Universality of Large-Scale Geometry For Critical Random-Cluster Models

Ulrik Thinggaard Hansen* and Ioan Manolescu†

November 2, 2026

Abstract

The understanding of the behaviour of critical lattice models in two dimensions has made immense progress in the past couple of decades, with conformal invariance as its conjectured end goal. We consider here the critical random-cluster model (or FK-percolation) on a large class of graphs — namely bi-periodic isoradial graphs — with cluster-weight $q \in [1, 4]$. Following up on the establishment of rotational invariance at criticality in [10], we prove that the large scale behaviour of these models is asymptotically the same on all such graphs. These include the critical homogeneous models on the square, triangular and hexagonal lattices.

1 Introduction

1.1 Motivation

The successful import of the renormalisation group formalism from quantum field theory was a landmark achievement in statistical mechanics[1] that, while remaining elusive from a rigorous point of view, yielded two concrete predictions:

1. Critical statistical mechanics models should be universal in the sense that their scaling limits should not depend on the choice of underlying lattice.

2. The scaling limits should enjoy conformal symmetries.

In particular, since the planar conformal group is infinite dimensional, scaling limits of 2D statistical mechanics models should enjoy incredibly rich symmetries.

With Schramm's introduction of the SLE in [28], this conjectural picture took on a much more tangible form, and the subsequent twenty years have seen a staggering amount of progress, establishing the conjecture of conformal invariance in a number of different models (see [1, 2, 4, 6, 16, 17, 23, 29, 30, 31] for a non-exhaustive list).

*Institute for Stochastics, KIT, ulrik.hansen@kit.edu

†Department of Mathematics, UniFr, ioan.manolescu@unifr.ch

[1]We refer to [15] for a historical exposition.

However, with the exception of [2, 16, 17], these approaches are characterised by exploiting very particular features of discrete holomorphicity in the models and as such, do not offer robust methods for treating other models.

Another approach is to focus on universality. That is, rather than taking a single model and trying to compute its scaling limit directly, take a class of related models and show that all of their possible scaling limits must agree.

Such an approach has also seen some development in the last fifteen years, starting with bond percolation in [20], extending to more general random-cluster models in [11] and lately, being used to establish not only universality across a class of square lattices, but rotational invariance of any possible scaling limit in [10].

In this paper, we combine the results of [10] with some of the conceptual techniques from [20] in order to extend the universality result of [10] to a larger class of graphs. As a special case, we establish large scale rotational invariance of the critical random-cluster model on any bi-periodic isoradial lattice.

1.2 Main Result

In order to state our results properly, we give a brief introduction to some of the objects with which we will be working, all of which will be treated in detail in Section 2.

We say that an embedded planar graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ is **isoradial** if every face of $\mathbb{G}$ is inscribed in a circle of radius 1 and the centre of the corresponding circle lies within that face. We say that $\mathbb{G}$ is **bi-periodic** if there exist two linearly independent vectors $v, w \in \mathbb{Z}^2$ such that $\mathbb{G}$ is invariant under translation by v and w.

For a finite graph $G = (V, E)$, the **random-cluster model** is the probability measure on $\{0,1\}^E$ given by

$$\phi_{\boldsymbol{p},q,G}[\{\omega\}] = \frac{1}{Z_{\boldsymbol{p},q,G}} q^{\kappa(\omega)} \prod_{e \in E} p_e^{\omega(e)} (1-p_e)^{1-\omega(e)},$$

where $\boldsymbol{p} = (p_e)_{e \in E} \in (0,1)^E$ is a vector of probabilities, $q \in (0, \infty)$ is called the cluster-weight, κ is the number of components in ω and Z is a normalising constant called the **partition function**.

To each edge e of an isoradial graph, we shall attach a certain weight p_e depending only on the local geometry of the graph in a way that shall be made precise later on. For this choice of weights, we shall simply refer to the corresponding finite volume measures as ϕ_G, the so-called critical model with isoradial weights. By usual weak compactness elements, this gives rise to corresponding infinite-volume measures $\phi_{\mathbb{G}}$ on infinite isoradial graphs $\mathbb{G}$, which turn out to be unique in the regime of interest for us, namely $q \in [1,4]$. We will also write $\phi_{\delta\mathbb{G}}$ for this same measure on the graph $\mathbb{G}$ rescaled by $\delta > 0$.

An example of an isoradial graph is the square lattice $\sqrt{2}\mathbb{Z}^2$; note the rescaling factor $\sqrt{2}$ that makes the faces of $\mathbb{Z}^2$ be inscribed in circles of radius 1 rather than $1/\sqrt{2}$. Then, $\phi_{\sqrt{2}\mathbb{Z}^2}$ is the critical random-cluster measure [3] with edge-parameters $\frac{\sqrt{q}}{1+\sqrt{q}}$ for all edges, which has been extensively studied in recent years [9, 12, 13, 14, 18, 27].

It is among this class of random-cluster models that universality is expected when $q \in [1,4]$ – that is, when the model on $\mathbb{Z}^2$ undergoes a continuous phase transition (see [9, 14] with more

recent proofs in [27, 18]). That the phase transition is of the same type on all bi-periodic isoradial graphs was proved in [11].

To make a precise statement, we fix a metric that is naturally associated with re-scaled percolation configurations. For two oriented loops, $\gamma, \gamma' : S^1 \to \mathbb{C}$, we define their natural distance

$$d(\gamma, \gamma') = \inf_{\varphi} \|\gamma \circ \varphi - \gamma'\|_\infty,$$

where φ ranges over all orientation-preserving homeomorphisms of S^1 into itself.

For two sets of oriented loops $\mathcal{F}$ and $\mathcal{G}$, we define their **Camia-Newman distance** d_{CN} such that $d_{\mathrm{CN}}(\mathcal{F}, \mathcal{G}) \leq \varepsilon$ if and only if for every loop $\gamma \in \mathcal{F}$ with $\|\gamma\|_\infty \leq \frac{1}{\varepsilon}$, there exists a loop $\gamma' \in \mathcal{G}$ such that $d(\gamma, \gamma') < \varepsilon$, and vice-versa[2].

For convenience, we will also sometimes consider the global[3] Camia-Newman metric $d_{\text{global CN}}$ which is defined similarly, except that we do not restrict our attention to loops with $\|\gamma\|_\infty \leq \frac{1}{\varepsilon}$.

Finally, for two probability measures ϕ, ϕ' on oriented families of loops, write

$$d_{\mathrm{CN}}(\phi, \phi') := \inf\big\{\varepsilon > 0 \mid \exists \text{ coupling } \mathbb{P} \text{ of } \phi \text{ and } \phi' \text{ with } \mathbb{P}[d_{\mathrm{CN}}(\omega, \omega') > \varepsilon] < \varepsilon\big\}.$$

In our applications, the loop configurations will be the set of cluster interfaces surrounding primal and dual clusters of a percolation configuration ω on some isoradial graph $\mathbb{G}$. The exact definition is given in Section 2, see also Figure 4. For now, the only thing to mention is that these loops will always be oriented so that the dual edges lie on their left.

The identification between percolation and loop configurations is one-to-one, and we shall use the Camia-Newman distance between percolation configurations by extension.

Theorem 1.1. *Fix $q \in [1, 4]$. For any bi-periodic isoradial graph $\mathbb{G}$, there exists $c > 0$ such that, for all $\delta \leq 1$,*

$$d_{\mathrm{CN}}\big(\phi_{\delta\mathbb{G}}, \phi_{\delta\sqrt{2}\mathbb{Z}^2}\big) \leq \delta^c. \tag{1}$$

The assumption of bi-periodicity of $\mathbb{G}$ above may be weakened, in the spirit of [20], but not completely removed. We believe the result to hold for general isoradial graphs, but our proof does rely on bi-periodicity.

Two issues arise when working with non-periodic graphs: the potential existence of edges of length arbitrarily close to 0 or 2 (this is excluded by the *bounded angles property* of [20]) and the potential lack of regularity within $\mathbb{G}$ (such as the existence of a square grid structure, which was termed the *square grid property* in [20]). To extend Theorem 1.1 to general isoradial graphs, one would need to tackle these difficulties. The authors believe they are best tackled together and that they are surmountable, but do require new ideas.

Corollary 1.2. *For any $q \in [1, 4]$ and any bi-periodic isoradial graph $\mathbb{G}$, any sub-sequential limit of $\phi_{\delta\mathbb{G}}$ as $\delta \searrow 0$ is invariant under rotations.*

[2]It is worth noting that the Camia-Newman distance is essentially only a function of loops which are big. Indeed, $d_{\mathrm{CN}}(\gamma, \gamma')$ is bounded above by the sums of the diameters of the loops and the Hausdorff distance between the traces. As such, small loops will not contribute as long as there are any loops of similar diameter nearby in the other configuration.

[3]Note that the global metric takes values in $[0, \infty]$ but otherwise satisfies the usual axioms of a metric. We will only use it on pairs of configurations where the distance is finite.

Proof. Consider a bi-periodic, isoradial graph $\mathbb{G}$ and a rotation $\widetilde{\mathbb{G}}$ of this same graph. As they are both bi-periodic and isoradial, (1) implies

$$d_{\mathrm{CN}}\big(\phi_{\delta\mathbb{G}}, \phi_{\delta\widetilde{\mathbb{G}}}\big) \xrightarrow[\delta\to 0]{} 0.$$

This implies the result. □

It was recently proved in [1] that the critical random-cluster model on $\mathbb{Z}^2$ with cluster weight $q = 4$ converges to the conformal loop ensemble CLE_4. By Theorem 1.1, this extends readily to $\phi_{\mathbb{G}}$ for any bi-periodic isoradial graph $\mathbb{G}$.

Corollary 1.3. *Fix $q = 4$ and $\mathbb{G}$ a bi-periodic isoradial graph. Then $\phi_{\delta\mathbb{G}}$ converges, as $\delta \searrow 0$, to the full-plane nested CLE_4. This includes the homogeneous models on the triangular and hexagonal lattices with critical parameters $1/2$ and $4/5$, respectively.*

The full-plane CLE_4 is the local limit of the process in increasing domains; see [21]. The convergence above holds for the d_{CN} distance.

We do not claim the convergence in domains of $\mathbb{G}$. Indeed, to deduce the convergence in domains from that in the full plane, one requires a form of continuity of the measure in the shape of the domain in which it is defined. Such a continuity for $\phi_{\mathbb{Z}^2}$ is the byproduct of the convergence in [1]. For isoradial graphs, we do not have an a priori proof of such continuity, which is why we restrict Corollary 1.3 to the full plane.

For $q = 2$, the scaling limit of the random-cluster model (also known in this case as the FK-Ising model) was determined on isoradial graphs and beyond in [7, 5] using direct analysis of each model. For $q \in [1, 2) \cup (2, 4)$, the scaling limit on $\mathbb{Z}^2$ and isoradial graphs alike remains conjectural.

Structure of the Paper

Section 2 is dedicated to the necessary background material for the paper. In particular, we review isoradial graphs and the associated critical random-cluster models as well the star-triangle transform and its consequences. The latter is our main tool for comparing a priori different models.

Section 3 then focuses on a particular composite type of transformation, which we will call a cube flip; it is essentially a large-scale star-triangle transformation. It is reminiscent of the interchanging of square lattices in [10], but in finite volume. The estimates of this section are crucial, but follow arguments similar to [10]. To avoid repeating tedious details, proofs here will be sketched, but may be followed without prior knowledge of [10].

Section 4 turns to the proof of Theorem 1.1. First, the graph is transformed so as to be composed of large cubes; these transformations are mesoscopic, and therefore incur an error which is easily controlled. Then, rapid fire applications of the cube flip of Section 3 transform this graph into a square lattice, and the estimates obtained there easily bound the ensuing errors.

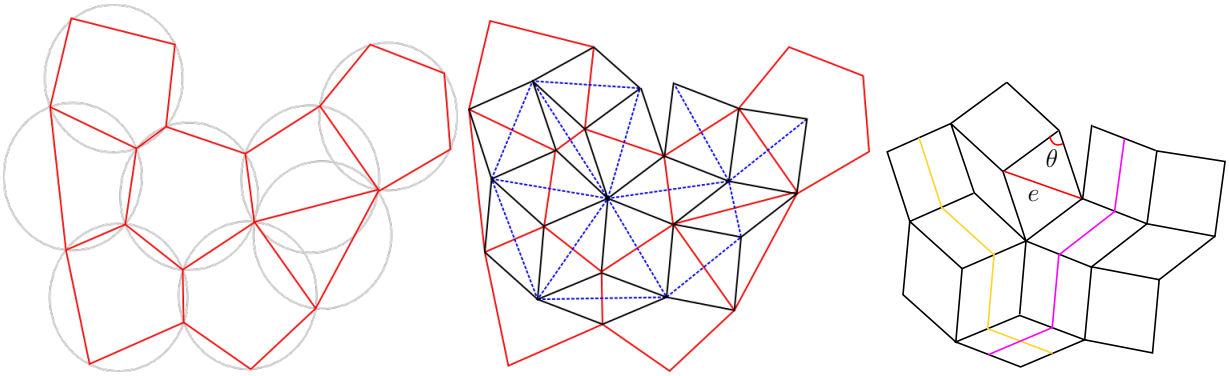


Figure 1: An isoradial domain and its various representations. To the left, we see the graph in red, in the middle, its dual graph superimposed in blue and the associated rhombic tiling in black and to the right, we see the rhombic tiling, the angle associated to an edge and two tracks of the graph.

2 Background

2.1 Isoradial Graphs

In this section, we introduce the concepts related to isoradial graphs which we will need throughout the paper. First, let us establish a bit of standard notation: $\mathbf{Z}^2$ will henceforth always refer to the square lattice in its usual isoradial embedding, which is the rotation of $\sqrt{2}\mathbb{Z}^2$ by $\frac{\pi}{4}$. Set Λ_N to the Euclidean box $[-N, N]^2$; when working with a fixed graph $\mathbb{G}$, we will sometimes identify Λ_N to the subgraph of $\mathbb{G}$ it contains.

One may note that any isoradial embedding of $\mathbb{G}$ also gives rise to an isoradial embedding of its dual graph $\mathbb{G}^* = (\mathbb{V}^*, \mathbb{E}^*)$ by embedding dual vertices as the centres of the circles inscribing the faces of $\mathbb{G}$.

To any isoradial graph $\mathbb{G}$, we associate a rhombic tiling of the plane $\mathbb{G}^\diamond = (\mathbb{V}^\diamond, \mathbb{E}^\diamond)$ as follows: $\mathbb{V}^\diamond = \mathbb{V} \cup \mathbb{V}^*$ and we draw an edge between every centre $v^* \in \mathbb{V}^*$ of a face and each of the vertices $v \in \mathbb{V}$ incident to that face. It is easy to check that $\mathbb{G}^\diamond$ is an embedded planar graph and that each of its faces is a rhombus of side-length 1. Note that $(\mathbf{Z}^2)^\diamond = \mathbb{Z}^2$.

Thus, to every isoradial graph, we can attach a certain rhombic tiling of the plane and any rhombic tiling gives rise to exactly two different isoradial graphs with the only ambiguity being which vertices to consider primal and which to consider dual. See Figure 1.

In this way, each edge $e \in \mathbb{E}$ is contained in a unique rhombus $R(e) \in (\mathbb{V}^\diamond)^*$ and we define the angle θ_e to be the interior angle of $R(e)$ at one of the vertices opposite of e (which is well-defined by symmetry).

For a given rhombic tiling $\mathbb{G}^\diamond$, define a **track** of $\mathbb{G}^\diamond$ to be a sequence of parallel edges $(e_n)_{n\in\mathbb{Z}} \subseteq \mathbb{E}^\diamond$ such that for every n, e_n and e_{n+1} are edges of the same rhombus $R \in (\mathbb{V}^\diamond)^*$ while e_n and e_{n+2} are not. See Figure 1. The transverse angle of the track $(e_n)_{n\in\mathbb{Z}}$ is the angle between each of the e_n and the horizontal axis; we will generally take it to be in $[0, \pi)$.

We say that two distinct tracks $t = (e_n)_{n\in\mathbb{N}}$ and $t' = (f_n)_{n\in\mathbb{N}}$ **intersect** if there exist n, m such that e_n, f_m are part of the same rhombus R. We say that two tracks t and t' are **parallel** if they do not intersect.

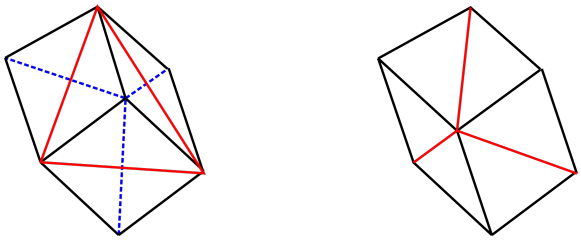


Figure 2: To the left we see a triangle in red, its dual star in blue and the associated hexagon of rhombi in $\mathbb{G}^\diamond$. To the right, the hexagon has been rotated by π, replacing the triangle by a star.

We will often identify a track with the mid-points z_n of the corresponding edges e_n. We will also, without comment, refer to the tracks of $\mathbb{G}^\diamond$ as the tracks of $\mathbb{G}$ (and thus, $\mathbb{G}$ and $\mathbb{G}^*$ have the same tracks).

We say that a track $t = (z_n)_{n\in\mathbb{Z}}$ has an **asymptotic direction** if there exists a vector v_t and a $C > 0$ such that, uniformly in n and m,

$$|z_n - z_m - v_t(n-m)| \leq C$$

Proposition 2.1. *Suppose that $\mathbb{G}$ is bi-periodic. Then, every track of $\mathbb{G}$ has an asymptotic direction.*

Remark 2.2. It is worth noting that two tracks with non-parallel asymptotic directions necessarily intersect. As two tracks can intersect at most once, for bi-periodic graphs, tracks with parallel asymptotic directions cannot intersect. Therefore, the above proposition says that for any bi-periodic isoradial graph $\mathbb{G}$, the tracks of $\mathbb{G}$ can be partitioned into a finite number of track types labelled by their asymptotic directions.

Proof. Let $t = (z_n)_{n\in\mathbb{Z}}$ be a track of $\mathbb{G}$ and let v and w be the vectors of translation corresponding to the actions of $(1,0)$ and $(0,1)$ respectively. By the pigeonhole principle, there exist distinct $n, m \in \mathbb{Z}$ such that z_n and z_m represent the same class in $\mathbb{G}/(\mathbb{Z}v \oplus \mathbb{Z}w)$. Accordingly, we have $(j,k) \in \mathbb{Z}^2 \setminus \{0\}$ such that $z_n - z_m = jv + kw$. It follows that t is invariant under translation by $jv + kw$. It is straightforward to check that $\frac{jv+kw}{n-m}$ is an asymptotic direction of t. □

Finally, note that any biperiodic graph also has the bounded angles property [20] — since the set of transverse angles of its tracks is finite — and the square grid property [20].

2.2 The Star-Triangle Transform and the Track-Exchange Operator

The core ingredient in our treatment of universality among random-cluster models is the star-triangle transform (also known as the Yang-Baxter relation), which is a bit of integrability found in a staggering number of places throughout the planar statistical mechanics literature. It was first discovered by Kennelly in 1899 in the context of electrical networks [22] before appearing in Onsager's famous exact solution of the 2D Ising model [25].

In the present section, we shall focus on the purely graphical aspect of the star-triangle transformation, as that is going to guide the surgery procedures of the paper.

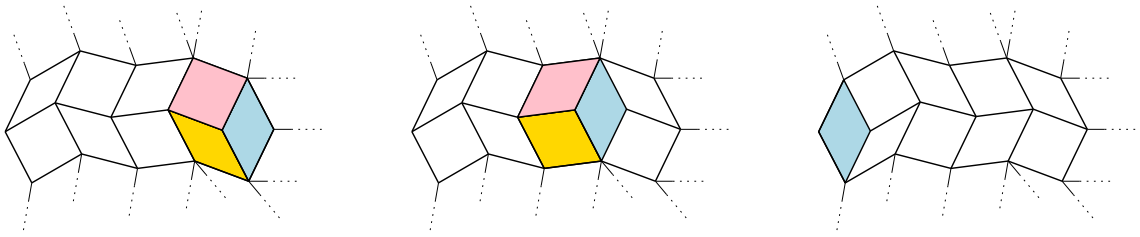


Figure 3: A finite track-exchange in a wider isoradial graph $\mathbb{G}$. Successive star-triangle transforms applied in the coloured hexagon move the blue rhombus to the left, thus exchanging the two horizontal rows of rhombi.

We call an isomorphic copy of the complete graph on three vertices a **triangle**. The set of dual edges to a triangle is called a **star**. One might note that, in the rhombic tiling description of an isoradial graph, this corresponds to three neighbouring rhombi forming a hexagon. See Figure 2.

For any isoradial graph $\mathbb{G}$ and any star or triangle in $\mathbb{G}$, we can define the star-triangle transform $\mathbf{S}(\mathbb{G})$, which replaces the star by a triangle or vice versa. From the rhombic tiling point of view, this corresponds to rotating the three rhombi forming the hexagon by π. Thus, the star-triangle transform preserves isoradiality of the graph.

Now, if two tracks t and s run next to each other before intersecting, we may apply multiple consecutive star triangle transforms to locally change the relative position of the two tracks. Such a sequence of transformations is called a finite **track-exchange**; see Figure 3. As in [10], infinite track-exchanges may also be defined; see also Remark 2.9 below.

2.3 The Random-Cluster Model

We now introduce the random-cluster model on isoradial graphs in more detail. We also recall its elementary properties that we are going to need later on. We refer to the monographs [8, 19] for thorough introductions to the random-cluster model.

Definition 2.3. For a finite subgraph $G = (V, E)$ of an isoradial graph $\mathbb{G}$, $q \geq 1$, vector $\boldsymbol{p} = p_e \in [0,1]^E$ and partition ξ of the vertex boundary of G, the **random-cluster measure** on G with boundary condition ξ is the measure $\phi^{\xi}_{\boldsymbol{p},q,G}$ on $\{0,1\}^E$ given by

$$\phi^{\xi}_{\boldsymbol{p},q,G}[\{\omega\}] = \frac{1}{Z^{\xi}_{\boldsymbol{p},q,G}} q^{\kappa^{\xi}(\omega)} \prod_{e \in E} p_e^{\omega(e)} (1-p_e)^{1-\omega(e)},$$

where κ^{ξ} is the number of connected components in the graph

$$G^{\xi}_{\omega} = (V, E_{\omega}) / \sim_{\xi},$$

$E_{\omega} := \{e \in E |\ \omega(e) = 1\}$, $\sim_{\xi}$ is the equivalence relation with equivalence classes given by ξ and $Z^{\xi}_{\boldsymbol{p},q,G}$ is a normalising constant called the partition function.

E_{ω} is called the set of **open edges** and $E \setminus E_{\omega}$ is called the set of **closed edges**.

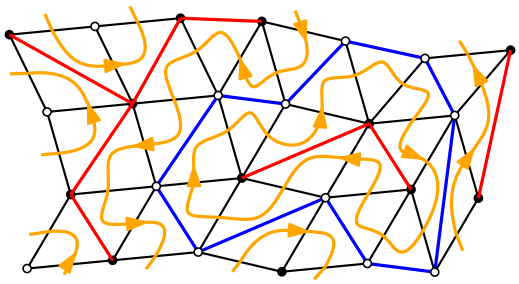


Figure 4: An isoradial graph ($\mathbb{G}^\diamond$ is in black), with a primal configuration (in red), its dual configuration (in blue) and the loop representation (in orange).

Proposition 2.4 (Domain Markov Property)**.** *If $G_1 = (V_1, E_1) \subseteq G_2 = (V_2, E_2)$ are two finite subgraphs of an isoradial graph $\mathbb{G}$, we write $\omega_1 := \omega|_{E_1}$ and $\omega_2 := \omega|_{E_2 \setminus E_1}$ and have*

$$\phi^{\xi}_{\boldsymbol{p},q,G_2}[\omega_1 \in A|\ \omega_2] = \phi^{\xi_{\omega_2}}_{\boldsymbol{p},q,G_1}[A],$$

where v and w belong to the same class of ξ_{ω_2} if and only if they are connected by a path in $((V_2 \setminus V_1), E_{\omega_2})/\sim_\xi$.

Definition 2.5. For an isoradial graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$, we say that a probability measure on $\phi_{\boldsymbol{p},q}$ on $\{0,1\}^{\mathbb{E}}$ is an **infinite-volume random-cluster measure** on $\mathbb{G}$ if, for any finite subgraph $G = (V, E)$ of $\mathbb{G}$, we write $\omega_1 = \omega|_E$ and $\omega_2 = \omega|_{E^c}$ and almost surely have

$$\phi_{\boldsymbol{p},q}[\omega_1 \in A|\ \omega_2] = \phi^{\xi_{\omega_2}}_{\boldsymbol{p},q,G}[A],$$

where v and w belong to the same element of ξ_{ω_2} if and only if they are connected by a path in $(\mathbb{V} \setminus V, E_{\omega_2})$.

For an isoradial graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$ and $\omega \in \{0,1\}^{\mathbb{E}}$ a FK-percolation configuration. We define the *dual configuration* ω^* on $\mathbb{G}^*$ (defined in Section 2.1) by

$$\omega^*(e^*) = 1 - \omega(e) \qquad \forall e \in \mathbb{E},$$

were e^* is the edge dual to e.

The medial graph of $\mathbb{G}$ is the dual of $\mathbb{G}^\diamond$. Its vertices are the intersection points between the primal and dual edges of $\mathbb{G}$; vertices of the medial graph are connected if they are centres of adjacent rhombi. The loop representation of ω is the partition of the edges of the medial graph into loops and bi-infinite paths in such a way that no path or loop intersects any open edge of ω or open dual edge of ω^*. These may be represented as in Figure 4 using bent edges and small turns at the vertices of the medial graph.

The loops are oriented to have the dual open edges on their left. Throughout this paper, the loop representation of ω contains no infinite paths; it is formed entirely of finite loops.

2.4 Isoradial weights and the Star-Triangle Transform

In this paper, we shall only be interested in so-called critical models with isoradial weights. These weights are chosen so as to be compatible with the isoradial embedding of the graph and the star-triangle transform. For a given star-triangle transform $\mathbf{S}$ and $\omega \in \{0,1\}^E$, we can associate a random configuration $\mathbf{S}(\omega) \in \{0,1\}^{\mathbf{S}(E)}$ as follows:

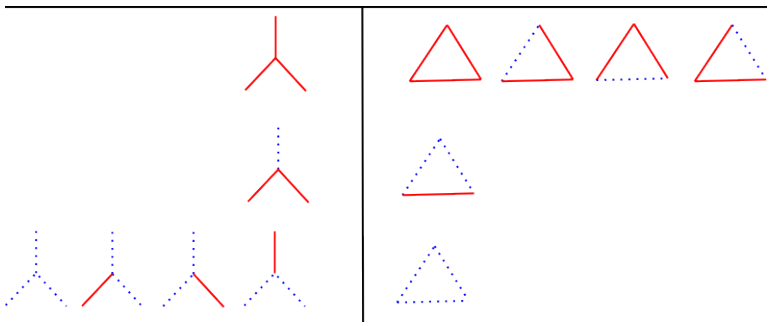


Figure 5: The pairing of percolation configurations under the star-triangle transform. We have omitted symmetric cases.

(*i*) $\omega(e) = \mathbf{S}(\omega)(e)$ for every e not affected by $\mathbf{S}$.

(*ii*) For the edges affected by $\mathbf{S}$, we pair configurations as indicated in Figure 5. In cases of ambiguity, we select the result with probabilities proportional to the random-cluster weights[4]. Solid red lines are open edges and dashed blue ones are closed edges.

Remark 2.6. We may note that the connection field of ω is preserved, meaning that for $v, w \in \mathbb{V} \cap \mathbf{S}(\mathbb{V})$, v is connected to w in ω if and only if they are in $\mathbf{S}(\omega)$.

Define the isoradial weights

$$f^q(\theta) = \begin{cases} \frac{\sqrt{q}\sin(r(\pi-\theta))}{\sin(r\theta)+\sqrt{q}\sin(r(\phi-\theta))} & q < 4 \\ \frac{2\pi-2\theta}{2\pi-\theta} & q = 4 \end{cases},$$

where $r = \frac{1}{\pi}\cos^{-1}(\sqrt{q}/2)$. The following is essentially due to Kenyon [24], but the precise statement is to be found in [11].

Proposition 2.7 ([11], Proposition 2.3)**.** *For every $q \in [1,4]$, any finite part G of an isoradial graph $\mathbb{G}$ and any boundary condition ξ, the star-triangle transform preserves the random-cluster measure with parameters $p_e^{\mathbb{G}} = f^q(\theta_e)$. I.e., for these edge weights, we have*

$$\mathbf{S}(\phi_G^\xi) = \phi_{\mathbf{S}(G)}^\xi.$$

Remark 2.8. The star-triangle transform neither removes nor adds new boundary vertices so that it does, indeed, make sense to regard the same ξ on both sides of the equation.

Throughout the rest of the paper, we will assume that $q \in [1,4]$ is fixed. In this range of parameters, it was shown in [11] that for bi-periodic isoradial graphs $\mathbb{G}$, there is a unique infinite-volume measure, which we denote by $\phi_{\mathbb{G}}$. When $\mathbb{G}$ is not bi-periodic, we will require that it has a unique infinite volume measure — we expect this to be the case for all isoradial graphs in this regime.

[4]One may note that the relative number of components in each case of ambiguity is independent of the boundary condition

Remark 2.9. As shown in [10], if an infinite isoradial graph $\mathbb{G}$ has a unique infinite volume measure $\phi_{\mathbb{G}}$ and contains two parallel tracks s and t that run next to each other (i.e. no rhombus is contained strictly between s and t), one may perform an **infinite track-exchange** of these two tracks.

More precisely, if we write $\mathbb{G}'$ for the isoradial graph obtained by exchanging the rhombi of t and s, there exists a random map $\mathbf{T}$ between configurations on $\mathbb{G}$ and configurations on $\mathbb{G}'$ such that

$$\omega \sim \phi_{\mathbb{G}} \quad \Longrightarrow \quad \mathbf{T}(\omega) \sim \phi_{\mathbb{G}'},$$

where $\phi_{\mathbb{G}'}$ is the unique infinite volume measure on $\mathbb{G}'$. Furthermore, $\mathbf{T}$ preserves the connectivity between vertices common to $\mathbb{G}$ and $\mathbb{G}'$ – that is, all vertices except those contained between s and t. Indeed, $\mathbf{T}$ may be viewed as a bi-infinite sequence of successive star-triangle transformation. Direct constructions of $\mathbf{T}$ are also possible, see for instance [26].

Let us briefly remind the reader of the RSW theory, which has played a central role in the study of planar critical percolation models.

Definition 2.10. Let Q be rotation and translation of $[0, 4n] \times [0, n]$ for $n \geq 1$. Let $\mathscr{C}(Q)$ be the event that there exists a path of open edges crossing Q between its two short sides. We say that an isoradial graph $\mathbb{G}$ has the **strong RSW property** with parameter $c > 0$ if, for every Q as above and $n \in \mathbb{N}$, every event B depending on the edges at distance at least n from Q, we have

$$\phi_{\mathbb{G}}[\mathscr{C}(Q) \,|\, B] \in [c, 1 - c]$$

We write $\mathbb{G} \in \mathcal{RSW}(c)$ if the above holds for $\mathbb{G}$.

The results in [11] imply that for any bi-periodic isoradial graph $\mathbb{G}$, $\phi_{\mathbb{G}}$ admits a strong RSW property. One standard consequence of RSW is a mixing property. See e.g. [12, Proposition 2.9].

Proposition 2.11 (Mixing)**.** *For every $c > 0$, there exists $C > 0$ such that for every $\mathbb{G} \in \mathcal{RSW}(c)$ and events A depending only on edges in Λ_r and B depending only on edges outside of Λ_R, we have*

$$\left|\phi_{\mathbb{G}}[A \cap B] - \phi_{\mathbb{G}}[A]\phi_{\mathbb{G}}[B]\right| \leq C \left(\tfrac{r}{R}\right)^{C} \phi_{\mathbb{G}}[A]\phi_{\mathbb{G}}[B].$$

2.5 Extensions of isoradial domains

An isoradial domain $\mathcal{D}$ is a finite, simply connected set of rhombi of a rhombic tiling $\mathbb{G}^{\diamond}$. We call the isoradial graph $\mathbb{G}$ an *extension* of $\mathcal{D}$. It is called a *natural extension* if it is such that there exists a unique infinite volume measure $\phi_{\mathbb{G}}$ on $\mathbb{G}$. In particular, biperiodic graphs are natural extensions.

Proposition 2.12. *Let $\mathcal{D}$ be an isoradial domain and $\mathbb{G}$ and $\mathbb{G}'$ be two natural extensions of $\mathcal{D}$. Then, the marginals of $\phi_{\mathbb{G}}$ and $\phi_{\mathbb{G}'}$ on $\mathcal{D}$ are the same. In fact, even the distribution of the boundary conditions on $\mathcal{D}$ induced by the outside configuration is the same under both measures.*

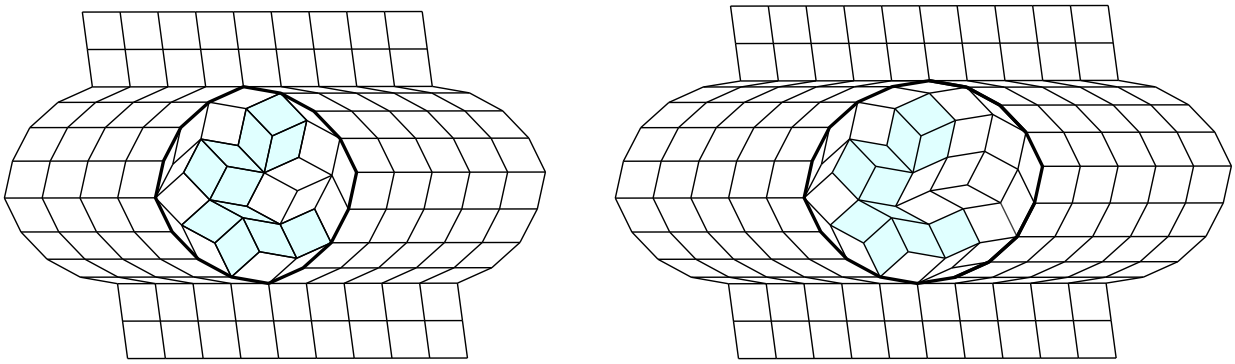


Figure 6: *Left:* The default extension $\mathbb{G}_0$ of an isoradial domain $\mathcal{D}$. The convexification $\mathcal{D}'$ is surrounded by the bold line. *Right:* A different extension $\mathbb{G}_\Lambda$ of $\mathcal{D}$. The convexified part Λ' is also highlighted; outside of it we use the same track angles as for $\mathbb{G}_0$.

Proof. The proof has two parts. We first define a particular extension $\mathbb{G}_0$ of $\mathcal{D}$; we will construct it by adding rhombi to $\mathcal{D}$ sequentially. There are two stages to the construction: first we *convexify* $\mathcal{D}$, and then we complete it into an isoradial graph in the full plane.

Write $t_1, \dots, t_k$ for the tracks traversing $\mathcal{D}$. Recall that any rhombus of an isoradial graph corresponds to the intersection of two tracks. Thus, the potential rhombi formed by the tracks $t_1, \dots, t_k$ is that of the pairs $\{t_i, t_j\}$ with $i \neq j$. Some of these intersections may appear in $\mathcal{D}$, but not necessarily all of them. If $\{t_i, t_j\}$ corresponds to no rhombus of $\mathcal{D}$, we say that the intersection is *missing*. We say that $\{t_i, t_j\}$ is *accessible* if it corresponds to a rhombus which shares two edges with existing adjacent rhombi on the boundary of $\mathcal{D}$.

Consider the lexicographical order on the pairs $\{t_i, t_j\}$ with $i < j$. Let $\{t_i, t_j\}$ denote the smallest pair so that $\{t_i, t_j\}$ is both missing and accessible. If such a pair exists, add the rhombus and restart the procedure. By repeatedly doing this, we arrive at a stage where no such pair of tracks exists. Let $\mathcal{D}'$ denote the isoradial domain obtained by this procedure.

We claim that the union of the rhombi of $\mathcal{D}'$ form a convex set, and that $\mathcal{D}'$ contains all intersections between tracks of distinct transverse angles. Indeed, if no rhombus as above may be added, the edges forming the boundary of $\mathcal{D}'$ have angles in $[0, 2\pi)$ that form an increasing sequence along the boundary in counterclockwise order (if starting from the lowest angle). In other words, the boundary can be parametrised by a curve which only ever turns to the left, and so, $\mathcal{D}'$ is convex. Moreover, it implies that any track enters and exits $\mathcal{D}'$ at opposing sides. For this to be true for all finitely many angles appearing, each angle must intersect every other angle. See Figure 6 for an illustration.

Finally, we construct $\mathbb{G}_0$. Consider two angles $\theta_1 \neq \theta_2$ so that none of the tracks $t_1, \dots, t_N$ has these particular transverse angles. Complete $\mathcal{D}'$ with tracks of transverse angles θ_1 and θ_2 so that the tracks $t_1, \dots, t_k$ intersect all those with angle θ_1, but none of those with angle θ_2.

For simplicity of exposition, we will consider $\theta_1 = 0$ henceforth. We call the tracks of transverse angle θ_1 *vertical* and those of transverse angle θ_2 *horizontal*.

Notice that $\mathbb{G}_0$ has the structure of a square lattice outside of $\mathcal{D}'$. As a result, the RSW estimates of [11, Thm. 1.1] apply for rectangles sufficiently far from $\mathcal{D}'$, which implies that $\mathbb{G}_0$ has a unique infinite volume FK-percolation measure. Thus, $\mathbb{G}_0$ is a natural extension of $\mathcal{D}$. We will consider $\mathbb{G}_0$ to be the *default* extension of $\mathbb{D}$.

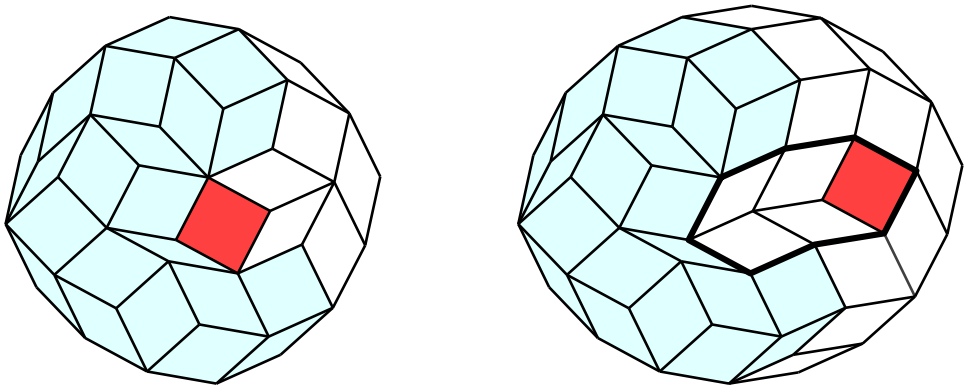


Figure 7: The procedure used in Step 1 of the transformation of $\mathbb{G}_\Lambda$ into $\mathbb{G}_0$. *Left:* the part $\mathcal{D}'$ of $\mathbb{G}_0$ with $\mathcal{A}$ in blue and the rhombus R_j marked in red. *Right:* the convex part Λ' of $\mathbb{G}_\Lambda$ with the rhombus $\tilde{R}$ in red. This is the vertex which needs to be moved to the position occupied by R; we do so by applying star triangle transformations in the delimited region. Here, two transformations suffice.

In the second part of the proof, we show that any other natural extension of $\mathcal{D}$ produces the same measure on $\mathcal{D}$ as $\phi_{\mathbb{G}_0}$. Fix such a natural extension $\mathbb{G}_1$. Let ω_0 be a configuration on $\mathcal{D}$ and $\varepsilon > 0$. Let Λ be a large enough connected part of $\mathbb{G}_1$ so that

$$\max_{\xi,\xi'} \left| \phi_\Lambda^\xi[\omega = \omega_0 \text{ on } \mathcal{D}] - \phi_\Lambda^{\xi'}[\omega = \omega_0 \text{ on } \mathcal{D}] \right| \leq \varepsilon. \tag{2}$$

Define now $\mathbb{G}_\Lambda$ as the completion of Λ obtained as above, by first convexifying Λ to obtain Λ', then completing it with tracks of transverse angles θ_1 and θ_2. In light of (2), we find

$$|\phi_{\mathbb{G}_1}[\omega = \omega_0 \text{ on } \mathcal{D}] - \phi_{\mathbb{G}_\Lambda}[\omega = \omega_0 \text{ on } \mathcal{D}]| \leq \varepsilon. \tag{3}$$

Next, we apply star-triangle transformations to $\mathbb{G}_\Lambda$ so as to create an arbitrarily large portion around $\mathcal{D}$ identical to $\mathbb{G}_0$. These transformations will not affect $\mathcal{D}$. This is done in two steps.

Step 1: We first show that by applying star-triangle transformations to $\mathbb{G}_\Lambda$ outside of $\mathcal{D}$, we may reconstruct $\mathcal{D}'$. See Figure 7 for an illustration of this step.

Write $R_1, R_2, \ldots$ for the rhombi of $\mathcal{D}' \setminus \mathcal{D}$ in the order in which they were added when convexifying $\mathcal{D}$. Let j be maximal so that $R_1, \ldots, R_{j-1}$ appear in Λ' at the same position as in $\mathcal{D}'$. Finally, write $\mathcal{A} = \mathcal{D} \cup R_1 \cup \cdots \cup R_{j-1}$.

If $\mathcal{A} = \mathcal{D}'$, we are done, and we may move on to Step 2. Otherwise, observe that R_j is the intersection of two tracks, s and t, which cross $\mathcal{D}$. These two tracks necessarily intersect in Λ' (as Λ' is convex), and we denote by $\tilde{R}$ the rhombus of $\mathbb{G}_\Lambda$ corresponding to their intersection.

Since $\mathcal{A}$ is connected, there exists a finite region in $\mathbb{G}_\Lambda$ delimited by s, t and $\mathcal{A}$. Write $\mathcal{B}$ for the set of rhombi contained in this region, including those contained on s and t, but not those contained in $\mathcal{A}$.

Note that since $R_j \neq \tilde{R}$ by assumption, $|\mathcal{B}| > 1$. Therefore, there necessarily exists a hexagon that, when flipped, reduces the number of rhombi in $\mathcal{B}$ by at least one. This transformation involves the rhombi in $\mathcal{B}$, and therefore does not affect $\mathcal{A}$, and in particular not $\mathcal{D}$.

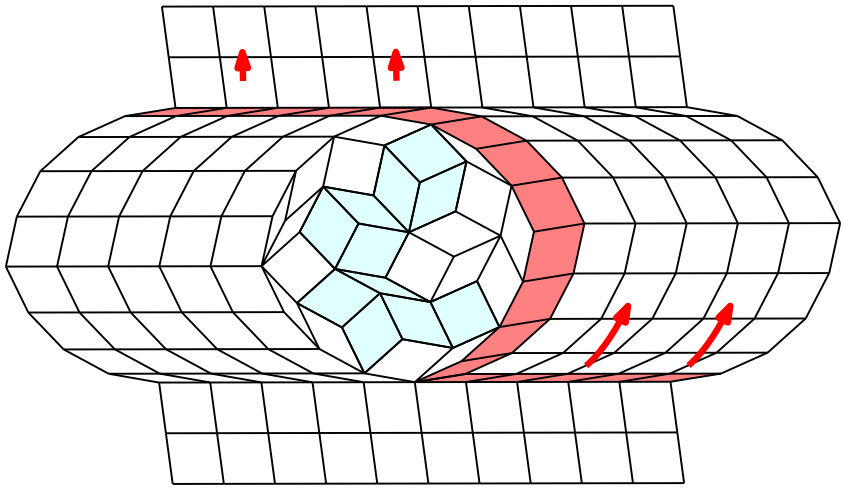


Figure 8: The procedure used in Step 2 of the transformation of $\mathbb{G}_\Lambda$ into $\mathbb{G}_0$. The graph $\mathbb{G}_\Lambda$ may now be considered to contain $\mathcal{D}'$ within the convex region Λ'. The tracks that intersect Λ' but not $\mathcal{D}'$ (here, there is only one, marked in red) need to be sent to infinity. We do this by star-triangle transformation and track-exchanges. Here the red track may be moved upwards indefinitely.

Applying such transformations repeatedly, we reduce $\mathcal{B}$ to a single rhombus, which in turn yields a new graph $\tilde{\mathbb{G}}_\Lambda$ with $\tilde{\mathcal{A}} = \mathcal{A} \cup \{R_j\}$. By repeated applications of this argument, $\mathcal{A}$ may be increased until it is equal to $\mathcal{D}'$.

Step 2: We may now assume that $\mathbb{G}_\Lambda$ contains $\mathcal{D}'$, and we wish to apply star-triangle transformations to it, outside of $\mathcal{D}'$, so as to recreate an arbitrarily large portion of $\mathbb{G}_0$. See Figure 8 for an illustration of this step.

For ease of exposition, we will assume that we may perform infinitely many star-triangle transformations and infinite track-exchanges, and will transform $\mathbb{G}_\Lambda$ into $\mathbb{G}_0$ exactly. This may be rendered formal and re-interpreted to show how $\mathbb{G}_\Lambda$ may be transformed into $\mathbb{G}_0$ on an arbitrarily large region by a finite number of transformations.

We refer to the tracks of $\mathbb{G}_\Lambda$ that cross Λ' but not $\mathcal{D}'$ as *extra tracks*. If no such tracks exist, then $\mathcal{D}' = \Lambda'$, and $\mathbb{G}_\Lambda = \mathbb{G}_0$. Assume next that at least one extra track exists. We will eliminate these tracks by "pushing them to infinity" using star-triangle transformations and track-exchanges.

As extra tracks cross neither $\mathcal{D}'$ nor any of the horizontal tracks, they either run above or below $\mathcal{D}'$. We will limit ourselves to eliminating those running above $\mathcal{D}'$. Among these, consider the extra track t with minimal transverse angle; this is the extra track which is topmost to the left of Λ' and bottommost to the right. We write $s_1, s_2, \ldots$ for the horizontal tracks above Λ', in vertically increasing order.

Applying star-triangle transformations on the right of t, this track may be moved to run adjacent to s_1 by a procedure which is very similar to the one used in step 1. Then, applying track-exchanges between (the image of) t and the horizontal tracks $s_1, s_2 \ldots$, t may be sent upwards to infinity. This concludes Step 2.

Observe now that none of the star-triangle transformations and track-exchanges applied

above to $\mathbb{G}_\Lambda$ affects the marginal of the measure in $\mathcal{D}$ or even the boundary conditions on that domain (i.e. the connections between points on the boundary). Moreover, since these transformations modify $\mathbb{G}_\Lambda$ so as to contain an arbitrary large region around $\mathcal{D}$ that is identical to $\mathbb{G}_0$, and since the latter has a unique infinite-volume measure, we conclude that

$$\phi_{\mathbb{G}_\Lambda}[\omega = \omega_0 \text{ on } \mathcal{D}] = \phi_{\mathbb{G}_0}[\omega = \omega_0 \text{ on } \mathcal{D}].$$

Combining the above with (3) and using the fact that ε is arbitrary, we conclude that

$$\phi_{\mathbb{G}_1}[\omega = \omega_0 \text{ on } \mathcal{D}] = \phi_{\mathbb{G}_0}[\omega = \omega_0 \text{ on } \mathcal{D}],$$

as desired. The same applies to the boundary conditions induced on $\mathcal{D}$ by the configuration outside. □

2.6 Irreducibility of pavings of convex domains

We close this section with a useful combinatorial observation.

Lemma 2.13. *Let $\mathcal{D}$ and $\mathcal{D}'$ be two isoradial domains that occupy the same convex set of $\mathbb{R}^2$. Write $2M$ for the length of the boundary of $\mathcal{D}$ and $\mathcal{D}'$. There exists a sequence $\mathbf{S}_1, \dots, \mathbf{S}_L$ of at most M^3 star triangle transformations that may be applied in order to $\mathcal{D}$ and such that*

$$(\mathbf{S}_L \circ \cdots \circ \mathbf{S}_1)(\mathcal{D}) = \mathcal{D}'. \tag{4}$$

Proof. It suffices to prove (4) for a specific choice of $\mathcal{D}'$, with $L \leq M^3/2$. This specific choice of $\mathcal{D}'$ is defined as follows:

Write $\mathcal{C}$ for the convex set occupied by $\mathcal{D}$ and $e_1, \dots, e_{2M}$ for the edges of the boundary of $\mathcal{C}$, in counterclockwise order, so that $e_1, \dots, e_M$ have increasing angles $\theta_1, \dots \theta_M$ in $[0, \pi)$. For any paving of $\mathcal{C}$ with rhombi, write t_i for the track entering through e_i. Such a track necessarily exits through $e_{M+k+i_0-(i-i_0)}$, where i_0 is the first index with $\theta_{i_0} = \theta_i$ and $k + i_0$ is the last index with $\theta_{k+i_0} = \theta_i$.

We define the standard paving $\mathcal{D}'$ of $\mathcal{C}$ inductively. Set i_0 equal to smallest index j such that $\theta_j = \theta_M$ and set $k = M - i_0$. We first require that t_M runs along the left boundary of $\mathcal{C}$ from e_M to e_{2M-k}. This fixes $M - k - 1$ rhombi in $\mathcal{C}$. The rest of $\mathcal{C}$ is surrounded by a convex perimeter with angles $\theta_1, \dots, \theta_{M-1}, \theta_1, \dots, \theta_{M-1}$. Proceeding by induction, this constructs the standard paving.

Now that $\mathcal{D}'$ is defined, we explain how to pass from $\mathcal{D}$ to $\mathcal{D}'$ using star-triangle transformations. This is again done by induction. Consider the track t_M of $\mathcal{D}$. If this is in the same position as in $\mathcal{D}'$, fix this region and proceed with transforming the rest of the domain, which corresponds to the case $M - 1$.

If t_M in $\mathcal{D}$ differs from its position in $\mathcal{D}'$, "push" it to the left. Indeed, our assumption implies the existence of at least one rhombus in $\mathcal{D}$ which lies strictly to the left of t_M. One may then find such a rhombus R that has two sides adjacent to t_M. For instance, one may orient the tracks $t_1, \dots, t_{M-k}$ from the left of t_M towards the right, inducing a partial order on rhombi to the left of t_M. Then, pick R to be any maximal rhombus of this partial ordering. See Figure 2.6. Applying a star triangle-transform to R and its two neighbouring rhombi on t_M, we obtain a graph $\mathbf{S}(\mathcal{D})$ with R to the right t_M. Repeating this procedure pushes t_M to the left, in the same position as in $\mathcal{D}'$.

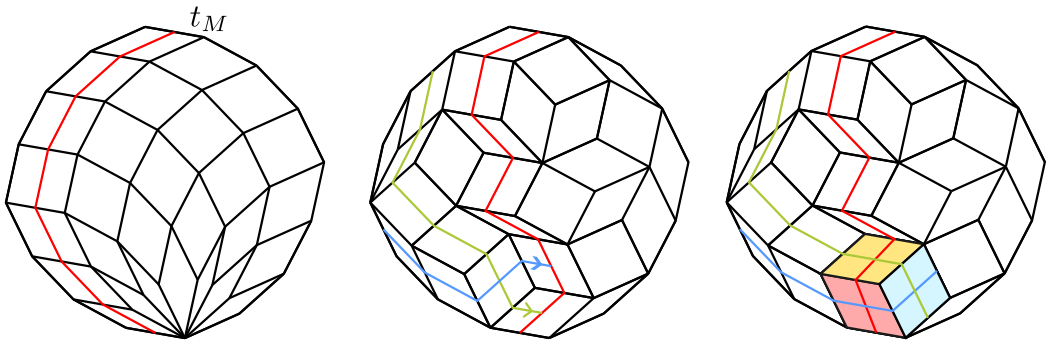


Figure 9: A convex set $\mathcal{C}$ with its standard paving on the left. In the middle, another paving of $\mathcal{C}$, where t_M is not the leftmost track. This implies the existence of a hexagon which may be flipped with the star-triangle transform and moved to the right of t_M, such as the intersection of the blue and green track, which is maximal for the ordering of rhombi given by the orientations of the tracks $t_1, \dots t_{M-1}$ from left to right.

The number of star-triangle transformations needed to push t_M to the left is bounded by the overall number of rhombi in $\mathcal{D}$, which is at most $M(M-1)/2$ - corresponding to the case where all angles are distinct. Since this procedure is then applied to t_{M-1} in the remaining part, then to t_{M-2} etc, we conclude that the total number of star-triangle transformations needed to pass from $\mathcal{D}$ to $\mathcal{D}'$ is bounded above by

$$\sum_{l=1}^{M} \frac{l(l-1)}{2} \leq \frac{M^3}{2}.$$

This concludes the proof. □

3 The Cube Flip Lemma

In the context of isoradial graphs, the star-triangle transformation may be viewed as the permutation of three adjacent rhombi which together form a hexagon. Here, we consider a large-scale version of this transformation and describe its effect on the configuration.

Given three unit vectors v, w and z not contained in the same closed half-plane, we define the *cube* $\mathcal{C}(v, w, z, N)$ to be the isoradial domain obtained by gluing translates of three $N \times N$ rhombi, each paved with N^2 unit rhombi with sides $v \times w$, $w \times z$ and $z \times v$, respectively, as in Figure 10, left diagram. Furthermore, we define the **flipped cube** $\mathcal{C}^{\mathfrak{F}}(v, w, z, N) = \mathcal{C}(-v, -w, -z, N)$. If $\mathbb{G}$ is an extension of $\mathcal{C}(v, w, z, N)$, write $\mathbb{G}^{\mathfrak{F}}$ for the same graph, but with $\mathcal{C}(v, w, z, N)$ replaced by the flipped cube $\mathcal{C}^{\mathfrak{F}}(v, w, z, N)$.

As v, w, z and N will mostly be fixed in this section, we will omit them from notation when no ambiguity is possible. The purpose of this section is to prove the following:

Theorem 3.1 (Cube flip)**.** *For any three unit vectors z, v, w not contained in the same closed half-plane, there exists a constant $c > 0$ with the following property: For any $N \geq 1$ and any natural extension $\mathbb{G}$ of $\mathcal{C}(v, w, z, N)$, there exists a coupling $\mathbb{P}$ between $\omega_1 \sim \phi_{\mathbb{G}}$ and $\omega_2 \sim \phi_{\mathbb{G}^{\mathfrak{F}}}$ such that*

(i) $\omega_1(e) = \omega_2(e)$ *for any edge e outside of $\mathcal{C}$;*

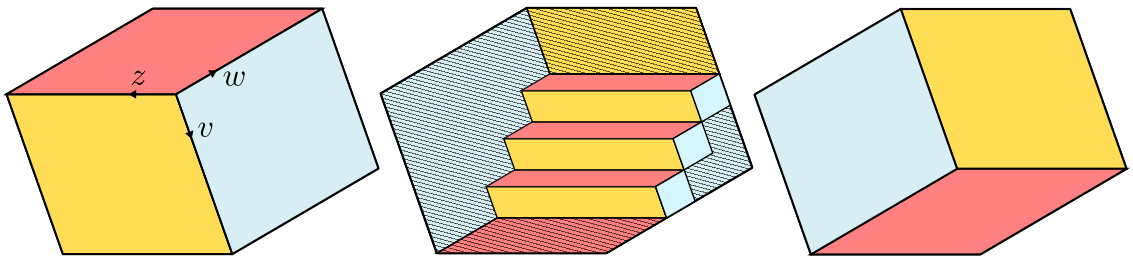


Figure 10: *Left:* a cube with the block of old rhombi in yellow, the new ones in red and the transverse ones in blue. *Right:* the flipped cube, with the three blocks exchanged. *Middle:* a stage $\mathbb{G}_{kN}$ in the transformation from the initial to the flipped cube. The frozen block is hashed and the active one appears bright. To get to $\mathbb{G}_{(k+1)N}$, all new rhombi in the active block are pushed down by track-exchanges. The blue transverse rhombi are also pushed to the left side.

(ii) the connections induced by $\omega_1 \cap \mathcal{C}$ on the vertices of $\partial\mathcal{C}$ are the same as those induced by $\omega_2 \cap \mathcal{C}^{\mathfrak{F}}$;

(iii) $\mathbb{P}\big[d_{\text{global CN}}(\omega_1, \omega_2) \geq N^{1-c}\big] \leq \frac{1}{N^c}$.

The proof of the above requires a slight re-working of the results of [10]. Indeed, the flipping of the cube from $\mathcal{C}$ to $\mathcal{C}^{\mathfrak{F}}$ may be done by a series of star-triangle transformations that resemble the track-exchanges of [10]. Some additional details need to be considered due to the boundary effects that appear when flipping the cube, rather than infinite horizontal tracks. Given the complexity of the argument in [10], we will give a short sketch below. The rest of the section is dedicated to the proof of Theorem 3.1.

3.1 RSW theory in the cube

In the proof of Theorem 3.1 , the RSW theory will be used inside $\mathcal{C}$. Notice that $\mathcal{C}(v, w, z, N)$, or rather its infinite version $\mathcal{C}(v, w, z, \infty)$, is not biperiodic, and therefore, the results of [11, 20] do not apply as such. Nevertheless, the RSW theory inside $\mathcal{C}(v, w, z, \infty)$ is deduced easily from [11].

Lemma 3.2 (RSW in infinite cubes)**.** *For every v, w, z as above, $\mathcal{C}(v, w, z, \infty)$ satisfies the strong RSW property $\mathcal{RSW}(c)$ for some $c > 0$.*

The constant $c > 0$ above may be shown to be independent of v, w, z, but this is not important for our purposes.

A piece of notation will be useful: we write $\mathbb{L}(v, w, z)$ for the isoradial embedding of the square lattice formed of one half-plane paved by rhombi of sides v, z and the other half-plane paved by rhombi of sides w, z; the common boundary of the half-planes is the line with direction z passing through 0.

Proof. Let Q be a rectangle as in Definition 2.10 and consider the three pieces Q^1, Q^2, Q^3 of Q lying inside $\mathbb{L}(v, w, z)$, $\mathbb{L}(w, z, v)$ and $\mathbb{L}(z, v, w)$ respectively. These are finite, potentially empty, subgraphs of square lattices, which have the RSW property by [11].

Now, deterministically, some Q^j must contain a translated, rotated and scaled copy of $[0, 4] \times [0, \frac{1}{2}]$, and thus, one may cross Q simply by crossing Q^j. This proves the desired by Proposition 2.12 and the RSW property of the associated square lattice. □

For convenience, we will avoid proving Theorem 3.1 for general extensions of the cube and focus on the special case of the infinite cube.

Lemma 3.3 (Reduction to the infinite cube)**.** *If Theorem 3.1 holds with* $\mathbb{G} = \mathcal{C}(v, w, z, \infty)$*, then it holds for any natural extension* $\mathbb{G}$*.*

Proof. Fix v, w, z and N as in the statement of Theorem 3.1. Assume Theorem 3.1 holds for $\mathbb{G} = \mathcal{C}(v, w, z, \infty)$, where we flip only $\mathcal{C}(v, w, z, N)$. Finally, fix a natural extension $\mathbb{G}$ of $\mathcal{C}(v, w, z, \infty)$.

Let $\omega_1 \sim \phi_{\mathcal{C}(v,w,z,\infty)}$ and $\omega_2 \sim \phi_{(\mathcal{C}(v,w,z,\infty))^{\mathfrak{F}}}$ (where, in the latter, we only flip $\mathcal{C}(v, w, z, N)$) be coupled so as to satisfy (i), (ii), and (iii) of Theorem 3.1. By Proposition 2.12, ω_1 may also be coupled to $\omega \sim \phi_{\mathbb{G}}$ in such a way that

- $\omega(e) = \omega_1(e)$ for e in $\mathcal{C}(v, w, z, N)$,
- for all $v, w \in \mathcal{C}(v, w, z, N)$, we have $v \overset{\omega}{\longleftrightarrow} w$ if and only if $v \overset{\omega_1}{\longleftrightarrow} w$ and
- ω is conditionally independent of ω_2 given ω_1.

Define now a configuration $\tilde{\omega}$ on $\mathbb{G}^{\mathfrak{F}}$ by

$$\tilde{\omega}(e) = \begin{cases} \omega(e) & e \notin \mathcal{C}(v, w, z, N), \\ \omega_2(e) & e \in \mathcal{C}(v, w, z, N). \end{cases}$$

Then, $\tilde{\omega} \sim \phi_{\mathbb{G}^{\mathfrak{F}}}$ by Proposition 2.12 and the fact that the coupling between ω_1 and ω_2 satisfies (ii). Furthermore, the coupling of ω and $\tilde{\omega}$ satisfies (i), (ii) and (iii) because that of ω_1 and ω_2 does. □

3.2 Nail lattices and homotopy distance

Fix unit vectors v, w, z not belonging to the same half-plane and $N \geq 1$ large. With Lemma 3.3 in hand, we now turn to proof of Theorem 3.1 for $\mathbb{G} = \mathcal{C}(v, w, z, \infty)$. Write $\mathcal{C} = \mathcal{C}(v, w, z, N)$.

We will encode the distance d_{CN} via the notion of nails and homotopy classes — see [10, Sec. 3.5]. We give a brief description below for completeness.

Several scales will come into play, all depending on N, the size of the cube. In principle, they are independent, but we will choose them as follows for some small $\varepsilon > 0$:

- $R = N^{1-\varepsilon^3}$ is the distance between the nails;
- $r_{\max} = N^{1-\varepsilon^2}$ is the maximal distance between a nail and the corresponding lattice point;
- $r_{\min} = N^{1-\varepsilon}$ is the scale of mesoscopic clusters;
- $r_{\mathrm{IIC}} = N^{1-C\varepsilon}$ is the scale at which we will compare extrema of nails with the IIC; C is a large constant chosen as in [10, eq. (18)].

Once these quantities are fixed, we define the following notions: We will henceforth assume that ε is sufficiently small and N is large, in particular so that $r_{\mathrm{IIC}} \gg 1$.

The lattice $\mathscr{L}$ is the set of points of $R\,\mathbb{Z}^2$ contained in $\mathcal{C}$, at a distance at least R from $\mathcal{C}^c$. For a configuration ω on $\mathbb{G}$ and a finite cluster C of ω, define its extremal values in the directions $-v$ and $-z$ by

$$\mathrm{T}(\mathsf{C}) = \max_{y\in\mathsf{C}}\langle -v, y\rangle; \quad \mathrm{B}(\mathsf{C}) = \min_{y\in\mathsf{C}}\langle -v, y\rangle; \quad \mathrm{L}(\mathsf{C}) = \min_{y\in\mathsf{C}}\langle -z, y\rangle; \quad \mathrm{R}(\mathsf{C}) = \max_{y\in\mathsf{C}}\langle -z, y\rangle.$$

The points realising these minimas and maximas are called the extrema of C. If multiple points satisfy one of the conditions, we choose one according to an arbitrary rule[5].

We say that C is a *nail* at some point $x \in \mathscr{L}$ if

$$\big|\mathrm{T}(\mathsf{C}) - \langle -v, x\rangle\big| \le r_{\max}; \qquad \big|\mathrm{B}(\mathsf{C}) - \langle -v, x\rangle\big| \le r_{\max};$$
$$\big|\mathrm{L}(\mathsf{C}) - \langle -z, x\rangle\big| \le r_{\max} \qquad \text{and} \qquad \big|\mathrm{R}(\mathsf{C}) - \langle -z, x\rangle\big| \le r_{\max}.$$

A lattice of nails $\mathsf{N} = (\mathsf{N}(x))_{x\in\mathscr{L}}$ is a choice of one nail $\mathsf{N}(x)$ for each lattice point $x \in \mathscr{L}$. Note that there need not be a unique lattice of nails a given configuration - and there might not even be one at all.

Given a configuration and a lattice of nails $\mathsf{N} = (\mathsf{N}(x))_{x\in\mathscr{L}}$, we will encode the large loops and paths of $\omega \cap \mathcal{C}$ as follows: For each $x \in \mathscr{L}$, let u_x be a point in $\mathsf{N}(x)$. For an edge xy of $\mathscr{L}$, let $[u_x, u_y]$ be the oriented segment of the straight line between u_x and u_y.

For a simple oriented curve γ with image in $\mathbb{R}^2 \setminus \{u_x : x \in \mathscr{L}\}$ let $[u_{x_1}u_{y_1}], \dots, [u_{x_K}u_{y_K}]$ be the sequence of oriented segments crossed by γ; the segments are oriented from left to right as γ crosses them. The *word* associated to γ is the subsequence of $(x_1, y_1), \dots, (x_K, y_K)$ obtained by inductively removing all consecutive pairs of the form (x, y), (y, x). When γ is an oriented loop, we consider the word modulo cyclic permutations. A word and, by extension, the path or loop it encodes is called *trivial* if the associated word contains at most one letter.

Recall that the ω is encoded by an infinite family of oriented loops. Write $[\omega]_{\mathsf{N}}$ for the multi-set of non-trivial words of all loops associated to ω that intersect $\mathcal{C}$. We call $[\omega]_{\mathsf{N}}$ the homotopy information of ω with respect to the lattice of nails N.

The following lemma shows how the homotopy information gives information controls the distance d_{CN}.

Lemma 3.4. *Let $\mathbb{P}$ be a coupling between $\omega_1 \sim \phi_{\mathbb{G}}$ and $\omega_2 \sim \phi_{\mathbb{G}^\delta}$ with $\omega_1|_{\mathcal{C}^c} = \omega_2|_{\mathcal{C}^c}$. Then, for any $\varepsilon > 0$ there exists $\delta > 0$, independent of N, such that*

$$\mathbb{P}\big[d_{\text{global CN}}(\omega_1, \omega_2) \le N^{1-\delta}\big] \le \mathbb{P}\big[\omega_1,\ \omega_2 \text{ contain lattices of nails } \mathsf{N}_1,\ \mathsf{N}_2 \text{ with } [\omega_1]_{\mathsf{N}_1} = [\omega_2]_{\mathsf{N}_2}\big] + \tfrac{1}{N^\delta}.$$

Proof. The result is a straightforward extension of [10, Prop. 3.9]. □

3.3 Coupling for the cube flip

The cube flip may be viewed as a composition of N^2 finite track-exchanges, where each finite track-exchange is the sequential application of N star-triangle transformations involving the

[5]The rule is not actually important; in [10] the choice was leftmost, rightmost, bottommost and topmost, respectively, but in retrospect, any choice would do. The actual definition of extremum in [10] is slightly more complicated, as it involves the extremal cell of the cluster. We will omit this for clarity here.

two same tracks of transverse angles v and w and N tracks of transverse angle z. The track-exchanges will be grouped into $2N-1$ batches, each batch containing simultaneous, non-overlapping applications of several finite track-exchanges.

To fix ideas, we will call the rhombi of with sides v and z the *old* rhombi (yellow in Figure 10), those of sides w and z the *new* rhombi (red in Figure 10) and the others the *transverse* rhombi (blue in Figure 10). We will perform the cube flip by exchanging old and new rhombi via track-exchanges. These are described on $\mathbb{G}$, but only affect the region in $\mathcal{C}$. We will group them as follows:

Fix $\mathbb{G}_0 = \mathbb{G}$. Suppose $\mathbb{G}_t$ was constructed. To obtain $\mathbb{G}_{t+1}$, exchange any horizontal track of old rhombi with a horizontal track of new rhombi if the two tracks are adjacent, and the track of new rhombi is on the top. Let $\mathbf{S}_t$ denote the finite composition of star-triangle transformations forming these multiple track-exchanges.

The rhombi in $\mathcal{C}$ are gradually permuted, but remain in the same hexagonal region of the plane, which we will always call $\mathcal{C} \subseteq \mathbb{G}_t$ for clarity. In any $\mathbb{G}_t$, $\mathcal{C}$ contains an *active block* formed of alternating horizontal tracks of old and new rhombi and up to four *frozen blocks*, formed entirely of old, new or transverse rhombi. The latter type of rhombi form two blocks, one to the left and one to the right of the active block. The old and new blocks lie below and above the active block, and exchange positions after N transformations. In $\mathbb{G}_{2N}$, all old rhombi live above all new rhombi, and the process terminates. We have $\mathbb{G}_{2N} = \mathbb{G}^{\mathfrak{F}}$.

These transformations will also act on configurations to produce a Markov chain of configurations ω_t sampled according to $\phi_{\mathbb{G}_t}$. To define the chain, we add a piece of notation.

A cluster C in some configuration ω_t on some lattice $\mathbb{G}_t$ is called *mesoscopic* if

$$\mathsf{C} \subseteq \mathcal{C} \quad \text{and} \quad r_{\min} \leq \mathrm{T}(\mathsf{C}) - \mathrm{B}(\mathsf{C}) \leq r_{\max}.$$

Write $\mathrm{Meso}(\omega_t)$ for the set of mesoscopic clusters of ω_t and

$$\mathcal{H}(\omega) = \big(\mathrm{T}(\mathsf{C}), \mathrm{B}(\mathsf{C}), \mathrm{L}(\mathsf{C}), \mathrm{R}(\mathsf{C})\big)_{\mathsf{C} \in \mathrm{Meso}(\omega)}$$

for the extremal coordinates of all mesoscopic clusters.

We now give the explicit construction of the chain $(\omega_t)_{t=0,\dots,2N}$. Sample $\omega_0 \sim \phi_{\mathbb{G}_0}$. For $t \geq 0$ even, given $\omega_0, \dots, \omega_t$, we first sample an intermediate configuration $\omega_{t+1/2} \sim \phi_{\mathbb{G}_t}$ as follows:

Consider the boxes $\Lambda_{r_{\mathrm{IIC}}}(z_j)$ around the extrema z_j of mesoscopic clusters. We say that an *overlap error* occurs (in ω_t) if any two such boxes overlap, if any such box intersects more than one cluster of diameter $r_{\min}/4$, if any such box intersects more than one block, or if any such box is at a distance less than $r_{\min}$ from $\mathcal{C}^c$. In this case, we set $\omega_{t+1/2} = \omega_t$. Otherwise, set $\omega_{t+1/2} = \omega_t$ for all edges outside of the boxes $\Lambda_{r_{\mathrm{IIC}}}(z_j)$ and resample the configuration in each box according to the measure

$$\phi_{\mathbb{G}_t}\Big[\,.\,\Big|\, \omega = \omega_t \text{ outside } \bigcup_j \Lambda_{r_{\mathrm{IIC}}}(z_j) \text{ and } \mathcal{H}(\omega) = \mathcal{H}(\omega_t)\Big].$$

Now, independently of the definition of $\omega_{t+1/2}$, set

$$\omega_{t+1} = \mathbf{S}_t(\omega_{t+1/2}) \quad \text{and} \quad \omega_{t+2} = \mathbf{S}_{t+1}(\omega_t).$$

The resampling procedure performed between ω_t and $\omega_{t+1/2}$ is described in more detail in [10, Sec. 3.2]. It is also explained why we resample the environment around the extrema only at even steps.

With these definitions, the main step of the proof is the following statement, corresponding to [10, Prop 3.7]: Denote by $\mathrm{Meso}^+(\omega_0)$ the set of clusters C of ω_0 such that

$$\mathrm{dist}(\mathsf{C}, \mathcal{C}^c) \geq r_{\max} \quad \text{and} \quad 2r_{\min} \leq \mathrm{T}(\mathsf{C}) - \mathrm{B}(\mathsf{C}) \leq \frac{r_{\max}}{2}.$$

For a cluster C_t of ω_t, we denote by C_{t+1} its coupled cluster in ω_{t+1} with the convention that $\mathsf{C}_{t+1} = \emptyset$ if the cluster ceases to exist. We call a cluster C of ω_0 *stable* if for all $t \in [0, 2N]$, we have $\mathsf{C}_t \neq \emptyset$ and furthermore,

$$|\mathrm{A}(\mathsf{C}_0) - \mathrm{A}(\mathsf{C}_t)| \leq \tfrac{1}{10} r_{\max} \qquad \text{for all } \mathrm{A} \in \{\mathrm{T}, \mathrm{B}, \mathrm{L}, \mathrm{R}\}.$$

The above corresponds to [10, Def. 3.5]. The constant $1/10$ above is arbitrary, and should just be taken sufficiently small. Note that, as opposed to [10, Def. 3.5], we do not mention the drifts $\mathrm{Drift}_{\mathrm{vert}}$ for $A = \mathrm{T}, \mathrm{B}$ and $\mathrm{Drift}_{\mathrm{lat}}$ in the definition of stability. Indeed, these were proved in [10, Lem. 4.4] to be zero.

Proposition 3.5. *For $\varepsilon > 0$ small enough, there exists $\delta = \delta(\varepsilon) > 0$ so that*

$$\mathbb{P}\big[\textit{all clusters } \mathsf{C} \in \mathrm{Meso}^+(\omega_0) \textit{ are stable}\big] \geq 1 - \tfrac{1}{N^\delta}. \tag{5}$$

Proof. The proof of (5) is the same as that of [10, Prop 3.7], with the added input that the drift in both vertical and lateral directions is zero [10, Lem. 4.4].

The only minor difference with [10, Prop 3.7] is the presence of transfer blocks. Indeed, in [10, Prop 3.7], horizontal tracks are infinite, and no transfer blocks exist; only the old and new blocks do.

We have incorporated in the notion of overlap errors the situation when an extremum box intersects the interface between the active and transverse blocks. As such, controlling the expectation of increments for extrema close to such interfaces is not necessary. Also, the probability of an overlap error occurring at a given step may be bounded exactly as in [10, Prop 3.7], with the addition of a factor accounting for the case where the overlap error occurs due to the transverse block.

Another small detail to take into account is in the definition of $\tau_{\mathrm{start}}(\mathsf{C})$ and $\tau_{\mathrm{end}}(\mathsf{C})$ for clusters $\mathsf{C} \in \mathrm{Meso}^+(\omega_0)$. These are roughly the times when the cluster enters and exits, respectively, the active block. In our setting, a point in the plane may enter or exit the active block by transitioning to the transverse blocks. As such, one needs to take both coordinates of the cluster into account when defining $\tau_{\mathrm{start}}(\mathsf{C})$ and $\tau_{\mathrm{end}}(\mathsf{C})$, rather than just the vertical coordinate, as was the case in [10, Prop 3.7].

These are only technical difficulties which we will not detail. We will however note that, in this adaptation, the fact that the drifts are null makes all estimates easier. □

Finally, following [10], we state a simple consequence of the strong RSW property:

Lemma 3.6. *For $\varepsilon > 0$ small enough, there exists $\delta = \delta(\varepsilon) > 0$ so that*

$$\mathbb{P}\big[\omega_0 \textit{ contains a lattice of nails in } \mathrm{Meso}^+(\omega_0)\big] \geq 1 - \tfrac{1}{N^\delta}.$$

Proof. The proof is identical to [10, Lem. 3.11]. □

We are now ready to prove Theorem 3.1 for $\mathbb{G} = \mathcal{C}(v, w, z, \infty)$.

Proof of Theorem 3.1 for $\mathbb{G} = \mathcal{C}(v, w, z, \infty)$. Combining Proposition 3.5 and Lemma 3.6 we find that, under the coupling $\mathbb{P}$ above, for $\varepsilon > 0$ small enough, there exists $\delta > 0$ such that

$$\mathbb{P}\big[\omega_0 \text{ contains a lattice of nails that survives to a lattice of nails of } \omega_{2N}\ \big] \geq 1 - \tfrac{2}{N^\delta}.$$

The construction of the coupling is such that,

- $\omega_0 \sim \phi_{\mathbb{G}}$, $\omega_{2N} \sim \phi_{\mathbb{G}^\mathfrak{s}}$, since, by induction, the laws of the different ω_t are $\phi_{\mathbb{G}_t}$;
- $\omega_0 = \omega_{2N}$ outside of $\mathcal{C}$;
- all connections between vertices outside $\mathcal{C}$ are identical in ω_0 and ω_{2N}: indeed, the transformations $\mathbf{S}_t$ preserve such connections, and the resampling between ω_t and $\omega_{t+1/2}$ is only performed in situations where no such connection may be broken or created (see the definition of overlap error);
- if N_0 is a lattice of nails in ω_0 that survives to clusters N_{2N} in ω_{2N}, which also form a lattice of nails, then
$$[\omega_0]_{\mathsf{N}_0} = [\omega_{2N}]_{\mathsf{N}_{2N}}.$$

Combining the last display above with Lemma 3.4 provides the desired conclusion after possibly altering δ. □

4 Universality On Bi-Periodic Isoradial Graphs

In this section, we show Theorem 1.1. To be precise, we will prove it in the following form:

Theorem 4.1. *Let* $\mathbb{G}$ *be a bi-periodic isoradial graph. There exists* $c > 0$ *such that, for any* $N \geq 1$*, there exists a coupling* $\mathbb{P}$ *between* $\omega_1 \sim \phi_{\mathbb{G}}$ *and* $\omega_2 \sim \phi_{\mathbf{Z}^2}$ *such that*

$$\mathbb{P}\big[d_{\text{global CN}}(\omega_1|_{\Lambda_N}, \omega_2|_{\Lambda_N}) > N^{1-c}\big] < N^{-c}.$$

In the above, by $\omega|_{\Lambda_N}$, we mean the family of loops of ω contained in Λ_N. The proof of the above occupies the rest of the section. Before diving in, let us mention how it implies Theorem 1.1.

Proof of Theorem 1.1 given Theorem 4.1. Fix $\delta > 0$. Let $\mathbb{P}$ be the coupling of Theorem 4.1 for $N = \delta^{-\frac{2}{2-c}}$, where $c > 0$ is the constant given by Theorem 4.1 — we will assume for simplicity that N is an integer, otherwise one should round up; we also assume without loss of generality that $c < 1/4$.

Observe that[6],

$$d_{\text{global CN}}(\delta\omega_1|_{\Lambda_{\delta N}}, \delta\omega_2|_{\Lambda_{\delta N}}) = \delta \cdot d_{\text{global CN}}(\omega_1|_{\Lambda_N}, \omega_2|_{\Lambda_N}).$$

Then Theorem 4.1 and some standard algebraic manipulations imply

$$\mathbb{P}[d_{\text{global CN}}(\delta\omega_1|_{\Lambda_{\delta^{-c/2}}}, \delta\omega_2|_{\Lambda_{\delta^{-c/2}}}) > \delta^{c/4}] < \delta^{c/2},$$

which in turn yields (1) for the exponent $c/4$. □

[6]This linearity of the distance is not true for the usual Camia-Newman distance, because rescaling the configuration also changes which loops to consider when determining the distance.

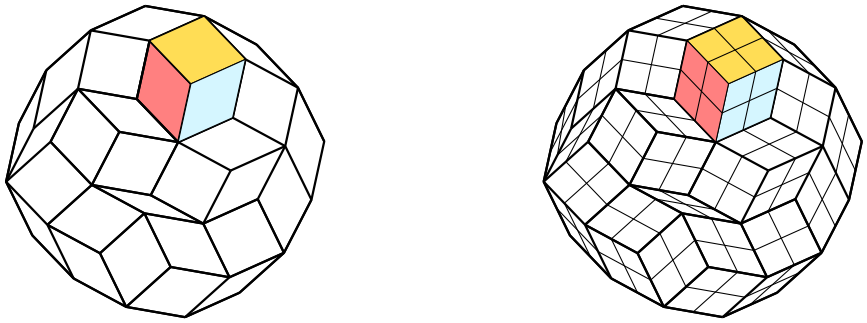


Figure 11: An isoradial domain $\mathcal{D}$ and the corresponding domain in the proper two-batch $\mathbb{G}^{(2)}$ (scaled by a factor of $\frac{1}{2}$). Note that stars and triangles of the original rhombic tiling correspond to cubes in $\mathbb{G}^{(2)}$ to which one may apply the Cube Flip Lemma (Theorem 3.1).

4.1 Batched graphs and general strategy

For an isoradial graph $\mathbb{G}$, we call a family of tracks $t_1, \dots, t_K$ a K-batch if

i) the tracks t_j do not mutually intersect. Furthermore, they are indexed such that for each j, the tracks $t_1, \dots, t_{j-1}$ lie on one side of t_j and $t_{j+1}, \dots, t_K$ on the other;

ii) any rhombus between t_1 and t_K belongs to exactly one of the tracks t_j.

We call $\mathbb{G}$ *K-batched* if its tracks may be partitioned into K-batches. We further call it *properly K-batched* if the tracks of each K-batch have the same transverse angle.

If $\mathbb{G}$ is an isoradial graph, its associated proper K-batch, which we write $\mathbb{G}^{(K)}$, is the isoradial graph obtained by replacing each track of $\mathbb{G}$ with K tracks of the same transverse angle that run next to each other. See Figure 11.

Our proof of Theorem 4.1 is split in two steps. First, we will prove that it applies to any properly K-batched bi-periodic isoradial graph, where $K = N^{1-\delta}$ for some well chosen $\delta > 0$. This is done by applying Theorem 3.1. We will then prove that if it applies to $\mathbb{G}^{(K)}$, it also does to $\mathbb{G}$. This step requires almost no probabilistic estimates.

4.2 Proof of Theorem 4.1 for K-batched graphs

Theorem 4.2. *Fix a bi-periodic isoradial graph $\mathbb{G}$. There exist $C > 0$ and $c \in (0,1)$ such that, for any $K \geq 1$ and $N \geq K$, there exists a coupling $\mathbb{P}$ between $\omega_1 \sim \phi_{\mathbb{G}^{(K)}}$ and $\omega_2 \sim \phi_{\mathbf{Z}^2}$ such that*

$$\mathbb{P}\left[d_{\text{global CN}}(\omega_1|_{\Lambda_N}, \omega_2|_{\Lambda_N}) \geq C\left(\tfrac{N}{K}\right)^3 K^{1-c}\right] \leq C\left(\tfrac{N}{K}\right)^3 K^{-c}.$$

Remark 4.3. In the above, C and c are allowed to depend on $\mathbb{G}$. A quick inspection of the proof shows that they actually only depend on the bounded angles property [20] satisfied by $\mathbb{G}$. The bi-periodicity of $\mathbb{G}$ is only needed here to ensure that $\mathbb{G}$ has this bounded angles property.

We believe that the constants may be taken independently of the bounded angles property, and thus that Theorem 4.2 applies uniformly to all isoradial graphs $\mathbb{G}$. This, however, requires controlling cube flips with degenerating angles, which we do not do here.

Proof. Fix N, K and a domain $\mathcal{D}_0$ of $\mathbb{G}$ that contains $\Lambda_{N/K}$ and let $\mathcal{D}$ denote its convexification. Then, the perimeter $2M$ of $\mathcal{D}$ satisfies $M \leq C_0 N/K$ for some constant C_0 that depends on $\mathbb{G}$. Write $e_1, \dots, e_{2M}$ for the edges of the boundary of $\mathcal{D}$, in counterclockwise order, so that $e_1, \dots, e_M$ have increasing angles in $[0, \pi)$. Write $\theta_1, \dots, \theta_M$ for the angles of $e_1, \dots, e_M$.

Next, add rows of $3M$ rhombi left and right of $\mathcal{D}$, each with edges with angles 0 and θ_i with $i = 1, \dots, M$ except when $\theta_i = 0$. This corresponds to prolonging the tracks intersecting $\mathcal{D}$ horizontally, as in Figures 6 and 12. Next, add $3M$ tracks running above the constructed graph, with transverse angle $\pi/2$, and the same below. Finally, let $\mathcal{G}$ be the convexification of the graph thus obtained.

Note that the edges of the perimeter of $\mathcal{G}$ have angles:

(I) $\theta_1, \dots, \theta_M$, each appearing twice;

(II) 0, appearing $12M$ additional times;

(III) $\pi/2$, appearing an additional $12M$ times.

The angles in category (I) different from 0 and $\frac{\pi}{2}$ are grouped into four sectors, which we denote by NE, NW, SW and SE, respectively. The tracks entering SW exit in NE and those entering NW exist in SW. Also note that the distance between $\Lambda_{N/K}$ and $\mathcal{G}^c$ is at least $3M$.

Let $\tilde{\mathcal{G}}$ be the paving of the interior of $\mathcal{G}$ by rhombi constructed as follows:

- all tracks entering SW run along the bottom and right side of the boundary up to their exit in NE;
- all tracks entering NW run along the top boundary, then left of the tracks in the point above, up to the time they intersect them and exit in SE;
- all other parts of $\tilde{\mathcal{G}}$ are paved by square rhombi - that is, rhombi with angles 0 and $\pi/2$.

See Figure 12 for an example.

Observe that, in $\tilde{\mathcal{G}}$, $\Lambda_{N/K}$ is entirely contained in the region paved with square rhombi. Indeed, the tracks with traverse angles other than 0 and $\pi/2$ run along the boundary of $\tilde{\mathcal{G}}$ and hence, remain within distance $2M$ of it.

By Lemma 2.13, there exists a sequence $\mathbf{S}_1, \dots, \mathbf{S}_L$ of star-triangle transformations with

$$L \leq C_1 M^3$$

that transforms $\mathcal{G}$ into $\tilde{\mathcal{G}}$ where $C_1 = 13^3$. Here, we have used the counts in (I), (II) and (III) to determine the length of the perimeter.

Let $\mathbb{L}$ be a natural extension of $\mathcal{G}$. Then, by applying cube flips $\mathbf{S}_1^{(K)}, \dots, \mathbf{S}_L^{(K)}$ to $\mathbb{L}^{(K)}$, corresponding to the star-triangle transformations $\mathbf{S}_1, \dots, \mathbf{S}_L$, we transform $\mathbb{L}^{(K)}$ into a graph $\tilde{\mathbb{L}}^{(K)}$ which contains $\tilde{\mathcal{G}}^{(K)}$. Let $\mathbb{P}$ be the coupling between $\omega_1 \sim \phi_{\mathbb{L}^{(K)}}$ and $\omega_2 \sim \phi_{\tilde{\mathbb{L}}^{(K)}}$ obtained by these cube flips. By Theorem 3.1, we find that

$$\mathbb{P}\big[d_{\text{global CN}}(\omega_1, \omega_2) \geq C_1 M^3 K^{1-c}\big] \leq C_1 M^3 \tfrac{1}{K^c}.$$

Here, we have used the triangle inequality for the global Camia-Newman distance and a union bound.

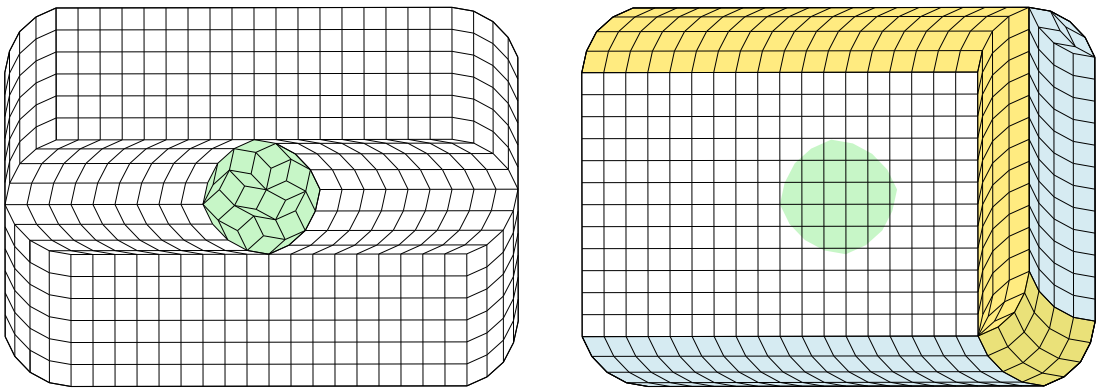


Figure 12: Left: a convex domain $\mathcal{D}$ of an isoradial graph (green), with the completion $\mathcal{G}$ around it. Right: the convex region occupied by $\mathcal{G}$, paved with the rhombic tiling $\tilde{\mathcal{G}}$. The tracks between SW and NE are blue, those between NW and SE are orange. Note that the green region is entirely covered by a copy of $\mathbb{Z}^2$. A sequence of at most $C_1 M^3$ star-triangle transformations allows to pass from $\mathcal{G}$ to $\tilde{\mathcal{G}}$.

Finally, since $\mathbb{L}^{(K)}$ in Λ_N is identical to $\mathbb{G}^{(K)}$, $\tilde{\mathbb{L}}^{(K)}$ in Λ_N is identical to $\mathbf{Z}^2$, and since $\mathbb{L}^{(K)}$ and $\tilde{\mathbb{L}}^{(K)}$ are natural extensions, Proposition 2.12 implies that we may extend $\mathbb{P}$ to include $\omega_1' \sim \phi_{\mathbb{G}^{(K)}}$ and $\omega_2' \sim \phi_{\mathbf{Z}^2}$ so that

$$d_{\text{global CN}}(\omega_1|_{\Lambda_N}, \omega_1'|_{\Lambda_N}) = 0 \quad \text{and} \quad d_{\text{global CN}}(\omega_2|_{\Lambda_N}, \omega_2'|_{\Lambda_N}) = 0.$$

Combining the last two displays, and using the bound on the value of M, we find

$$\mathbb{P}\big[d_{\text{global CN}}(\omega_1'|_{\Lambda_N}, \omega_2'|_{\Lambda_N}) \geq C_2\big(\tfrac{N}{K}\big)^3 K^{1-c}\big] \leq C_2\big(\tfrac{N}{K}\big)^3 K^{-c},$$

for some $\mathbb{G}$-dependent constant C_2, as required. □

4.3 Proof of Theorem 4.1 for general graphs

This section contains the proof of Theorem 4.1. We start with a proposition.

Proposition 4.4. *Fix $\mathbb{G}$ a bi-periodic isoradial graph. For any $c > 0$, there exists $\delta > 0$ such that, for all $N \geq 1$, if we set $K = \lceil N^{1-c} \rceil$, there exists a coupling $\mathbb{P}$ between $\omega_1 \sim \phi_{\mathbb{G}}$ and $\omega_2 \sim \phi_{\mathbb{G}^{(K)}}$ such that*

$$\mathbb{P}\big[d_{\text{global CN}}(\omega_1|_{\Lambda_N}, \omega_2|_{\Lambda_N}) > N^{1-\delta}\big] < N^{-\delta}.$$

Proof. We ignore the integer parts for simplicity. We will transform $\mathbb{G}$ into $\mathbb{G}^{(K)}$ using star-triangle transformations and track-exchanges.

A collection of tracks $(t_n)_{n\in\mathbb{Z}}$ is said to be a **track type** if

(*i*) the t_n are mutually parallel;

(*ii*) no other track in $\mathbb{G}$ is parallel to t_n for any n;

(*iii*) every other track s intersects the t_n in order (i.e. there exists an orientation of s such that t_n is intersected before t_{n+1} for all n).

Since $\mathbb{G}$ is bi-periodic, its tracks may be partitioned into a finite number of track types. Write $(t_n^{(j)})_{n\in\mathbb{Z}}$ with $j = 1, \dots, J$ for these families.

We start off with the case when all tracks of the same type have the same transverse angle. Let us start by batching the tracks of the type $(t_n^{(1)})_{n\in\mathbb{Z}}$. There exists a finite sequence of transformations $\mathbf{S}_1, \dots, \mathbf{S}_T$ with the following properties:

- each $\mathbf{S}_t$ is the composition of infinitely many, non-overlapping star-triangle transformations;
- the tracks $t_{nK}^{(1)}$ with $n \in \mathbb{Z}$ are not involved in any star-triangle transformation of any $\mathbf{S}_t$;
- $\mathbf{S}_T \circ \cdots \circ \mathbf{S}_1(\mathbb{G})$ is bi-periodic, with the tracks $(t_n^{(1)})_{n\in\mathbb{Z}}$ being batched into K-batches.

The $\mathbf{S}_t$'s may be constructed by taking limits of the construction in [20, Lemma 7.1]. That the limit exists is a direct consequence of the fact that the star-triangle transformations are non-overlapping and local.

For a configuration $\omega_0 \sim \phi_{\mathbb{G}}$ set

$$\mathbb{G}_1 = \mathbf{S}_T \circ \cdots \circ \mathbf{S}_1(\mathbb{G}) \quad \text{and} \quad \omega_1 = \mathbf{S}_T \circ \cdots \circ \mathbf{S}_1(\omega_0).$$

Note that the connections in ω_0 and ω_1 between points on the tracks $(t_{nK}^{(1)})_{n\in\mathbb{Z}}$ are identical. Indeed, the transformations involved in $\mathbf{S}_1, \dots, \mathbf{S}_T$ do not affect the points on these track, and thus preserve connections between them. Furthermore, $\omega_1 \sim \phi_{\mathbb{G}_1}$.

Using the same procedure, define sequentially $\mathbb{G}_j$ from $\mathbb{G}_{j-1}$ so that each of the families of tracks $(t_n^{(1)}), \dots, (t_n^{(j)})$ are batched into K-batches in $\mathbb{G}_j$. To get $\mathbb{G}_j$ from $\mathbb{G}_{j-1}$, we apply star-triangle transformations that do not affect any of the tracks $(t_{nK}^{(j)})_{n\in\mathbb{Z}}$. Thus, the connections in ω_{j-1} and ω_j between points on the tracks $(t_{nK}^{(j)})_{n\in\mathbb{Z}}$ are identical.

The following lemma is a direct consequence of the RSW property and the preservation of connections between points of $(t_{nK}^{(j)})_{n\in\mathbb{Z}}$. Its proof is postponed to the end of the section.

Lemma 4.5. *For any fixed $c > 0$, there exists $\delta > 0$ such that, for all N large enough and K and $1 \le j \le J$ as above,*

$$\mathbb{P}\big[d_{\text{global CN}}(\omega_{j-1}|_{\Lambda_N}, \omega_j|_{\Lambda_N}) > N^{1-\delta}\big] < N^{-\delta}.$$

We return to the proof of Proposition 4.4. By repeatedly applying the lemma and using a union bound, we conclude that

$$\mathbb{P}\big[d_{\text{global CN}}(\omega_0|_{\Lambda_N}, \omega_J|_{\Lambda_N}) > J \cdot N^{1-\delta}\big] < J \cdot N^{-\delta}.$$

Recall that $\omega_0 \sim \phi_{\mathbb{G}}$ and observe that $\omega_J \sim \phi_{\mathbb{G}_J} = \phi_{\mathbb{G}^{(K)}}$. This provides the desired conclusion by adjusting the constants in the case where all track of any given type have the same transverse angle.

We now describe how to treat the case where tracks of one type may have different transverse angles. Due to bi-periodicity, the transverse angles of any track type form a periodic sequence; write ℓ_i for the period of the transverse angles of the tracks of type i. Define $\mathbb{G}'$ as the graph obtained by batching $\ell_i K$ tracks of track type i for each i.

The same reasoning as above provides a coupling between $\omega \sim \phi_{\mathbb{G}}$ and $\omega' \sim \phi_{\mathbb{G}'}$ with

$$\mathbb{P}\big[d_{\text{global CN}}(\omega|_{\Lambda_N}, \omega'|_{\Lambda_N}) > C \cdot N^{1-\delta}\big] < C \cdot N^{-\delta}, \tag{6}$$

where C is a constant that may depend on $\mathbb{G}$, but not on N or $c > 0$. Note, however, that now, $\mathbb{G}'$ may be different from $\mathbb{G}^{(K)}$, as it may not be properly K-batched.

To get from $\mathbb{G}'$ to $\mathbb{G}^{(K)}$ we may need to permute tracks of the same type to obtain a properly K-batched graph $\mathbb{G}''$, then move certain batches of $\mathbb{G}''$ to obtain $\mathbb{G}^{(K)}$. More precisely we claim the following: There exists a sequence $\mathbf{T}_1^{(1)}, \dots, \mathbf{T}_{L_1}^{(1)}, \dots, \mathbf{T}_1^{(J)}, \dots, \mathbf{T}_{L_J}^{(J)}$, each of which is the composition of infinitely many, non-overlapping track-exchanges and a sequence $\mathbf{S}_1^{(1)}, \dots, \mathbf{S}_{T_1}^{(1)}, \dots, \mathbf{S}_1^{(J)}, \dots, \mathbf{S}_{T_J}^{(J)}$, each of which is the composition of infinitely many non-overlapping star-triangle transformations such that,

$$(\mathbf{S}_{T_J}^{(J)} \circ \cdots \circ \mathbf{S}_1^{(1)} \circ \mathbf{T}_{L_J}^{(J)} \circ \cdots \circ \mathbf{T}_1^{(1)})(\mathbb{G}') = \mathbb{G}^{(K)},$$

and the transformations $\mathbf{T}_j^{(i)}$ and $\mathbf{S}_j^{(i)}$ do not affect the tracks $(t_{K\ell_i n}^{(i)})_{n\in\mathbb{Z}}$.

In particular, Lemma 4.5 applies to each transformation of the type $\mathbf{T}_{L_i}^{(i)} \circ \cdots \circ \mathbf{T}_1^{(i)}$ or $\mathbf{S}_{T_i}^{(i)} \circ \cdots \circ \mathbf{S}_1^{(i)}$ and extends the coupling $\mathbb{P}$ to include a configuration $\omega'' \sim \phi_{\mathbb{G}^{(K)}}$ with

$$\mathbb{P}\big[d_{\text{global CN}}(\omega'|_{\Lambda_N}, \omega''|_{\Lambda_N}) > C' \cdot N^{1-\delta}\big] < C' \cdot N^{-\delta}, \tag{7}$$

where C' is a constant that may depend on $\mathbb{G}$, but not on N or $c > 0$. Combining (6) and (7) provides the desired conclusion after adjustment of the constants. □

Proof of Theorem 4.1, general case. We may restrict our attention to N sufficiently large. Let $c, C > 0$ be given by Theorem 4.2.

Consider N large and let $K = N^{1-c/6}$. Applying Theorem 4.2 to $\mathbb{G}^{(K)}$ and using Proposition 4.4, we obtain a coupling $\mathbb{P}$ between $\omega_0 \sim \phi_{\mathbb{G}}$, $\omega_1 \sim \phi_{\mathbb{G}^{(K)}}$ and $\omega_2 \sim \phi_{\mathbf{Z}^2}$ such that

$$\begin{aligned}
&\mathbb{P}\big[d_{\text{global CN}}(\omega_0|_{\Lambda_N}, \omega_2|_{\Lambda_N}) > CN^{1-c/2} + N^{1-\delta}\big] \\
&\quad \le \mathbb{P}\big[d_{\text{global CN}}(\omega_0|_{\Lambda_N}, \omega_1|_{\Lambda_N}) > N^{1-\delta}\big] + \mathbb{P}\big[d_{\text{global CN}}(\omega_1|_{\Lambda_N}, \omega_2|_{\Lambda_N}) \ge CN^{1-c/2}\big] \\
&\quad \le N^{-\delta} + CN^{-c/3},
\end{aligned}$$

where $\delta > 0$ is given by Lemma 4.5. Adjusting the constants provides the desired conclusion. □

Proof of Lemma 4.5. Fix $c > 0$, N, K and j as in the statement. We will choose $\delta > 0$ below. We will only work with the tracks of type j, which we denote $(t_n)_{n\in\mathbb{Z}}$.

Let $\texttt{Strip}_n$ denote the region enclosed by t_{nK} and $t_{(n+1)K}$ in Λ_N. Observe that Λ_N is then partitioned by $\texttt{Strip}_{-M}, \dots, \texttt{Strip}_M$ for some $M \le CN/K = CN^c$, where C is a constant potentially depending on $\mathbb{G}$.

Using the fact that connections between the points of t_{nK} are identical in ω_{j-1} and ω_j, we get a bijection between loops intersecting these tracks in the two configurations. In order for $d_{\text{global CN}}(\omega_{j-1}|_{\Lambda_N}, \omega_j|_{\Lambda_N}) > N^{1-\delta}$ to occur, one of two things must happen:

a) there exists some $-M \le n \le M$ and a loop γ of ω_{j-1} intersecting t_{nK} such that its coupled loop γ' of ω_j satisfies $d(\gamma, \gamma') > N^{1-\delta}$ or

b) there exists some $-M \leq n \leq M$ and a loop γ of ω_{j-1} or ω_j, completely contained in $\texttt{Strip}_n$ that has a diameter at least $N^{1-\delta}$.

Both points imply the existence of a primal path in either ω_{j-1} or ω_j, contained in some $\texttt{Strip}_n$ with $-M \leq n \leq M$ and of diameter at least $\frac{1}{2}N^{1-\delta}$. Write `Long` for this event. By the strong RSW property and Proposition 2.11 applied to $\phi_{\mathbb{G}_{j-1}}$ and $\phi_{\mathbb{G}_j}$ and using a union bound, we find that

$$\mathbb{P}[\texttt{Long}] \leq C_0 \cdot M \cdot \exp(-c_0 N^{1-\delta}/K),$$

where C_0 and $c_0 > 0$ are constants potentially depending on $\mathbb{G}$. Since $K = N^{1-c}$ and $M \leq CN^c$, we conclude that, by taking $\delta > 0$ small enough, we have $\mathbb{P}[\texttt{Long}] \leq N^{-\delta}$. This provides the desired conclusion. □

Statement on the use of AI, This work was entirely developed and written without the use of AI.

Acknowledgements. The second author acknowledges financial support from the Swiss National Science Foundation (SNSF).

References

[1] Emile Averous, Hugo Duminil-Copin, Tiancheng He, Piet Lammers, and Ioan Manolescu. Conformal invariance of the planar four-state Potts model and its random-cluster representation. 2026. preprint arXiv:2610.08507.

[2] Emile Averous, Hugo Duminil-Copin, Tiancheng He, Piet Lammers, and Ioan Manolescu. Conformal invariance of the six-vertex model's height function. 2026. preprint arXiv:2610.07454.

[3] Vincent Beffara and Hugo Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probab. Theory Relat. Fields*, 153:511–542, 2012.

[4] Federico Camia and Charles M. Newman. Two-dimensional critical percolation: the full scaling limit. *Comm. Math. Phys.*, 268(1):1–38, 2006.

[5] Dmitry Chelkak. Ising model and s-embeddings of planar graphs. *Annales scientifiques de l'École Normale Supérieure*, 57(5):1271–1346, 2024.

[6] Dmitry Chelkak, Hugo Duminil-Copin, Clément Hongler, Antti Kemppainen, and Stanislav Smirnov. Convergence of Ising interfaces to Schramm's SLE curves. *C. R. Math. Acad. Sci. Paris*, 352(2):157–161, 2014.

[7] Dmitry Chelkak and Stanislav Smirnov. Universality in the 2D Ising model and conformal invariance of fermionic observables. *Invent. Math.*, 189(3):515–580, 2012.

[8] Hugo Duminil-Copin. Lectures on the Ising and Potts models on the hypercubic lattice. *PIMS-CRM Summer School in Probability*, 2019.

[9] 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 $q > 4$. *Ann. Sci. Éc. Norm. Supér. (4)*, 54(6):1363–1413, 2021.

[10] Hugo Duminil-Copin, Karol K. Kozlowski, Dmitry Krachun, Ioan Manolescu, and Mendes Oulamara. Rotational invariance in critical planar lattice models. *ArXiV pre-print 2012.11672*, 2021.

[11] Hugo Duminil-Copin, Jhih-Huang Li, and Ioan Manolescu. Universality for the random-cluster model on isoradial graphs. *Electron. J. Probab.*, 23:Paper No. 96, 70, 2018.

[12] Hugo Duminil-Copin and Ioan Manolescu. Planar random-cluster model: scaling relations. *Forum Math. Pi*, 10:Paper No. e23, 83, 2022.

[13] Hugo Duminil-Copin, Ioan Manolescu, and Vincent Tassion. Planar random-cluster model: fractal properties of the critical phase. *Probab. Theory Relat. Fields*, pages 401–449, 2021.

[14] Hugo Duminil-Copin, Vladas Sidoravicius, and Vincent Tassion. Continuity of the phase transition for planar random-cluster and Potts models with $1 \leq q \leq 4$. *Communications in Mathematical Physics*, 349(1):47–107, 2017.

[15] Michael E. Fisher. Renormalization group theory: its basis and formulation in statistical physics. *Rev. Modern Phys.*, 70(2):653–681, 1998.

[16] Alessandro Giuliani, Vieri Mastropietro, and Fabio Lucio Toninelli. Height fluctuations in interacting dimers. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 53(1):98 – 168, 2017.

[17] Alessandro Giuliani, Vieri Mastropietro, and Fabio Lucio Toninelli. Non-integrable dimers: Universal fluctuations of tilted height profiles. *Commun. Math. Phys.*, 377:1883–1959, 2020.

[18] Alexander Glazman and Piet Lammers. Delocalisation and continuity in 2d: Loop O(2), six-vertex, and random-cluster models. *Communications in Mathematical Physics*, 406, 2025.

[19] Geoffrey R. Grimmett. The random-cluster model. *volume 333 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences].*, 2006.

[20] Geoffrey R. Grimmett and Ioan Manolescu. Bond percolation on isoradial graphs: criticality and universality. *Probab. Theory Related Fields*, 159(1-2):273–327, 2014.

[21] Antti Kemppainen and Wendelin Werner. The nested simple conformal loop ensembles in the riemann sphere. *Probability Theory and Related Fields*, (165):835–866, 2016.

[22] Arthur E. Kennelly. The equivalence of triangles and three-pointed stars in conducting networks. *Electrical World and Engineer*, 34(12):413–414, 1899.

[23] Richard Kenyon. Conformal invariance of domino tiling. *Ann. Probab.*, 28(2):759–795, 2000.

[24] Richard Kenyon. The Laplacian and Dirac operators on critical planar graphs. *Invent. Math.*, 150(2):409–439, 2002.

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

[26] R. Peled and D. Romik. Bijective combinatorial proof of the commutation of transfer matrices in the dense $O(1)$ loop model. *Séminaire Lotharingien de Combinatoire*, 73(B73b), 2015.

[27] Gourab Ray and Yinon Spinka. A short proof of the discontinuity of phase transition in the planar random-cluster model with $q > 4$. *Communications in Mathematical Physics*, 378, 2020.

[28] Oded Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000.

[29] Oded Schramm and Scott Sheffield. Harmonic explorer and its convergence to SLE_4. *Ann. Probab.*, 33(6), 2005.

[30] Oded Schramm and Scott Sheffield. Contour lines of the two-dimensional discrete Gaussian free field. *Acta Math.*, 202(1):21–137, 2009.

[31] Stanislav Smirnov. Critical percolation in the plane: conformal invariance, Cardy's formula, scaling limits. *C. R. Acad. Sci. Paris Sér. I Math.*, 333(3):239–244, 2001.