising Model Correlations”. PhD thesis.
- [HS13] Clément Hongler and Stanislav Smirnov. “The Energy Density in the Planar Ising Model”. In: *Acta Mathematica* 211.2 (2013), pp. 191–225.
- [Hul38] Lamek Hulthén. *Über das Austauschproblem eines Kristalles*. 1938.
- [IKR89] A. G. Izergin, V. E. Korepin, and N. Y. Reshetikhin. “Conformal Dimensions in Bethe Ansatz Solvable Models”. In: *Journal of Physics A: Mathematical and General* 22.13 (1989), p. 2615.

- [Kad+67] Leo P. Kadanoff, Wolfgang Götze, David Hamblen, Robert Hecht, E. A. Lewis, V. V. Palciauskas, Martin Rayl, J. Swift, David Aspnes, and Joseph Kane. “Static Phenomena near Critical Points: Theory and Experiment”. In: *Reviews of Modern Physics* 39.2 (1967), pp. 395–431.
- [Kad66] Leo P. Kadanoff. “Scaling Laws for Ising Models near T_c ”. In: *Physique Physique Fizika* 2.6 (1966), pp. 263–272.
- [Kau49] Bruria Kaufman. “Crystal Statistics. II. Partition Function Evaluated by Spinor Analysis”. In: *Physical Review* 76.8 (1949), pp. 1232–1243.
- [KB90] A. Klumper and M. T. Batchelor. “An Analytic Treatment of Finite-Size Corrections in the Spin-1 Antiferromagnetic XXZ Chain”. In: *Journal of Physics A: Mathematical and General* 23.5 (1990), p. L189.
- [KBP91] A. Klumper, M. T. Batchelor, and P. A. Pearce. “Central Charges of the 6- and 19-Vertex Models with Twisted Boundary Conditions”. In: *Journal of Physics A: Mathematical and General* 24.13 (1991), p. 3111.
- [Ken00] Richard Kenyon. “Conformal Invariance of Domino Tiling”. In: *The Annals of Probability* 28.2 (2000), pp. 759–795.
- [Ken99] Arthur Edwin Kennelly. “The Equivalence of Triangles and Three-Pointed Stars in Conducting Networks”. In: *Electrical world and engineer* 34.12 (1899), pp. 413–414.
- [Koz18] Karol K. Kozłowski. “On Condensation Properties of Bethe Roots Associated with the XXZ Chain”. In: *Communications in Mathematical Physics* 357.3 (2018), pp. 1009–1069.
- [KPS14] Filip Kos, David Poland, and David Simmons-Duffin. “Bootstrapping Mixed Correlators in the 3D Ising Model”. In: *Journal of High Energy Physics* 2014.11 (2014), p. 109.
- [KRV20] Antti Kupiainen, Rémi Rhodes, and Vincent Vargas. “Integrability of Liouville Theory: Proof of the DOZZ Formula”. In: *Annals of Mathematics* 191.1 (2020), pp. 81–166.
- [KT23] Laurin Köhler-Schindler and Vincent Tassion. “Crossing Probabilities for Planar Percolation”. In: *Duke Mathematical Journal* 172.4 (2023), pp. 809–838.
- [KT25] Laurin Köhler-Schindler and Vincent Tassion. “An Introduction to Russo-Seymour-Welsh Theory”. In: *Stochastic Geometry: Percolation, Tesselations, Gaussian Fields and Point Processes*. Ed. by Hermine Biermé. Cham: Springer Nature Switzerland, 2025, pp. 1–34.
- [KWZ93] A. Klumper, T. Wehner, and J. Zittartz. “Conformal Spectrum of the Six-Vertex Model”. In: *Journal of Physics A: Mathematical and General* 26.12 (1993), pp. 2815–2827.
- [Lam23] Piet Lammers. *A Dichotomy Theory for Height Functions*. 2023. arXiv: [2211.14365](https://arxiv.org/abs/2211.14365) [math].
- [Law14] Gregory F. Lawler. *Conformally Invariant Processes in the Plane*. Providence: American Mathematical Society, 2014.
- [Lis21] Marcin Lis. “On Delocalization in the Six-Vertex Model”. In: *Communications in Mathematical Physics* 383.2 (2021), pp. 1181–1205.
- [Lis22] Marcin Lis. “Spins, Percolation and Height Functions”. In: *Electronic Journal of Probability* 27 (2022), pp. 1–21.

- [LO24] Piet Lammers and Sébastien Ott. “Delocalisation and Absolute-Value-FKG in the Solid-on-Solid Model”. In: *Probability Theory and Related Fields* 188.1 (2024), pp. 63–87.
- [LSM61] Elliott Lieb, Theodore Schultz, and Daniel Mattis. “Two Soluble Models of an Antiferromagnetic Chain”. In: *Annals of Physics* 16.3 (1961), pp. 407–466.
- [LSW02] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. *On the Scaling Limit of Planar Self-Avoiding Walk*. 2002. arXiv: [math/0204277](https://arxiv.org/abs/math/0204277).
- [LSW04] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. “Conformal Invariance of Planar Loop-Erased Random Walks and Uniform Spanning Trees”. In: *The Annals of Probability* 32.1B (2004), pp. 939–995.
- [LT24] Piet Lammers and Martin Tassy. “Macroscopic Behavior of Lipschitz Random Surfaces”. In: *Probability and Mathematical Physics* 5.1 (2024), pp. 177–267.
- [LW80] Elliott H. Lieb and F. Y. Wu. “Two-Dimensional Ferroelectric Models”. In: *12th School of Modern Physics on Phase Transitions and Critical Phenomena*. 1980.
- [Mas04] Vieri Mastropietro. “Ising Models with Four Spin Interaction at Criticality”. In: *Communications in Mathematical Physics* 244.3 (2004), pp. 595–642.
- [MS16] Jason Miller and Scott Sheffield. “Imaginary Geometry I: Interacting SLEs”. In: *Probability Theory and Related Fields* 164.3 (2016), pp. 553–705.
- [MS20] Jason Miller and Scott Sheffield. “Liouville Quantum Gravity and the Brownian Map I: The QLE(8/3, 0) Metric”. In: *Inventiones mathematicae* 219.1 (2020), pp. 75–152.
- [MW14] Barry M. McCoy and Tai Tsun Wu. *The Two-Dimensional Ising Model*. second. Dover Publications, 2014.
- [MW68] B. M. McCoy and Tai Tsun Wu. “Hydrogen-Bonded Crystals and the Anisotropic Heisenberg Chain”. In: *Il Nuovo Cimento B (1965-1970)* 56.2 (1968), pp. 311–315.
- [Nie82] Bernard Nienhuis. “Exact Critical Point and Critical Exponents of O(n) Models in Two Dimensions”. In: *Physical Review Letters* 49.15 (1982), pp. 1062–1065.
- [Ons44] Lars Onsager. “Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition”. In: *Physical Review* 65.3-4 (1944), pp. 117–149.
- [Pol68] A M Polyakov. “Microscopic Description of Critical Phenomena”. In: *Soviet Physics JETP* 28.3 (1968), pp. 533–539.
- [Pol70a] A M Polyakov. “Properties of Long and Short Range Correlations in the Critical Region”. In: *Soviet Physics JETP* 30.1 (1970), pp. 151–157.
- [Pol70b] Alexander M. Polyakov. “Conformal Symmetry of Critical Fluctuations”. In: *JETP Letters* 12 (1970), pp. 381–383.
- [PP64] A Z Patashinskii and V L Pokrovskii. “Second Order Phase Transitions in a Bose Fluid”. In: *Soviet Physics JETP* 19.3 (1964), pp. 677–691.
- [PS] H. Pinson and T. Spencer. “Universality in 2D Critical Ising Model.” Private communication.
- [RS22] Gourab Ray and Yinon Spinka. “Finitary Codings for Gradient Models and a New Graphical Representation for the Six-Vertex Model”. In: *Random Structures & Algorithms* 61.1 (2022), pp. 193–232.

- [Rus78] Lucio Russo. “A Note on Percolation”. In: *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 43.1 (1978), pp. 39–48.
- [RV14] Rémi Rhodes and Vincent Vargas. “Gaussian Multiplicative Chaos and Applications: A Review”. In: *Probability Surveys* 11 (2014), pp. 315–392.
- [Sch00] Oded Schramm. “Scaling Limits of Loop-Erased Random Walks and Uniform Spanning Trees”. In: *Israel Journal of Mathematics* 118.1 (2000), pp. 221–288.
- [Seg88] G. B. Segal. “The Definition of Conformal Field Theory”. In: *Differential Geometrical Methods in Theoretical Physics*. Ed. by K. Bleuler and M. Werner. Dordrecht: Springer Netherlands, 1988, pp. 165–171.
- [She05] Scott Sheffield. “Random Surfaces”. In: *Astérisque* 304 (2005).
- [She07] Scott Sheffield. “Gaussian Free Fields for Mathematicians”. In: *Probability Theory and Related Fields* 139.3 (2007), pp. 521–541.
- [She09] Scott Sheffield. “Exploration Trees and Conformal Loop Ensembles”. In: *Duke Mathematical Journal* 147.1 (2009), pp. 79–129.
- [Smi01] Stanislav Smirnov. “Critical Percolation in the Plane: Conformal Invariance, Cardy’s Formula, Scaling Limits”. In: *Comptes Rendus de l’Académie des Sciences - Series I - Mathematics* 333.3 (2001), pp. 239–244.
- [Smi10] Stanislav Smirnov. “Conformal Invariance in Random Cluster Models. I. Holomorphic Fermions in the Ising Model”. In: *Annals of Mathematics* 172.2 (2010), pp. 1435–1467.
- [SS05] Oded Schramm and Scott Sheffield. “Harmonic Explorer and Its Convergence to SLE₄”. In: *The Annals of Probability* 33.6 (2005), pp. 2127–2148.
- [STF79] E. K. Sklyanin, L. A. Takhtadzhyan, and L. D. Faddeev. “Quantum Inverse Problem Method. I”. In: *Theoretical and Mathematical Physics* 40.2 (1979), pp. 688–706.
- [SW12] Scott Sheffield and Wendelin Werner. “Conformal Loop Ensembles: The Markovian Characterization and the Loop-Soup Construction”. In: *Annals of Mathematics* 176.3 (2012), pp. 1827–1917.
- [SW78] P.D. Seymour and D.J.A. Welsh. “Percolation Probabilities on the Square Lattice”. In: *Annals of Discrete Mathematics*. Vol. 3. Elsevier, 1978, pp. 227–245.
- [Tes01] J. Teschner. “Liouville Theory Revisited”. In: *Classical and Quantum Gravity* 18 (2001), R153–R222.
- [Tes19] J. Teschner. “Liouville theory revisited”. In: *Integrability: From Statistical Systems to Gauge Theory. Lecture Notes of the Les Houches Summer School* 106 (2019), pp. 60–120.
- [Tes95] Jörg Teschner. “On the Liouville Three-Point Function”. In: *Physics Letters B* 363.1 (1995), pp. 65–70.
- [Til07] Béatrice de Tilière. “Scaling Limit of Isoradial Dimer Models and the Case of Triangular Quadri-Tilings”. In: *Annales de l’Institut Henri Poincaré (B) Probability and Statistics* 43.6 (2007), pp. 729–750.
- [WH31] N. Wiener and E. Hopf. “Über eine Klasse Singulärer Integralgleichungen.” In: *Sitzungsberichte der Berliner Akademie der Wissenschaften* (1931), pp. 696–706.

- [Wid65a] B. Widom. “Equation of State in the Neighborhood of the Critical Point”. In: *The Journal of Chemical Physics* 43.11 (1965), pp. 3898–3905.
- [Wid65b] B. Widom. “Surface Tension and Molecular Correlations near the Critical Point”. In: *The Journal of Chemical Physics* 43.11 (1965), pp. 3892–3897.
- [Wil71a] Kenneth G. Wilson. “Renormalization Group and Critical Phenomena. I. Renormalization Group and the Kadanoff Scaling Picture”. In: *Physical Review B* 4.9 (1971), pp. 3174–3183.
- [Wil71b] Kenneth G. Wilson. “Renormalization Group and Critical Phenomena. II. Phase-space Cell Analysis of Critical Behavior”. In: *Physical Review B* 4.9 (1971), pp. 3184–3205.