

Near-critical Ornstein–Zernike theory for the planar random-cluster model

Lucas D’Alimonte*, Ioan Manolescu†

October 15, 2025

Abstract

We develop an Ornstein–Zernike theory for the two-dimensional random-cluster model with $1 \leq q < 4$ that also applies in its near-critical regime. In particular, we prove an asymptotic formula for the two-point function which holds uniformly for $p < p_c$ and blends the subcritical and near-critical behaviours of the model.

The analysis is carried out by studying the renewal properties of a subcritical percolation cluster, *at the scale of the correlation length*. More precisely, we explore sequentially the cluster in a given direction, by slices of thickness comparable to the correlation length. We show that this exploration satisfies the properties of a *killed Markov renewal process* — a class of processes that may be analysed independently and have Brownian behaviour. In addition to the two-point function estimate, we derive other consequences of the Ornstein–Zernike theory such as an invariance principle for the rescaled cluster and the strict convexity of the inverse correlation length — all at the scale of the correlation length, uniformly in $p < p_c$.

Finally, our approach differs from that of earlier papers of Campanino, Ioffe, Velenik and others, with the cluster being dynamically explored rather than constructed from its diamond decomposition.

Contents

1	Introduction	2
1.1	Definition of the model	3
1.2	Results	5
1.3	Overview of the proof	6
2	Background on the random-cluster model	7
2.1	Standard properties of the random-cluster model	7
2.2	Near-critical theory	8

*LPSM, Sorbonne Université, dalimonte@lpsm.paris

†Université de Fribourg, ioan.manolescu@unifr.ch

3	Exploring the cluster as a Markov renewal process	12
3.1	Cone-connections	15
3.2	The number of active segments is subcritical	23
3.3	Mixing when $N_k = 1$	31
3.4	Proof of Theorem 3.2	35
4	Local limit theorem for the Markov renewal process	41
4.1	Probability of hitting a half-space	41
4.2	Endpoint concentration when conditioned on survival	44
4.3	Invertibility of the drift and proof of Theorem 1.1	47
4.4	Strict convexity of Wulff shape	49
4.5	Invariance principle	51

1 Introduction

Rigorous understanding of the correlation structure of random fields is one of the main objectives of modern statistical mechanics. This article focuses on one of the most studied such models: the so-called random-cluster measure (also known as FK-percolation). The classical Ornstein–Zernike (OZ) formula describes the behaviour of the two-point function $G(x) := \phi_{p,q}[0 \leftrightarrow x]$ when x goes to infinity, where $\phi_{p,q}$ is the random-cluster measure of parameters p and q , and p is such that the model is in its subcritical (or disordered) phase. It quantifies the behaviour of G beyond its exponential decay and is sharp when x tends to infinity.

The OZ asymptotic was first conjectured in two seminal works by Ornstein and Zernike in 1914 [OZ14], and Zernike [Zer16]. Trying to correct a formula describing the phenomenon of *opalescence* in a crystal, they provided a non-rigorous computation of what would later become the OZ formula. The first rigorous implementation of this reasoning is due to Abraham and Kunz [AK77] and Paes–Leme [PL78] in the context of classical lattice gases theory, by means of a graphical representation of the partition function. Soon after, the Ornstein–Zernike asymptotic was proved for a variety of models *in a perturbative regime*, see [BF85, MZ96].

The next breakthrough consisted in a rigorous derivation of the Ornstein–Zernike asymptotic in the whole regime of exponential decay of the correlation functions; this was done in the case of the self-avoiding walk in the direction of an axis in [CC86] and later on in any direction in [Iof98]. For percolation models, the case of Bernoulli percolation was first treated for an axis direction in [CCC91] and later on in any direction in [CI02]. The case of subcritical Ising models was treated in [CIV03] via the random-line graphical representation of the two-point function. Finally, the analysis was carried out for all the subcritical random-cluster measures in [CIV08].

In recent developments, the theory has been extended to Ising models with long-range interactions [AOV24], but was also shown to fail in some long-range Ising models, when the coupling constants decay too slowly [AIOV21, AOV23].

All works cited above aim to understand correlations at a fixed, subcritical temperature. In this article, we are interested in how the OZ formula behaves when the temperature parameter (p in the random-cluster model) approaches its critical value. Our study is limited to the random-cluster model on the two-dimensional square lattice $\mathbb{Z}^2$, with $q \in [1, 4)$ and $p < p_c(q) = \frac{\sqrt{q}}{1+\sqrt{q}}$. The near-critical properties of the model were recently studied in [DCM22].

Our main output is an asymptotic formula for the function $x \mapsto G(x)$ *uniformly both in p and x* , in the regime $p < p_c$ and $\|x\|$ above the correlation length. When $p = p_c$ and $q \in [1, 4]$ the phase transition is known to be continuous [DCST17] and the function G decays algebraically in $\|x\|$. The main feature of our analysis is that it captures how the two-point function switches from a disordered behaviour (characterised by its exponential decay) to a near-critical behaviour (with the same algebraic decay as at criticality). To do this, we needed to replace the arguments of [CIV08] relying on the finite energy property by RSW-type constructions at the scale of the correlation length.

Besides this new result, another interest of the present paper is to revisit the approach developed in [CIV08] and related papers. Indeed, while our argument uses key elements of [CIV08], several differences may be highlighted. We circumvent the use of the “skeleton calculus” and employ a direct exploration of the subcritical cluster in a given direction. The argument also bypasses the perturbation theory used in [CIV08] to deduce the strict convexity properties of the correlation length and of the associated Wulff shape. Instead, we deduce these properties from rather standard large deviations estimates for the associated random exploration.

Our argument is limited to the two-dimensional model as this is the only case where the near-critical behaviour of the random-cluster model has been established. We believe that the same strategy may be implemented in general dimension for fixed $p < p_c$, where all RSW arguments may be replaced by finite-energy constructions, to reprove the results of [CIV08].

1.1 Definition of the model

In this section we briefly define the model, with some additional features deferred to Section 2. We refer to the monographs [Gri06, DC20] for further background.

We slightly abuse notation by writing $\mathbb{Z}^2 = (\mathbb{Z}^2, E(\mathbb{Z}^2))$ for the two-dimensional square lattice, with vertex set $\mathbb{Z}^2$ and edges between vertices at Euclidean distance 1. For $G = (V(G), E(G))$ a subgraph of $\mathbb{Z}^2$, the space of percolation configurations on G is $\Omega^G := \{0, 1\}^{E(G)}$. For an edge $e \in E(G)$ and $\omega \in \Omega^G$, we say that e is *open* if $\omega(e) = 1$ and *closed* otherwise. A percolation configuration will be identified both with the set of its open edges as well as with the sub-graph of G with vertex set $V(G)$ and edge-set formed of the open edges of ω . The connected components of ω are called *clusters*.

The boundary ∂G of G is the set of vertices of G incident to at least one edge of $E(\mathbb{Z}^2) \setminus E(G)$. A *boundary condition* η is a partition of ∂G ; we say that the vertices of ∂G that belong to the same component of η are *wired* together. To a boundary condition η and a percolation configuration $\omega \in \{0, 1\}^{E(G)}$, associate the graph ω^η which

is obtained from ω by identifying all the mutually wired vertices of ∂G .

A percolation configuration ξ on $\mathbb{Z}^2$ induces certain boundary conditions on ∂G : two vertices are wired together if they are connected in $\xi \setminus E(G)$. We shall make a slight notational abuse by identifying the percolation configuration ξ with the boundary condition it induces on ∂G , and keeping the notation ω^ξ when ξ is a percolation configuration on $\mathbb{Z}^2 \setminus G$. Two boundary conditions play a special role: the *free* boundary condition, denoted by 0, in which no boundary vertices are wired together, and the *wired* boundary condition, denoted by 1, in which all boundary vertices are wired together.

The random-cluster measure on a finite subgraph G of $\mathbb{Z}^2$, with boundary conditions η and parameters $p \in (0, 1)$ and $q \geq 1$ is defined as follows. For a percolation configuration ω on G write $o(\omega)$ for number of open edges of ω , and $k(\omega^\eta)$ the number of clusters of ω^η . Set

$$\phi_{G,p,q}^\eta[\omega] = \frac{1}{Z_{G,p,q}^\eta} \left(\frac{p}{1-p} \right)^{o(\omega)} q^{k(\omega^\eta)},$$

where $Z_{G,p,q}^\eta$ is called the *partition function* of the model, and is the unique constant guaranteeing that $\phi_{G,p,q}^\eta$ is a probability measure.

Random-cluster measures may also be defined on the full graph $\mathbb{Z}^2$ either through the DLR formalism or by taking limits of measures on increasing finite subgraphs of $\mathbb{Z}^2$. Due to monotonicity properties, it is classical that the free and wired measures $\phi_{G,p,q}^0$ and $\phi_{G,p,q}^1$ admit limits as G increases to $\mathbb{Z}^2$. These will be denoted by $\phi_{p,q}^0$ and $\phi_{p,q}^1$, respectively, and are instances of *infinite-volume measures*, which are invariant under translations and ergodic.

It was proved in [BDC12] that the model exhibits a phase transition at the self-dual parameter $p_c(q) = \frac{\sqrt{q}}{1+\sqrt{q}}$. That is, for any $p < p_c(q)$ (in the so-called *subcritical regime*), there exists a unique infinite-volume measure ($\phi_{p,q}^0 = \phi_{p,q}^1$) and all clusters are almost surely finite under this measure. When $p > p_c(q)$, the infinite-volume measure is again unique and contains almost surely a unique infinite cluster. The phase transition was shown to be continuous for $q \in [1, 4]$ [DCST17], in the sense that there exists a unique infinite-volume measure also at $p = p_c(q)$. Additionally, strong RSW-type estimates were established at criticality for these values of q [DCST17, DCT20, DCMT21]. For $q > 4$, the phase transition is discontinuous [DCGH⁺21, RS20], which is to say that $\phi_{p_c(q),q}^0 \neq \phi_{p_c(q),q}^1$, with the former having a subcritical behaviour and the latter a super-critical one.

As already mentioned, we are treating here the case of a continuous phase transition, for which the correlation length diverges as $p \nearrow p_c$. Furthermore, for reasons explained in Remark 1.2, the case $q = 4$ is excluded from our analysis.

Henceforth, when $q \in [1, 4)$ is fixed, we will omit it from the notation. As the infinite-volume measure is unique for any $p \in [0, 1]$, we denote it by ϕ_p .

For the exposition of our results, we need to introduce two classical quantities in the study of subcritical and near-critical random-cluster model. Fix $q \in [1, 4)$ and $p < p_c$. The *correlation length* of the model is defined as the following limit, the existence of

which is based on super-multiplicativity arguments. For $\vec{v} \in \mathbb{R}^2 \setminus \{0\}$, set¹

$$\xi_p(\vec{v}) = \left(\lim_{n \rightarrow \infty} -\frac{1}{n} \log \phi_p[0 \leftrightarrow \lfloor n\vec{v} \rfloor] \right)^{-1}.$$

Write $\|\cdot\|$ for the Euclidian norm on $\mathbb{R}^2$ and let $\mathbb{S}^1 = \{\vec{v} \in \mathbb{R}^2 : \|\vec{v}\| = 1\}$. Notice that, for any $\vec{v} \neq 0$,

$$\|\vec{v}\| \xi_p(\vec{v}) = \xi_p(\vec{v}/\|\vec{v}\|),$$

so we will mostly consider the case $\vec{v} \in \mathbb{S}^1$. The results of [BDC12] imply that whenever $p < p_c$, $\xi_p(\vec{v}) > 0$ for any $\vec{v} \in \mathbb{S}^1$.

For $R \geq 0$, let $\Lambda_R := \{-R, \dots, R\}^2$ and $\Lambda_R(x)$ for its translate by x . Define the *critical one-arm probability* at distance R as

$$\pi_1(R) = \phi_{p_c}^0[0 \leftrightarrow \partial\Lambda_R].$$

1.2 Results

We are now ready to state our main results. For positive quantities f, g we write $f \lesssim g$ to mean that there exists a constant $C > 0$ such that $f \leq Cg$. The constants C may be chosen uniformly in certain parameters in the definitions of f and g which will be explicitly stated – these will generally be $p < p_c$, n and t below. Write $f \asymp g$ when $f \lesssim g$ and $g \lesssim f$. We will often use the expression *uniformly in* $p < p_c$ to mean uniformly for $p \in [\varepsilon, p_c)$ for some $\varepsilon > 0$ fixed throughout the paper. The choice of ε is arbitrary, but may affect the value of certain constants.

Theorem 1.1 (Ornstein–Zernike asymptotics at the scale of the correlation length). *Fix $q \in [1, 4)$. Then, uniformly in $\vec{v} \in \mathbb{S}^1$, $n \geq \xi_p(\vec{v})$ and $p < p_c$,*

$$\phi_p[0 \leftrightarrow \lfloor n\vec{v} \rfloor] \asymp \pi_1(\xi_p(\vec{v}))^2 \sqrt{\frac{\xi_p(\vec{v})}{n}} e^{-\frac{n}{\xi_p(\vec{v})}}. \quad (1)$$

Compared to the classical OZ results of [CIV08], the advantage of our result is that it is uniform in p . Indeed, for $q \in [1, 4)$, $\xi_p(\vec{v}) \rightarrow \infty$ as $p \nearrow p_c$. Previous results often employ local surgeries based on finite energy at scales lower than $\xi_p(\vec{v})$. As such they do not apply when $\xi_p(\vec{v})$ diverges. Our approach blends the critical and subcritical behaviour, each accounting for one of the terms in the right-hand side of (1): the term $e^{-\frac{n}{\xi_p(\vec{v})}} (\frac{n}{\xi_p(\vec{v})})^{-1/2}$ is reminiscent of classical OZ-type formulas and is entirely a subcritical phenomenon, while the term $\pi_1(\xi_p(\vec{v}))^2$ appears due to the near-critical behaviour of the model. Finally, when $n \leq \xi_p(\vec{v})$, the right-hand side may be replaced by $\pi_1(n)^2$, as proved in [DCM22].

Theorem 1.1 is a consequence of a more detailed description of the cluster of 0 when conditioned to be connected to a half-plane at a distance n in the direction $\vec{w}$, dual to $\vec{v}$; see Theorem 3.2 below. We would like to emphasise that Theorem 3.2 also has other consequences, such as:

¹ $\lfloor n\vec{v} \rfloor$ denotes the vertex of $\mathbb{Z}^2$ which is the closest to $n\vec{v}$. In case of equality, we arbitrarily choose the top-left most one.

- Corollary 4.2 which states that the probability for 0 to be connected to the half-space $\{\langle x, \vec{w} \rangle \geq n\}$ is uniformly comparable to $\pi_1(\xi_p(\vec{w}))$ times a pure exponential of $n/\xi_p(\vec{w})$.
- Theorem 4.9 which states that the inverse correlation length ξ_p is strictly convex in terms of the direction, as is its convex dual, known as the *Wulff shape*.
- Theorem 4.10 which establishes an invariance principle for the cluster of 0 when conditioned to be connected to $[n\vec{v}]$ or to a half-plane.

Contrary to previous approaches, Theorem 3.2 provides all these consequences simultaneously.

Remark 1.2. We conclude this introduction by discussing the requirement that $q \in [1, 4)$. When $q > 4$, it has been shown in [DCGH⁺21] that the correlation length is uniformly bounded in the whole subcritical regime and that the phase transition is of first order. Thus, a OZ result valid uniformly for $p < p_c$ would also hold for the free measure at p_c . At the critical point, interfaces between the free and wired phases are expected to have random-walk behaviour, while primal clusters under the free measure, when conditioned to hit a far-away half-space, resemble the space between two random walks conditioned not to intersect — the upcoming paper [DGO25] aims to prove this. This behaviour is inconsistent with the OZ behaviour of Theorem 1.1; see also [D'A24, Corollary 1.8] for a more detailed discussion of the asymptotics of the two-point function in this case.

For $q = 4$, the correlation length does diverge as $p \nearrow p_c$, as for $q < 4$. However, contrary to the case $q < 4$, we expect that the ratio between the typical width of a long cluster and the correlation length diverges as $p \nearrow p_c$, possibly logarithmically. The core difference between the cases $q < 4$ and $q = 4$ is that for $q < 4$ the RSW theory crossing estimate holds up to the boundary, with adverse boundary conditions. This is ultimately reflected in Lemma 2.1, which we do not believe holds for $q = 4$; see Section 2 for details.

That being said, our strategy may be used for any $q \geq 1$ and $p < p_c$ (and even any dimension $d \geq 2$) to obtain OZ results which are uniform in n and $\vec{v}$, but not in $p < p_c$. Thus, the present paper offers an alternative approach to the theory of [CIV08].

1.3 Overview of the proof

Fix some $\vec{w} \in \mathbb{S}^1$. We aim to describe the cluster of the origin, conditioned to hit a distant half-plane in the direction $\vec{w}$. The idea is to slice the plane with lines orthogonal to $\vec{w}$ placed at regular intervals of length comparable to $\xi_p(\vec{w})$. We call these lines *hyperplanes*, as they should have co-dimension 1 in the more general d -dimensional setting. For simplicity of exposition, imagine that $\vec{w} = e_1$ is the horizontal unit vector. This is not limiting, as the symmetries of the lattice are never used.

The cluster of 0 is then explored from left to right in a Markovian way: explore the connected component in the half-space left of the k^{th} hyperplane and write X_k its highest point on the hyperplane². When no such intersection exists, write $X_k = \dagger$ and

²The choice of X_k as highest on the hyperplane is arbitrary. Any choice determined by the explored cluster should have the same properties

say that X_k dies. We will show that the process $(X_k)_k$ has a certain renewal structure, even when conditioned on surviving for n steps. As such, the process is decomposed into irreducible pieces and behaves essentially as a random walk, with all consequences following by standard tools.

The ideas are similar to the previous works such as [CIV08], but we believe are rephrased in a novel way. The decomposition of the cluster in irreducible pieces, and its random walk interpretation, appeared already in these works. However, the present work does not use the “diamond decomposition” and explores the cluster in a dynamical way, that is less dependent on the point it ultimately reaches.

Acknowledgements The authors are grateful to Sébastien Ott for helpful discussions and insights into the Ornstein–Zernike theory. This work was partially supported by the Swiss National Science Foundation.

2 Background on the random-cluster model

In this section we provide some more elements from the standard theory of the random-cluster model, which readers familiar with the model may skip. We also review some recent results of [DCM22] about the near-critical regime of the two-dimensional random-cluster model. Fix $q \geq 1$ for the whole section.

2.1 Standard properties of the random-cluster model

Monotonicity properties. The set of percolation configurations on $\mathbb{Z}^2$ is equipped with the partial order defined by $\omega \leq \omega'$ if for any $e \in E(\mathbb{Z}^2)$, $\omega(e) \leq \omega'(e)$. An event A is said to be *increasing* if for any two percolation configurations $\omega \leq \omega'$, $\omega \in A \Rightarrow \omega' \in A$. Set $q \geq 1$, $p \in [0, 1]$ and G a finite subgraph of $\mathbb{Z}^2$. The *FKG inequality* asserts that for any increasing events A, B , and any boundary condition η on G ,

$$\phi_{G,p}^\eta[A \cap B] \geq \phi_{G,p}^\eta[A] \phi_{G,p}^\eta[B]. \quad (\text{FKG})$$

The random-cluster measure also possesses the following monotonicity property. If $\xi \leq \xi'$ are boundary conditions on G (which is to say that ξ' is a coarser partition than ξ), then for any increasing event A ,

$$\phi_{G,p}^\xi[A] \leq \phi_{G,p}^{\xi'}[A]. \quad (\text{MON})$$

Both properties above extend to infinite-volume limits.

Domain Markov property. Let G be finite some subgraph of $\mathbb{Z}^2$. Fix $q \geq 1$ and $p \in (0, 1)$. Let $G' = (V', E')$ be a subgraph of G . Then for any boundary condition η on G , any percolation configuration $\xi \in \{0, 1\}^{E \setminus G'}$,

$$\phi_{G,p}^\eta[\cdot | G' \mid \omega_{E \setminus E'} = \xi] = \phi_{G',p}^{\xi \eta}[\cdot],$$

where ξ^η is the boundary condition induced on the complement of G' by ξ together with the boundary condition η .

Duality. Consider the dual graph $(\mathbb{Z}^2)^*$ with vertex set $V(\mathbb{Z}^2) + (1/2, 1/2)$ and edge set $\{i + (1/2), j + (1/2)\}$, for i, j such that $\{i, j\} \in E(\mathbb{Z}^2)$. To any percolation configuration on $\mathbb{Z}^2$ we associate its *dual configuration*, defined on the graph $(\mathbb{Z}^2)^*$ by setting $\omega^*(e) = 1 - \omega(e^*)$, where e^* is the unique edge of $(\mathbb{Z}^2)^*$ that crosses e . It is classical that when $\omega \sim \phi_p^0$ then $\omega^* \sim \phi_{p^*}^1$, where p and p^* are linked by the following duality relation:

$$pp^* = q(1 - p)(1 - p^*).$$

Note that the value $p_{\text{sd}} := \frac{\sqrt{q}}{1 + \sqrt{q}}$ is the unique solution of $p = p^*$, and coincides with the critical parameter p_c as first proved in [BDC12].

2.2 Near-critical theory

The near-critical regime of the random-cluster model is the set of parameters n and p for which n is sufficiently small — or p sufficiently close to p_c — so that the system behaves almost critically at scale n . It is expected, and is indeed the case for the two-dimensional random-cluster model, that the system behaves roughly critically in the near-critical regime, and sub- or super-critically outside of it.

The rigorous understanding of the near-critical regime of percolation models in two dimensions started with Kesten's seminal work on Bernoulli percolation [Kes87]. Kesten's results were adapted to the random-cluster model on $\mathbb{Z}^2$ with $q \in [1, 4]$ in [DCM22]. Here we will mention only the consequences of these works that are relevant to us.

For $q \geq 1$ fixed and $p < p_c$, define the *characteristic length*

$$L(p) = \inf \{n \geq 0 : \phi_p[\text{Cross}(\Lambda_n)] \notin [\delta, 1 - \delta]\},$$

where $\text{Cross}(\Lambda_n)$ is the event that Λ_n contains an open path between its left and right sides, and $\delta > 0$ is some small fixed quantity.

It was proved in [DCM22] that, for any $\vec{v} \in \mathbb{S}^1$,

$$L(p) \asymp \xi_p(\vec{v}) \tag{2}$$

uniformly in $p < p_c$, for $q \in [1, 4]$. Moreover, $L(p) \rightarrow \infty$ as $p \nearrow p_c$.

The equivalence (2) also holds for $q > 4$, but only because both $L(p)$ and $\xi_p(\vec{v})$ are uniformly bounded away from 0 and ∞ .

Henceforth we will use $L(p)$ rather than $\xi_p(\vec{v})$ to designate a quantity of the order of the correlation length, for instance when referring to the interspacing of the hyperplanes used to define the process $(X_n)_n$. We do so to emphasise that its use is not related to the direction $\vec{v}$ and is only important up to a bounded multiplicative constant.

Fix $\vec{w} \in \mathbb{S}^1$; for all practical purposes, think of $\vec{w}$ as the unit vector in the horizontal direction. Define the half-spaces

$$\mathcal{H}_{\leq t}^{\vec{w}} = \{x \in \mathbb{R}^2 : \langle x, \vec{w} \rangle \leq t L(p)\} \quad \text{and} \quad \mathcal{H}_{\geq t}^{\vec{w}} = \{x \in \mathbb{R}^2 : \langle x, \vec{w} \rangle \geq t L(p)\}.$$

Call $\partial \mathcal{H}_{\leq t}^{\vec{w}} = \partial \mathcal{H}_{\geq t}^{\vec{w}} = \{x \in \mathbb{R}^2 : \langle x, \vec{w} \rangle = t L(p)\}$ a hyperplane and set $\mathcal{H}_{< t}^{\vec{w}} = \mathcal{H}_{\leq t}^{\vec{w}} \setminus \partial \mathcal{H}_{\leq t}^{\vec{w}}$. In the following we will use an arbitrary integer approximation of these objects, which

we do not detail. We will mostly work with $\vec{w}$ fixed, and will omit it from the notation whenever no ambiguity is possible.

A particular consequence of the discussion above is the following bound on the speed of exponential decay. The particular form of the statement below is due to its specific use in our proofs. A *potential past* is any connected set of edges $A \subset \mathcal{H}_{\leq 0}$. Write $\partial_{\leq 0}A$ for all edges adjacent to A , contained in $\mathcal{H}_{\leq 0}$.

Lemma 2.1. *Fix $q \in [1, 4)$. There exists a constant $c > 0$ such that the following holds. For any $p < p_c$, any $n \geq L(p)$ and any potential past A ,*

$$\phi_{\Lambda_n}^1[\Lambda_{L(p)} \xleftrightarrow{A^c} \partial\Lambda_n \mid \omega \equiv 1 \text{ on } A \text{ and } \omega \equiv 0 \text{ on } \partial_{\leq 0}A] \leq e^{-cn/L(p)}. \quad (3)$$

The novelty in the statement above is the conditioning on the edges in A and $\partial_{\leq 0}A$; indeed the result for the unconditioned measure is a simple consequence of [DCM22]. Note that the connection in (3) needs to occur *outside* of A . Thus, it would be tempting to assume that the probability in (3) decreases with A , but the potential boundary conditions A induces may actually increase the probability of the connection. Indeed, if we set $D = \Lambda_n \setminus (A \cup \partial_{\leq 0}A)$, we work above in the measure ϕ_D^ξ , where ξ are the boundary conditions on D that are wired the endpoints of edges of A on $\partial\mathcal{H}_{\leq 0}$, wired on $\partial\Lambda_n$, but free on the endpoints of $\partial_{\leq 0}A$.

Finally, the assumption $q < 4$ is important for the above. For $q > 4$ the result is known to fail, and is expected to fail for $q = 4$. In the latter case, if $A = \mathcal{H}_{\leq 0}$, the cluster of the origin may “creep along” $\partial\mathcal{H}_{\leq 0}$ at a slightly smaller cost than if $A = \emptyset$.

We turn our attention to scales below the characteristic length, called the *critical window*. It was proved in [DCM22, Thm. 2.1] that the RSW property still holds in this regime. Let $\text{Circ}(r, R)$ be the event that Λ_R contains an open circuit surrounding Λ_r . Write $\text{Circ}^*(r, R)$ for the event that the dual configuration contains such a circuit, which for the primal model translates to Λ_r not being connected to $\partial\Lambda_R$.

Proposition 2.2 (RSW in the critical window). *Fix $q \in [1, 4]$. There exists $c > 0$ such that for any $p < p_c$ and $n \leq L(p)$ and any boundary condition ξ on Λ_{2n} ,*

$$c \leq \phi_{\Lambda_{2n}, p}^\xi[\text{Circ}(n, 2n)] \leq 1 - c \quad \text{and} \quad c \leq \phi_{\Lambda_{2n}, p}^\xi[\text{Circ}^*(n, 2n)] \leq 1 - c. \quad (\text{RSW})$$

Finally, the most significant contribution of [DCM22] was to prove the stability of the arm event probabilities in the near-critical regime.

Theorem 2.3. *Fix $q \in [1, 4]$. Then*

$$\phi_{\Lambda_{2R}, p}^\xi[0 \leftrightarrow \partial\Lambda_R] \asymp \pi_1(R), \quad (4)$$

uniformly in $p < p_c$, $R \leq L(p)$ and any boundary conditions ξ on $\partial\Lambda_{2R}$

We finish the section with the proof of Lemma 2.1. This is a relatively standard consequence of the exponential decay above $L(p)$ [DCM22, Prop. 2.13], but the authors are unaware of a specific reference. We will omit some details in the proof below.

Proof of Lemma 2.1. Fix $q \in [1, 4)$ and $p < p_c$; write $L = L(p)$. Constants below are independent of p . Figure 1 contains helpful illustrations for the proof.

Before we start, we mention that it is an immediate consequence of [DCM22, Prop. 2.13] that, for all $n \geq L$,

$$\phi_{\Lambda_{2n}}^1[\Lambda_n \leftrightarrow \Lambda_{2n}] \leq e^{-cn/L}, \quad (5)$$

for some universal constant $c > 0$. In particular, this implies (3) for $A = \emptyset$.

We turn to the general proof. For A a potential past set $\text{Past}(A)$ for the event that all edges of A are open, but all those of $\partial_{\leq 0}A$ are closed. Also, write A^* for all edges dual to edges of A or $\partial_{\leq 0}A$.

For $k \geq 1$, set³

$$u_k = \inf_A \phi_{\Lambda_{2^k L}}^1[\Lambda_{2^{k-2}L}(-2^{k-1}L, 0) \xrightarrow{\omega^* \cup A^*} \Lambda_{2^{k-2}L}(2^{k-1}L, 0) \mid \text{Past}(A)],$$

where the infimum is over all potential pasts A . We will start by proving that there exists some universal constant $c_1 > 0$ such that, for any $k \geq 1$,

$$u_k \geq 1 - e^{-c_1 2^k}. \quad (6)$$

We do this by proving that u_1 is bounded away from 0 uniformly in p and that

$$u_{k+1} \geq (1 - e^{-c_0 2^k})^2 (1 - (1 - u_k)^2) \quad (7)$$

for all $k \geq 1$, where $c_0 > 0$ is a universal constant.

To prove the bound on u_1 , consider a potential past A . If $A \cap \Lambda_{2L}$ has a diameter smaller than L , there exists a rectangle of width $L/2$ and length $2L$ connecting the boxes $\Lambda_{L/2}(-L, 0)$ to $\Lambda_{L/2}(L, 0)$. Using (RSW), we prove that the two boxes are connected in ω^* with uniformly positive probability. If $A \cap \Lambda_{2L}$ has a diameter larger than L , there exists tubes of length at most $4L$ and width at most $L/2$ that connect $\Lambda_{L/2}(-L, 0)$ and $\Lambda_{L/2}(L, 0)$ to A . Then⁴, applying [DCMT21, Thm.7], we find

$$\phi_{\Lambda_{2L}}^1[\Lambda_{L/2}(-L, 0) \xrightarrow{\omega^*} A^* \mid \text{Past}(A)] \geq 1 \text{ and } \phi_{\Lambda_{2L}}^1[\Lambda_{L/2}(L, 0) \xrightarrow{\omega^*} A^* \mid \text{Past}(A)] \geq 1.$$

Applying (FKG), we conclude that, in all cases,

$$\phi_{\Lambda_{2L}}^1[\Lambda_{L/2}(-L, 0) \xrightarrow{\omega^* \cup A^*} \Lambda_{L/2}(L, 0) \mid \text{Past}(A)] \geq 1.$$

We turn to the recurrence relation (7). Fix $k \geq 1$ and A a potential past. Write $N = 2^k L$. Consider the two non-overlapping vertical translates of Λ_N by $(0, -N)$ and $(0, N)$, respectively, and the corresponding translates H_+ and H_- of the events in the definition of u_k :

$$H_{\pm} = \left\{ \Lambda_{N/4}(-N/2, \pm N) \xrightarrow{(\omega^* \cup A^*) \cap \Lambda_N(0, \pm N)} \Lambda_{N/4}(N/2, \pm N) \right\}.$$

³Recall that $\Lambda_r(a, b)$ is the translate of Λ_r by the vector $(a, b) \in \mathbb{R}^2$.

⁴It is essential here that $q < 4$; this estimate is expected to fail at $q = 4$.




Figure 1: *Left*: in the proof of (7) we consider two translates H_+ and H_- of the event in u_k , one in the upper half-plane, one in the lower one. Here the red path ensures that H_+ occurs. The blue paths (also part of $\omega^* \cup A^*$) produce the events B_L and B_R . When at least one of H_+ and H_- occur, and both B_L and B_R , then the two horizontal translates of $\Lambda_{N/2}$ connect to each-other. *Right*: To produce a path separating $\Lambda_{2^k L}$ from $\partial\Lambda_{2^{k+1}L}$, it suffices for two vertical translates of the events in u_{k-1} to occur, and for the blue paths to occur in $\omega^* \cup A^*$. As for B_L and B_R , the latter occur with exponentially high probability in 2^k .

See Figure 1, left diagram.

Under $\phi_{\Lambda_{2N}}^1[\cdot | \text{Past}(A)]$, that the events H_+ and H_- occur may be bounded by independent Bernoulli variables of parameter u_k . Thus,

$$\phi_{\Lambda_{2N}}^1[H_+ \cup H_- | \text{Past}(A)] \geq 1 - (1 - u_k)^2.$$

Now, let B_R be the event that, in $\omega^* \cap \Lambda_{2N} \cap \mathcal{H}_{>0}$, there exists a circuit around $\Lambda_{N/4}(N/2, N)$, contained in $\Lambda_{N/2}(N/2, N)$, connected to a circuit around $\Lambda_{N/4}(N/2, -N)$, contained in $\Lambda_{N/2}(N/2, -N)$. Then, by (FKG) and (5),

$$\phi_{\Lambda_{2N}}[B_R | \text{Past}(A)] \geq \phi_{\Lambda_{2N} \cap \mathcal{H}_{\geq 0}}^1[B_R] \geq 1 - e^{-c_0 2^k}.$$

for some universal positive constant $c_0 > 0$.

Similarly, define B_L be the event that, in $(\omega^* \cup A^*) \cap \Lambda_{2N} \cap \mathcal{H}_{<0}$, there exists a circuit surrounding $\Lambda_{N/4}(-N/2, N)$, contained in $\Lambda_{N/2}(-N/2, N)$, connected to a circuit surrounding $\Lambda_{N/4}(-N/2, -N)$, contained in $\Lambda_{N/2}(-N/2, -N)$. Due to (MON), when sampled according to $\phi_{\Lambda_{2N}}^1[\cdot | \text{Past}(A)]$, $\omega^* \cup A^*$ in $\mathcal{H}_{\leq 0}$ dominates the dual configuration sampled according to $\phi_{\Lambda_{2N} \cap \mathcal{H}_{\leq 0}}^1$. As such

$$\phi_{\Lambda_{2N}}[B_L | \text{Past}(A)] \geq \phi_{\Lambda_{2N} \cap \mathcal{H}_{\leq 0}}^1[B_L] \geq 1 - e^{-c_0 2^k}.$$

Notice that, when B_L and B_R and at least one of H_+ or H_- occur, $\Lambda_{N/2}(-N, 0)$ is connected to $\Lambda_{N/2}(N, 0)$ in $\omega^* \cup A^*$. By (FKG) we conclude that

$$\begin{aligned} \phi_{\Lambda_{2N}}^1[\Lambda_{N/2}(-N, 0) \xleftrightarrow{\omega^* \cup A^*} \Lambda_{N/2}(N, 0) \mid \text{Past}(A)] \\ \geq \phi_{\Lambda_{2N}}^1[(H_+ \cup H_-) \cap B_L \cap B_R \mid \text{Past}(A)] \geq (1 - e^{-c_0 2^k})^2 (1 - (1 - u_k)^2). \end{aligned} \quad (8)$$

This concludes the proof of (7). Finally, the lower bound on u_1 and (7) readily imply (6).

We turn to proving (3). Let $k = \lfloor \log_2(n/L(p)) \rfloor - 1$. It suffices to consider the cases $k \geq 2$; (RSW) allow us to extend (3) to smaller values of n by potentially altering the constant. Then

$$\phi_{\Lambda_n}^1[\Lambda_L \leftrightarrow \partial\Lambda_n \mid \text{Past}(A)] \leq 1 - \phi_{\Lambda_{2^{k+1}L}}^1[\omega^* \cup A^* \text{ has circuit around } \Lambda_{2^k L} \mid \text{Past}(A)] \quad (9)$$

The event in the right-hand side above may be constructed by two instances of the events in the definition of u_{k-1} (translated by $(0, 3 \cdot 2^{k-1}L)$ and $(0, -3 \cdot 2^{k-1}L)$, respectively) and a series of connection events relating only to the left and right half-planes. Rather than defining these precisely, we direct the reader to Figure 1, right diagram, and to the explanation behind (8), which prove that

$$\phi_{\Lambda_{2^{k+1}L}}^1[\omega^* \cup A^* \text{ has circuit around } \Lambda_{2^k L} \mid \text{Past}(A)] \geq (1 - e^{-c_2 2^k})^2 u_{k-1}^2.$$

Combining the above with (6) and inserting the result in (9) we obtain (3) with a constant $c > 0$ that is independent of p and A . $\square$

3 Exploring the cluster as a Markov renewal process

The key idea of our approach is to show that the subcritical cluster of the origin admits what we call a “killed renewal structure”, and that the renewal structure persists even when conditioned to hit a far-away half-plane. We start by defining a class of processes which will have a killed random-walk like behaviour.

Definition 3.1. A stochastic process $(X_t, Y_t)_{t \geq 1} \in (\mathbb{R} \cup \{\dagger\} \times \{0, 1\})^{\mathbb{N}}$ is called a *killed Markov renewal process* (KM RP in short) with respect to some filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$ if

- It is adapted to $(\mathcal{F}_t)_{t \in \mathbb{N}}$;
- If $X_t = \dagger$, then $Y_t = 0$ and $X_{t+1} = \dagger$;
- If we set $T_0 = 0$ and $T_k = \inf\{t > T_{k-1} : Y_t = 1\}$, then, for any $k \geq 1$, the process $(X_{t+T_k} - X_{T_k}, Y_{t+T_k})_{t \in \mathbb{N}}$ conditionally on $\mathcal{F}_{T_k}$ and $X_{T_k} \neq \dagger$ has a fixed law $\mathcal{L}$.

We say that the process has *exponential tails* if there exists $c > 0$ such that, for all $k, j, n \geq 1$,

$$\mathbb{P}[X_{n+T_k} \neq \dagger \text{ and } T_{k+1} - T_k > n \mid \mathcal{F}_{T_k}, T_k < \infty] \leq \exp(-cn) \quad \text{and} \quad (10)$$

$$\mathbb{P}[X_{T_k+j} \neq \dagger \text{ and } |X_{T_k+j} - X_{T_k}| \geq n \mid \mathcal{F}_{T_k}, T_k < \infty] \leq \exp(-cn/j). \quad (11)$$

Moreover, for the initial step, for any $j, n \geq 1$,

$$\mathbb{P}[X_n \neq \dagger \text{ and } T_1 > n \mid X_1 \neq \dagger] \leq e^{-cn} \quad \text{and} \quad (12)$$

$$\mathbb{P}[X_j \neq \dagger \text{ and } |X_j| \geq n \mid X_1 \neq \dagger] \leq e^{-cn/j}. \quad (13)$$

We say that the model exhibits a *mass-gap* $c > 0$ if the above displays also hold under the additional conditioning $X_{T_k+N} \neq \dagger$ for any $N \geq 1$.

Finally, we call the process *aperiodic* if the greatest common divisor of the support of the increment $T_2 - T_1$ is one.

For such a process and $k \geq 1$, call the laws of $X_{T_{k+1}} - X_{T_k}$ and $T_{k+1} - T_k$ conditionally on $T_k \neq \infty$ the X -step and the T -step (or vertical and horizontal steps, respectively). Notice that these only depend on the law $\mathcal{L}$ mentioned above, and therefore do not depend on k . Define the X -step mean and variance as

$$\mu_X := \mathbb{E}[X_{T_{k+1}} - X_{T_k} \mid T_{k+1} < \infty] \quad \text{and} \quad \sigma_X := \text{Var}[X_{T_{k+1}} - X_{T_k} \mid T_{k+1} < \infty]$$

and the killing rate as

$$\kappa := \mathbb{P}[T_{k+1} = \infty \mid T_k < \infty].$$

We will now describe how such general processes are related to the cluster of 0 under different conditionings.

Fix $\vec{w} \in \mathbb{S}^1$. For all practical purposes, think of $\vec{w}$ as the unit vector in the horizontal direction, pointing to the right; this is only so that the vocabulary and illustrations below make sense, but has no bearing on the arguments. Recall the definitions of the half-spaces $\mathcal{H}_{\leq t}^{\vec{w}}$ and $\mathcal{H}_{\geq t}^{\vec{w}}$; the dependence on $\vec{w}$ will be omitted whenever possible.

Write $\mathbf{C}$ for the cluster of 0 and $\mathbf{C}_{\leq t}$ for the cluster of 0 in $\omega \cap \mathcal{H}_{\leq t}$. Note that this is contained in, but not always equal to $\mathbf{C} \cap \mathcal{H}_{\leq t}$. For any $t \in \mathbb{Z}$, set

$$X_t := \max \{h \in \mathbb{R} : tL(p) \cdot \vec{w} + h \cdot \vec{w}^\perp \in \mathbf{C}_{\leq t}\}, \quad (14)$$

where $\vec{w}^\perp \in \mathbb{S}^1$ is a unit vector orthogonal to $\vec{w}$ (the direction of $\vec{w}^\perp$ is irrelevant for now; think of it as pointing upwards). If the set above is empty, which is to say that $\mathbf{C}$ does not intersect $\mathcal{H}_{\geq t}$, set $X_t = \dagger$. Thus X_t is the “highest” coordinate of the intersection of $\mathbf{C}_{\leq t}$ with $\partial\mathcal{H}_{\leq t}$.

The main objective of this section is the following result.

Theorem 3.2. *Fix $q \in [1, 4)$, $p < p_c$ and $\vec{w} \in \mathbb{S}^1$; sample $\mathbf{C}$ according to ϕ_p and define $(X_t)_{t \geq 0}$ as above. There exists an enlarged probability space supporting a random process Y_t such that $(X_t/L(p), Y_t)$ is an aperiodic killed renewal Markov process with a mass-gap, killing rate and X -step variance, all bounded away from 0 uniformly in p and $\vec{w}$. Moreover, the killing rate is also bounded uniformly away from 1.*

Finally, the initial step survival rate satisfies

$$\phi_p[X_1 \neq \dagger] \asymp \phi_p[T_1 < \infty] \asymp \pi_1(L(p)), \quad (15)$$

uniformly in p and $\vec{w}$.

Remark 3.3. By the construction below, we will ensure that, for all renewal times T_k , the cluster $\mathsf{C}_{\geq T_k} := \mathsf{C} \setminus \mathsf{C}_{\leq T_k}$ after T_k is contained in a cone of fixed aperture and apex $(L(T_k - 1), X_{T_k})$. Since the inter-renewal times have exponential tails, this cone-containment property ensures that the full cluster remains close to the linear interpolation of the points $(Lt, X_t)_{t \geq 0}$.

Theorem 3.2 suffices to prove Theorem 1.1, as well as a large number of other properties of subcritical clusters. We chose to formulate it using the concept of KMRP so as to separate the model-dependent part of the argument from the generic analysis of a class of processes with random-walk behaviour. The latter is contained in Section 4. While the formalism is new in this context, we do not claim the theorem to be entirely original. Indeed, the only formal novelty compared to [CIV08] is the uniformity in $p < p_c$.

We mentioned uniform lower bounds on σ_X , but no upper bounds, as the uniform exponential tails induce uniform upper bounds on σ_X, μ_X . Note that we do not claim that $\mu_X = 0$. This is the case when $\vec{w}$ is aligned to the coordinate axis due to symmetry, but is not generally true. As such, under the conditioning on $\{X_n \neq \dagger\}$, the cluster does not “aim” for the point $n\vec{w}$, but rather for a point $n\vec{v}$ for some $\vec{v}$ depending on $\vec{w}$ (with $\langle \vec{v}, \vec{w} \rangle = 1$). The relation between $\vec{w}$ and $\vec{v}$ will yield the strict convexity of the Wulff shape and ultimately will prove how $\vec{w}$ needs to be chosen to deduce Theorem 1.1 for some direction $\vec{v}$ – see Section 4.4 for details.

The proof of Theorem 3.2 relies on a geometric analysis of C under the survival event $\{X_N \neq \dagger\}$. The analysis is performed at a scale $L(p)$, in particular showing that there exists a density of “times” at which the cluster is confined to boxes of size $L(p)$. Those boxes play the role of “pre-renewal times”; indeed, it will be shown that at each pre-renewal time, there is a uniformly positive probability that the “future” cluster is sampled independently from its past, thus generating an actual renewal. The pre-renewal times play a similar role to the cone points in [CIV08].

Remark 3.4. Any KMRP with exponential tails and a mass-gap has a pure exponential rate of survival

$$\mathbb{P}[X_n \neq \dagger] \asymp \exp(-n/\zeta),$$

as will be proved in Section 4. In our context, $\zeta = \zeta(p, \vec{w})$ depends on p and $\vec{w}$ but is uniformly bounded away from 0 and ∞ . The constants in $\asymp$ above are not uniform in p ; they will be shown to be of order $\pi_1(L(p))$ and are due to the requirement of survival up to the first renewal time. Finally, the mass-gap states that the exponential rate of survival with no renewals is strictly smaller than ζ .

Remark 3.5. Theorem 3.2 provides a rigorous construction of “the infinite cluster conditioned to survive in the direction $\vec{w}$ ”. Indeed, it is a general property of KMRPs with a mass-gap that the distribution of $T_{k+1} - T_k$ under the measure $\mathbb{P}[\cdot | \mathcal{F}_{T_k}, T_k < \infty, X_n \neq \dagger]$ converges as $n \rightarrow \infty$ to a measure fully supported on $\mathbb{N}^*$, with exponential tail. Remark 4.3 allows to identify this distribution as an exponential tilt of the distribution of $T_{k+1} - T_k$ conditioned on $T_1 < \infty$.

The same type of convergence applies to the law of $(X_{T_k+1} - X_{T_k}, \dots, X_{T_{k+1}} - X_{T_k})$, with the limit being supported on finite sequences. If we call $\mathcal{L}_{\text{irred}}$ this limiting law, we may construct an infinite sequence (X_t) by concatenating i.i.d. samples from $\mathcal{L}_{\text{irred}}$. An infinite cluster may also be constructed in a similar way, by sampling i.i.d. pieces of cluster D_k — see the proof of Theorem 3.2 for the exact definition of D_k — according to a similar limiting law. Its statistics (for instance the mean value or variance of the T -step) are different from the ones of $\mathcal{L}$ conditioned to survive for only one step, as $\mathcal{L}_{\text{irred}}$ is an exponential tilt of $\mathcal{L}$.

Remark 3.6. For KMRP with exponential tails but no mass-gap a *condensation* phenomenon can occur, in which the process survives by making a very large step of linear order rather than many small steps of constant order. This was discussed in the context of long range Ising models in [AIOV21, AOV23].

The rest of the section is dedicated to proving Theorem 3.2. For the rest of the section, $q \in [1, 4)$ and $p < p_c$ are fixed and we omit them from notation. A direction $\vec{w} \in \mathbb{S}^1$ is also fixed and omitted from notation. For simplicity, we will assume $\vec{w}$ to be the horizontal unit vector $(1, 0)$. Unless otherwise stated, constants and equivalences below will be uniform in p and $\vec{w}$; they will be referred to as *universal*.

3.1 Cone-connections

For $\alpha > 0$, define the cone in direction $\vec{w}$ with aperture $2 \arctan \alpha$ as

$$\mathcal{Y}_\alpha = \{z \in \mathbb{R}^2 : |\langle z, \vec{w}^\perp \rangle| \leq \alpha \langle z, \vec{w} \rangle\}.$$

Also set $\partial \mathcal{Y}_\alpha$ to be the set of vertices in $\mathcal{Y}_\alpha$ that have at least one neighbour outside $\mathcal{Y}_\alpha$.

Recall the notion of potential past A and the event $\text{Past}(A)$ that requires that the edges of A are all open, while those of $\partial_{\leq 0} A$ are closed.

Proposition 3.7 (Connections in cones). *There exist constants $\alpha, c > 0$ such that for any $k, n \geq 1$, and any potential past A ,*

$$\phi[\Lambda_L \leftrightarrow (\mathcal{Y}_\alpha - (kL, 0))^c \mid \Lambda_L \leftrightarrow \mathcal{H}_{\geq n} \text{ and } \text{Past}(A)] \leq e^{-ck}. \quad (16)$$

The above states that, when long connections in the direction $\vec{w}$ occur, the whole connecting cluster is contained in a cone with high probability. Note that the aperture of the cone is fixed, and is not claimed to be arbitrarily small. Indeed, Proposition 3.7 may ultimately be shown to hold for any $\alpha > 0$, but only after proving the strict convexity of the Wulff shape. For now, it suffices to consider a fixed, large α .

Before proving Proposition 3.7, we formulate a useful mixing property. The proof is surprisingly intricate, but follows the argument of [DCM22, Prop. 2.9].

Lemma 3.8. *For any $\alpha > 0$ and events A and B depending on the edges of $\mathcal{H}_{\leq 0}$ and $\mathcal{Y}_\alpha \cap \mathcal{H}_{\geq 1}$, respectively,*

$$\phi[A \cap B] \asymp \phi[A]\phi[B], \quad (17)$$

where the constants in $\asymp$ depend on α but not on p , A or B .

Moreover, there exists $c > 0$ such that, for any A as above and B depending only on edges in $\mathcal{Y}_\alpha + (kL, 0)$ for some $k \geq 1$

$$\left| \frac{\phi[A \cap B]}{\phi[A]\phi[B]} - 1 \right| \leq e^{-ck}. \quad (18)$$

Note that the events A and B are not assumed to be increasing or decreasing. This is why the proof below is more complicated than one might expect.

Proof. Fix $\alpha > 0$. We will focus on proving (17); the proof of (18) is explained at the end.

Write $\mathcal{D} = \mathcal{Y}_{2\alpha} \cap \mathcal{H}_{\geq 1/2}$ and $\mathcal{E} = \mathcal{Y}_\alpha \cap \mathcal{H}_{\geq 1}$. We first prove that there exists a constant $\eta < 1$ depending on α but not on p , such that

$$\phi_{\mathcal{D} \setminus \mathcal{E}}^1[\partial\mathcal{D} \leftrightarrow \mathcal{E}] \leq \eta \quad \text{and} \quad \phi_{\mathcal{H}_{\geq 0} \setminus \mathcal{D}}^1[\partial\mathcal{H}_{\geq 0} \leftrightarrow \mathcal{D}] \leq \eta. \quad (19)$$

We focus on the second inequality and we write S for the set of vertices in $\mathcal{H}_{\geq 0} \setminus \mathcal{D}$ that are at equal L^∞ -distance from $\partial\mathcal{H}_{\geq 0}$ and $\mathcal{D}$ (up to an error of 1). Batch the points in S according to their distance to $\mathcal{D}$: a batch S_k for $k \geq 0$ is the set of points $x \in S$ such that

$$k \leq \text{dist}(x, \mathcal{D}) < (k+1)L.$$

For the second event in (19) to occur, at least one of the batches S_k needs to be connected to $\mathcal{D}$ — see Figure 2. For k large enough, we have

$$\phi_{\mathcal{H}_{\geq 0} \setminus \mathcal{D}}^1[S_k \leftrightarrow \mathcal{D}] \leq 2\phi_{\Lambda_{c_1 k L}}^1[\Lambda_{L/c_1} \leftrightarrow \partial\Lambda_{c_1 k L}] \leq e^{-c_0 k},$$

for some universal constant $c_1 > 0$. The final inequality is due to (3).

Summing over $k \geq K_0$ we find

$$\phi_{\mathcal{H}_{\geq 0} \setminus \mathcal{D}}^1[\exists k \geq K_0 : S_k \leftrightarrow \mathcal{D}] \leq C_0 \sum_{k \geq K_0} e^{-c_0 k} \leq \frac{1}{2},$$

provided that K_0 is a large constant, which may depend on α , but not on L . We may now use (RSW) to deduce the second inequality in (19). The first inequality is proved in the same way.

With (19) in hand, the rest of the proof follows that of [DCM22, Prop. 2.9]. Fix an event B depending on the edges in $\mathcal{E}$ and ζ a boundary condition on $\partial\mathcal{H}_{\geq 0}$. We will show that

$$(1 - \eta)^2 \phi_{\mathcal{H}_{\geq 0}}^0[B] \leq \phi_{\mathcal{H}_{\geq 0}}^\xi[B] \leq (1 - \eta)^{-1} \phi_{\mathcal{H}_{\geq 0}}^0[B]. \quad (20)$$

We start off by proving that, for any set of edges Ω with $\mathcal{D} \subset \Omega \subset \mathcal{H}_{\geq 0}$ and any boundary conditions ξ on $\partial\Omega$,

$$\phi_\Omega^0[B] \geq (1 - \eta) \phi_\Omega^\xi[B]. \quad (21)$$




Figure 2: *Left*: the dotted line represents the points at equal distance from $\mathcal{D}$ and $\partial\mathcal{H}_{\geq 0}$; the solid black lines represent wired boundary conditions. The blue segment is the top half of a batch S_k . For it to be connected to $\mathcal{D}$, the annulus surrounding it should be crossed, which occurs with an exponentially small probability in k . For small values of k , use (RSW) to complete the dual path separating $\partial\mathcal{H}_{\geq 0}$ from $\mathcal{D}$.

Right: For the second part of the argument, consider a domain Ω containing $\mathcal{D}$ and prove that the probability of B with arbitrary boundary conditions on Ω is not much larger than that with free boundary conditions — we use here the first estimate of (19). One bound of (20) follows directly. The opposite bound is proved by considering Ω to be the complement of the cluster of $\partial\mathcal{H}_{\geq 0}$ and using the second bound of (19) to show that $\mathcal{D} \subset \Omega$ with positive probability.

See Figure 2, right-diagram, for an illustration. Indeed, consider the increasing coupling $\mathbf{P}$ between $\omega \sim \phi_{\Omega}^0$ and $\omega' \sim \phi_{\Omega}^{\xi}$ obtained by revealing first the cluster of $\partial\Omega$ in ω' , then the rest of the configurations. For details about this type of coupling see [DCM22]; the important feature to keep in mind here is that $\omega = \omega'$ for all edges not connected to $\partial\Omega$ in ω' . Then

$$\begin{aligned}
 \phi_{\Omega}^{\xi}[B] - \phi_{\Omega}^0[B] &\leq \mathbf{P}[\omega \neq \omega' \text{ on } \mathcal{E} \text{ and } \omega' \in B] \\
 &\leq \mathbf{P}[\partial\Omega \xleftrightarrow{\omega'} \mathcal{E} \text{ and } \omega' \in B] \\
 &\leq \phi_{\mathcal{D} \setminus \mathcal{E}}^1[\partial\mathcal{D} \leftrightarrow \mathcal{E}] \cdot \phi_{\Omega}^{\xi}[B] \\
 &\leq \eta \cdot \phi_{\Omega}^{\xi}[B],
 \end{aligned} \tag{22}$$

where the third inequality is obtained by observing that

$$\phi_{\Omega}^{\xi}[\Omega \leftrightarrow \mathcal{E} \mid B] \leq \phi_{\mathcal{D} \setminus \mathcal{E}}^1[\partial\mathcal{D} \leftrightarrow \mathcal{E}],$$

and the last is due to the first bound in (19). The above immediately implies (21). We insist here that B is not assumed increasing and that the argument works because in the second line of (22), the event depends only on ω' . This is why the same reasoning may not be used to prove an opposite bound.

Applying (21) to $\Omega = \mathcal{H}_{\geq 0}$ we obtain the upper bound in (20). Another consequence of the above is that for all sets Ω with $\mathcal{D} \subset \Omega \subset \mathcal{H}_{\geq 0}$,

$$\phi_{\Omega}^0[B] \geq (1 - \eta) \phi_{\mathcal{H}_{\geq 0}}^0[B], \quad (23)$$

since the probability on the right-hand side above is an average of quantities of the form $\phi_{\Omega}^{\xi}[B]$.

We turn to the second bound in (21). Sample ω from $\phi_{\mathcal{H}_{\geq 0}}^{\zeta}$ and let $\mathbf{C}$ be the cluster of $\partial\mathcal{H}_{\geq 0}$. Then

$$\phi_{\mathcal{H}_{\geq 0}}^{\zeta}[B] \geq \sum_{\Omega} \phi_{\Omega}^0[B] \phi_{\mathcal{H}_{\geq 0}}^{\zeta}[\mathcal{H}_{\geq 0} \setminus \mathbf{C} = \Omega],$$

where the sum is over all sets of edges Ω containing $\mathcal{D}$. Using (23), we conclude that

$$\phi_{\mathcal{H}_{\geq 0}}^{\zeta}[B] \geq (1 - \eta) \phi_{\mathcal{H}_{\geq 0}}^0[B] \phi_{\mathcal{H}_{\geq 0}}^{\zeta}[\mathcal{H}_{\geq 0} \not\leftrightarrow \mathcal{D}] \geq (1 - \eta)^2 \phi_{\mathcal{H}_{\geq 0}}^0[B],$$

with the second inequality following from (19). This proves the second inequality in (20).

Finally, note that (20) implies that, for any event A depending only on edges in $\mathcal{H}_{\leq 0}$,

$$(1 - \eta)^4 \leq \frac{\phi_p[B | A]}{\phi_p[B]} \leq (1 - \eta)^{-2},$$

which yields the desired result.

To prove (18) follow the same lines, but observe that $\eta \leq e^{-ck}$ for some $c > 0$, which implies the result. $\square$

We are finally ready for the proof of Proposition 3.7.

Proof of Proposition 3.7. The proof is complicated by the fact that we consider a general potential past A and work in A^c . We start with the simpler⁵ case of $A = \mathcal{H}_{\leq 0}$, which is to say that we work under $\phi_{\mathcal{H}_{\geq 0}}^1$. We will then explain how to deduce the general case.

We start by explaining how to choose α . First notice that due to direct (RSW) constructions, we have that for any $\ell \geq 0$

$$\phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \Lambda_L(\ell L, 0)] \geq \exp(-C_0 \ell), \quad (24)$$

for some universal constant $C_0 > 0$. Let $c > 0$ be the constant given by (3) and set $\alpha = 2C_0/c$.

Fix k and n . All constants below are allowed to depend on α , but not on p , k or n .

Case 1: measure in the half-space. Consider the cones

$$\mathcal{Y}_{\text{out}} = \mathcal{Y}_{\alpha} - (kL, 0) \quad \text{and} \quad \mathcal{Y}_{\text{in}} = \mathcal{Y}_{\alpha/2} - (kL/2, 0);$$

see Figure 3, left diagram. To bound the probability in (16), we will distinguish two scenarios:

⁵The important feature here is that $\phi_{\mathcal{H}_{\geq 0}}^1$ is invariant under vertical translations; the same proof also applies to $\phi_{\mathcal{H}_{\geq 0}}^0$ and the infinite volume measure ϕ . When $\partial\mathcal{H}_{\geq 0}$ is not aligned to the axis of the lattice, the measure is not exactly translationally invariant, but dominates any vertical translate of $\phi_{\mathcal{H}_{\geq -1}}^1$, which suffices.

- (1.a) either Λ_L is connected to $\mathcal{H}_{\geq n}$ inside $\mathcal{Y}_{\text{in}}$, but is also connected to $\mathcal{Y}_{\text{out}}^c$,
(1.b) or Λ_L is connected to $\mathcal{H}_{\geq n}$, but not inside in $\mathcal{Y}_{\text{in}}$.

Write S_a and S_b for the two events above. We will bound their probabilities under $\phi_{\mathcal{H}_{\geq 0}}^1[\cdot | \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}]$ separately.

(1.a) To bound the probability of S_a we condition on the configuration inside of $\mathcal{Y}_{\text{in}}$. If S_a occurs, then there exists a point $x \in \partial\mathcal{Y}_{\text{out}}$ that is connected to $\mathcal{Y}_{\text{in}}$. Applying the same argument of summation by batches as used for (19) we conclude that

$$\phi_{\mathcal{H}_{\geq 0}}^1[S_a \mid \omega = \omega_0 \text{ on } \mathcal{Y}_{\text{in}}] \leq Ce^{-ck} \quad (25)$$

for universal constants $c, C > 0$ and any ω_0 on $\mathcal{Y}_{\text{in}}$ containing a connection between Λ_L and $\mathcal{H}_{\geq n}$. Now,

$$\phi_{\mathcal{H}_{\geq 0}}^1[S_a \mid \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] = \sum_{\omega_0} \phi_{\mathcal{H}_{\geq 0}}^1[S_a \mid \omega = \omega_0 \text{ on } \mathcal{Y}_{\text{in}}] \phi_{\mathcal{H}_{\geq 0}}^1[\omega = \omega_0 \text{ on } \mathcal{Y}_{\text{in}} \mid \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}],$$

where the sum is over all ω_0 containing a connection between Λ_L and $\mathcal{H}_{\geq n}$. Injecting (25) into the above we conclude that

$$\phi_{\mathcal{H}_{\geq 0}}^1[S_a \mid \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] \leq Ce^{-ck}.$$

(1.b) To bound the probability of S_b under the conditional measure, we will directly compare $\phi_{\mathcal{H}_{\geq 0}}^1[S_b]$ to $\phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}]$.

For $x \in L \cdot (\mathbb{N} \times \mathbb{Z})$, write $\mathcal{Y}(x)$ for the horizontal translate of $\mathcal{Y}_{\alpha/2}$ whose boundary passes through x . Let $\mathcal{B}_x$ be the event that $\Lambda_L(x)$ is connected to both Λ_L and $\mathcal{H}_{\geq n}$ inside $\mathcal{Y}(x)$, but that these two connections are produced by disjoint clusters of $\omega \cap \mathcal{Y}(x) \cap \Lambda_L(x)^c$. By considering the right-most translate of $\mathcal{Y}_{\alpha/2}$ that contains a connection between Λ_L and $\mathcal{H}_{\geq n}$, we conclude that if S_b occurs, then there exists at least one $x \in L \cdot (\mathbb{N} \times \mathbb{Z}) \setminus \mathcal{Y}_{\text{in}}$ for which $\mathcal{B}_x$ occurs. See Figure 3, middle diagram.

We will bound the probability of each event $\mathcal{B}_x$ with x as above. Fix one such $x = (x_1, x_2)$ and assume $x_2 \geq 0$; the case $x_2 \leq 0$ is identical. Note that $\mathcal{B}_x$ is measurable in terms of the configuration inside $\mathcal{Y}(x)$. By the same argument as in the proof of Lemma 3.8, conditionally on $\mathcal{B}_x$ (or more generally on any event in $\mathcal{Y}(x)$), $\Lambda_L(x)$ is connected to infinity by a dual path in $\mathcal{Y}(x)^c$ with positive probability. By performing an additional surgery inside $\Lambda_L(x)$ using (RSW) and doing a similar argument in the lower half-plane (after exploring all the clusters of $\Lambda_L(x)$ in $\mathcal{Y}(x)$), we conclude that

$$\begin{aligned} \phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{B}_x] &\leq \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L(x) \text{ connected in } \mathcal{Y}(x) \text{ to } \Lambda_L \text{ and to } \mathcal{H}_{\geq n} \text{ by disjoint clusters}] \\ &\leq \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L(x) \xleftrightarrow{\mathcal{Y}(x)} \Lambda_L] \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L(x) \xleftrightarrow{\mathcal{Y}(x)} \mathcal{H}_{\geq n}]. \end{aligned} \quad (26)$$

The second inequality is due to (MON) and requires conditioning on the cluster producing one of the connections and bounding the probability of the other connection.




Figure 3: *Left:* The occurrence of S_a conditionally on $\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}$ induces an “unforced” connection between $\mathcal{Y}_{\text{in}}$ and $\mathcal{Y}_{\text{out}}^c$, which has an exponential cost in k . *Middle:* To bound S_b , consider the right-most translate of $\mathcal{Y}_{\text{in}}$ that contains a connection between Λ_L and $\mathcal{H}_{\geq n}$. Its boundary intersects a “pivotal” box $\Lambda_L(x)$, that is a box that is connected to Λ_L and $\mathcal{H}_{\geq n}$ inside $\mathcal{Y}(x)$ by distinct clusters. *Right:* When working with an arbitrary potential past A , the connection between Λ_L and $\mathcal{H}_{\geq n}$ may use the left half-plane. Nevertheless, there exists a vertical translate of Λ_L that is connected to $\mathcal{H}_{\geq n}$ inside $\mathcal{H}_{\geq 0}$. We will consider the closest such translate to 0.

We now bound each term in the right-hand side of (26). First, by vertical translation,

$$\phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L(x) \xleftrightarrow{\mathcal{Y}(x)} \mathcal{H}_{\geq n}] \leq \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L(x_1, 0) \leftrightarrow \mathcal{H}_{\geq n}].$$

Moreover,

$$\begin{aligned} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \xleftrightarrow{\mathcal{Y}(x)} \Lambda_L(x)] &\leq \phi_{\mathcal{H}_{\geq 0}}^0[\Lambda_L \xleftrightarrow{\mathcal{Y}(x)} \Lambda_L(x)] \\ &\leq \exp(-c(x_1 + x_2)/L) \\ &\leq \exp(-C_0 x_1/L - \frac{c}{2}(x_1 + x_2)/L) \\ &\leq \phi_{\mathcal{H}_{\geq 0}}^1[0 \leftrightarrow \Lambda_L(x_1, 0)] \exp(-\frac{c}{2}\|x\|/L), \end{aligned}$$

where c is a universal constant and C_0 is given by (24). The first inequality comes from the mixing property (17), the second one from (2), and the third one is due to the choice of α and the fact that $x_2 \geq \alpha x_1$.

Inserting the last two displays into (26) we conclude that

$$\begin{aligned} \phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{B}_x] &\leq \exp(-\frac{c}{2}\|x\|/L) \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \Lambda_L(x_1, 0)] \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L(x_1, 0) \leftrightarrow \mathcal{H}_{\geq n}] \\ &\leq \exp(-\frac{c}{2}\|x\|/L) \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}], \end{aligned}$$

The second inequality is due to **(FKG)** and an application of **(RSW)** to glue the two connections. Summing the above over all $x \in L \cdot (\mathbb{N} \times \mathbb{Z}) \setminus \mathcal{Y}_{\text{in}}$, we find that

$$\phi_{\mathcal{H}_{\geq 0}}^1[S_b] \leq e^{-ck} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}],$$

for a potentially modified value of c .

Combining points (a) and (b) we conclude that

$$\phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow (\mathcal{Y}_\alpha - (kL, 0))^c \mid \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] \leq Ce^{-ck},$$

for universal constants $c, C > 0$. The constant C may be removed by modifying c .

A particular consequence of the above and Lemma 3.8 is that

$$\begin{aligned} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] &\leq \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \xleftrightarrow{\mathcal{Y}_\alpha} \mathcal{H}_{\geq n}] \\ &\leq \phi_{\mathcal{H}_{\geq 0}}^0[\Lambda_L \xleftrightarrow{\mathcal{Y}_\alpha} \mathcal{H}_{\geq n}] \leq \phi[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n} \mid \text{Past}(A)], \end{aligned} \quad (27)$$

for any potential past A . This will be useful for Case 2, below.

Case 2: general potential pasts A . We are now ready to prove the general statement. Fix a potential past A . Write $D = (A \cup \partial_{\leq 0} A)^c$ and ξ for the boundary conditions induced on D by the conditioning $\text{Past}(A)$, so that $\phi[\cdot \mid \text{Past}(A)] = \phi_D^\xi$.

For a point $x \in \{0\} \times L\mathbb{Z}$, write $\tilde{\mathcal{C}}_x$ for the cluster of $\omega \cap \mathcal{H}_{\geq 0}$ that connects $\Lambda_L(x)$ to $\mathcal{H}_{\geq n}$ — if several such clusters exists, let $\tilde{\mathcal{C}}_x$ denote their union and if no such cluster exists set $\tilde{\mathcal{C}}_x = \emptyset$.

Let ω be a configuration in which Λ_L is connected to $\mathcal{H}_{\geq n}$. Write $X(\omega)$ be the point $x \in \{0\} \times L\mathbb{Z}$ of minimal norm for which $\tilde{\mathcal{C}}_x \neq \emptyset$ — when two such points exist, choose one arbitrarily.

The fact that ω contains a connection between Λ_L and $\mathcal{H}_{\geq n}$ guarantees that $X(\omega)$ is well defined and that Λ_L is connected to $\Lambda_{\|X(\omega)\| - L}^c$ in $\tilde{\mathcal{C}}_{X(\omega)}^c$; see Figure 3, right diagram. It may be noted that we do not claim $\tilde{\mathcal{C}}_{X(\omega)}$ to be connected to Λ_L ; it may be that Λ_L is connected to $\mathcal{H}_{\geq n}$ via a different cluster.

As a consequence, if x is a possible realisation of $X(\omega)$ and $\mathbf{C}$ is a possible realisation of $\tilde{\mathcal{C}}_{X(\omega)}$, we have that

$$\begin{aligned} \phi_D^\xi[X(\omega) = x, \tilde{\mathcal{C}}_x = \mathbf{C} \text{ and } \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] &\leq \phi_D^\xi[\tilde{\mathcal{C}}_x = \mathbf{C}] \phi_D^\xi[\Lambda_L \xleftrightarrow{\mathbf{C}^c} \partial\Lambda_{\|x\| - L} \mid \tilde{\mathcal{C}}_x = \mathbf{C}] \\ &\leq \phi_D^\xi[\tilde{\mathcal{C}}_x = \mathbf{C}] \phi_{D \cap \Lambda_{\|x\| - L}}^\zeta[\Lambda_L \leftrightarrow \partial\Lambda_{\|x\| - L}] \\ &\leq \phi_D^\xi[\tilde{\mathcal{C}}_x = \mathbf{C}] e^{-c\|x\|}. \end{aligned} \quad (28)$$

Above, ζ are the boundary conditions on $D \cap \Lambda_{\|x\| - L}$ that are wired on $\partial\Lambda_{\|x\| - L}$ and identical to ξ elsewhere. Indeed, the conditioning on $\tilde{\mathcal{C}}_x = \mathbf{C}$ produces a mixture of free and wired boundary conditions on $\mathbf{C}^c$, together with the additional conditioning that no other connections between boxes $\Lambda_L(x')$ and $\mathcal{H}_{\geq n}$ occur in $\mathcal{H}_{\geq 0}$ for $x' \in \{0\} \times L\mathbb{Z}$ with $\|x'\| \leq \|x\|$. The latter is a decreasing conditioning and the boundary conditions induced

by the conditioning on $\mathbf{C}$ only have wired parts outside of $\Lambda_{\|x\|-L}$ — this is due to the choice of $X(\omega)$. The last inequality is due to (3). The bound (28) will be used below.

Fix some small constant $c_0 > 0$ that depends only on α ; we will see below how to choose it. If ω is a configuration containing connections between Λ_L and both $\mathcal{H}_{\geq n}$ and the outside of $\mathcal{Y}_{\text{out}} = \mathcal{Y}_\alpha - (kL, 0)$, then at least one of the cases below occurs.

$$(2.a) \quad \|X(\omega)\| \geq \alpha c_0 kL;$$

$$(2.b) \quad \|X(\omega)\| \leq \alpha c_0 kL \text{ but } \tilde{\mathbf{C}}_{X(\omega)} \text{ is not contained in } \mathcal{Y}_\alpha - (2c_0 kL, 0);$$

$$(2.c) \quad \tilde{\mathbf{C}}_{X(\omega)} \text{ is contained in } \mathcal{Y}_\alpha - (2c_0 kL, 0), \text{ but } \Lambda_L \text{ is connected to } \mathcal{Y}_{\text{out}}^c.$$

We will bound the probability of each of the events above separately.

To bound the probability of (2.a), observe that, if x is a potential realisation of $X(\omega)$, summing (28) over $\mathbf{C}$ implies that

$$\begin{aligned} \phi_D^\xi[X(\omega) = x \text{ and } \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] &\leq \phi_D^\xi[\Lambda_L(x) \xrightarrow{\mathcal{H}_{\geq 0}} \mathcal{H}_{\geq n}] e^{-c\|x\|} \\ &\leq \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L(x) \leftrightarrow \mathcal{H}_{\geq n}] e^{-c\|x\|}. \end{aligned}$$

Summing over x with $\|x\| \geq \alpha c_0 kL$ we conclude that

$$\phi_D^\xi[\|X(\omega)\| \geq \alpha c_0 kL \text{ and } \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] \leq e^{-c c_0 k} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}], \quad (29)$$

for some universal constant $c > 0$.

The probability of (2.b) may be bounded in a similar way. For x a potential realisation of $X(\omega)$, summing (28) over realisations of $\tilde{\mathbf{C}}_x$ as in (2.b), we find

$$\begin{aligned} \phi_D^\xi[X(\omega) = x, \tilde{\mathbf{C}}_x \text{ intersects } (\mathcal{Y}_\alpha - (2c_0 kL, 0))^c \text{ and } \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] \\ \leq \phi_D^\xi[\tilde{\mathbf{C}}_x \text{ intersects } (\mathcal{Y}_\alpha - (2c_0 kL, 0))^c] e^{-c\|x\|} \\ \leq \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n} \text{ and } \Lambda_L \leftrightarrow (\mathcal{Y}_\alpha - (c_0 kL, 0))^c] e^{-c\|x\|} \\ \leq e^{-c c_0 k + c\|x\|} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}], \end{aligned}$$

with the second inequality due to (MON) — note that the event considered is increasing — and the last one given by (16) applied to $\phi_{\mathcal{H}_{\geq 0}}^1$, which we already proved. Summing over the possible values of $X(\omega)$ we find that

$$\begin{aligned} \phi_D^\xi[\|X(\omega)\| \leq \alpha c_0 kL, \tilde{\mathbf{C}}_{X(\omega)} \not\subset (\mathcal{Y}_\alpha - (2c_0 kL, 0)) \text{ and } \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] \\ \leq e^{-c c_0 k} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}]. \end{aligned} \quad (30)$$

Finally, to bound the probability of (2.c), notice that for this event to occur, there need to exist a connection between $\Lambda_{2\alpha c_0 kL}$ and $\mathcal{Y}_{\text{out}}^c$, outside of $\tilde{\mathbf{C}}_{X(\omega)}$. Using the same

argument of conditioning on $\tilde{\mathbf{C}}_{X(\omega)}$ as in (28),

$$\begin{aligned}
& \phi_D^\xi[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}, \tilde{\mathbf{C}}_{X(\omega)} \subset \mathcal{Y}_\alpha - (2c_0kL, 0) \text{ and } \Lambda_L \leftrightarrow \mathcal{Y}_{\text{out}}^c] \\
& \leq \sum_x \phi_D^\xi[\Lambda_L(x) \xrightarrow{\mathcal{H}_{\geq 0}} \mathcal{H}_{\geq n}] \phi_{D \cap \Lambda_{2\alpha c_0kL}^c}^\zeta[\Lambda_{2\alpha c_0kL} \leftrightarrow \mathcal{Y}_{\text{out}}^c] \\
& \leq c_0k e^{C_{c_0k}} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] \phi_D^\xi[\Lambda_{2\alpha c_0kL} \leftrightarrow \mathcal{Y}_{\text{out}}^c] \\
& \leq e^{-c'k} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}].
\end{aligned} \tag{31}$$

where the sum is over all $x \in \{0\} \times L\mathbb{Z}$ with $\|x\| \leq 2\alpha c_0kL$ and ζ are the boundary conditions on $D \cap \Lambda_{2\alpha c_0kL}^c$ that are wired on $\partial\Lambda_{2\alpha c_0kL}$ and identical to ξ everywhere else. The second inequality uses (MON) and the finite energy property, with C some universal constant. The last inequality is due to (3) and is valid for c_0 below some universal constant, and may be ensured by the choice of c_0 ; $c' > 0$ is a universal constant.

Summing (29), (30) and (31), we obtain

$$\begin{aligned}
\phi_D^\xi[\Lambda_L \leftrightarrow \mathcal{Y}_{\text{out}}^c \text{ and } \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] & \leq e^{-c''k} \phi_{\mathcal{H}_{\geq 0}}^1[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] \\
& \leq e^{-c''k} \phi_D^\xi[\Lambda_L \leftrightarrow \mathcal{H}_{\geq n}].
\end{aligned}$$

for some universal constant $c'' > 0$. We used (27) in the last line. This proves (16). $\square$

Henceforth $\alpha \geq 4$ is fixed to satisfy Proposition 3.7 and we write $\mathcal{Y} := \mathcal{Y}_\alpha$.

3.2 The number of active segments is subcritical

For integers $t \geq 0, k \in \mathbb{Z}$, define $\mathcal{L}_{t,k} := \{tL\} \times [kL, (k+1)L)$, so that the hyperplane $\partial\mathcal{H}_{\leq t}$ is the disjoint union of the line segments $\mathcal{L}_{t,k}$ when k runs over $\mathbb{Z}$. Also, write $x_{t,k} = (tL, (k+1/2)L)$ for the midpoint of $\mathcal{L}_{t,k}$.

Recall that in Theorem 3.2 we consider a configuration sampled according to ϕ and $\mathbf{C}_{\leq t}$ is the cluster of 0 in $\mathcal{H}_{\leq t}$. A line segment $\mathcal{L}_{t,k}$ is said to be *active at time t* if $\mathbf{C}_{\leq t} \cap \mathcal{L}_{t,k} \neq \emptyset$. Write N_t for the number of active segments at time t . A time t for which $N_t = 1$ will be called a *pre-renewal time*. The goal of this section is the following.

Proposition 3.9 (Density of pre-renewal times). *There exists $c > 0$ such that for any $n, t, r \geq 1$,*

$$\begin{aligned}
\phi[N_s > 1, \forall s \in \{t+1, \dots, t+r\} \mid \mathbf{C}_{\leq t}, N_t = 1, X_n \neq \dagger] & \leq \exp(-cr) \quad \text{and} \\
\phi[N_s > 1, \forall s \in \{1, \dots, r\} \mid X_n \neq \dagger] & \leq \exp(-cr).
\end{aligned} \tag{32}$$

This tells us that pre-renewal times arrive quickly, even when conditioning on a survival event far in the future. Proposition 3.9 will follow from the result below.

Proposition 3.10. *There exist constants $C \geq 1, \mu < 1$ and $K > 0$ such that for any $t \geq 0$ and $n \geq t + C$,*

$$\phi[N_{t+C} \mid \mathbf{C}_{\leq t}, X_n \neq \dagger] \leq \mu N_t + K. \tag{33}$$

The rest of this section is dedicated to the proofs of the two results above. In spirit, the results of that section are very close to limit theorems for the distribution of the number of children of a subcritical branching process conditioned to survive for a long time (see [AN04]). Our proof is a bit more complex due to the lack of control on the offspring distribution and most importantly due to the positive correlations between the different activated segments. As mentioned, Proposition 3.10 is the key step, and we start with it. First we explain how to choose the constant C .

Lemma 3.11. *For any $\mu > 0$, there exists some constant $C > 0$ such that, for any potential past A*

$$\phi[\#j \text{ such that } \mathcal{L}_{0,0} \leftrightarrow \mathcal{L}_{C,j} \mid \text{Past}(A)] \leq \mu. \quad (34)$$

Proof of Lemma 3.11. Fix $\mu > 0$, A a potential past and let C be some positive integer. Then, for $j \in \mathbb{Z}$, for $\mathcal{L}_{0,0}$ to be connected to $\mathcal{L}_{C,j}$, Λ_L needs to be connected to $\partial\Lambda_{L \cdot \max\{C, |j| - 1\}}$. Thus, by (3),

$$\phi[\mathcal{L}_{0,0} \leftrightarrow \mathcal{L}_{C,j} \mid \text{Past}(A)] \leq \exp(-c \max\{C, |j| - 1\}).$$

Summing the above and taking C large enough implies (34). $\square$

Proof of Proposition 3.10. Fix the parameters n and t as in the statement. Also fix some realisation of $C_{\leq t}$. We will always work conditionally on this realisation of the “past cluster”, and write $\phi_{C_{\leq t}}$ for this conditional measure. It is the translate of a measure of the type $\phi[\cdot \mid \text{Past}(A)]$, and Lemma 3.11 applies to it. All notions of connections and clusters below refer to the configuration in $C_{\leq t}^c$ only. All constants and equivalences below are uniform in the choices of $p, \vec{w}, n, t$ and $C_{\leq t}$.

Fix C given by Lemma 3.11 with the choice of an arbitrary $\mu < 1/2$. Write $\mathbf{j}$ for the index of the top-most active box $\mathcal{L}_{t,j}$ such that $\mathcal{L}_{t,j} \cap C_{\leq t}$ is connected to $\mathcal{H}_{\geq n}$. If no such connection exists, write $\mathbf{j} = \emptyset$. Also write K for the minimal value $K \geq 0$ such that the cluster of $\mathcal{L}_{t,\mathbf{j}}$ is contained in $\mathcal{Y} + x_{t,\mathbf{j}} - (KL, 0)$ (recall that $x_{t,j}$ denotes the mid-point of $\mathcal{L}_{t,j}$).

Our goal is to bound

$$\phi_{C_{\leq t}}[N_{t+C} \mid \mathbf{j} \neq \emptyset] = \sum_j \phi_{C_{\leq t}}[\mathbf{j} = j \mid \mathbf{j} \neq \emptyset] \phi_{C_{\leq t}}[N_{t+C} \mid \mathbf{j} = j],$$

and we will do so by bounding each of the terms $\phi_{C_{\leq t}}[N_{t+C} \mid \mathbf{j} = j]$ individually.

We first argue that

$$\phi_{C_{\leq t}}[K > k \mid \mathbf{j} = j] \leq e^{-c_0 k} \quad (35)$$

for some constant $c_0 > 0$ and all $k \geq 3$.

We say $\mathcal{L}_{t,j}$ is a top-most seed if $\mathcal{L}_{t,j} \cap C_{\leq t}$ is connected to $\partial\Lambda_{2L}(x_{t,j})$ but not to $\mathcal{L}_{t,j+1}$ in $C_{\leq t}^c \cap \Lambda_{2L}(x_{t,j})$. We also write $s_t = \phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{L}_{0,0} \leftrightarrow \mathcal{H}_{\geq t}]$. We start off by estimating the probability of $\mathbf{j} = j$.

Lemma 3.12 (Gluing lemma). *Uniformly in n, t and r and past clusters $\mathcal{C}_{\leq t}$ satisfying $\phi[\mathcal{C}_{\leq t} \leftrightarrow \mathcal{H}_n] < \frac{1}{2}$,*

$$\phi_{\mathcal{C}_{\leq t}}[\mathbf{j} = j] \asymp s_{n-t} \phi_{\mathcal{C}_{\leq t}}[j \text{ is a top-most seed}]. \quad (36)$$

We defer the proof of the lemma to later in the section and finish that of Proposition 3.10. Due to the condition on $\mathcal{C}_{\leq t}$ appearing in Lemma 3.12, we will distinguish two cases depending on $\mathcal{C}_{\leq t}$.

Case $\mathcal{C}_{\leq t}$ is such that $\phi_{\mathcal{C}_{\leq t}}[\mathcal{C}_{\leq t} \leftrightarrow \mathcal{H}_n] < \frac{1}{2}$. Due to (36) we have

$$\begin{aligned} \phi_{\mathcal{C}_{\leq t}}[\mathbf{K} > k \mid \mathbf{j} = j] &\leq \frac{\phi_{\mathcal{C}_{\leq t}}[\mathbf{K} \geq k \text{ and } \mathbf{j} = j]}{s_{n-t} \phi_{\mathcal{C}_{\leq t}}[j \text{ is top-most seed}]} \\ &\leq \frac{\phi_{\mathcal{C}_{\leq t}}[j \text{ is top-most seed, } \Lambda_{2L}(x_{t,j}) \leftrightarrow \mathcal{H}_{\geq n} \text{ and } \Lambda_{2L}(x_{t,j}) \leftrightarrow (\mathcal{Y} + x_{t,j} - (kL, 0))^c]}{s_{n-t} \phi_{\mathcal{C}_{\leq t}}[j \text{ is top-most seed}]} \\ &\leq \frac{\phi_{\mathcal{C}_{\leq t}}[\Lambda_{2L}(x_{t,j}) \leftrightarrow \mathcal{H}_{\geq n} \text{ and } \Lambda_{2L}(x_{t,j}) \leftrightarrow (\mathcal{Y} + x_{t,j} - (kL, 0))^c \mid \omega \equiv 1 \text{ on } \Lambda_{2L}(x_{t,j})]}{s_{n-t}} \\ &\leq e^{-c_0 k}, \end{aligned}$$

The second inequality is simply an inclusion of events: for $\mathbf{j} = j$ to occur, j needs to be a top-most seed and $\Lambda_{2L}(x_{t,j})$ should be connected to $\mathcal{H}_{\geq n}$. The third is a consequence of (FKG) and the last one is due⁶ to Proposition 3.7. This proves (35).

Now, for j fixed, by the almost sure finiteness of the cluster of 0,

$$\phi_{\mathcal{C}_{\leq t}}[N_{t+C} \mid \mathbf{j} = j] = \sum_{k \geq 1} \phi_{\mathcal{C}_{\leq t}}[N_{t+C} \mid \mathbf{K} = k, \mathbf{j} = j] \phi_{\mathcal{C}_{\leq t}}[\mathbf{K} = k \mid \mathbf{j} = j]. \quad (37)$$

To bound $\phi_{\mathcal{C}_{\leq t}}[N_{t+C} \mid \mathbf{K} = k, \mathbf{j} = j]$ explore first the connected component $\tilde{\mathcal{C}}$ of $\mathcal{L}_{t,j}$ and observe that it is connected to at most $2\alpha(k+C)$ intervals $\mathcal{L}_{t+C,\ell}$, since it is contained in $\mathcal{Y} + x_{t,j} - (kL, 0)$.

Conditionally on $\tilde{\mathcal{C}}$, all other active intervals $\mathcal{L}_{t,\ell}$ are connected in $\tilde{\mathcal{C}}^c$ to a random number of intervals $\mathcal{L}_{t+C,\ell'}$ with an average at most μ . Indeed, the conditioning on $\tilde{\mathcal{C}}$ induces free boundary conditions on $\tilde{\mathcal{C}}^c$, and by (MON) the estimate (34) applies.

We conclude that $\phi_{\mathcal{C}_{\leq t}}[N_{t+C} \mid \mathbf{K} = k, \mathbf{j} = j] \leq \mu N_t + 2\alpha(k+C)$. Inserting this into (37) and using (35) we conclude that

$$\begin{aligned} \phi_{\mathcal{C}_{\leq t}}[N_{t+C} \mid \mathbf{j} = j] &\leq \mu N_t + 2\alpha \sum_{k \geq 0} (k+C) \phi_{\mathcal{C}_{\leq t}}[\mathbf{K} = k \mid \mathbf{j} = j] \\ &\leq \mu N_t + 2\alpha C_0 \sum_{k \geq 0} (k+C) e^{-c_0 k} \end{aligned}$$

⁶The careful reader might note that Proposition 3.7 does not rigorously apply in this setting, because of the additional conditioning on the configuration in $\Lambda_{2L}(x_{t,j})$ which does not correspond to a potential past. Nevertheless, a direct application of (RSW) proves that the bound (16) still applies in this case, with a potential bounded multiplicative factor.

which produces the desired result with $K = 2\alpha C_0 \sum_k (k + C) e^{-c_0 k} < \infty$.

Case $C_{\leq t}$ is such that $\phi[C_{\leq t} \leftrightarrow \mathcal{H}_n] \geq \frac{1}{2}$. Here we may simply write

$$\phi_{C_{\leq t}}[N_{t+C} \mid X_n \neq \dagger] \leq 2\phi_{C_{\leq t}}[N_{t+C}] \leq 2\mu N_t,$$

with the last bound obtained from summing (34) over all active segments. As μ was chosen strictly below $1/2$, Proposition 3.10 follows with $\mu' = 2\mu < 1$. $\square$

Remark 3.13. A minor modification of the proof above also implies the existence of a universal constant c such that, for all t and $n \geq t + C$

$$\phi[N_{t+C} = 1 \mid C_{\leq t}, X_n \neq \dagger] \geq e^{-cN_t}. \quad (38)$$

This observation shall be useful later on.

Recall that we still need to prove the “gluing estimate” (36). This is a relatively standard, but tedious use of the RSW property (RSW).

Proof of Lemma 3.12. Fix all quantities as in the statement of the lemma and Proposition 3.10. All constants appearing below will be independent of these quantities.

The upper bound on $\phi_{C_{\leq t}}[\mathbf{j} = j]$ is immediate, since $\{\mathbf{j} = j\}$ requires j to be a top-most seed and for $\Lambda_{4L}(x_{t,j})$ to be connected to $\mathcal{H}_{\geq n}$. These two events depend on the inside of $\Lambda_{2L}(x_{t,j})$ and the outside $\Lambda_{4L}(x_{t,j})$, respectively. The mixing property proved in (27) allows one to factorise the probabilities of these two events, up to a universally bounded multiplicative constant.

We now focus on the lower bound. Fix $R = L/10$ and write $x = x_{t,j}$ for the center of the interval $\mathcal{L}_{t,j}$. Write $\mathcal{A}$ for the points of the cluster $C_{\leq t}$ on $\mathcal{L}_{t,j}$. Also write $\mathcal{B}$ for the points of $C_{\leq t}$ on $\mathcal{L}_{t,\ell}$ with $\ell > j$; these all lie on $\partial\mathcal{H}_{\geq t}$, above $\mathcal{L}_{t,\ell}$. Let x_+ and x_- be the top-most and bottom-most points of $\mathcal{A}$ respectively. Write y_- for the bottom-most point of $\mathcal{B}$. Note that x_+ and y_- are linked by a dual path of the boundary of $C_{\leq t}$.

We start off by exploring certain interfaces starting on the boundary of $C_{\leq t}$ up to certain stopping times; write Exp for the set of edges thus explored. If *successful* (see below for a definition), the exploration will produce two *exposed* arcs, which is to say a primal arc χ_1 connected to $\mathcal{A}$ and a dual arc χ_0 connected in the dual configuration the free path of $\partial C_{\leq t}$ between x_+ and y_- , such that both χ_1 and χ_0 may be linked to the right side of $\partial\Lambda_{2L}(x)$ by disjoint tubes contained in $\Lambda_{2L}(x) \setminus (C_{\leq t} \cup \text{Exp})$, of width $R/4$ and length at most $100R$. See the left diagram of Figure 4 for an example.

Our procedure will be such that Exp is contained within distance $5R$ of $\mathcal{L}_{t,j}$ and

$$\phi_{C_{\leq t}}[\text{exploration is successful}] \gtrsim \phi_{C_{\leq t}}[j \text{ top-most seed}]. \quad (39)$$

Once this is done, a simple argument based on (16), (27), (RSW) and the condition on $C_{\leq t}$ will show that

$$\phi_{C_{\leq t}}[\mathbf{j} = j \mid \text{exploration is successful}] \gtrsim s_{n-t}. \quad (40)$$




Figure 4: *Left*: An example of a successful exploration (purple), with the arcs χ_0 and χ_1 being the top and bottom sides of the interface and the tubes depicted in grey. The solid blue lines form $\mathcal{A}$, while $\mathcal{B}$ is the solid black segment above y_- . An exploration is needed when either $\mathcal{A}$ or the dual arc between x_+ and y_- are inaccessible. *Right*: For a successful exploration, χ_1 may be connected to $\mathcal{H}_{\geq n}$ by a primal cluster that does not intersect $\mathcal{B}$ with probability at least s_{n-t} . After exploring the cluster $\mathcal{C}_\chi$ of χ_1 (red) $\mathcal{B}$ may be separated from $\mathcal{H}_{\geq n}$ with uniformly positive probability due to our assumption on $\mathcal{C}_{\leq t}$.

The combination of (39) and (40) yields the desired lower bound. Below we focus on proving these two equations. The former is the major difficulty of the proof, and is obtained via a tedious, but almost deterministic construction. For illustration, note that, when $\mathcal{A}$ contains a segment of length R and x_+ and y_- are at a distance at least R of each other, no exploration is necessary, and (39) is trivial.

Explorations and proof of (39). The exploration is done in two steps, each defining one of χ_0 and χ_1 .

Step 1: Write $\partial_\infty \Lambda_{2R}(x_+)$ for the arc of $\partial \Lambda_{2R}(x_+)$ separating x_+ from ∞ in $\mathcal{C}_{\leq t}^c$. This contains at least $\partial \Lambda_{2R}(x_+) \cap \mathcal{H}_{\geq t}$. Split it into the top and bottom part, which lie above and below the point $x_+ + (2R, 0)$, respectively.

Write Γ for the top boundary of the cluster of $\mathcal{A}$ in $\mathcal{C}_{\leq t}^c$. This is a path that will be indexed by $[0, 1]$ starting from the point x_+ . Since the cluster of $\mathcal{A}$ is a.s. finite, Γ eventually ends at x_- . Below we will explore Γ up to a convenient stopping time τ . For now set τ to be the first time Γ hits $\partial_\infty \Lambda_{2R}(x_+)$; the actual stopping time will be different, but we avoid introducing new notation.

We distinguish several cases:

- (1) x_+ and y_- are at a distance larger than $2R$ of each-other. Then update τ to be equal to 0;
- (2) x_+ and y_- are at a distance smaller than $2R$ of each-other, $\Gamma_{[0, \tau]}$ does not connect $\mathcal{A}$ to $\mathcal{B}$ and Γ_τ is in the bottom part of $\partial_\infty \Lambda_{2R}(x_+)$;

- (3) x_+ and y_- are at a distance smaller than $2R$ of each-other, $\Gamma_{[0,\tau]}$ does not connect $\mathcal{A}$ to $\mathcal{B}$ and Γ_τ is in the top part of $\partial_\infty \Lambda_{2R}(x_+)$;
- (4) $\Gamma_{[0,\tau]}$ connects $\mathcal{A}$ to $\mathcal{B}$.

In case (4), the exploration is considered unsuccessful. Note however that, when j is a top-most seed, Γ exits $\Lambda_{2L}(x)$ without connecting $\mathcal{A}$ to $\mathcal{B}$, and therefore excludes case (4). Thus,

$$\phi_{C_{\leq t}}[\text{cases (1), (2) or (3)}] \geq \phi_{C_{\leq t}}[j \text{ top-most seed}].$$

Next we define χ_0 or χ_1 for each of the cases (1)-(3) differently. In case (1), set χ_0 to be the part of the boundary of $C_{\leq t}$ between x_+ and y_- . Also write $z = x_+$ and $\tilde{C} = C_{\leq t}$.

In the second and third case set $z = \Gamma_\tau$ and $\tilde{C} = C_{\leq t} \cup \Gamma([0, \tau])$. In the second case, write χ_0 for the top part of $\Gamma([0, \tau])$; while in the third case we denote by χ_1 the bottom part of $\Gamma([0, \tau])$.

Notice that in the first two cases, χ_0 is a dual arc accessible in $\tilde{C}^c$, in the sense defined above. In the third case χ_1 is a primal arc accessible in $\tilde{C}^c$. Regardless, our present exploration up to this point only produces one of the two accessible arcs required. In the second step, we will perform further explorations which will produce the second arc. This step depends on which case occurred in Step 1.

Step 2 in cases (1) and (2): Define the region Λ' as $\Lambda' = \Lambda_R(x_+)$ in case (1) and, in case (2), as the union of $\Lambda_R(u)$ for u in the bottom part of $\partial_\infty \Lambda_{2R}(x_+)$. Write $\partial_\infty \Lambda'$ for the arc of $\partial \Lambda'$ separating z from ∞ in $\tilde{C}^c$. Split $\partial_\infty \Lambda'$ at $x_+ + (R, 0)$ in case (1) and at $x_+ + (3R, 0)$ in case (2) into its top and bottom sections – see Figure 5.

Consider now τ' as the first time after τ when Γ touches $\partial_\infty \Lambda'$. We distinguish two sub-cases

- (a) $\Gamma_{\tau'}$ is in the top part of $\partial_\infty \Lambda'$ or
- (b) $\Gamma_{\tau'}$ is in the bottom part of $\partial_\infty \Lambda'$.

In case (a), write χ_1 for the wired arc of $\Gamma_{[0,\tau']}$. Then χ_1 is a primal arc and χ_0 is a dual arc, both accessible in $(\tilde{C} \cup \Gamma_{[0,\tau']})^c$.

In case (b) a more complicated construction is needed. Write $\tilde{\Gamma}$ for the exploration path starting at x_+ , leaving vertices connected to $\mathcal{A}$ in $C_{\leq t}^c$ on its left (including those of $\mathcal{A}$, but not other vertices of $C_{\leq t}$), and all other vertices on the right. See Figure 5 for an illustration.

Write τ'' for the first time this path touches $\partial_\infty \Lambda'$; Figure 5 shows why such a time exists when j is a top-most seed, and why $\tilde{\Gamma}_{\tau''}$ is “below” $\Gamma_{\tau'}$ on $\partial_\infty \Lambda'$. It follows that the wired arc χ_1 of $\tilde{\Gamma}_{[0,\tau'']}$ is accessible in $(\tilde{C} \cup \tilde{\Gamma}_{[0,\tau'']})^c$.

Notice that we do not claim that χ_1 is accessible when $\Gamma_{[\tau,\tau']}$ has been explored. Thus, one should make a choice at the stopping time τ whether to continue exploring Γ up to time τ' or whether to explore $\tilde{\Gamma}$ up to time τ'' .




Figure 5: The different scenarios that may occur when exploring Γ , and potentially $\tilde{\Gamma}$. The exploration up to τ and the resulting connection is drawn in red, the subsequent construction is in purple. The red and purple arrow refer connections to the paths from $\Lambda_R(x + (5R, \pm R))$ to $\mathcal{H}_{\geq n}$ and ∞ , which occur with probability comparable to s_{n-t} . From left to right: scenarios 1(a), 2(a), 2(b) and 3. Note that in scenario 2(b), the exploration path Γ never exposes an accessible primal arc. That is why we do not explore Γ between τ and τ' (grey) but rather explore $\tilde{\Gamma}$ up to τ'' (purple).

Notice however that either

$$\begin{aligned} \phi[\text{case (a)} \mid \text{case (1) or (2), } \Gamma_{[0,\tau]} \text{ and } j \text{ top-most seed}] &\geq 1/2 \text{ or} \\ \phi[\text{case (b)} \mid \text{case (1) or (2), } \Gamma_{[0,\tau]} \text{ and } j \text{ top-most seed}] &\geq 1/2. \end{aligned}$$

If (a) is more probable, then explore Γ up to time τ' . If the resulting exploration corresponds to case (a), define χ_1 as above and set $\text{Exp} = \Gamma_{[0,\tau']}$. Otherwise say the exploration is unsuccessful.

If (b) is more probable, explore $\tilde{\Gamma}$ up to time τ'' . If the wired part of $\tilde{\Gamma}_{[0,\tau']}$ is indeed accessible, denote it by χ_1 and set $\text{Exp} = \Gamma_{[0,\tau]} \cup \tilde{\Gamma}_{[0,\tau']}$. Otherwise say the exploration is unsuccessful.

With this construction, if the exploration is successful, the arcs χ_0 and χ_1 are indeed both accessible in $(C_{\leq t} \cup \text{Exp})^c$. Furthermore, our analysis proves that

$$\phi[\text{exploration is successful} \mid \text{case (1) or (2) and } j \text{ top-most seed}] \geq 1/2,$$

which implies (39) in cases (1) and (2).

Step 2 in case (3): This case is treated similarly to case (2), except that χ_1 is now defined and we aim to define χ_0 . Set Λ' to be the union of $\Lambda_R(u)$ for u in the top part of $\partial_\infty \Lambda_{2R}(x_+)$. Define $\partial_\infty \Lambda'$ and its top and bottom sections similarly to how this was done in case (2).

Define sub-cases (a) and (b) depending on where Γ exits Λ' ; set τ' to be the exit time. If it is more likely to exit on the bottom part, define χ_0 as the dual (top) side

of $\Gamma_{[\tau, \tau']}$. If it is more likely to exit on the top part, start an exploration $\tilde{\Gamma}$ from x_+ leaving primal open edges on the left and dual open ones on the right (this exploration runs along $\partial\mathcal{C}_{\leq t}$ up to y_- , and continues at least up to $\partial\Lambda_{2L}(x)$ when j is a top-most seed). Define then χ_0 as the dual side of $\tilde{\Gamma}$ up to the first exit time $\partial_\infty\Lambda'$.

The same type of analysis as above shows that

$$\phi_{\mathcal{C}_{\leq t}}[\text{exploration is successful} \mid \text{case (3) and } j \text{ top-most seed}] \geq 1/2,$$

which implies (39) in case (3).

Conclusion: proof of (40). Fix now a successful realisation of Exp, with the primal and dual arcs denoted χ_1 and χ_0 , respectively. To ensure $\mathbf{j} = j$, it suffices to connect χ_1 to $\mathcal{H}_{\geq n}$ and produce a dual path starting from χ_0 that prevents any connection between $\mathcal{B}$ and $\mathcal{H}_{\geq n}$. See Figure 4, right diagram.

By (16), (27) and (RSW), with probability comparable to s_{n-t} , the cluster $\mathcal{C}_\chi$ of χ_1 in $(\mathcal{C}_{\leq t} \cup \text{Exp})^c$ intersects $\mathcal{H}_{\geq n}$, but does not connect to any point in $\mathcal{B}$.

Furthermore, conditionally on $\mathcal{C}_\chi$, the measure in $(\mathcal{C}_{\leq t} \cup \text{Exp} \cup \mathcal{C}_\chi)^c$ is dominated by the measure $\phi_{\mathcal{C}_{\leq t}}$. Indeed, the exploration of $\mathcal{C}_\chi$ produces free boundary conditions in the remaining part of the space. Thus,

$$\phi_{\mathcal{C}_{\leq t}}[\mathcal{B} \not\leftrightarrow \mathcal{H}_{\geq n} \mid \text{Exp}, \mathcal{C}_\chi] \geq \phi_{\mathcal{C}_{\leq t}}[\mathcal{C}_{\leq t} \not\leftrightarrow \mathcal{H}_{\geq n}] \geq 1/2,$$

Averaging over $\mathcal{C}_\chi$ and using the previous observation, we conclude (40). $\square$

We finally turn to the proof of Proposition 3.9.

Proof of Proposition 3.9. Let C , K and μ be given by Proposition 3.10. Set

$$\ell = \frac{2K}{1-\mu} \quad \text{so that} \quad \mu\ell + K = \frac{1+\mu}{2}\ell.$$

The proof is in two steps. First we prove that, for any $n, t \geq 1$,

$$\phi[N_s > \ell, \forall s \in \{t+1, \dots, t+r\} \mid \mathcal{C}_{\leq t}, N_t \leq \ell, X_n \neq \dagger] \leq \exp(-cr) \quad \text{and} \quad (41)$$

$$\phi[N_s > \ell, \forall s \in \{1, \dots, r\} \mid X_n \neq \dagger] \leq \exp(-cr), \quad (42)$$

for some universal constant $c > 0$. Then we will argue that

$$\phi[N_{t+C} = 1 \mid \mathcal{C}_{\leq t}, N_t \leq \ell, X_n \neq \dagger] \geq 1. \quad (43)$$

Indeed, (41), (42) and (43) imply Proposition 3.9, and we focus on the proof of these three bounds below.

We start off with the proof of (41). Let $\alpha = \frac{2}{1+\mu} > 1$. Fix $t \geq 1$ and work under the conditional measure $P := \phi[\cdot \mid \mathcal{C}_{\leq t}, X_n \neq \dagger]$. For $s \geq 0$, set $M_s = \alpha^s N_{t+Cs}$ and $\tau = \inf\{s \geq 1 : N_{t+Cs} \leq \ell\}$. Then, (33) and the choice of ℓ imply that $M_{s \wedge \tau}$ is a positive super-martingale.

It follows that

$$P[\tau > r] \leq P[M_{r \wedge \tau} \geq \alpha^r \ell] \leq \frac{1}{\alpha^r \ell} E[M_{r \wedge \tau}] \leq \frac{1}{\alpha^r \ell} E[M_0].$$

This directly proves (41), since $E[M_0] \leq \ell$.

The same reasoning applies to (42), provided that we have a uniform bound on the expectation $\phi[N_1 \mid X_n \neq \dagger]$. Such a bound is easily obtained from (16) since

$$\phi[N_1 \geq 2\alpha k \mid X_n \neq \dagger] \leq \phi[\Lambda_L \leftrightarrow (\mathcal{Y}_\alpha - (kL, 0))^c \mid \Lambda_L \leftrightarrow \mathcal{H}_{\geq n}] \leq e^{-ck}.$$

Finally, (43) is a direct consequence of (38). $\square$

3.3 Mixing when $N_k = 1$

We now prove that the process $(X_k)_k$ satisfies a mixing property at every time t for which $N_t = 1$.

Write $\mathbf{C}_{\geq t} = \mathbf{C} \setminus \mathbf{C}_{\leq t}$ and $\mathbf{C}_{s \leq \cdot \leq t} = \mathbf{C}_{\leq t} \setminus \mathbf{C}_{\leq s}$ for $0 \leq s \leq t$. We will argue that, conditionally on the past $\mathbf{C}_{\leq t}$, the future $\mathbf{C}_{\geq t+1/2}$ may be sampled independently of the past with positive probability. To state our result well, some vocabulary needs to be introduced. First of all, for any $n \geq 0$, we call $\mathcal{L}_n := \phi_{\mathcal{H}_{\geq 0}}^1[\cdot \mid \mathcal{L}_{0,0} \leftrightarrow \mathcal{H}_{\geq n}]$.

A realisation χ of $\mathbf{C}_{\geq t+1/2} - (Lt, X_t)$ is called a *cluster-future*. It may be formed of several connected components, each containing at least some edge in $\mathcal{H}_{\geq 1/2}$ and at least one point on $\partial\mathcal{H}_{\geq 1/2}$; see Figures 6 and 7 for illustrations. The section $\mathbf{C}_{t \leq \cdot \leq t+1/2} = \mathbf{C} \setminus (\mathbf{C}_{\leq t} \cup \mathbf{C}_{\geq t+1/2})$ connecting the past cluster to the (translate of the) cluster-future is called the *link*. We now define a particular class of *well-behaved* cluster-futures.

Consider the measure $\phi_{\mathcal{H}_{\geq 0}}^1$ and write $\mathbf{C}$ for the cluster of $\mathcal{L}_{0,0}$. Fix a cluster-future χ and define $\mathcal{E}(\chi)$ as the event that the rectangle $[\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$ contains horizontal open crossings, that the lowest one is contained in $[\frac{3}{8}L, \frac{1}{2}L] \times [-L/2, -L/3]$ and the highest one is contained in $[\frac{3}{8}L, \frac{1}{2}L] \times [L/3, L/2]$, and that they are connected via an open path in $[\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$. Moreover, $\mathcal{E}(\chi)$ requires that, if we consider the cluster Ξ of χ in $\mathcal{H}_{\geq 3/8}$, we have $\Xi \setminus \chi \subset [\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$ and it crosses horizontally $[\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$. See Figure 6 for an illustration.

For a constant $c_{\text{WB}} > 0$, we call χ *c_{WB}-well-behaved* if χ is contained in $\mathcal{Y}$ and

$$\phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{E}(\chi) \mid \mathbf{C}_{\geq 1/2} = \chi] \geq c_{\text{WB}}.$$

The following is a direct consequence of Proposition 3.7 and (RSW).

Corollary 3.14. *There exists a universal constant $c_{\text{WB}} > 0$ such that, for all $n, s \geq 1$,*

$$\mathcal{L}_n[\mathbf{C}_{\geq 1/2} \text{ c}_{\text{WB}}\text{-well-behaved}] \geq c_{\text{WB}} \quad \text{and} \quad \mathcal{L}_n[\mathbf{C} \subset \mathcal{Y} - (sL, 0)] \geq 1 - e^{-c_{\text{WB}} s}.$$

Moreover, the same holds under $\phi[\cdot \mid \text{Past}(A)]$ and $\mathbf{C} \cap \mathcal{H}_{\geq n} \neq \emptyset$ for any potential past A .




Figure 6: The cluster future χ is formed of the pink connected components, surrounded by dual paths. The thick red crossings in the rectangle $[\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$ are required for $\mathcal{E}(\chi)$. Moreover, the orange part ensures that χ is linked to these crossings in such a way that $\mathcal{E}(\chi)$ occurs.

Proof. The second inequality is indeed a direct consequence of Proposition 3.7. We focus on the first and treat the case of $\mathcal{L}_n$ for simplicity.

By Proposition 3.7 and (RSW) we find

$$\phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{C}_{\geq 1/2} \subset \mathcal{Y}, \mathcal{E}(\mathcal{C}_{\geq 1/2}) \text{ and } \mathcal{L}_{0,0} \leftrightarrow \mathcal{H}_{\geq n}] \geq \phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{L}_{0,0} \leftrightarrow \mathcal{H}_{\geq n}] \quad (44)$$

Indeed, we may first explore the top and bottom boundaries of $\mathcal{C}$, starting on $\partial\mathcal{H}_{\geq 0}$, up to the first time they reach $\mathcal{H}_{\geq 1}$. By (RSW), we may direct these to satisfy the requirements of $\mathcal{E}(\cdot)$ and to arrive at a macroscopic distance from each other. Conditionally on their realisation, using (RSW) again, we may connect them by a primal path in $[\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$, and connect them by dual paths to the top and bottom of $[\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$. Finally, we connect them to $\mathcal{H}_{\geq t}$ while ensuring that their cluster remains in $\mathcal{Y}$. Conditionally on the previous events, the latter has probability of the order of $\phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{L}_{0,0} \leftrightarrow \mathcal{H}_{\geq n}]$ due to Proposition 3.7 and Lemma 3.8.

From (44), we conclude that

$$\sum_{\chi} \phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{E}(\chi) | \mathcal{C}_{\geq 1/2} = \chi] \mathcal{L}_n[\mathcal{C}_{\geq 1/2} = \chi] = \mathcal{L}_n[\mathcal{C}_{\geq 1/2} \subset \mathcal{Y}, \mathcal{E}(\mathcal{C}_{\geq 1/2})] \geq 1.$$

where the sum is over all possible cluster-futures $\chi \subset \mathcal{Y}$. The result follows readily. $\square$

Henceforth we fix c_{WB} given by Corollary 3.14 and omit it from notation. We may now state the crucial mixing estimate for the process $(X_t)_{t \geq 0}$, and more generally for the exploration $(\mathcal{C}_{\leq t})_{t \geq 0}$ of the cluster of 0.

Proposition 3.15. *There exists $\eta > 0$ such that for any $1 \leq t \leq n$ and any well-behaved cluster-future χ*

$$\begin{aligned} & \phi[\mathbf{C}_{\geq t+1/2} = \chi + (Lt, X_t) \text{ and } \mathbf{C}_{\geq t} \subset \mathcal{Y} + (L(t-1), X_t) \mid \mathbf{C}_{\leq t}, N_t = 1, X_n \neq \dagger] \\ & \geq \eta \mathcal{L}_{n-t}[\mathbf{C}_{\geq 1/2} = \chi], \end{aligned} \quad (45)$$

where $\chi + (Lt, X_t)$ is the translation of χ by (Lt, X_t) .

Moreover, for any $1 \leq s < t \leq n$, any possible realisation ζ of $\mathbf{C}_{s+\frac{1}{2} \leq \cdot \leq t} - (Ls, X_s)$ such that $N_t = 1$, and any realisation χ of $\mathbf{C}_{\geq t} - (Lt, X_t)$ contained in $\mathcal{Y} - (\frac{t-s}{2}L, 0)$,

$$\begin{aligned} & \sum_{\chi} \left| \phi[\mathbf{C}_{\geq t} = \chi + (Lt, X_t) \mid \mathbf{C}_{\leq t}, \mathbf{C}_{s+\frac{1}{2} \leq \cdot \leq t} = \zeta + (Ls, X_s), X_n \neq \dagger] \right. \\ & \quad \left. - \mathcal{L}_{n-s}[\mathbf{C}_{\geq t-s} = \chi + (L(t-s), X_{t-s}) \mid \mathbf{C}_{\frac{1}{2} \leq \cdot \leq t-s} = \zeta] \right| \leq e^{-\eta(t-s)}. \end{aligned} \quad (46)$$

The first inequality together with Corollary 3.14 state that, at every pre-renewal time, there is a positive probability for the cluster-future to be sampled independently of the past, thus creating an actual point of renewal. It is tempting to think that this implies exponential tails for the spacing between renewal times, thus proving Theorem 3.2.

Unfortunately this is not the case. Indeed, (45) implies a first-moment bound on the number of renewals, but may not be used to control correlations between renewal times. Hence the need for (46), which states that, if the clusters sampled according to $\phi[\cdot \mid \mathbf{C}_{\leq s}, X_n \neq \dagger]$ and $\mathcal{L}_{n-s}$ are coupled for $t-s$ steps, then they remain coupled for the rest of the process with probability exponentially close to 1. Crucially, the value of $\eta > 0$ may be chosen uniformly in p , the direction $\vec{w}$ as well as in t, s and n and the realisation of the cluster up to t .

Finally, note that in (45) we require that both the translated cluster-future and the link be contained in $\mathcal{Y} + (L(t-1), X_t)$. This will eventually ensure the cone-containment property of the cluster; it is a technical detail which may be ignored in a first instance.

Remark 3.16. Eventually (45) will be used to state that at any pre-renewal time, $\mathbf{C}_{\geq t+1/2}$ is sampled with positive probability independently from the past. We do not claim this about the entirety of $\mathbf{C}_{\geq t}$; indeed, (45) is not expected to apply to the link. See also Figure 7.

This is a crucial difference with [CIV08], where the link is trivial due to the use of cone points, and independent sampling estimates apply to the whole of $\mathbf{C}_{\geq t}$.

The rest of the section is dedicated to proving Proposition 3.15.

Proof of Proposition 3.15. We start with the proof of (45), which is the most complicated property.

Working under $\phi_{\mathcal{H}_{\geq 0}}^1$, write Γ_N and Γ_S for the highest and lowest, respectively, open horizontal crossings of $[\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$ and Γ_E for the right-most open path connecting these. When all three paths exist, write UnExp for the region of $[\frac{3}{8}L, \frac{1}{2}L] \times [-L, L]$ between Γ_N and Γ_S and left of Γ_E . Furthermore, write $\mathbf{C}^\Gamma$ the cluster of Γ_N in the region




Figure 7: *Top:* A realisation χ of $\mathbf{C}_{\geq 1/2}$ under $\mathcal{L}_{n-t}$ in red, and the link in purple. *Bottom:* The same realisation, shifted by (Lt, X_t) as $\mathbf{C}_{\geq t+1/2}$. The measure here is $\phi[\cdot | \mathbf{C}_{\leq t}, X_n \neq \dagger]$. Note that the link is different in the two diagrams.

$\mathcal{H}_{\geq \frac{3}{8}L} \setminus \text{UnExp}$. When the paths Γ_N , Γ_S or Γ_E do not exist, set $\mathbf{C}^\Gamma = \emptyset$. Finally, define $\mathbf{C}_{\geq 1/2}^\Gamma$ in a same way as $\mathbf{C}_{\geq 1/2}$.

Fix a well-behaved cluster-future χ that reaches $\mathcal{H}_{\geq n-t}$. The event $\mathcal{E}(\chi)$ may be determined by $\mathbf{C}^\Gamma$. Indeed, it requires that all three paths Γ_N , Γ_S and Γ_E exists, that the first two be contained in specific regions and that $\mathbf{C}_{\geq 1/2}^\Gamma = \chi$. Recall the notation $s_n = \phi_{\mathcal{H}_{\leq 0}}^1[\mathcal{L}_{0,0} \leftrightarrow \mathcal{H}_{\geq n}]$. We have

$$\begin{aligned} \mathcal{L}_{n-t}[\mathbf{C}_{\geq 1/2} = \chi] &\leq \frac{1}{s_{n-t}} \phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{E}(\chi) \text{ and } \mathbf{C}_{\geq 1/2} = \chi] \\ &\leq \frac{1}{s_{n-t}} \phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{E}(\chi) \text{ and } \mathbf{C}_{\geq 1/2}^\Gamma = \chi], \end{aligned} \quad (47)$$

where the first line is a consequence of the definition of well-behaved cluster-futures and the second is due to the inclusion of events.

Similarly, when working under $\phi[\cdot | \mathbf{C}_{\leq t}]$, with a past $\mathbf{C}_{\leq t}$ such that $N_t = 1$, define $\tilde{\mathbf{C}}^\Gamma$ in the same way as $\mathbf{C}^\Gamma$, but shifted by (Lt, X_t) . Then

$$\begin{aligned} &\phi[\mathbf{C}_{\geq t+1/2} = \chi + (Lt, X_t) \text{ and } \mathbf{C}_{\geq t} \subset \mathcal{Y} + (L(t-1), X_t) | \mathbf{C}_{\leq t}, X_n \neq \dagger] \\ &\geq \frac{1}{\phi[X_n \neq \dagger | \mathbf{C}_{\leq t}]} \sum_{\chi^\Gamma} \phi[\tilde{\mathbf{C}}^\Gamma = \chi^\Gamma + (Lt, X_t) | \mathbf{C}_{\leq t}] \cdot \phi_D^\tau[\mathbf{C}_{\leq t} \Leftrightarrow \Gamma_E + (Lt, X_t)], \end{aligned}$$

where the sum is over all possible realisations χ^Γ of $\mathbf{C}^\Gamma$ such that $\mathbf{C}_{\geq 1/2}^\Gamma = \chi$ and such that $\mathcal{E}(\chi)$ occurs. Above D is the complement of $\mathbf{C}_{\leq t} \cup (\chi^\Gamma + (Lt, X_t))$ and τ are the boundary conditions induced by conditioning on $\tilde{\mathbf{C}}^\Gamma = \chi^\Gamma + (Lt, X_t)$, that is, those induced by $\mathbf{C}_{\leq t}$ on its boundary and which are wired on the portion on $\partial(\chi^\Gamma + (Lt, X_t))$ formed of

$\Gamma_N + (Lt, X_t)$, $\Gamma_S + (Lt, X_t)$ and $\Gamma_E + (Lt, X_t)$, and free elsewhere. Finally, we write $C_{\leq t} \Leftrightarrow \Gamma_E + (Lt, X_t)$ for the event that $C_{\leq t}$ and $\Gamma_E + (Lt, X_t)$ are connected, but that they are not connected to any other point of $\mathcal{H}_{\geq t+1/2}$ or $(\mathcal{Y} + (L(t-1), X_t))^c$.

We now argue that

$$\begin{aligned} \phi_D^T[C_{\leq t} \Leftrightarrow \Gamma_E + (Lt, X_t)] &\geq \phi_D^T[C_{\leq t} \leftrightarrow \partial\mathcal{H}_{\geq t+\frac{1}{4}}] \geq \phi[C_{\leq t} \leftrightarrow \partial\mathcal{H}_{\geq t+\frac{1}{4}} \mid C_{\leq t}] \\ \text{and } \phi[X_n \neq \dagger \mid C_{\leq t}] &\leq \phi[C_{\leq t} \leftrightarrow \partial\mathcal{H}_{\geq t+\frac{1}{4}} \mid C_{\leq t}]^{s_{n-t}}. \end{aligned}$$

The first inequality is proved as Lemma 3.12 and uses (RSW). Indeed, under the event $C_{\leq t} \leftrightarrow \partial\mathcal{H}_{\geq t+\frac{1}{4}}$, we may explore either the top-most or bottom-most interface stemming from $C_{\leq t}$ so that, with positive probability, it exposes a wired arc. Then, it suffices to connect the wired arc to Γ_E via a primal path and to connect the top and bottom of $C_{\leq t}$ to the top of Γ_N and the bottom of Γ_S by dual paths, all contained in $\mathcal{H}_{\leq t+1/2} \cap (\mathcal{Y} + (L(t-1), X_t))$ — (RSW) and the wide opening of $\mathcal{Y}$ allow one to construct these connections with positive probability. The second and third inequalities follow from Lemma 3.8.

We conclude that

$$\begin{aligned} \phi[C_{\geq t+1/2} = \chi + (Lt, X_t) \text{ and } C_{\geq t} \subset \mathcal{Y} + (L(t-1), X_t) \mid C_{\leq t}, X_n \neq \dagger] \\ \geq \frac{1}{s_{n-t}} \sum_{\chi^\Gamma} \phi[\tilde{C}^\Gamma = \chi^\Gamma + (Lt, X_t) \mid C_{\leq t}] \\ \geq \frac{1}{s_{n-t}} \sum_{\chi^\Gamma} \phi_{\mathcal{H}_{\geq 0}}^1[C^\Gamma = \chi^\Gamma] \\ = \frac{1}{s_{n-t}} \phi_{\mathcal{H}_{\geq 0}}^1[\mathcal{E}(\chi) \text{ and } C_{\geq 1/2}^\Gamma = \chi], \end{aligned} \tag{48}$$

where the sum is over all χ^Γ as above. The second inequality is due to the mixing property of Lemma 3.8 and the fact that χ^Γ is contained in $\mathcal{Y} \cap \mathcal{H}_{\geq 3/8}$. Finally, (47) and (48) together imply (45).

The second property (46) follows by the reasoning leading to the mixing property (18). Indeed, even though we are conditioning on the edges of the “middle part” given by ζ , observe that the conditioning actually helps the existence of dual paths separating $\mathcal{H}_{\geq 0}$ and $\mathcal{Y} + (\frac{t-s}{2}L, 0)$ in ζ^c . Thus, the proof of the mixing property can be repeated *mutatis mutandis* and yields (46), and we do not give further details. $\square$

3.4 Proof of Theorem 3.2

Fix a direction $\vec{w}$. Recall the notations C for the cluster of 0 and N_t for the number of active segments on $\partial\mathcal{H}_{\leq t}$. Write $S_0 = 0$ and for $k \geq 0$ set

$$S_{k+1} = \inf\{t \geq S_k + 2 : N_t = 1\}$$

Note that we impose that $S_{k+1} - S_k \geq 2$; this is purely for technical reasons. Recall that the times $(S_k)_{k \geq 1}$ are called the pre-renewal times of $(X_t)_t$. Let K be the first index for which $S_K = \infty$.

Let $D_k = C_{S_{k-1} \leq \cdot \leq S_k}$ for $k = 1, \dots, K-1$. We also set $D_K = C_{\geq S_{K-1}}$; this is the only piece D_k which does not end with a pre-renewal time. First we will describe how to sample the pieces $(D_k)_{k \geq 0}$ sequentially, which in turn constructs the sequence $(S_k)_{k \geq 0}$. By constructing the pieces $(D_k)_{k \geq 0}$, we construct the cluster C and therefore the sequence $(X_t)_{t \geq 0}$. We will later construct the variables $(Y_t)_{t \geq 0}$ and prove the various properties of Definition 3.1.

For this proof, write $\mathbb{P}$ for the probability measure used to sample C according to the procedure described below. To directly prove that the process has a mass-gap, we will sample C under the conditional measure $\phi[\cdot | 0 \leftrightarrow \mathcal{H}_{\geq n}]$, or equivalently $\phi[\cdot | X_n \neq \dagger]$, for some arbitrary $n \geq 0$. To sample under the unconditional measure, it suffices to set $n = 0$.

Sequential sampling of C . Sample the initial step D_1 as $\tilde{C}_{\leq S_1}$ where $\tilde{C}$ is a sample of the cluster of 0 under $\phi[\cdot | X_n \neq \dagger]$.

Fix $k \geq 1$ and assume the pieces $D_1, \dots, D_k$ already defined. Write ζ for the realisation of $C_{\leq S_k}$ and assume that $S_k < \infty$ (otherwise the sampling procedure is finished). We now describe how to sample D_{k+1} conditionally on $\zeta = C_{\leq S_k}$. For simplicity write $s := S_k$; the value of S_k is determined by the conditioning, so may be treated as a constant.

Proposition 3.15 states that the information on $C_{\leq s}$ needed to sample the future is limited to a random number j of past steps, which is to say $C_{s-j \leq \cdot \leq s}$, with j having exponential tails. Our first task is to formalise this.

For $j > 0$ and χ a potential realisation of $C_{\geq s} - (Ls, X_s)$ contained in $\mathcal{Y} - (\frac{jL}{2}, 0)$ set

$$q_j(\zeta, \chi) = \min_{\xi} \phi[C_{\geq s} = \chi + (Ls, X_s) | C_{\leq s} = \xi, X_n \neq \dagger],$$

where the minimum is taken over all possible realisations ξ of $C_{\leq s}$ with $\xi_{s-j \leq \cdot \leq s} = \zeta_{s-j \leq \cdot \leq s}$. For $j > s$ we simply take $\xi = \zeta$; but the cone condition on χ still depends on j . For all futures χ not contained in the appropriate cones, we set $q_j(\zeta, \chi) = 0$.

When $j = 0$, $q_0(\zeta, \chi) = q_0(\chi)$ requires a special definition. For any well-behaved cluster-future χ set

$$q_0(\chi) = \min_{\xi} \phi[C_{\geq s+1/2} = \chi + (Ls, X_s), C_{\geq s} \subset \mathcal{Y} + (L(s-1), X_s) | C_{\leq s} = \xi, X_n \neq \dagger],$$

where the minimum is over all potential realisations ξ of $C_{\leq s}$ producing $N_s = 1$. For all other cluster futures χ set $q_0(\zeta, \chi) = 0$.

Also, define

$$Z_j(\zeta) := \sum_{\chi} q_j(\zeta, \chi),$$

where the sum is over all potential realisation χ of $C_{\geq s} - (Ls, X_s)$ (or all cluster-futures when $j = 0$). By definition, the quantities $q_j(\zeta, \chi)$ and $Z_j(\zeta)$ are increasing in $j \geq 1$.

The combination of Propositions 3.15 and 3.7 states that, uniformly in ζ as above

$$Z_0(\zeta) \geq 1 \quad \text{and} \quad 1 - Z_j(\zeta) \leq e^{-\eta j}. \quad (49)$$

Define a random variable $\tilde{M}_k$ taking values in $\mathbb{N}$ with

$$\mathbb{P}[\tilde{M}_k \leq j \mid \mathbf{C}_{\leq s} = \zeta] = Z_j(\zeta).$$

Sample $\tilde{M}_k$, then, conditionally on $\tilde{M}_k$ we will sample a potential future for $\mathbf{C}$ as follows.

If $\tilde{M}_k = 0$ set

$$\mathbb{P}[\tilde{\mathbf{C}}_{\geq s+1/2} = \chi + (Ls, X_s) \mid \mathbf{C}_{\leq s}, \tilde{M}_k = 0] = \frac{1}{Z_0(\zeta)} q_0(\zeta, \chi), \quad (50)$$

for any possible cluster-future χ . Then sample $\tilde{\mathbf{C}}_{s \leq \cdot \leq s+1/2}$ according to the conditional measure

$$\phi[\cdot \mid \mathbf{C}_{\leq s} = \zeta, \mathbf{C}_{\geq s+1/2} = \tilde{\mathbf{C}}_{\geq s+1/2} \text{ and } \mathbf{C}_{\geq s} \subset \mathcal{Y} - (L(s-1), X_s)].$$

Note here that $\mathbf{C}_{\geq s+1/2}$ is sampled independently of $\mathbf{C}_{\leq s}$, but that the link is allowed to depend on it.

For $j \geq 1$ and χ a potential realisation of $\mathbf{C}_{\geq s} - (Ls, X_s)$, define $\mathbb{P}[\tilde{\mathbf{C}}_{\geq s} = \chi + (Ls, X_s) \mid \mathbf{C}_{\leq s}, \tilde{M}_k = j]$ inductively by

$$\begin{aligned} & \mathbb{P}[\tilde{\mathbf{C}}_{\geq s} = \chi + (Ls, X_s) \mid \mathbf{C}_{\leq s}, \tilde{M}_k = j] \\ &= \frac{1}{\mathbb{P}[\tilde{M}_k = j \mid \mathbf{C}_{\leq s}]} \left(q_j(\zeta, \chi) - \mathbb{P}[\tilde{\mathbf{C}}_{\geq s} = \chi + (Ls, X_s) \text{ and } \tilde{M}_k < j \mid \mathbf{C}_{\leq s}] \right). \end{aligned} \quad (51)$$

It is immediate from the construction above that the sampling of $\tilde{\mathbf{C}}_{\geq s}$ follows the law of $\mathbf{C}_{\geq s}$ under $\phi[\cdot \mid \mathbf{C}_{\leq s} = \zeta]$.

If $\tilde{\mathbf{C}}_{\geq t}$ contains a pre-renewal, set $\mathbf{D}_{k+1} = \mathbf{D}_1(\tilde{\mathbf{C}}_{\geq t}) - (Ls, X_s)$, which is to say the piece of $\tilde{\mathbf{C}}_{\geq t}$ up to its first pre-renewal. Otherwise, set $\mathbf{D}_{k+1} = \tilde{\mathbf{C}}_{\geq s} - (Ls, X_s)$ and set $S_{k+1} = \infty$; the sampling of the sequence $(\mathbf{D}_j)_{j \geq 1}$ is finished.

At this stage, we have constructed $\mathbf{C}$ as a concatenation of pieces $(\mathbf{D}_k)_{k=0, \dots, K}$. The sequence $(X_t)_{t \geq 0}$ is implicitly defined, with $X_1, \dots, X_{S_k}$ depending on $\mathbf{D}_1, \dots, \mathbf{D}_k$ only. Note that each piece $\mathbf{D}_{k+1}$ needs to be translated by (LS_k, X_{S_k}) when attached to $\mathbf{C}_{\leq S_k}$.

Write $\text{len}(\mathbf{D}_k)$ is the horizontal “length” of $\mathbf{D}_k$, that is the maximal t for which $\mathbf{D}_k$ intersects $\mathcal{H}_{\geq S_k+t}$ — unless $\mathbf{D}_k$ is the last step of the process, this is equal to $S_{k+1} - S_k$. For $k > K$ we formally set $\mathbf{D}_k = \emptyset$ and $\text{len}(\mathbf{D}_k) = 0$.

Then, Proposition 3.9 states the existence of a universal constant $c > 0$ such that

$$\mathbb{P}[\text{len}(\mathbf{D}_k) \geq r \mid \mathbf{C}_{\leq S_k}] \leq \exp(-cr) \quad (52)$$

for all $k, r, n \geq 0$.

Memory variables and filtration. By (49),

$$\mathbb{P}[\tilde{M}_k > j \mid \mathbf{C}_{\leq t}] \leq \exp(-\eta j), \quad (53)$$

for all $j \geq 0$ and some universal constant $\eta > 0$. We may always decrease η , so we will henceforth assume that $0 < \eta \leq c$, where $c > 0$ is the constant appearing in (52).

By (53), we may bound the variables $\tilde{M}_k$ from above by i.i.d. modified geometric variables⁷ M_k with parameter $1 - e^{-\eta} > 0$. To be precise, at each sampling step S_k , sample a geometric variable M_{k+1} starting at 0, independent of $\mathbf{C}_{\leq t}$, then sample $\tilde{M}_{k+1} \leq M_{k+1}$ and $\mathbf{D}_{k+1}$ according to the procedure described above.

Set $Y_t = 1$ if $t = S_k$ for some k and $M_{k+j} \leq j$ for all $j \geq 0$. We have now defined the sequences $(X_t)_{t \geq 0}$ and $(Y_t)_{t \geq 0}$. Observe that Y_t depends on the “future” variables M_{k+j} (if t is such that $S_k = t$).

To render dependencies clear, let us properly define the filtration $(\mathcal{F}_t)_{t \geq 0}$ associated to the KMRP process. For $t \geq 0$, let $\mathcal{F}_t$ be the σ -algebra generated by $\mathbf{C}_{\leq t}$, all variables M_k and $\tilde{M}_k$ with $S_k < t$ and all variables Y_s with $s \leq t$.

Finally, the dependence between the memory variables M_k and the piece-lengths $\text{len}(\mathbf{D}_{k+1})$ will be of interest. Notice that, due to the relationship between η in (53) and c in (52), for any $k, r, j \geq 0$,

$$\begin{aligned} \mathbb{P}[\text{len}(\mathbf{D}_{k+1}) \geq r + j \mid \mathbf{C}_{\leq S_k}, M_k = j] &\leq \frac{\mathbb{P}[\text{len}(\mathbf{D}_{k+1}) \geq r + j \mid \mathbf{C}_{\leq S_k}]}{\mathbb{P}[M_k = j]} \\ &= \frac{\mathbb{P}[\text{len}(\mathbf{D}_{k+1}) \geq r + j \mid \mathcal{F}_{S_k}]}{\mathbb{P}[M_k = j]} \leq e^{-cr}. \end{aligned} \quad (54)$$

The equality above comes from the fact that the additional conditioning on $\mathcal{F}_k$ has no bearing on $\mathbf{D}_{k+1}$. It follows that the variables $(\text{len}(\mathbf{D}_{k+1}) - M_k)_{k \geq 0}$ may be bounded by i.i.d. geometric variable $(\text{ExtraLen}_k)_{k \geq 0}$ of some universal parameter, which are also independent of the variables $(M_k)_{k \geq 0}$.

Renewal structure. We now argue that a time t such that $Y_t = 1$ is indeed a renewal time. Fix t such that $t = S_k$ and $Y_t = 1$. Then, under $\mathcal{F}_t$, the variables $(M_{k+j})_{j \geq 0}$ are independent geometric variables conditioned on $M_{k+j} \leq j$ for each j . In particular, they are independent of $\mathcal{F}_t$, and therefore so is the sequence $(Y_s)_{s \geq t}$.

Additionally, $\mathbf{D}_{k+1}$ only depends on $\mathcal{F}_t$ via its link and $\mathbf{D}_{k+j}$ is independent of $\mathbf{C}_{\leq S_{k+j}-j-1}$ and in particular of $\mathbf{C}_{\leq S_{k+1}}$ (here we use that $S_{j+1} - S_j \geq 2$ for all j). It follows that $\mathbf{C}_{\geq S_{k+1}/2}$ is independent of $\mathcal{F}_t$. As a consequence $(X_{k+j} - X_k)_{j \geq 1}$ is independent of $\mathbf{C}_{\leq k}$.

The law $\mathcal{L}$ appearing in the definition of a KMRP is explicitly constructed when sampling $\tilde{\mathbf{C}}_{\geq s}$ in (50) and (51) under the unconditioned measure ϕ .

Exponential tails and mass-gap. We now discuss the exponential tails of (10) and (11); those of (12) and (13) will be discussed below.

Fix $\ell \geq 1$, condition on $\mathcal{F}_{T_\ell}$ so that $T_\ell < \infty$ and let $k \geq \ell$ be the index such that $S_k = T_\ell$. Due to our construction, the whole of $\mathbf{C}_{\geq T_\ell + \frac{1}{2}}$ is contained in the cone $\mathcal{Y} + (LS_k, X_{S_k})$, which implies (11). We turn to (10).

Define $\kappa > k$ to be the first index for which $M_{\kappa+j} \leq j$ for all $j \geq 0$. In other words, κ is such that $T_{\ell+1} = S_\kappa$. Then, for $m \geq 1$,

$$\mathbb{P}[X_{m+T_\ell} \neq \dagger \text{ but } T_{\ell+1} - T_\ell > m \mid \mathcal{F}_{T_\ell}] = \mathbb{P}\left[\sum_{j=k+1}^{\kappa} \text{len}(\mathbf{D}_j) > m \mid \mathcal{F}_{T_\ell}\right]. \quad (55)$$

⁷That is, variables M_k with $\mathbb{P}[M_k \geq j] = e^{-\eta j}$ for $j \geq 0$.

We will now argue that $\kappa - k$ has exponential tails under $\mathbb{P}[\cdot | \mathcal{F}_{T_\ell}]$ and that the quantity above is exponentially decreasing in m .

The tails for κ are relatively standard and appear in [OV18]. We give a full proof for completeness.

We say a memory variable M_i reaches an index $j < i$ if $i - M_i \leq j$. With this formulation, $\kappa > k$ is the first index such that none of the M_i with $i \geq \kappa$ reaches indices strictly smaller than κ .

Set $J_0 = k$ and, for $i \geq 0$, let

$$J_{i+1} = \max\{j > J_i : j - M_j \leq J_i\}$$

to be the index of the last memory variable reaching J_i . The above is well defined up to the first i for which the set is empty. Then $\kappa = J_i + 1$ and we formally define $J_{i+1} = \emptyset$.

Notice that, for all $i \geq 0$

$$\mathbb{P}[J_{i+1} = \emptyset | \mathcal{F}_{T_\ell}, J_1, \dots, J_i] = \prod_{j \geq 1} \mathbb{P}[M_{j+J_i} < j | \mathcal{F}_{T_\ell}, J_1, \dots, J_i] > 0.$$

Indeed, the variables $(M_{j+J_i})_{j \geq 1}$ under the measure above are independent geometrics, each conditioned not to reach J_{i-1} (nor $J_0 - 1$). As such the probability above may be computed explicitly and be shown to be uniformly positive.

The same computation shows that

$$\mathbb{P}[J_{i+1} - J_i \geq m | \mathcal{F}_{T_\ell}, J_1, \dots, J_i] \leq \sum_{j \geq m} \mathbb{P}[M_{j+J_i} \geq j | \mathcal{F}_{T_\ell}, J_1, \dots, J_i] \leq e^{-\eta m}.$$

Thus, $\kappa - k - 1$ is the sum of variables $(J_{i+1} - J_i)_{i \geq 0}$ which have exponential tails, up to the first time $J_{i+1} = \emptyset$. It follows that $\kappa - k$ also has exponential tails.

Conditionally on $J_1, J_2, \dots$, the variables $(M_j)_{j > k}$ are independent, but have certain conditionings. Indeed, each variable M_{J_i} is conditioned to be larger than $J_i - J_{i-1}$, but strictly smaller than $J_i - J_{i-2}$, while all other memory variables are bounded above. We conclude that, under $\mathbb{P}[\cdot | \mathcal{F}_{T_\ell}]$,

$$\sum_{j=k+1}^{\kappa} M_{j-1} - (\kappa - k) = \sum_{j=k+1}^{\kappa} M_{j-1} - \sum_{i \geq 1} (J_i - J_{i-1}) \quad (56)$$

has exponential tails. Indeed, the above is the sum of $(\kappa - k)$ random variables which may be bounded by i.i.d. geometrics independent of $(\kappa - k)$.

Finally, by (54), we conclude that, under $\mathbb{P}[\cdot | \mathcal{F}_{T_\ell}]$,

$$\sum_{j=k+1}^{\kappa} (\text{len}(\mathcal{D}_j) - M_{j-1}) \quad (57)$$

also has exponential tails.

Inserting the exponential tails for $\kappa - k$, (56) and (57) into (55) we conclude the existence of a universal constant $c > 0$ such that

$$\mathbb{P}\left[\sum_{j=k+1}^{\kappa} \text{len}(\mathbf{D}_j) > m \mid \mathcal{F}_{T_\ell}\right] \leq e^{-cm} \quad \text{for all } m \geq 1.$$

Initial step. It remains to prove (15) as well as the exponential bounds (12) and (13) for the initial step. It is a direct consequence of (4)

$$\phi_p[X_1 \neq \dagger] = \phi_p[0 \leftrightarrow \partial\mathcal{H}_{\geq 1}] \leq \pi_1(L(p)).$$

Conversely, (RSW) and (4) allow us to prove that

$$\phi_p[X_1 \neq \dagger \text{ and } N_1 = 1] \geq \pi_1(L(p)).$$

Then, by the properties of the memory variables,

$$\phi_p[T_1 < \infty] \geq \phi_p[T_1 = 1] \geq \phi_p[X_1 \neq \dagger \text{ and } N_1 = 1] \geq \pi_1(L(p)).$$

This proves (15).

Proposition 3.7 implies an exponential tail for $\min\{j \geq 0 : \mathbf{C} \subset Y - (Lj, 0)\}$ when $\mathbf{C}$ is sampled under $\phi_p[\cdot \mid 0 \leftrightarrow \partial\mathcal{H}_{\geq n}]$, whence (13) follows readily.

Finally, (32) states that, under $\phi_p[\cdot \mid 0 \leftrightarrow \partial\mathcal{H}_{\geq n}]$, S_1 has a uniform exponential tail. The rest of the proof of (12) and is identical to that of (10) above.

Killing rate and step variance. From (RSW) we conclude directly that, for any time T_k with $k \geq 1$,

$$\phi[0 \leftrightarrow \mathcal{H}_{\geq T_k+1} \mid \mathcal{F}_{T_k}] \leq 1 - c < 1,$$

for some universal constant $c > 0$. This implies that the killing rate is bounded away from 0 uniformly in all parameters. Furthermore, a quick analysis shows that

$$\phi[T_{k+1} = T_k + 2 \text{ and } X_{T_{k+1}} \geq X_{T_k} + 1 \mid \mathcal{F}_{T_k}] \geq 1,$$

which implies a uniform lower bound on the X -step variance and a uniform upper bound on the killing rate.

Aperiodicity. Recall that our construction formally prevents us to have $T_{k+1} = T_k + 1$. However, direct (RSW) constructions prove that

$$\phi[T_{k+1} = T_k + 2 \mid \mathcal{F}_{T_k}] \geq 1 \quad \text{and} \quad \phi[T_{k+1} = T_k + 3 \mid \mathcal{F}_{T_k}] \geq 1,$$

which implies the aperiodicity of the process. $\square$

4 Local limit theorem for the Markov renewal process

This section contains general results about killed Markov renewal processes formulated for the process $(X_t, Y_t)_{t \geq 1}$ of Theorem 3.2. While the details of this specific process are not important, we find it easier to formulate the results in this context.

4.1 Probability of hitting a half-space

First, we show that the probability that a KMRP with a mass-gap survives for n steps behaves like a pure exponential. We highlight that this result is purely stated for KMRPs. In particular, its proof is not based on percolation arguments but rather on analytic ones, and is somewhat orthogonal to the rest of the paper.

Theorem 4.1. *Let $(X_t, Y_t)_{t \geq 1}$ be an aperiodic KMRP with a mass-gap and killing rate $\kappa > 0$. Then there exists $\zeta \leq \kappa^{-1}$ such that*

$$\mathbb{P}[X_n \neq \dagger] \asymp \mathbb{P}[X_1 \neq \dagger] e^{-n/\zeta} \quad \text{for all } n \geq 1, \quad (58)$$

where the constants in the equivalence and ζ are bounded away from 0 and ∞ depending only on the killing rate and the mass-gap.

Expressed in terms of the random-cluster model, Theorem 4.1 states the following.

Corollary 4.2. *For any $\vec{w} \in \mathbb{S}^1$ and $p < p_c$, there exists $\zeta(p, \vec{w}) > 0$ such that*

$$\phi_p[0 \leftrightarrow \mathcal{H}_{\geq n}^{\vec{w}}] \asymp \pi_1(L(p)) e^{-\frac{n}{\zeta(p, \vec{w})L(p)}}, \quad (59)$$

uniformly in $n \geq L(p)$, p and $\vec{w}$. Moreover, ζ is bounded uniformly away from 0 and ∞ .

The above does not imply that $\xi_p(\vec{w}) = \zeta(p, \vec{w})L(p)$, since $\xi_p(\vec{w})$ is not defined in terms of hitting a half-space, but rather a specific point. We will see how to deduce $\xi_p(\vec{v})$ in the next section.

Proof of Theorem 4.1. When studying $\mathbb{P}[X_n \neq \dagger]$, the first step needs to be treated with particular care. Indeed,

$$\begin{aligned} \mathbb{P}[X_n \neq \dagger] &= \sum_{1 \leq k \leq n} \mathbb{P}[T_1 = k] \mathbb{P}[X_n \neq \dagger \mid \mathcal{F}_k, T_1 = k] + \mathbb{P}[T_1 > n \text{ but } X_n \neq \dagger] \quad (60) \\ &= \sum_{1 \leq k \leq n} \mathbb{P}[T_1 = k] \mathbb{P}[X_n \neq \dagger \mid \mathcal{F}_k, T_1 = k] + O(e^{-cn} \mathbb{P}[X_n \neq \dagger]), \end{aligned}$$

where the second line is due to the mass-gap $c > 0$. Furthermore, also due to the mass-gap, $\mathbb{P}[T_1 = k] = O(e^{-ck} \mathbb{P}[X_k \neq \dagger])$.

The main analysis is dedicated to

$$s_{n-k} := \mathbb{P}[X_n \neq \dagger \mid \mathcal{F}_k, T_1 = k],$$

which is a quantity that depends only on $n - k$ and the law $\mathcal{L}$ of Definition 3.1. For $n \geq 1$, define the auxiliary quantities

$$\begin{aligned} a_n &:= \mathbb{P}[T_2 - T_1 = n \mid \mathcal{F}_{T_1}, T_1 < \infty] \\ c_n &:= \mathbb{P}[X_{n+T_1} \neq \dagger \text{ but } T_2 > n + T_1 \mid \mathcal{F}_{T_1}, T_1 < \infty] \quad \text{and} \\ p_n &:= \frac{1}{\kappa} \mathbb{P}[\exists k \geq 0, T_k - T_1 = n \text{ and } T_{k+1} = \infty \mid \mathcal{F}_{T_1}, T_1 < \infty] \\ &= \mathbb{P}[\exists k \geq 0, T_k - T_1 = n \mid \mathcal{F}_{T_1}, T_1 < \infty], \end{aligned}$$

where we recall that κ denotes the killing rate. The equality is due to the renewal property.

By analysing the value of the last renewal time, we find

$$s_n = \sum_{1 \leq k \leq n} p_k c_{n-k} + c_n = \sum_{1 \leq k \leq n} p_k c_{n-k} + O(e^{-cn} s_n). \quad (61)$$

due to the mass-gap.

We will prove that there exists $\zeta > 0$ such that

$$p_n \asymp e^{-n/\zeta} \quad (62)$$

for n large enough. Before doing so, observe that (62) combined with (60), (61) and the consequences of the mass-gap, yields

$$s_n \asymp e^{-n/\zeta} \quad \text{and} \quad \mathbb{P}[X_n \neq \dagger] \asymp \mathbb{P}[X_1 \neq \dagger] e^{-n/\zeta}. \quad (63)$$

The rest of the proof is dedicated to (62). By analysing the first renewal time, we find

$$p_n = \sum_{k=1}^n a_k p_{n-k} + \mathbf{1}_{n=0}. \quad (64)$$

Denote by P and A the generating series of the sequences (p_n) and (a_n) , respectively. Note that these are both power series with positive coefficients. Writing R_p and R_a for their respective radii of convergence, the mass-gap states that $R_p < R_a$. From (64) we deduce the “killed renewal equation”

$$P(z) = \frac{1}{1 - A(z)}. \quad (65)$$

The equality holds in the whole disk of convergence of P .

Observe that A is a power series with positive coefficients that satisfies $A(1) = 1 - \kappa < 1$. Also observe that, due to (65), it is the case that $A(R_p) = 1$.

Write $\mathbb{D}_{R_p} \subset \mathbb{C}$ for the open disc centered at 0 of radius R_p . We now are going to argue that R_p is the only 0 in $\mathbb{D}_{R_p}$ of the series $A - 1$ and that it is simple. Indeed, it is easy to see that if $|z| < 1$, then

$$\left| \sum_{n \geq 1} a_n (R_p z)^n \right| \leq \sum_{n \geq 1} a_n (R_p |z|)^n < 1.$$

Furthermore, we argue that the aperiodicity of the process implies that $A - 1$ does not have an additional 0 on $\partial\mathbb{D}_{R_p} \setminus \{R_p\}$. Indeed, if it were the case, then there would exist some $\theta \in (0, 2\pi)$ such that $\sum_{k=0}^{+\infty} a_k R_p^k e^{ik\theta} = 1$. By the equality case in the triangular inequality, one would then have that for all the $k \in \text{Supp}(T_2 - T_1)$, the $e^{ik\theta}$ are aligned; this contradicts the aperiodicity of T . Finally, since $R_p < R_a$, A has a positive derivative at R_p , which implies that R_p is a simple zero of $A - 1$.

Summarising the above, we proved that the function

$$g(z) := \frac{1 - A(z)}{R_p - z}$$

does not vanish on $\mathbb{D}_{R_p}$. We wish to apply Wiener's $1/f$ theorem (see [Zyg02, Theorem 5.2]) to the function g . To that end, we need to check that the corresponding series is summable at R_p . It is a simple observation that, for all $|z| < R_p$,

$$g(z) = \sum_{n=0}^{+\infty} \left(\frac{z}{R_p}\right)^n \left(1 - \sum_{k=1}^n R_p^k a_k\right) = \sum_{n=0}^{+\infty} \left(\frac{z}{R_p}\right)^n \left(\sum_{k=n+1}^{+\infty} R_p^k a_k\right).$$

Due to the positivity of the terms a_k , when inserting $z = R_p$ in the above, we find

$$\sum_{n=0}^{+\infty} \sum_{k=n+1}^{+\infty} R_p^k a_k = \sum_{k=1}^{+\infty} R_p^k k a_k = R_p A'(R_p) < \infty.$$

By Wiener's $1/f$ theorem [Zyg02, Theorem 5.2], $1/g$ may be expanded as a power series $\sum_n b_n z^n$ in $\mathbb{D}_{R_p}$ and this series is absolutely summable at the point $z = R_p$. By (65), for any $z \in \mathbb{D}_{R_p}$,

$$\frac{1}{g(z)} = (R_a - z)P(z) = \sum_{n \geq 0} z^n (R_p p_n - p_{n-1}).$$

As the above expression can be evaluated at $z = R_p$, by Abel's radial theorem

$$\lim_{z \rightarrow R_p^-} \frac{1}{g(z)} = \sum_{n \geq 0} (R_p^{n+1} p_n - R_p^n p_{n-1}) = \lim_{n \rightarrow \infty} R_p^{n+1} p_n. \quad (66)$$

Notice however that

$$\lim_{z \rightarrow R_p^-} \frac{1}{g(z)} = \lim_{z \rightarrow R_p^-} \frac{R_p - z}{1 - A(z)} = \frac{1}{A'(R_p)} > 0.$$

The last two displays allow us to conclude that $p_n = \frac{1}{R_p A'(R_p)} R_p^{-n} (1 + o(1))$ when $n \rightarrow \infty$. The explicit value of $\lim_n p_n R_p^n$ is irrelevant for us; one simply needs to observe that it is strictly positive and finite.

Setting $\zeta = (\log R_p)^{-1}$, we obtain that $p_n \asymp e^{-n/\zeta}$ as required. $\square$

Proof of Corollary 4.2. Fix $\vec{w}$ and $p < p_c$. All constants below will be uniform in $\vec{w}$, p and $N \geq L(p)$. We have

$$\phi_p[X_{\lfloor N/L(p)+1 \rfloor} \neq \dagger] \leq \phi_p[0 \leftrightarrow \mathcal{H}_{\geq N}^{\vec{w}}] \leq \phi_p[X_{\lfloor N/L(p) \rfloor} \neq \dagger].$$

Apply Theorem 4.1 to conclude that

$$\phi_p[0 \leftrightarrow \mathcal{H}_{\geq N}^{\vec{w}}] \asymp \phi_p[X_1 \neq \dagger] \exp\left(-\frac{n}{L(p)\zeta(p, \vec{w})}\right)$$

for some $\zeta(p, \vec{w})$. Moreover, by (15), $\phi_p[X_1 \neq \dagger] \asymp \pi_1(L(p))$, which concludes the proof of the corollary. $\square$

Remark 4.3. This computation allows us to identify the law of the inter-renewal times when the process is conditioned on surviving at least n steps. Indeed, observe that for any $\ell \in \{1, \dots, n-1\}$,

$$\begin{aligned} \mathbb{P}[T_2 - T_1 = \ell \mid \exists k \geq 0, T_k - T_1 = n, \mathcal{F}_{T_1}, T_1 < \infty] &= \frac{s_{n-\ell}}{s_n} \mathbb{P}[T_2 - T_1 = \ell \mid T_1 < \infty] \\ &\asymp e^{\ell/\zeta} \mathbb{P}[T_2 - T_1 = \ell \mid T_1 < \infty]. \end{aligned}$$

We conclude that when n tends to infinity, the law of the inter-renewal times converges towards an exponentially tilted sample of $\mathcal{L}(T_2 - T_1 = \cdot \mid T_1 < \infty)$. The first and last steps of the process, however, have a different distribution. The same computation applied to finite-dimensional marginals yields that the joint distribution of the different inter-renewal times (except the first and last ones) converges towards i.i.d. samples of the exponential tilt of $\mathcal{L}(T_2 - T_1 = \cdot \mid T_1 < \infty)$.

Finally, the distribution of the X -steps also converges towards i.i.d. samples of some probability distribution. We call $\mathcal{L}_{\text{irred}}$ the limit of the law of $(T_2 - T_1, X_2 - X_1)$ under the conditionings $\{X_n \neq \dagger\}$ with $n \rightarrow \infty$. The mass gap implies that this distribution has exponential tails both for the T and the X -steps.

4.2 Endpoint concentration when conditioned on survival

We now prove a local limit theorem for X_n under the conditioning $X_n \neq \dagger$. Let $g_\sigma(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-x^2/2\sigma^2}$ be the Gaussian density with variance σ^2 .

Proposition 4.4. Fix $\vec{w} \in \mathbb{S}^1$ and $p < p_c$. There exists $\mu = \mu(p, \vec{w})$ and $\sigma = \sigma(p, \vec{w})$ such that, for any $k \in \mathbb{Z}$, when $n \rightarrow \infty$,

$$\left| \sqrt{n} \phi_p[\lfloor X_n/L(p) \rfloor = \lfloor n \cdot \mu \rfloor + k \mid X_n \neq \dagger] - g_\sigma\left(\frac{k}{\sqrt{n}}\right) \right| \rightarrow 0. \quad (67)$$

with the asymptotics being uniform in p and $\vec{w}$. Furthermore $|\mu(p, \vec{w})|$ is uniformly bounded away from ∞ , and $\sigma(p, \vec{w})$ is uniformly bounded away from 0 and ∞ .

Proof. We argued in Remark 4.3 that when conditioned on $X_n \neq \dagger$, the process has steps sampled from some i.i.d. distribution. It thus converges towards a so-called *compound Markov process*, using the terminology of [Bor22]. Local limit-type theorems for compound

Markov processes are classical; for instance, Proposition 4.4 directly follows from [Bor22, Theorem 2.1.2].

The uniform bounds on the $\mu(p, \vec{w})$ and $\sigma(p, \vec{w})$ follow from those on the mass-gap, mean, variance and killing rate of the KMRP in Theorem 3.2. $\square$

Remark 4.5. Other properties typical of random walks may be extended to the process $(X_n)_n$ such as the existence of a uniform constant $C > 0$ such that for any $k \in \mathbb{Z}$,

$$\phi_p \left[\lfloor X_n/L(p) \rfloor = \lfloor n \cdot \mu \rfloor + k \mid X_n \neq \dagger \right] \leq C \phi_p \left[\lfloor X_n/L(p) \rfloor = \lfloor n \cdot \mu \rfloor \mid X_n \neq \dagger \right] \quad (68)$$

and a large deviation estimate

$$\phi_p \left[\left| X_n/L(p) - \lfloor n \cdot \mu \rfloor \right| \geq \alpha n \mid X_n \neq \dagger \right] = e^{-I(\alpha)n+o(n)} \quad (69)$$

where I is a differentiable function on an interval $(-\varepsilon, \varepsilon)$ for some $\varepsilon > 0$ with $I(0) = 0$ and $I(\alpha) > 0$ for $\alpha \neq 0$. All constants appearing in the above are uniform in n , $\vec{w}$ and $p < p_c$. The rate function I does depend on the direction $\vec{w}$. For proofs of such statements, see [Bor22].

As a consequence, we deduce a preliminary form of the OZ-formula. Recall the quantity $\zeta(p, \vec{w})$ defined in Corollary 4.2.

Corollary 4.6. Fix $\vec{w} \in \mathbb{S}^1$ and $p < p_c$. There exists $\mu = \mu(p, \vec{w})$ such that

$$\phi_p[0 \leftrightarrow x] \asymp \frac{\pi_1(L(p))^2}{\sqrt{n}} e^{-\frac{n}{\zeta(p, \vec{w})}}, \quad (70)$$

uniformly in p , $\vec{w}$, $n \geq 1$ and any $x \in \mathbb{Z}^2 \setminus \Lambda_{L(p)}$ with $\|x - nL(p)(\vec{w} + \mu \cdot \vec{w}^\perp)\| \leq L(p)$.

Remark 4.7. The proof below also allows one to deduce from (69) a large deviation estimate for the hitting position

$$\phi_p \left[0 \leftrightarrow \lfloor nL(p)(\vec{w} + (\mu(p, \vec{w}) + \alpha) \cdot \vec{w}^\perp) \rfloor \mid 0 \leftrightarrow \mathcal{H}_{\geq n}^{\vec{w}} \right] = \pi_1(L(p))^2 \cdot e^{-I(\alpha)n+o(n)} \quad (71)$$

for any α close enough to 0 and some differentiable rate function I with $I(0) = 0$ and $I(\alpha) > 0$ for $\alpha \neq 0$.

Proof of Corollary 4.6. Fix p , $\vec{w} \in \mathbb{S}^1$ and $x \in \Lambda_{L(p)/2}(nL(p)(\vec{w} + \mu(p, \vec{w}) \cdot \vec{w}^\perp))$ for some $n \geq 1$. For $n < 3$ it was proved in [DCM22] that $\phi_p[0 \leftrightarrow x] \asymp \pi_1(L(p))^2$. We will henceforth assume that $n \geq 3$. The upper and lower bound on $\phi_p[0 \leftrightarrow x]$ will be treated differently. All constants below are uniform in p and $\vec{w}$.

We start with the lower bound. Write Γ for the top-most path connecting $C_{\leq 1}$ to $\mathcal{H}_{\geq n-1}$. This path is explorable starting from the top of $C_{\leq 1} \cap \partial\mathcal{H}_{\leq 1}$ and X_{n-1} is its endpoint on $\mathcal{H}_{\geq n-1}$. Conditioning on a realisation $\Gamma = \gamma$ with $\lfloor X_{n-1}/L(p) \rfloor = \lfloor (n-1) \cdot \mu(p, \vec{w}) \rfloor$, by an applying RSW-type construction in $\Lambda_{2L(p)}(x)$ (see Figure 8) we find that

$$\phi_p[0 \leftrightarrow x \mid \Gamma = \gamma] \geq c \phi_{\Lambda_{L(p)}, p}^0[x \leftrightarrow \partial\Lambda_{L(p)/2}] \asymp \pi_1(L(p)),$$




Figure 8: If we assume that $X_{n-1} \neq \dagger$ and that it is at a distance of order $L(p)$ from x , then connecting x to 0 requires an additional cost of $\pi_1(L(p))$. The picture appears skewed since the horizontal direction is $\vec{w} + \mu(p, \vec{w}) \cdot \vec{w}^\perp$, while the hyperplanes $\partial\mathcal{H}_{\leq t}$ are orthogonal to $\vec{w}$.

where (RSW) and (4) state that all constants may be chosen uniform in $\vec{w}$ and p . Summing over γ as above we conclude that

$$\phi_p[0 \leftrightarrow x] \geq \pi_1(L(p)) \phi_p[\lfloor X_{n-1}/L(p) \rfloor] = \lfloor (n-1) \cdot \mu_{\vec{w}} \rfloor \asymp \frac{\pi_1(L(p))^2}{\sqrt{n}} e^{-\frac{n}{\zeta(p, \vec{w})}}, \quad (72)$$

where we used (67) and (59) in the last equivalence.

We turn to the upper bound, for which we decompose

$$\begin{aligned} \phi_p[0 \leftrightarrow x] &\leq \phi_p[0 \leftrightarrow \Lambda_{L(p)}(x) \text{ and } x \leftrightarrow \partial\Lambda_{L(p)}(x)] \\ &\leq \phi_p[0 \leftrightarrow \Lambda_{L(p)}(x)] \phi_{\Lambda_{L(p)}, p}^1[0 \leftrightarrow \partial\Lambda_{L(p)}(x)], \end{aligned} \quad (73)$$

where the second inequality is due to the fact that the two events are measurable in terms of what happens outside and inside of $\Lambda_{L(p)}(x)$, respectively. The second probability may be bounded above by a universal multiple $\pi_1(L(p))$ due to (4).

To bound the first quantity, decompose depending on the value of the last renewal time T_ℓ before n , and keep in mind that cluster after this renewal time is contained in $\mathcal{Y} + (L(p)T_k, X_k)$. Thus, applying (10), we find

$$\begin{aligned} \phi_p[0 \leftrightarrow \Lambda_{L(p)}(x)] &\leq \sum_{k \geq 0} \phi_p[|X_{n-k} - nL(p)\mu\vec{w}^\perp| \leq \alpha(k+1), Y_{n-k} = 1 \text{ and } Y_{n-k+1} = \dots = Y_n = 0] \\ &\leq \sum_{k \geq 0} \phi_p[|X_{n-k} - nL(p)\mu\vec{w}^\perp| \leq \alpha(k+1)] e^{-n(c+1/\zeta(p, \vec{w}))} \\ &\leq \alpha\pi_1(L(p)) e^{-\frac{n}{\zeta(p, \vec{w})}} \cdot \sum_{k \geq 0} \frac{k+1}{\sqrt{n-k}} e^{-cn} \\ &\leq \frac{\pi_1(L(p))}{\sqrt{n}} e^{-\frac{n}{\zeta(p, \vec{w})}}, \end{aligned}$$

with the before-last inequality due to (68). Inserting these bounds in (73) yields the converse of (72). $\square$




Figure 9: *Left:* when $\mu(\vec{w})$ and δ have same sign, the point $x = n(\vec{w} + \mu(\vec{w})\vec{w}^\perp)$ is contained in $\mathcal{H}_{\geq n \cos \delta}^{\rho_\delta(\vec{w})}$. *Right:* when $\mu(\vec{w}) > 0$ and $\delta < 0$, the same point x is contained in $\mathcal{H}_{\geq n \cdot \frac{\cos(\theta-\delta)}{\cos \theta}}^{\rho_\delta(\vec{w})}$. In both images, the grey direction is given by $\vec{v}(\vec{w})$.

4.3 Invertibility of the drift and proof of Theorem 1.1

In light of the previous section, write

$$\vec{v}(\vec{w}) := \frac{\vec{w} + \mu(p, \vec{w}) \cdot \vec{w}^\perp}{\|\vec{w} + \mu(p, \vec{w}) \cdot \vec{w}^\perp\|} \in \mathbb{S}^1.$$

Taking $n \rightarrow \infty$ in (70), we conclude that

$$\xi_p(\vec{v}(\vec{w})) = L(p)\zeta(p, \vec{w})\|\vec{w} + \mu(p, \vec{w}) \cdot \vec{w}^\perp\| = \frac{L(p)\zeta(p, \vec{w})}{\langle \vec{v}(\vec{w}), \vec{w} \rangle}. \quad (74)$$

The following proposition shows that $\vec{w} \mapsto \vec{v}(\vec{w})$ is a one-to-one function from $\mathbb{S}^1$ to itself. Theorem 1.1 will be a consequence of the surjectivity of this map. Moreover, we shall see in the next section that its injectivity has interesting consequences regarding strict convexity properties of the Wulff shape. Henceforth the choice of $\vec{w}^\perp$ is important, and we will assume that it is the rotation by $\pi/2$ of $\vec{w}$ in the positive direction.

Proposition 4.8. *For any $\vec{v} \in \mathbb{S}^1$, there exists exactly one $\vec{w} \in \mathbb{S}^1$ so that*

$$\vec{v}(\vec{w}) = \vec{v}.$$

Proof. The lemma decomposes in two disjoint statements, the surjectivity and injectivity of the function $\vec{v}$. We start with the former.

We will prove that the function $\vec{w} \mapsto \vec{v}(\vec{w})$ is continuous. We start off by proving that $\vec{w} \mapsto \zeta(p, \vec{w})$ is continuous. This is a simple geometric construction; see Figure 9 for an illustration.

Fix some $\vec{w} \in \mathbb{S}_1$ and assume for simplicity that $\mu(p, \vec{w}) \geq 0$. Write $\rho_\delta(\cdot)$ be the rotation by an angle δ . For $\delta > 0$, a connection between 0 and $x = nL(p)(\vec{w} + \mu \cdot \vec{w}^\perp)$ intersects $\mathcal{H}_{\geq n \cos \delta}^{\rho_\delta(\vec{w})}$. Together with (70) and (59) this implies that

$$\zeta(p, \rho_\delta(\vec{w})) \geq \cos \delta \cdot \zeta(p, \vec{w}).$$

For $\delta < 0$, write $\theta = \arctan(\mu(p, \vec{w}))$ – note that θ is uniformly bounded away from ∞ by Proposition 4.4. Then, a connection between 0 and $x = nL(p)(\vec{w} + \mu \cdot \vec{w}^\perp)$ ensures that the cluster of 0 intersects $\mathcal{H}_{\geq n \frac{\cos(\theta+\delta)}{\cos \theta}}^{\rho_\delta(\vec{w})}$. Thus

$$\zeta(p, \rho_\delta(\vec{w})) \geq \frac{\cos(\theta+\delta)}{\cos(\theta)} \cdot \zeta(p, \vec{w}).$$

Using the two displays above, and their versions with the roles of $\vec{w}$ and $\rho_\delta(\vec{w})$ inverted, we conclude that $\vec{w} \mapsto \zeta(p, \vec{w})$ is continuous.

The reasoning above shows more than just the continuity of ζ . Indeed, due to the specific definition of x we conclude that, for any $\varepsilon > 0$ and δ small enough

$$\phi_p \left[0 \leftrightarrow \frac{n}{\langle \vec{v}(\vec{w}), \rho_\delta(\vec{w}) \rangle} \vec{v}(\vec{w}) \mid 0 \leftrightarrow \mathcal{H}_{\geq n}^{\rho_\delta(\vec{w})} \right] \geq \exp(-\varepsilon n).$$

Together with the large deviation estimate (71) and the continuity of the rate function, this implies that $\vec{v}(\vec{w})$ may be rendered arbitrarily close to $\vec{v}(\rho_\delta(\vec{w}))$, provided δ is small enough. Thus we conclude the proof of the continuity of $\vec{w} \mapsto \vec{v}(\vec{w})$.

To conclude the surjectivity of $\vec{w} \mapsto \vec{v}(\vec{w})$ observe that, for $\vec{w}$ a coordinate vector, $\mu(p, \vec{w}) = 0$ due to the symmetry of the lattice. It follows that $\vec{v}(\vec{w}) = \vec{w}$ in this case. Furthermore, symmetry also implies that $\vec{v}$ commutes with any reflexion around a lattice axis. By the intermediate value theorem and the precedent observation, we conclude that the image of the function $\vec{w} \mapsto \vec{v}(\vec{w})$ covers the whole unit circle, as claimed.

We turn to the proof of the injectivity of $\vec{w} \mapsto \vec{v}(\vec{w})$. We will prove that for any $\vec{w}$ and $\delta > 0$ small enough, $\vec{v}(\rho_\delta(\vec{w}))$ is a small, strictly positive rotation of $\vec{v}(\vec{w})$.

Fix $\delta \in [-\pi/4, \pi/4]$. Applying the large deviation estimate (71) we conclude that, for all α small enough (independent of δ),

$$\begin{aligned} \exp \left(- (I(\alpha) + \zeta(\vec{w})^{-1})(n + o(n)) \right) &\asymp \phi_p \left[0 \leftrightarrow \lfloor nL(p)(\vec{w} + (\mu(p, \vec{w}) + \alpha) \cdot \vec{w}^\perp) \rfloor \right] \\ &\geq \phi_p \left[0 \leftrightarrow \mathcal{H}_{\geq n \langle \vec{w} + (\mu(p, \vec{w}) + \alpha) \cdot \vec{w}^\perp, \vec{\rho}_\delta(\vec{w}) \rangle}^{\vec{\rho}_\delta(\vec{w})} \right] \\ &\asymp \exp \left(- \frac{\cos \delta + (\mu(p, \vec{w}) + \alpha) \sin \delta}{\zeta(\rho_\delta(\vec{w}))} (n + o(n)) \right) \end{aligned} \quad (75)$$

where I is the large deviation rate function associated to the walk in direction $\vec{w}$. The constants in the asymptotics above depend on p , but not on n . Furthermore, we have asymptotic equality in (75) if and only if

$$\begin{aligned} \vec{w} + (\mu(p, \vec{w}) + \alpha) \cdot \vec{w}^\perp &= \rho_\delta(\vec{w}) + \mu(p, \rho_\delta(\vec{w})) \cdot \delta(\vec{w})^\perp, \\ \text{which translates to} \quad \vec{v}(\rho_\delta(\vec{w})) &= \frac{\vec{w} + (\mu(p, \vec{w}) + \alpha) \cdot \vec{w}^\perp}{\|\vec{w} + (\mu(p, \vec{w}) + \alpha) \cdot \vec{w}^\perp\|}. \end{aligned} \quad (76)$$

By the continuity of $\vec{w} \mapsto \vec{v}(\vec{w})$, if δ is small enough, the equality above holds for α in the range of validity of (75). We conclude that (76) holds for the unique α minimising

$$\frac{I(\alpha) + \zeta(\vec{w})^{-1}}{\cos \delta + (\mu(p, \vec{w}) + \alpha) \sin \delta}.$$

Using the fact that I is an even, convex function that is differentiable at 0 with $I(0) = I'(0) = 0$, we conclude that the minimiser occurs for $\alpha_{\min} = \alpha_{\min}(\delta) > 0$ within the range of (75).

Recalling that $\vec{v}(\vec{w})$ is the unit vector in the direction $\vec{w} + \mu(p, \vec{w})\vec{w}^\perp$ and that $\vec{v}(\rho_\delta(\vec{w}))$ is the one in the direction $\vec{w} + \mu(p, \vec{w}) \cdot \vec{w}^\perp + \alpha_{\min}(\delta) \cdot \vec{w}^\perp$, we conclude that the latter is indeed a small positive rotation of the former.

This readily implies the injectivity of the function $\vec{w} \mapsto \vec{v}(\vec{w})$. Indeed, consider its expression in terms of angle arguments

$$\begin{aligned}\varphi : [0, \pi/2] &\longrightarrow [-\pi, \pi) \\ \theta &\mapsto \arg(\vec{v}(\cos \theta, \sin \theta));\end{aligned}$$

recall that the symmetries of the lattice ensure that φ maps 0 and $\pi/2$ to themselves. Then our analysis says that, for any $\theta \in [0, \pi/2)$, there exists $\delta > 0$ such that $\varphi(\theta') > \varphi(\theta)$ for all $\theta < \theta' \leq \theta + \delta$. Together with the previously proved continuity of φ , this implies that φ is strictly increasing, hence injective. The injectivity extends to the whole range of angles $[0, 2\pi)$ by symmetries. $\square$

We can finally prove Theorem 1.1.

Proof of Theorem 1.1. Fix $q \geq 1$, $\vec{v} \in \mathbb{S}^1$, $p < p_c$ and $n \geq \xi_p(\vec{v})$. Define $\vec{w}$ as a vector such that $\vec{v}(\vec{w}) = \vec{v}$. Then (70) implies that

$$\phi_p[0 \leftrightarrow \lfloor n\vec{v} \rfloor] \asymp \frac{\pi_1(L(p))^2}{\sqrt{m}} e^{-\frac{m}{\zeta(p, \vec{w})}}, \quad (77)$$

where $m = \frac{n}{\langle \vec{v}, \vec{w} \rangle L(p)}$. Recall that the drift $\mu(p, \vec{w})$ is bounded uniformly, and therefore $\langle \vec{v}, \vec{w} \rangle$ is uniformly positive. Thus $m \asymp n/L(p)$ and therefore $\xi_p(\vec{v}) = \langle \vec{v}, \vec{w} \rangle L(p) \zeta(p, \vec{w})$. Inserting these observations in (77) we find

$$\phi_p[0 \leftrightarrow \lfloor n\vec{v} \rfloor] \asymp \pi_1(\xi_p(\vec{v}))^2 \sqrt{\frac{\xi_p(\vec{v})}{n}} e^{-\frac{n}{\xi_p(\vec{v})}},$$

as claimed. $\square$

4.4 Strict convexity of Wulff shape

We first introduce some notation and discuss classical properties. Fix $p < p_c$. For $\vec{w} \in \mathbb{R}^2 \setminus \{0\}$ define

$$(\xi_p^*(\vec{w}))^{-1} = - \lim_{n \rightarrow \infty} \frac{1}{n} \log \phi_p[0 \leftrightarrow \mathcal{H}_{\geq n}^{\vec{w}}],$$

with the existence of the limit being guaranteed by sub-additivity arguments. This definition is reminiscent of the one of the inverse correlation length, with the point-to-point connection being replaced by a point-to-hyperplane connection. For coherence, we set $\xi_p(0) = \xi_p^*(0) = \infty$. Observe that (59) states that, for $\vec{w} \in \mathbb{S}^1$,

$$\xi^*(\vec{w}) = \zeta(p, \vec{w}) L(p). \quad (78)$$

It is classical (see [MS95] or [CIV08] for a more modern exposition) that both $\xi_p(\cdot)^{-1}$ and $\xi_p^*(\cdot)^{-1}$ define norms on $\mathbb{R}^2$. Indeed they are both positive homogeneous, and FKG inequalities imply that they are convex.

Define their unit balls

$$\mathcal{U}_p = \{\vec{v} \in \mathbb{R}^2 : \xi_p(\vec{v}) \geq 1\} \quad \text{and} \quad \mathcal{W}_p = \{\vec{w} \in \mathbb{R}^2 : \xi_p^*(\vec{w}) \geq 1\}.$$

The latter is called the *Wulff shape* (see [BIV00] for an extensive review on the subject).

The goal of this section is to explain how our approach allows to re-prove some known results about $\mathcal{U}_p$ and $\mathcal{W}_p$. The method is different of [CIV08], as it does not rely on the analysis of the regularity properties of the associated Ruelle–Perron–Frobenius operator. Indeed we crucially use the invertibility of the function $\vec{v}$ defined in the last section.

Theorem 4.9. *For any $p < p_c$, $\mathcal{U}_p$ and $\mathcal{W}_p$ are strictly convex bounded sets of $\mathbb{R}^2$, symmetric with respect to the coordinate axis with differentiable boundaries. Moreover, they are convex dual to each other, i.e. they are linked by the following relation:*

$$\mathcal{U}_p = \{\vec{v} \in \mathbb{R}^2 : \langle \vec{v}, \vec{w} \rangle \leq \xi_p^*(\vec{w}), \forall \vec{w} \in \mathbb{S}^1\} \quad \text{and} \quad \mathcal{W}_p = \{\vec{w} \in \mathbb{R}^2 : \langle \vec{w}, \vec{v} \rangle \leq \xi_p(\vec{v}), \forall \vec{v} \in \mathbb{S}^1\}. \quad (79)$$

In the language of convex bodies, $\mathcal{U}_p$ is the convex body with *gauge function* given by ξ_p and the first equality in (79) states that its *support function* is ξ_p^* . The same holds for its dual $\mathcal{W}_p$, with the roles of ξ_p and ξ_p^* reversed; see [Web94] for details. The strict convexity of one shape is equivalent to the fact that the boundary of the other is differentiable [Web94]. The strict convexity of $\mathcal{U}_p$ and $\mathcal{W}_p$ is a consequence of our construction; their duality may be proved directly.

Proof. Classical results on convex duality [Web94, polar duality theorem, p. 238] imply that the two identities of (79) are equivalent. We prove the first one.

It is immediate that 0 satisfies the conditions of the right-hand side terms of (79).

Let $\vec{v} \in \mathbb{R}^2 \setminus \{0\}$ and $\vec{w} \in \mathbb{S}^1$ be a unit vector with $\langle \vec{v}, \vec{w} \rangle > 0$. The hyperplane with normal vector $\vec{w}$ containing the vertex $n\vec{v}$ is given by $\partial \mathcal{H}_{\geq n\langle \vec{v}, \vec{w} \rangle}^{\vec{w}}$. Thus,

$$\phi_p[0 \leftrightarrow n\vec{v}] \leq \phi_p[0 \leftrightarrow \mathcal{H}_{\geq n\langle \vec{v}, \vec{w} \rangle}^{\vec{w}}].$$

Considering the exponential rate of decay of these two quantities as $n \rightarrow \infty$, we find

$$\xi_p(\vec{v}) \langle \vec{v}, \vec{w} \rangle \leq \xi_p^*(\vec{w}). \quad (80)$$

The inequality holds trivially when $\langle \vec{v}, \vec{w} \rangle \leq 0$. Thus

$$\mathcal{U}_p \subset \{\vec{v} \in \mathbb{R}^2 : \langle \vec{v}, \vec{w} \rangle \leq \xi_p^*(\vec{w}), \forall \vec{w} \in \mathbb{R}^2\}.$$

For the converse inclusion, let $\vec{v}$ be a vector such that for any $\vec{w} \in \mathbb{S}^1$, $\langle \vec{v}, \vec{w} \rangle \leq \xi_p^*(\vec{w})$. Lemma 4.8 yields the existence of $\vec{w} \in \mathbb{S}^1$ such that $v(\vec{w}) = \vec{v}/\|\vec{v}\|$. Then (74) and (78) imply that

$$\xi_p(\vec{v}) = \frac{1}{\|\vec{v}\|} \xi_p(v(\vec{w})) = \frac{\xi_p^*(\vec{w})}{\|\vec{v}\| \langle \vec{v}(\vec{w}), \vec{w} \rangle} = \frac{\xi_p^*(\vec{w})}{\langle \vec{v}, \vec{w} \rangle} \geq 1.$$

We turn to the strict convexity and differentiability properties. As a standard consequence of the polar duality (79), for any pair $\vec{v}, \vec{w}$ with $\vec{v} = \vec{v}(\vec{w})$, $\vec{w}$ is the normal vector (pointing outwards) of a tangent to $\mathcal{U}_p$ at the point $\xi_p(\vec{v}) \cdot \vec{v}$ and $\vec{v}$ is the normal vector (pointing outwards) of a tangent to $\mathcal{W}_p$ at the point $\xi_p^*(\vec{w}) \cdot \vec{w}$.

By the bijectivity of $\vec{w} \mapsto \vec{v}(\vec{w})$ proved in Proposition 4.8, we conclude that neither $\mathcal{U}_p$ or $\mathcal{V}_p$ have facets; that is, they are strictly convex. As strict convexity of a convex body implies differentiability of the boundary of its convex dual, this concludes the proof of the theorem. $\square$

4.5 Invariance principle

Fix $p < p_c$ and $\vec{w} \in \mathbb{S}^1$. Define the linear interpolation $X(t)_{t \geq 0}$ of the discrete process $(X_n)_{n \geq 0}$. The representation of the cluster in terms of a random-walk like object allows to prove several invariance principles, similar to Donsker's Theorem for random walks. For instance, non-uniform invariance principles similar to Theorem 4.10 were derived in [CIV08] and [D'A24]. Our representation allows to prove the two following invariance principles.

Theorem 4.10.

1. Under the family of measures $\phi_p[\cdot \mid 0 \leftrightarrow \mathcal{H}_{\geq n}^{\vec{w}}]$, uniformly in p and $\vec{w}$, when $n \rightarrow \infty$,

$$\frac{1}{\sqrt{nL(p)}}(X_{nt} - tn\mu(p, \vec{w}))_{t \in (0,1)} \Rightarrow (B_t^{\sigma(p, \vec{w})})_{t \in (0,1)},$$

where $B_t^{\sigma(p)}$ denotes the one-dimensional Brownian motion started at 0 with variance $\sigma(p, \vec{w})$ given by Proposition 4.4.

2. Under the family of measures $\phi_p[\cdot \mid 0 \leftrightarrow n\vec{v}]$, uniformly in p and $\vec{v}$, when $n \rightarrow \infty$,

$$\frac{1}{\sqrt{nL(p)}}(X_{nt} - tn\mu(p, \vec{w}))_{t \in (0,1)} \Rightarrow (\mathbb{B}_t^{\sigma(p, \vec{w})})_{t \in (0,1)},$$

where $\vec{w}$ is the unique unit vector such that $\vec{v}(\vec{w}) = \vec{v}$ and $\mathbb{B}_t^{\sigma(p, \vec{w})}$ is the one-dimensional Brownian bridge started at 0 and ending at 0 with variance $\sigma(p, \vec{w})$.

In both statements, the convergence occurs in the space of continuous functions from $[0, 1]$ to $\mathbb{R}$, endowed with the topology of uniform convergence.

Proof. This follows directly from Theorem 3.2 together with classical considerations on Markov renewal processes. The first invariance principle directly follows from [Bor22, Theorem 1.5.3]. The second is a consequence of the first together with the local limit estimate of Proposition 4.4. As this is classical, we do not give further details, and refer to [D'A24] for the formal reasoning leading to the invariance principle. $\square$

By a standard union bound, it may be proved that, with probability tending to 1 under both $\phi_p[\cdot \mid 0 \leftrightarrow \mathcal{H}_{\geq n}^{\vec{w}}]$ and $\phi_p[\cdot \mid 0 \leftrightarrow n\vec{v}]$,

$$\max_k (T_{k+1} - T_k) \leq C \log(n/L(p)),$$

for $C > 0$ a universal constant.

Due to the cone-containment property (see Remark 3.3), the full cluster is at a distance $C'L(p) \max_k (T_{k+1} - T_k)$ from the points of renewal $(L(p)T_k, X_{T_k})_k$. Thus, with probability tending to 1 under both of the measures above,

$$d_{\mathcal{H}}(\mathbb{C}, (L(p)t, X_t)_{0 \leq t \leq n/L(p)}) \leq CL(p) \log(n/L(p)),$$

for some universal constant C , where $d_{\mathcal{H}}$ is the Hausdorff distance and $(L(p)t, X_t)_{0 \leq t \leq n/L(p)}$ is the family of points defined as in (14), expressed in the basis $(\vec{w}, \vec{w}^\perp)$.

Thus, one may easily deduce from Theorem 4.10 that the full cluster $\mathbb{C}$ sampled according to $\phi_p[\cdot \mid 0 \leftrightarrow \mathcal{H}_{\geq n}^{\vec{w}}]$ or $\phi_p[\cdot \mid 0 \leftrightarrow n\vec{v}]$, when properly rescaled, converges in law⁸ to the graph of a Brownian motion and bridge, respectively.

References

- [AIOV21] Yacine Aoun, Dmitry Ioffe, Sébastien Ott, and Yvan Velenik. Non-analyticity of the correlation length in systems with exponentially decaying interactions. *Comm. Math. Phys.*, 386(1):433–467, 2021.
- [AK77] D. B. Abraham and H. Kunz. Ornstein-Zernike theory of classical fluids at low density. *Phys. Rev. Lett.*, 39(16):1011–1014, 1977.
- [AN04] K.B. Athreya and P.E. Ney. *Branching Processes*. Dover Books on Mathematics. Dover Publications, 2004.
- [AOV23] Yacine Aoun, Sébastien Ott, and Yvan Velenik. On the two-point function of the Potts model in the saturation regime. *Comm. Math. Phys.*, 399(2):1103–1138, 2023.
- [AOV24] Yacine Aoun, Sébastien Ott, and Yvan Velenik. Ornstein-Zernike behavior for Ising models with infinite-range interactions. *Ann. Inst. Henri Poincaré Probab. Stat.*, 60(1):167–207, 2024.
- [BDC12] 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 Related Fields*, 153(3-4):511–542, 2012.
- [BF85] J. Bricmont and J. Fröhlich. Statistical mechanical methods in particle structure analysis of lattice field theories. I. General theory. *Nuclear Phys. B*, 251(4):517–552, 1985.
- [BIV00] T. Bodineau, D. Ioffe, and Y. Velenik. Rigorous probabilistic analysis of equilibrium crystal shapes. *J. Math. Phys.*, 41(3):1033–1098, 2000. Probabilistic techniques in equilibrium and nonequilibrium statistical physics.
- [Bor22] A. A. Borovkov. *Compound Renewal Processes*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2022.
- [CC86] J. T. Chayes and L. Chayes. Ornstein-Zernike behavior for self-avoiding walks at all noncritical temperatures. *Comm. Math. Phys.*, 105(2):221–238, 1986.
- [CCC91] M. Campanino, J. T. Chayes, and L. Chayes. Gaussian fluctuations of connectivities in the subcritical regime of percolation. *Probab. Theory Related Fields*, 88(3):269–341, 1991.
- [CI02] Massimo Campanino and Dmitry Ioffe. Ornstein-Zernike theory for the Bernoulli bond percolation on $\mathbb{Z}^d$. *Ann. Probab.*, 30(2):652–682, 2002.
- [CIV03] Massimo Campanino, Dmitry Ioffe, and Yvan Velenik. Ornstein-Zernike theory for finite range Ising models above T_c . *Probab. Theory Related Fields*, 125(3):305–349, 2003.

⁸We consider the convergence in law for sets of $\mathbb{R}^2$ under the Hausdorff distance.

- [CIV08] Massimo Campanino, Dmitry Ioffe, and Yvan Velenik. Fluctuation theory of connectivities for subcritical random cluster models. *Ann. Probab.*, 36(4):1287–1321, 2008.
- [D’A24] Lucas D’Alimonte. Entropic repulsion and scaling limit for a finite number of non-intersecting subcritical FK interfaces. *Electron. J. Probab.*, 29:1–53, 2024.
- [DC20] Hugo Duminil-Copin. Lectures on the Ising and Potts models on the hypercubic lattice. In *Random graphs, phase transitions, and the Gaussian free field*, volume 304 of *Springer Proc. Math. Stat.*, pages 35–161. Springer, Cham, 2020.
- [DCGH⁺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 $q > 4$. *Ann. Sci. Éc. Norm. Supér. (4)*, 54(6):1363–1413, 2021.
- [DCM22] Hugo Duminil-Copin and Ioan Manolescu. Planar random-cluster model: scaling relations. *Forum Math. Pi*, 10:Paper No. e23, 83, 2022.
- [DCMT21] Hugo Duminil-Copin, Ioan Manolescu, and Vincent Tassion. Planar random-cluster model: fractal properties of the critical phase. *Probability Theory and Related Fields*, 181(1):401–449, 2021.
- [DCST17] Hugo Duminil-Copin, Vladas Sidoravicius, and Vincent Tassion. Continuity of the phase transition for planar random-cluster and Potts models with $1 \leq q \leq 4$. *Comm. Math. Phys.*, 349(1):47–107, 2017.
- [DCT20] Hugo Duminil-Copin and Vincent Tassion. Renormalization of crossing probabilities in the planar random-cluster model. *Mosc. Math. J.*, 20(4):711–740, 2020.
- [DGO25] Moritz Dober, Alexander Glazman, and Sébastien Ott. Discontinuous transition in 2D Potts: I. Order-disorder interface convergence, 2025.
- [Gri06] G.R. Grimmett. *The Random-Cluster Model*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2006.
- [Iof98] D. Ioffe. Ornstein-Zernike behaviour and analyticity of shapes for self-avoiding walks on $\mathbb{Z}^d$. *Markov Process. Related Fields*, 4(3):323–350, 1998.
- [Kes87] Harry Kesten. Scaling relations for 2D-percolation. *Comm. Math. Phys.*, 109(1):109–156, 1987.
- [MS95] Salvador Miracle-Sole. Surface tension, step free energy, and facets in the equilibrium crystal. *Journal of Statistical Physics*, 79(1):183–214, 1995.
- [MZ96] R. A. Minlos and E. A. Zhizhina. Asymptotics of decay of correlations for lattice spin fields at high temperatures. I. The Ising model. *J. Statist. Phys.*, 84(1-2):85–118, 1996.
- [OV18] Sébastien Ott and Yvan Velenik. Potts models with a defect line. *Comm. Math. Phys.*, 362(1):55–106, 2018.
- [OZ14] L. Ornstein and F. Zernike. Accidental deviations of density and opalescence at the critical point of a single substance. *Koninklijke Nederlandse Akademie van Wetenschappen Proceedings Series B Physical Sciences*, 17:793–806, 1914.
- [PL78] Paulo Jorge Paes-Leme. Ornstein-Zernike and analyticity properties for classical lattice spin systems. *Ann. Physics*, 115(2):367–387, 1978.
- [RS20] 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(3):1977–1988, 2020.
- [Web94] R. Webster. *Convexity*. Oxford science publications. Oxford University Press, 1994.
- [Zer16] F. Zernike. The clustering-tendency of the molecules in the critical state and the extinction of light caused thereby. *Koninklijke Nederlandse Akademie van Wetenschappen Proceedings Series B Physical Sciences*, 18 II:1520–1527, 1916.
- [Zyg02] A. Zygmund. *Trigonometric series. Vol. I, II*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, third edition, 2002. With a foreword by Robert A. Fefferman.