


The scaling limits of the Minimal Spanning Tree and Invasion Percolation in the plane

Christophe Garban Gábor Pete Oded Schramm

Abstract

We prove that the Minimal Spanning Tree and the Invasion Percolation Tree on a version of the triangular lattice in the complex plane have unique scaling limits, which are invariant under rotations, scalings, and, in the case of the MST, also under translations. However, they are not expected to be conformally invariant. We also prove some geometric properties of the limiting MST. The topology of convergence is the space of spanning trees introduced by Aizenman, Burchard, Newman & Wilson (1999), and the proof relies on the existence and conformal covariance of the scaling limit of the near-critical percolation ensemble, established in our earlier works.



The MST in a box, and InvPerc started from the midpoint of the left boundary of the box until reaching the right boundary, on \mathbb{Z}^2 .

Contents

1	Introduction	3
1.1	The Minimal Spanning Tree MST	4
1.2	The Invasion Percolation Tree InvPerc	8
1.3	The scaling limit of the near-critical ensemble	9
1.4	Strategy of the proof and organization of the paper	12
2	Topological and measurability preliminaries	14
2.1	The space of essential spanning forests	14
2.2	The quad-crossing topology	18
2.3	Pivotal and pivotal measures	21
3	Enhanced networks and cut-off forests built from the near-critical ensemble	27
4	Approximation of MST_η by the cut-off trees $\text{MST}_\eta^{\bar{\lambda},\epsilon}$	39
4.1	Preparatory lemmas and the definition of $\text{MST}_\eta^{\bar{\lambda},\epsilon}$	39
4.2	Approximation as $\epsilon \rightarrow 0$ and $(\lambda, \lambda') \rightarrow (-\infty, \infty)$	43
5	Proof of the main result	47
5.1	Putting the pieces together for MST on tori \mathbb{T}_M^2	47
5.2	Extension to the full plane; invariance under translations, scalings and rotations	48
6	Geometry of the limit tree MST_∞	50
6.1	Degree types and pinching	50
6.2	A dimension bound for the trunk	57
7	Invasion percolation	59
8	Questions and conjectures	61
	References	61

1 Introduction

The Minimal Spanning Tree of weighted graphs is a classical combinatorial object [15, 39, 31], and is also very interesting from the viewpoint of probability theory and statistical physics: when the weights on the edges of a graph are chosen at random, using i.i.d. variables, then the resulting random tree turns out to be closely related to the near-critical regime of Bernoulli bond percolation on that graph.

In Bernoulli bond percolation at density $p \in [0, 1]$, each edge of the graph is kept open with probability p or becomes closed with probability $1 - p$, independently, and then one looks at the connected open components, called clusters. In site percolation, the vertices are chosen to be open or closed instead of the edges. These are among the most important spatial stochastic processes, due to their simultaneous simplicity and richness [16, 36, 32]. The main interest is in the phase transition near the critical density p_c , below which all clusters are small, above which a cluster (sometimes clusters) of positive density emerge. The theory of critical percolation in the plane has seen a lot of progress lately, starting with Smirnov's proof of conformal invariance of crossing probabilities for site percolation on the triangular lattice [51], and with the introduction of the Stochastic Loewner Evolution [48] that describes the conformally invariant curves that are the scaling limits of interfaces between open and closed clusters. These SLE curves can be used to understand critical percolation in depth [57], including the computation of critical exponents that had been predicted by physicists using non-rigorous conformal field theory techniques.

Beyond the static critical system, it is natural to consider dynamical versions: first, to slowly change p near p_c and observe how the phase transition exactly takes place — called near-critical percolation; second, to apply a stationary dynamics and observe how the critical system is changing in time — called dynamical percolation. Indeed, by “perturbing” critical percolation, the static results of the previous paragraph have also given way to an exhaustive study of dynamical and near-critical percolation [50, 26, 34, 28, 29]; see also the surveys [53, 30]. In particular, in [28, 29] we have proved the existence and conformal covariance of the scaling limit of the near-critical percolation ensemble, w.r.t. the quad-crossing topology introduced in [49]. Very roughly, this near-critical scaling limit is constructed from the critical scaling limit, plus independent randomness that governs how macroscopic clusters merge as we raise p .

It turns out that the macroscopic structure of the Minimal Spanning Tree (MST) and the Invasion Percolation Tree (InvPerc) can also be de-

scribed based on this merging process. Thus, building on [28, 29], in the present paper we prove the existence and some conformal properties of the scaling limits of **MST** and **InvPerc** on the triangular lattice, in the space of essential spanning forests introduced in [4]. In that paper, tightness results were proved, implying that subsequential scaling limits of the Minimal and Uniform Spanning Trees in the plane exist. Our proof of the uniqueness of the scaling limit has the important implication that the conjectural universality of critical percolation implies universality for many processes related to the near-critical ensemble, including **MST** and **InvPerc**. That this program of describing near-critical objects from the critical scaling limit may have a chance to work was suggested in [18]. Another motivation for our work is that it leads to interesting new objects: these two scaling limits are invariant under rotations and scalings, but, conjecturally, not under general conformal maps. Furthermore, the methods developed to establish these scaling limits also give information about the large-scale geometry of the discrete trees.

1.1 The Minimal Spanning Tree **MST**

For each edge of a finite graph, $e \in E(G)$, let $U(e)$ be an independent $\text{Unif}[0, 1]$ label. The **Minimal Spanning Tree**, denoted by **MST**, is the spanning tree T for which $\sum_{e \in T} U(e)$ is minimal. This is well-known to be the same as the union of lowest level paths between all pairs of vertices (i.e., the path between the two points for which the maximum label on the path is minimal). One can also use the so-called **reversed Kruskal algorithm** to construct **MST**: delete from each cycle the edge with the highest label U . This algorithm also shows that **MST** depends only on the ordering of the labels, not on the values themselves. Moreover, this algorithm also makes sense on any infinite graph, and produces what in general is called the Free Minimal Spanning Forest (**FMSF**) of the infinite graph. The Wired Minimal Spanning Forest (**WMSF**) is the one when we also remove the edge with the highest label (if such edge exists) from each cycle that “goes through infinity”, i.e., which is the union of two disjoint infinite simple paths starting from a vertex. For the case of Euclidean planar lattices, these two measures on spanning forests are known to be the same, again denoted by **MST**, and it almost surely consists of a single tree [7]. This measure can also be obtained as a **thermodynamical limit**: take any exhaustion by finite subgraphs $G_n(V_n, E_n)$, introduce a boundary condition by identifying some of the vertices on the boundary of G_n (i.e., elements of V_n that have neighbors in G outside of V_n), and then take the weak limit. On a general infinite graph, when no identifications are made in

the boundary, one gets the FMSF, and when all vertices are glued into a single vertex, one gets the WMSF. Studying these measures has a rich history on \mathbb{Z}^d , on point processes in \mathbb{R}^d , and on general transitive graphs; see [6, 47, 60, 5, 43, 55, 22, 56, 44, 42] and the references therein.

One can use the same $\text{Unif}[0, 1]$ labels that defined the MST to obtain a coupling of percolation for all densities $p \in [0, 1]$: an edge is “open at level p ” if $U(e) \leq p$. This way we get a **coupling** between the MST and the **percolation ensemble**. Moreover, as we explain in the next paragraph, the macroscopic structure of the MST is basically determined by the labels in the near-critical regime of percolation, and hence one may hope that the scaling limit of the MST is determined by the scaling limit of the near-critical ensemble.




Figure 1.1: The MST connects the percolation p -clusters without creating cycles, yielding the cluster-tree MST^p .

Consider the p -clusters (i.e., open components at level p) in the percolation ensemble on some large finite graph. Contract each component into a single vertex, keeping the edges (together with their labels) between the clusters, resulting in the “cluster graph”. It is easy to verify that making these contractions on the MST we get exactly the MST on the cluster graph. We denote this **cluster tree** by MST^p . See Figure 1.1. Now assume that p_1 is small enough so that even the largest p_1 -clusters are of small macroscopic size — then the tree MST^{p_1} will tell us the macroscopic structure of MST. On the other hand, if $p_2 > p_1$ is large enough, then most sites are in just one giant p_2 -cluster. Note that, for any $p > p_1$, we get the tree MST^p from MST^{p_1} by contracting the edges with labels in $(p_1, p]$. Thus, if we have the collection of *all* the p -clusters for all $p \in (p_1, p_2)$, then by following how they merge as we are raising p , we can reconstruct the tree MST^{p_1} . Now, one may hope that in order to tell the macroscopic structure of MST^{p_1} , it is enough to know only the *macroscopic* p -clusters for all $p \in (p_1, p_2)$ and follow how

those merge. The **near-critical window** of percolation is exactly the window (p_1, p_2) in which the above phase transition of the cluster sizes takes place, and the scaling limit of the near-critical ensemble is exactly the object that describes the macroscopic p -clusters in this window. Therefore, the above hope has the interpretation that the scaling limit of the near-critical ensemble should describe the scaling limit of the MST. This, of course, raises several questions: May the dust of microscopic p -clusters condensate into a new macroscopic p' -cluster at some $p' > p$, ruining the strategy of “following how macroscopic clusters merge”? Could MST^{p_1} go through microscopic p_1 -clusters in a way that significantly influences its macroscopic structure?

Our work addresses these questions in the case of **planar lattices**. The near-critical window for Bernoulli(p) percolation on the triangular lattice $\eta\mathbb{T}$ or the square lattice $\eta\mathbb{Z}^2$ with mesh $\eta > 0$ is given by

$$p = 1/2 + \lambda r(\eta) \text{ with } \lambda \in (-\infty, \infty) \text{ fixed and } \eta \rightarrow 0, \quad (1.1)$$

where $r(\eta) = \eta^2/\alpha_4(\eta, 1)$, with $\alpha_4(\eta, 1)$ being the alternating 4-arm probability of critical percolation [57]. It was proved on $\eta\mathbb{T}$ using SLE₆ computations [52] that $r(\eta) = \eta^{3/4+o(1)}$. As shown in [38], for $\lambda \ll -1$ we are at the subcritical end of the near-critical window, for $\lambda \gg 1$ we are at the supercritical end, and for any fixed $\lambda \in \mathbb{R}$, box-crossing probabilities are comparable to the critical case (just they are close to 0 for $\lambda \ll -1$, and close to 1 for $\lambda \gg 1$). That is, (1.1) is indeed the near-critical window. Then it was proved in [28, 29] that for any $\lambda \in \mathbb{R}$ there is a unique scaling limit as $\eta \rightarrow 0$; moreover, the entire coupled percolation ensemble, viewed near the critical point via the parametrization (1.1), where all the macroscopic changes happen, has a scaling limit as a Markov process in $\lambda \in \mathbb{R}$. It is important to keep in mind that even for any given $\lambda \neq 0$, this scaling limit is an interesting new object, known to be different from the critical scaling limit: the interfaces are singular w.r.t. SLE₆ [46]. (See also [10] and [29, Theorem 13.4] for the much simpler result that the full scaling limits are singular.)

Since we have a proof of the existence and properties of the scaling limit of the near-critical ensemble only for site percolation on the triangular lattice \mathbb{T} , if we want to use that to build the MST scaling limit, we will need a version of the MST that uses $\text{Unif}[0, 1]$ vertex labels $\{V(x)\}$ on \mathbb{T} . So, assign to each edge $e = (x, y)$ the vector label

$$U(e) := (V(x) \vee V(y), V(x) \wedge V(y)), \quad (1.2)$$

and consider the lexicographic ordering on these vectors to determine the MST. See Figure 1.2. With a slight abuse of terminology, this is what we

will call the MST on the lattice \mathbb{T} . Our strongest results will apply to this model, but some of them will also hold for subsequential limits of the usual MST on \mathbb{Z}^2 , known to exist by [4].




Figure 1.2: The minimal spanning tree associated to vertex labels of the triangular lattice \mathbb{T} , with a periodic boundary condition.

Let us make an important remark here. The use of the lexicographic ordering for the vector labels (1.2) is somewhat arbitrary, and starting from the same vertex labels, using a different way to get edge labels or using a different natural ordering, one could a priori get an MST with a very different global structure. In fact, this does happen if the vertex labels are assigned maliciously. Nevertheless, with the $\text{Unif}[0, 1]$ labels, for any rule to construct the MST on \mathbb{T} that ensures that *any two p -clusters are connected by a unique path of this MST* (which is exactly how our definition works), our approximation of the macroscopic structure of the MST using the near-critical ensemble will work with large probability, and hence the scaling limit will be the same.

We can now state our main theorem:

Theorem 1.1 (Limit of MST_η in \mathbb{C}). *As $\eta \rightarrow 0$, the spanning tree MST_η on $\eta\mathbb{T}$ converges in distribution, in the metric d_Ω of Definition 2.2 below, to a unique scaling limit MST_∞ that is invariant under translations, scalings, and rotations.*

The strategy of the proof will be described in Subsection 1.4. As a key step, we also prove convergence in any fixed torus \mathbb{T}_M^2 ; see Theorem 5.1. We work in tori to avoid the technicalities related to boundary issues, but with not too much additional work the extension to finite domains with free or wired boundary conditions would be certainly doable.

In Section 6, strengthening the results of [4], we study the geometry of the limiting tree MST_∞ . The degree of a vertex in a tree graph has the usual meaning, but the degree of a point in a spanning forest of the plane needs to

be defined carefully, which we will do in Subsection 6.1. To give an example, a pinching point on an MST_∞ path should not be called a branching point, but it still gives rise to a degree 4 point. Consequently, stating the results on the geometry of the limiting tree also needs some care, to be done precisely only in Theorem 6.2. Nevertheless, here are some of the earlier results and our new ones in rough terms. It was proved in [4] that there is an unspecified absolute bound k_0 such that almost surely all degrees in any subsequential limit of MST_η are at most k_0 . Furthermore, the set of branching points was shown to be almost surely countable. Here, we will prove that there are almost surely no pinching points, all degrees are bounded by 4, and the set of points with degree 4 is at most countable. We will also prove, in Subsection 6.2, that the Hausdorff dimension of the trunk is strictly below $7/4$. However, we do not have a guess for the exact dimension; the situation is similar to the somewhat related problem of finding the percolation chemical distance exponent [25].

To conclude this subsection, let us note that the recent works [1, 2] follow a strategy similar to ours, but in a very different setting: namely, in the **mean-field** case. It is well-known that there is a phase transition at $p = 1/n$ for the **Erdős-Rényi random graphs** $G(n, p)$. Similarly to the above case of planar percolation, it is a natural problem to study the geometry of these random graphs near the transition $p_c = 1/n$. It turns out in this case that the non-trivial rescaling is to work with $p = 1/n + \lambda/n^{4/3}$, $\lambda \in \mathbb{R}$. If $R_n(\lambda) = (C_n^1(\lambda), C_n^2(\lambda), \dots)$ denotes the sequence of clusters at $p = 1/n + \lambda/n^{4/3}$, ordered in decreasing order of size, say, then it is proved in [1] that as $n \rightarrow \infty$, the normalized sequence $n^{-1/3} R_n(\lambda)$ converges in law to a limiting object $R_\infty(\lambda)$ for a certain topology on sequences of compact spaces which relies on the Gromov-Hausdorff distance. This near-critical ensemble $\{R_\infty(\lambda)\}_{\lambda \in \mathbb{R}}$ has then been used in [2] to obtain a scaling limit as $n \rightarrow \infty$ (in the Gromov-Hausdorff sense) of the MST on the complete graph with n vertices. One could say that [29] is the Euclidean ($d = 2$) analogue of the mean-field case [1], and our present paper is the analogue of [2]. However, an important difference is that in the mean-field case one is interested in the intrinsic metric properties (and hence works with the Gromov-Hausdorff distance between metric spaces), while in the Euclidean case one is first of all interested in how the graph is embedded in the plane.

1.2 The Invasion Percolation Tree InvPerc

The connection between WMSF and critical percolation on infinite graphs can also be seen through invasion percolation. For a vertex x in an infinite

graph $G(V, E)$, and the labels $\{U(e)\}$, let $T_0 = \{x\}$, then, inductively, given T_n , let $T_{n+1} = T_n \cup \{e_{n+1}\}$, where e_{n+1} is the edge between T_n and $V \setminus T_n$ that has the smallest label U . The **Invasion Percolation Tree** of x is then $\text{InvPerc}(x) := \bigcup_{n \geq 0} T_n$. It is easy to see that, even deterministically, if $U : E(G) \rightarrow \mathbb{R}$ is an injective labelling of a locally finite graph, then $\text{WMSF} = \bigcup_{x \in V(G)} \text{InvPerc}(x)$.

Once the invasion tree enters an infinite p -cluster \mathcal{C} , it will not use edges outside it. Furthermore, it is not surprising (though non-trivial to prove, see [33]) that for any transitive graph G and any $p > p_c(G)$, the invasion tree eventually enters an infinite p -cluster. Therefore, $\limsup\{U(e) : e \in \text{InvPerc}(x)\} = p_c(G)$ for any $x \in V(G)$. This way, invasion percolation can be considered as a “self-organized criticality” version of critical percolation; finer results for the planar case are given in [20, 23, 24]. Moreover, InvPerc can be used to study Bernoulli percolation itself: e.g., for the well-behavedness of the supercritical phase on \mathbb{Z}^d , $d > 2$ [21], and for uniqueness monotonicity on non-amenable graphs [33]. Invasion percolation can be analyzed very well on regular trees [8], with a scaling limit that can be described using diffusion processes [9].

For planar lattices, since InvPerc_η is so intimately related to MST_η , it will be quite easy to modify the proof of Theorem 1.1 for the case of InvPerc ; see Section 7.

1.3 The scaling limit of the near-critical ensemble

We need to recall how the scaling limit of the near-critical ensemble is constructed in [28, 29], because the present paper is heavily built on this. To start with, we slightly change the near-critical parametrization given in (1.1):

Definition 1.2. *The near-critical ensemble $(\omega_\eta^\lambda)_{\lambda \in \mathbb{R}}$ will denote the following process:*

- (i) *Sample $\omega_\eta^{\lambda=0}$ according to \mathbf{P}_η , the law of critical percolation on $\eta\mathbb{T}$. We will sometimes represent this as a black-and-white coloring of the faces of the dual hexagonal lattice, with white hexagons standing for closed (empty) sites.*
- (ii) *As λ increases, closed sites (white hexagons) switch to open (black) at an exponential rate $r(\eta)$, as given after (1.1).*
- (iii) *As λ decreases, black hexagons switch to white at rate $r(\eta)$.*

Note that, for any $\lambda \in \mathbb{R}$, the near-critical percolation ω_η^λ corresponds exactly to a percolation configuration on $\eta\mathbb{T}$ with parameter

$$\begin{cases} p = p_c + (1 - p_c)(1 - e^{-\lambda r(\eta)}) & \text{if } \lambda \geq 0 \\ p = p_c e^{-|\lambda| r(\eta)} & \text{if } \lambda < 0. \end{cases}$$

For any site x , the value $\lambda(x) \in \mathbb{R}$ where x switches from closed to open will be called the near-critical percolation label of x .

The same definitions can be made on $\eta\mathbb{Z}^2$.

It is easy to understand intuitively why $r(\eta)$ is the right time rescaling to obtain the near-critical window. Assume that in the unit square there is no left-right crossing in $\omega_\eta^{\lambda=0}$. Then the expected number of those sites that are closed at $\lambda = 0$ but are **pivotal** for the left-right crossing (i.e., opening any of them would establish the crossing) and which actually become open in ω_η^λ is known to be of order λ . Therefore, for $\lambda > 0$ small, it is unlikely that a left-right crossing has been established if it was not already there, hence the system must have stayed very close to critical; on the other hand, one may expect that for $\lambda \gg 1$ a crossing is already quite likely, hence the system should already be quite supercritical. This was rigorously proved in [38]. Then, if one wants to describe the **scaling limit** of ω_η^λ as $\eta \rightarrow 0$, a natural idea that was detailed in [18] is that this should be possible by following which of those points get opened (for $\lambda > 0$) or get closed (for $\lambda < 0$) that were pivotal at $\lambda = 0$ for at least some small macroscopic distance $\epsilon > 0$. To this end, one should look at the counting measure on ϵ -pivotal points at criticality, normalized such that the measure stays non-trivial as $\eta \rightarrow 0$, and hope that these ϵ -pivotal measures have limits that are measurable w.r.t. the scaling limit of critical percolation itself. This is the main result of [28] (with a slight change of what ϵ -pivotal means). Then, the scaling limit of the near-critical ensemble may be described by taking Poisson point processes of switch times, with intensity measures being these ϵ -pivotal measures, and by updating the crossings of all the quads (certain generalized rectangles) according to these pivotal switches. This is done in [29]. Here there are roughly two main issues: firstly, it is not immediately clear how one can update the crossings of *all* the quads by pivotal switches that are happening at *all* spatial and time scales. For this, one should code the percolation configuration in a suitable manner that is minimal enough so that the updates can be done, but rich enough so that it contains all the relevant information. This coding and updating takes up a large part of [29], done through the so-called ϵ -networks that we will actually recall in Section 3. The second main issue is that one needs to prove that despite

all the switches that take place as λ increases, following the switches of all the initially ϵ -pivotal sites gives a good idea about the ϵ -pivotal switches at later times. For this, the key discrete result from [23, 29] is the following proposition, which we will often use also in the present paper:

Proposition 1.3 (Near-critical stability). *For any fixed $-\infty < \lambda < \lambda' < \infty$, in the near-critical ensemble on $\eta\mathbb{T}$, let $\mathcal{A}_k^{\lambda, \lambda'}(r, R)$ denote the following **near-critical polychromatic k -arm event**: there exist $k \geq 2$ disjoint paths in the lattice that connect the boundary pieces of the annulus $B_R(0) \setminus B_r(0)$, each called either “primal” or “dual”, and all the percolation ensemble labels along all the primal arms are at most λ' , while all the labels along the dual arms are at least λ . Note that $\lambda = \lambda'$ gives back the usual notion of primal and dual arms in the percolation configuration ω_η^λ . Then,*

$$\mathbf{P}[\mathcal{A}_k^{\lambda, \lambda'}(r, R)] \leq C_{\lambda, \lambda'} \alpha_k(r, R),$$

where $\alpha_k(r, R) = \alpha_k^\eta(r, R)$ is the polychromatic k -arm probability in critical percolation on the same lattice. Similarly, for the **monochromatic k -arm events**, $k \geq 1$, where all arms are primal,

$$\mathbf{P}[\mathcal{A}_k^{\lambda'}(r, R)] \leq C_{\lambda'} \alpha'_k(r, R),$$

where $\alpha'_k(r, R) = \alpha_k^{\eta'}(r, R)$ is the monochromatic k -arm probability at criticality. For $\alpha'_1(\cdot, \cdot)$, we will just use $\alpha_1(\cdot, \cdot)$.

The same statements hold for bond percolation on $\eta\mathbb{Z}^2$, just with dual arms being paths in the dual lattice, in the usual manner.

Remark 1.4. For fixed radii $0 < r < R < \infty$, the discrete multi-arm probabilities $\alpha_k^\eta(r, R)$ and $\alpha_k^{\eta'}(r, R)$ converge, as $\eta \rightarrow 0$, to their SLE₆ counterparts (see [57]). In the present paper, we will be interested in these quantities only up to constant factors, not in the details of their convergence, hence their dependence on η is not important and will be omitted from the notation. Formulas like $\alpha_k(\eta, 1)$ will also be understood on the discrete lattice, always with mesh η . We will also use the quasi-multiplicativity of multi-arm probabilities (both for the discrete and continuum versions): for any $k \geq 1$, there exists $c_k > 0$ such that

$$c_k \alpha_k(r_1, r_2) \alpha_k(r_2, r_3) \leq \alpha_k(r_1, r_3) \leq \alpha_k(r_1, r_2) \alpha_k(r_2, r_3),$$

for all $0 < r_1 < r_2 < r_3 < \infty$. Similarly for α'_k . See, e.g., [45, Subsection 4.5].

The proof of Proposition 1.3 for the alternating 4-arm event is given in [23, Lemma 6.3], or follows directly from [29, Lemma 8.4], which is more general in that it does not assume that the dynamics is monotone in λ . For general k , the case of $\lambda = \lambda'$ is known as Kesten's near-critical stability [38]. And just as in Kesten's approach, the proof for general k and general $\lambda < \lambda'$ is a simple modification of the proof for the alternating 4-arm event: the key point is that the pivotality of a site for a general k -event still depends on an alternating 4-arm event around that site, and hence the near-critical stability of the alternating 4-arm probability, proved using a recursion in [29], easily implies the stability of the general k -arm event, as well. We omit the details.

The above sketch of the contents of [28, 29] should make it clear that the scaling limit of the near-critical ensemble is constructed entirely from the critical scaling limit, plus independent randomness of the pivotal switch times. Moreover, all the proofs in [28, 29] are universal in the sense that they use lattice-independent discrete percolation technology that have been available since [38]. Altogether, once one proves Cardy's formula for critical percolation on $\eta\mathbb{Z}^2$, which would imply the same scaling limit as on $\eta\mathbb{T}$, we would also immediately get that the scaling limit for the entire near-critical ensemble is the same. This universal aspect remains true for the present paper.

1.4 Strategy of the proof and organization of the paper


First of all, in Subsection 2.1, we describe the topological space in which the convergence of our random trees will take place: the space of essential spanning forests in \mathbb{C} , introduced in [4]. There are possible alternatives to using this topology, such as the quad-crossing topology of [49] (suggested to us for this purpose by Nicolas Broutin) or the topology introduced in [48] for the scaling limit of the Uniform Spanning Tree. Especially the quad-crossing topology (recalled in Subsection 2.2) would seem natural, since the scaling limit of near-critical percolation is taken in this space. Nevertheless, we chose the topology of [4] for several reasons: that was the first paper dealing with subsequential scaling limits of MST_η , proving results that we are sharpening here; using this topology to describe paths in the spanning trees is not harder than using quad-crossings, while it also gives a natural way to glue the paths into more complicated trees; there is a simple explicit metric generating this topology. However, we will unfortunately need more topological preparations than just recalling these definitions, because the minimalist structure, based on just the pivotal measures of [28], which was enough to describe the scaling limit of the near-critical ensemble in [29], will

not be enough for the tree structures of the present paper. In particular, in Proposition 2.6, we will prove that that **set of colored pivotals** also has a limit as $\eta \rightarrow 0$.

In Section 3, we first recall the definition of the **networks** $N_\eta^{\bar{\lambda}, \epsilon}$ and $N_\infty^{\bar{\lambda}, \epsilon}$ introduced in [29], where $\bar{\lambda} = (\lambda, \lambda')$ is a pair of near-critical parameters with $\lambda < \lambda'$. These are graphs with vertex sets X given by those ϵ -pivotals in the configuration ω^λ on a torus \mathbb{T}_M^2 that experience a switch between level λ and λ' , and edges given roughly by the primal and dual connections in $\omega^\lambda \setminus X$. Then we need to add a bit more structure to these networks, creating the so-called **enhanced networks**: roughly, we will need to know which of these pivotals are connected together by an open cluster of $\omega^\lambda \setminus X$, and will need to know the colors of these pivotals in ω^λ . For this, we will use Proposition 2.6 mentioned in the previous paragraph and Proposition 3.6 saying that clusters of large diameter also have large volume (which excludes certain pathological geometric behaviour that would ruin the construction). From these enhanced networks, we will obtain finite labelled graphs whose vertices will basically be open λ -clusters that have ϵ -pivotals switching in the time interval (λ, λ') , with edges labelled by the times of the pivotal switches, showing how the λ -clusters merge. We will define the MST on this finite labelled graph, denoted by $MST_\eta^{\bar{\lambda}, \epsilon}$ in the discrete and $MST_\infty^{\bar{\lambda}, \epsilon}$ in the continuum case — these are basically the macroscopic approximations to the cluster trees that we discussed in Subsection 1.1. To be more precise, in Section 3 we define only some Minimal Spanning Forests, and we need a bit more work until in Lemma 4.4 we can actually define the trees. The fact that these approximating **cut-off trees** $MST_\eta^{\bar{\lambda}, \epsilon}$ and $MST_\infty^{\bar{\lambda}, \epsilon}$ are close to each other if the underlying near-critical ensembles $\omega_\eta^{[\lambda, \lambda']}$ and $\omega_\infty^{[\lambda, \lambda']}$ are close follows easily from [29].

In Section 4 we prove that the cut-off trees $MST_\eta^{\bar{\lambda}, \epsilon}$ are close to the true MST_η if $\lambda \ll -1$, $\lambda' \gg 1$, and $\epsilon > 0$ is small. Here the key technique is near-critical stability, Proposition 1.3.

Summarizing, we get that MST_η is close to $MST_\infty^{\bar{\lambda}, \epsilon}$. Since the latter does not depend on η , while the former does not depend on $\bar{\lambda}$ and ϵ , they both need to be close to an object that does not depend on any of these parameters: this will be the scaling limit MST_∞ . To give a succinct pictorial summary of this strategy:



This conclusion will be materialized in Section 5, together with the extension from the case of the tori \mathbb{T}_M^2 to the full plane, and with the proof of the claimed invariance properties.

As already advertised in Subsections 1.1 and 1.2, the results on the geometry of MST_∞ are discussed in Section 6, while Section 7 establishes the existence and invariance properties of InvPerc_∞ . We conclude the paper with some open problems in Section 8.

Acknowledgments. We thank Louigi Addario-Berry, Nicolas Broutin, Laure Dumaz, Grégory Miermont and David Wilson for stimulating discussions, Rob van den Berg for pointing out the connection between Proposition 3.6 and [35], and Alan Hammond and two amazing anonymous referees for very good comments on the manuscript.

Part of this work was done while all authors were at Microsoft Research, Redmond, WA, or GP was visiting CG at ENS Lyon. CG was partially supported by the ANR grant MAC2 10-BLAN-0123. GP was supported by an NSERC Discovery Grant at the University of Toronto, an EU Marie Curie International Incoming Fellowship at the Technical University of Budapest, and partially supported by the Hungarian National Research, Development and Innovation Office, NKFIH grant K109684, and by the MTA Rényi Institute “Lendület” Limits of Structures Research Group.

2 Topological and measurability preliminaries

2.1 The space of essential spanning forests

The following topological setup for discrete and continuum spanning trees was introduced in [4]. We are summarizing here the definitions and the notation, with small modifications; the main difference is roughly that Ω will also contain spanning trees of subsets of the complex plane, to accommodate the invasion percolation tree InvPerc and our approximating trees $\text{MST}^{\lambda, \epsilon}$.

We will work in a one-point compactification of $\mathbb{C} = \mathbb{R}^2$, denoted by $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$, with the Riemannian metric

$$\frac{4}{(1+x^2+y^2)^2}(dx^2+dy^2); \tag{2.1}$$

by stereographic projection, $\hat{\mathbb{C}}$ is isometric with the unit sphere. Note that this metric is equivalent to the Euclidean metric in bounded domains, while the distance between any two points outside the square of radius M around the origin in \mathbb{C} is at most $O(1/M)$. This will imply that convergence of

spanning trees in $\hat{\mathbb{C}}$ is the same as convergence within bounded subsets of \mathbb{C} . This is necessary, since convergence of random spanning trees cannot be uniform in \mathbb{C} : on $\eta\mathbb{Z}^2$, inside the infinitely many pieces $[i, i+1) \times [j, j+1)$, $i, j \in \mathbb{Z}$, one can find arbitrary topological behavior (e.g., macroscopically vanishing areas with arbitrarily large numbers of macroscopic branches emanating from them) that will be very far from the almost sure behavior of the continuum tree.

Spanning trees on infinite graphs are usually defined and studied as weak limits of spanning trees in finite subgraphs exhausting the infinite graph. For these finite graphs, one may consider different boundary conditions: most importantly, free or wired. As mentioned in the Introduction, for the MST on Euclidean planar lattices, all such boundary conditions give the same limit measure, and we will work in the tori \mathbb{T}_M^2 of side-length $2M$, which can be realized as the subdomains $[-M, M)^2$ of \mathbb{C} , or even as subgraphs of $\eta\mathbb{T}$ for suitable values of M , with a periodic boundary condition (which is sandwiched between the free and the wired conditions). See Figure 1.2 in the Introduction.

Definition 2.1. *A reference tree τ is a tree with a finite set of leaves (or external vertices), denoted by $\xi(\tau)$, with each edge considered to be a unit interval. A reparametrization is a continuous map $\phi : \tau \rightarrow \tau$ that fixes all the vertices and is monotone on the edges. An immersed tree indexed by τ is an equivalence class of continuous maps $f : \tau \rightarrow \hat{\mathbb{C}}$, where f_1 and f_2 are considered equivalent if there exist reparametrizations ϕ_1, ϕ_2 with $f_1 \circ \phi_1 = f_2 \circ \phi_2$. The collection of immersed trees indexed by τ is denoted by \mathcal{S}_τ , and we set*

$$\mathcal{S}^{(\ell)} := \bigcup_{\tau: |\xi(\tau)|=\ell} \mathcal{S}_\tau, \quad (2.2)$$

the set of immersed trees with ℓ leaves. A tree with leaves $x_1, \dots, x_\ell \in \hat{\mathbb{C}}$ will often be denoted by $T(x_1, \dots, x_\ell) \in \mathcal{S}^{(\ell)}$.

We will also consider trees immersed into the torus \mathbb{T}_M^2 with the flat Euclidean metric; the corresponding collection of immersed trees with ℓ leaves is denoted by $\mathcal{S}_M^{(\ell)}$.

One may consider trees immersed not just into $\hat{\mathbb{C}}$ or \mathbb{T}_M^2 , but into a graph $G(V, E)$ that is embedded into $\hat{\mathbb{C}}$ or \mathbb{T}_M^2 , and then the image of τ is required to be a subtree of $G(V, E)$, with its vertices mapped into V and any of its edges mapped to a union of edges from E .

Note that if a reference tree τ' is given by contracting some edges of some τ , denoted by $\tau' \prec \tau$, then $\mathcal{S}_{\tau'}$ is naturally a subset of \mathcal{S}_τ , represented




Figure 2.1: A reference tree τ with four leaves, with one immersion into \mathbb{Z}^2 and another into \mathbb{C} . The *image* in \mathbb{C} is not a tree, but this is allowed. In the scaling limit of any discrete random tree in $\hat{\mathbb{C}}$ one cannot see such self-intersections, but could see touch-points, and self-intersections might happen in scaling limits in higher dimensions.

by maps $f : \tau \rightarrow \hat{\mathbb{C}}$ that are constants on the contracted edges. By contractions in two non-isomorphic trees, τ_1 and τ_2 , we may reach the same tree $\tau' \prec \tau_i$, hence $\mathcal{S}^{(\ell)}$ may be viewed as covered by patches \mathcal{S}_τ that are sewn together along “smaller dimensional” patches $\mathcal{S}_{\tau'}$, similarly to a simplicial complex. (In particular, after these identifications, (2.2) stops being a disjoint union.)

We now equip each \mathcal{S}_τ with a very natural metric, extending the notion of uniform closeness up to reparametrization of curves: for two immersed trees $f_1, f_2 : \tau \rightarrow \hat{\mathbb{C}}$,

$$\text{dist}_\tau(f_1, f_2) = \inf_{\phi_1, \phi_2} \sup_{t \in \tau} \text{dist}_{\hat{\mathbb{C}}}(f_1 \circ \phi_1(t), f_2 \circ \phi_2(t)), \quad (2.3)$$

where the ϕ_i 's run over all reparametrizations of τ . This can be easily extended to immersed trees indexed by different reference trees: by the above remark about patches, for any pair of reference trees τ, τ' there exist sequences $\tau = \tau_0, \tau_1, \dots, \tau_m = \tau'$ such that $\tau_i \prec \tau_{i+1}$ or $\tau_i \succ \tau_{i+1}$ for all $i = 0, \dots, m-1$, and then for any $f : \tau \rightarrow \hat{\mathbb{C}}$ and $f' : \tau' \rightarrow \hat{\mathbb{C}}$ we can take

$$\text{dist}(f, f') = \inf \left\{ \sum_{i=0}^{m-1} \text{dist}_{\tau_i \vee \tau_{i+1}}(f_i, f_{i+1}) : \right. \\ \left. f_0 = f, f_m = f', \tau_i \xrightarrow{f_i} \hat{\mathbb{C}} \text{ for } i = 1, \dots, m-1 \right\},$$

where, with a rather obvious notation, $\tau_i \vee \tau_j = \tau_i$ if $\tau_i \succ \tau_j$. Note that for any τ, τ' there exists a τ'' with $\tau, \tau' \prec \tau''$ and $|\xi(\tau'')| \leq |\xi(\tau)| + |\xi(\tau')| -$

2, hence distances between immersed trees in $\mathcal{S}^{(\ell)}$ can be measured inside $\mathcal{S}^{(2\ell-2)}$, and $\text{dist}(f, f') \leq \text{dist}_{\tau''}(f, f') < \infty$.

With this metric, $\mathcal{S}^{(\ell)}$ is clearly a complete separable metric space, called the **space of ℓ -trees**. Of course, a Cauchy sequence of trees contained fully in \mathbb{C} might have a limit that has an edge going through ∞ . Similarly, $\mathcal{S}_M^{(\ell)}$ is complete and separable with the analogous metric, just using the Euclidean metric on \mathbb{T}_M^2 in (2.3).

Now that we have a definition for the space of finite trees immersed in $\hat{\mathbb{C}}$ or \mathbb{T}_M^2 , we can start defining what a spanning tree of $\hat{\mathbb{C}}$ or \mathbb{T}_M^2 should be: a set of finite trees that satisfy certain compatibility conditions.

The set of non-empty closed subsets of $\mathcal{S}^{(\ell)}$ in the above metric, equipped with the Hausdorff metric, is denoted by $\Omega^{(\ell)}$. We will consider graded sets

$$\mathcal{F} = (\mathcal{F}^{(\ell)})_{\ell \geq 1} \in \Omega^\times := \prod_{\ell \geq 1} \Omega^{(\ell)},$$

with the product topology. Clearly, Ω^\times is again complete, separable and metrizable; in one word, it is a Polish metric space.

Extending the map $\tau \mapsto \xi(\tau)$ giving the external vertices of an index tree, for any $\mathcal{F} \in \Omega^\times$ we can define

$$\xi(\mathcal{F}) := \bigcup \left\{ f(\xi(\tau)) : \tau \xrightarrow{f} \hat{\mathbb{C}} \in \mathcal{F}^{(\ell)}, \ell \geq 1 \right\} \subset \hat{\mathbb{C}},$$

which gives the set of external vertices occurring in \mathcal{F} . It is clearly a Borel measurable function, since for any open $U \subset \hat{\mathbb{C}}$, the preimage $\xi^{-1}(U)$ is a countable intersection (over $\ell \geq 1$) of open sets.

Let $\mathcal{S}_{B_1, \dots, B_\ell}$ be the set of immersed trees with endpoints $x_i \in B_i$, where each B_i is a closed subset of $\hat{\mathbb{C}}$. Note that this is a closed subset of $\mathcal{S}^{(\ell)}$, hence the map

$$\Omega^\times \longrightarrow \Omega_{B_1, \dots, B_\ell} \subseteq \Omega^{(\ell)}, \quad \mathcal{F} \mapsto \mathcal{F}^{(\ell)} \cap \mathcal{S}_{B_1, \dots, B_\ell}$$

is measurable. In words, extracting the subtrees of \mathcal{F} with leaves in prescribed closed sets (e.g., the branches of \mathcal{F} connecting two given points) is a measurable map.

Definition 2.2. A graded set $\mathcal{F} = (\mathcal{F}^{(\ell)})_{\ell \geq 1} \in \Omega^\times$ is called an **essential spanning forest** on its external vertices $\xi(\mathcal{F})$ if it satisfies the following properties:

- (i) for each $\ell \in \mathbb{N}^+$ and any ℓ -tuple $\{x_1, \dots, x_\ell\}$ of vertices in $\xi(\mathcal{F})$, there exists at least one immersed tree $T(x_1, \dots, x_\ell) \in \mathcal{F}^{(\ell)}$ with those leaves;

- (ii) for any immersed tree $T \in \mathcal{F}^{(\ell)}$, any subtree $T' \subset T$ (given by restricting the immersion to a combinatorial subtree of the index tree τ) is again in some $\mathcal{F}^{(\ell)}$;
- (iii) for any two trees $T_i \in \mathcal{F}^{(\ell_i)}$, $i = 1, 2$, there is a tree in some $\mathcal{F}^{(\ell)}$ that contains both T_i 's as subtrees and has no leaves beyond those of the T_i 's.

Note that (ii) implies that $\xi(\mathcal{F})$ contains all the vertices of all the embedded trees.

An essential spanning forest \mathcal{F} is called a **spanning tree** if $\xi(\mathcal{F}) \subset \mathbb{C}$ and every path $T(x, y) \in \mathcal{F}^{(2)}$ stays within a bounded region of \mathbb{C} . A spanning tree is called **quasi-local** if for any bounded $\Lambda \subset \mathbb{C}$ there exists a bounded domain $\bar{\Lambda}(\mathcal{F}, \Lambda) \subset \mathbb{C}$ such that every tree of \mathcal{F} with leaves in Λ is contained in $\bar{\Lambda}$.

The set of essential spanning forests in $\hat{\mathbb{C}}$ (with an arbitrary set of vertices $\xi(\mathcal{F})$) will be denoted by Ω . It is easy to check that Ω is a closed subset of the Polish space Ω^\times , hence itself is Polish. A simple explicit **metric**, denoted by d_Ω , is given by the restriction from Ω^\times to Ω of the sum over ℓ of the Hausdorff distance on $\mathcal{S}^{(\ell)}$ multiplied by the weight $2^{-\ell}$.

For the tori \mathbb{T}_M^2 , the spaces $\Omega_M^{(\ell)}$, Ω_M^\times , Ω_M are defined analogously, with the only difference being that any essential spanning forest here is a single tree. The metric d_{Ω_M} is defined the same way as d_Ω .

The only way in which two vertices may be disconnected in an essential spanning forest \mathcal{F} in $\hat{\mathbb{C}}$ is that all the paths between them go through ∞ ; therefore, either \mathcal{F} is a spanning tree, or no component of it is contained in a bounded domain of \mathbb{C} . This is the property that the adjective ‘‘essential’’ for these spanning forests refers to. (In the setting of discrete infinite graphs, this reduces to saying that all components of the forest are infinite trees.) Also, note that the above definition allows for having more than one path between two vertices, which will in fact happen in the scaling limit of the MST.

2.2 The quad-crossing topology

Let us quickly recall the notation and the basic results for the quad-crossing topology of percolation configurations, introduced in [49] and studied further in [28, 29].

Let $D \subset \hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ be open, or be equal to the torus \mathbb{T}_M^2 . A **quad** in the domain D can be considered as a homeomorphism Q from $[0, 1]^2$ into D . The space of all quads in D , denoted by \mathcal{Q}_D , can be equipped with the

following metric: $d_{\mathcal{Q}}(Q_1, Q_2) := \inf_{\phi} \sup_{z \in \partial[0,1]^2} |Q_1(z) - Q_2(\phi(z))|$, where the infimum is over all homeomorphisms $\phi : [0, 1]^2 \rightarrow [0, 1]^2$ which preserve the 4 corners of the square. A **crossing** of a quad Q is a connected closed subset of $[Q] := Q([0, 1]^2)$ that intersects both $\partial_1 Q = Q(\{0\} \times [0, 1])$ and $\partial_3 Q = Q(\{1\} \times [0, 1])$. We say that Q has a **dual crossing** between $\partial_1 Q$ and $\partial_3 Q$ by some closed subset $S \subseteq [Q]$ if there is no crossing in S between $\partial_2 Q = Q([0, 1] \times \{0\})$ and $\partial_4 Q = Q([0, 1] \times \{1\})$.

From the point of view of crossings, there is a natural partial order on \mathcal{Q}_D : we write $Q_1 \leq Q_2$ if any crossing of Q_2 contains a crossing of Q_1 . Furthermore, we write $Q_1 < Q_2$ if there are open neighborhoods \mathcal{N}_i of Q_i (in the uniform metric) such that $N_1 \leq N_2$ holds for any $N_i \in \mathcal{N}_i$. A subset $S \subset \mathcal{Q}_D$ is called **hereditary** if whenever $Q \in S$ and $Q' \in \mathcal{Q}_D$ satisfies $Q' < Q$, we also have $Q' \in S$. The collection of all closed hereditary subsets of \mathcal{Q}_D will be denoted by \mathcal{H}_D . Any discrete percolation configuration ω_η of mesh $\eta > 0$, considered as a union of the topologically closed percolation-wise open hexagons in the plane, naturally defines an element $S(\omega_\eta)$ of \mathcal{H}_D : the set of all quads for which ω_η contains a crossing. In particular, near-critical percolation at level $\lambda \in \mathbb{R}$, as defined in Definition 1.2, induces a probability measure on \mathcal{H}_D , which will be denoted by \mathbf{P}_η^λ .

By introducing a natural topology, \mathcal{H}_D can be made into a compact metric space. Indeed, let

$$\Xi_Q := \{S \in \mathcal{H}_D : Q \in S\} \quad \text{for any } Q \in \mathcal{Q}_D,$$

and let

$$\Xi_U := \{S \in \mathcal{H}_D : S \cap U = \emptyset\} \quad \text{for any open } U \subset \mathcal{Q}_D.$$

Then, define \mathcal{T}_D to be the minimal topology that contains every Ξ_Q and Ξ_U^c as open sets. It is proved in [49, Theorem 3.10] that for any nonempty open D , the topological space $(\mathcal{H}_D, \mathcal{T}_D)$ is compact, Hausdorff, and metrizable. Furthermore, for any dense $\mathcal{Q}_0 \subset \mathcal{Q}_D$, the events $\{\Xi_Q : Q \in \mathcal{Q}_0\}$ generate the Borel σ -field of \mathcal{H}_D . An arbitrary metric generating the topology \mathcal{T}_D will be denoted by $d_{\mathcal{H}}$. Now, since Borel probability measures on a compact metric space are always tight, we have subsequential scaling limits of \mathbf{P}_η^λ on \mathcal{H}_D , as $\eta = \eta_k \rightarrow 0$. Moreover, the following convergence of probabilities holds. For critical percolation, $\lambda = 0$, it is Corollary 5.2 of [49]; for general λ , the exact same proof works, using that the RSW estimates hold in near-critical percolation.

Lemma 2.3. *For any $\lambda \in \mathbb{R}$, any subsequential scaling limit $\mathbf{P}_{\eta_k}^\lambda \rightarrow \mathbf{P}_\infty^\lambda$, and any quad $Q \in \mathcal{Q}_D$, one has $\mathbf{P}_\infty^\lambda[\partial\Xi_Q] = 0$. Therefore, by the weak*

convergence of $\mathbf{P}_{\eta_k}^\lambda$ to $\mathbf{P}_\infty^\lambda$,

$$\mathbf{P}_{\eta_k}^\lambda[\Xi_Q] \rightarrow \mathbf{P}_\infty^\lambda[\Xi_Q].$$

For the case of site percolation on $\eta\mathbb{T}$, we know much more than just the existence of subsequential limits. As explained in [28, Subsection 2.3], the existence of a unique quad-crossing scaling limit for $\lambda = 0$ follows from the loop scaling limit result of [51, 19]. The case of general λ is Theorem 1.4 of [29]:

Theorem 2.4 (Near-critical scaling limit). *For any $\lambda \in \mathbb{R}$, there is a unique measure $\mathbf{P}_\infty^\lambda$ for percolation configurations ω_∞^λ in $(\mathcal{H}_D, \mathcal{T}_D)$ such that the weak convergence $\omega_\eta^\lambda \xrightarrow{d} \omega_\infty^\lambda$ holds.*

We have shown in [28] that the arm events between the boundary pieces of an annulus are measurable w.r.t. the quad-crossing topology, and the convergence of probabilities (analogous to Lemma 2.3) holds. Namely, for any topological annulus $A \subset D$ with piecewise smooth inner and outer boundary pieces $\partial_1 A$ and $\partial_2 A$ (and for the case of $D = \mathbb{T}_M^2$, we also require A to be null-homotopic), we define the **alternating 4-arm event** in A as $\mathcal{A}_4 = \bigcup_{\delta > 0} \mathcal{A}_4^\delta$, where \mathcal{A}_4^δ is the existence of quads $Q_i \subset D$, $i = 1, 2, 3, 4$, with the following properties (see the left side of Figure 2.2):

- (i) Q_1 and Q_3 are disjoint and are at distance at least δ from each other; the same for Q_2 and Q_4 ;
- (ii) for $i \in \{1, 3\}$, the sides $\partial_1 Q_i = Q_i(\{0\} \times [0, 1])$ lie inside $\partial_1 A$ and the sides $\partial_3 Q_i = Q_i(\{1\} \times [0, 1])$ lie outside $\partial_2 A$; for $i \in \{2, 4\}$, the sides $\partial_2 Q_i = Q_i([0, 1] \times \{0\})$ lie inside $\partial_1 A$ and the sides $\partial_4 Q_i = Q_i([0, 1] \times \{1\})$ lie outside $\partial_2 A$; all these sides are at distance at least δ from the annulus A and from the other Q_j 's;
- (iii) the four quads are ordered cyclically around A according to their indices;
- (iv) For $i \in \{1, 3\}$, we have $\omega \in \Xi_{Q_i}$, while for $i \in \{2, 4\}$, we have $\omega \in \Xi_{Q_i}^c$. In plain words, the quads Q_1, Q_3 are crossed, while the quads Q_2, Q_4 are dual crossed between the boundary pieces of A , with a margin δ of safety.

The definitions of general (**mono- or polychromatic**) k -arm events in A are of course analogous: for arms of the same color we require the corresponding quads to be completely disjoint, and we still require all the boundary pieces lying outside the annulus A to be disjoint.




Figure 2.2: Defining the alternating 4-arm event (in Subsection 2.2) and the color of a pivotal point (in Subsection 2.3) using quad-crossings. Quads with a primal crossing are in solid red, quads with a dual crossing are in dashed blue.

The next lemma is proved for critical percolation in [28, Lemma 2.9]. For near-critical percolation, the same proofs work, using the stability of multi-arm probabilities (see [29, Lemma 8.4, Proposition 11.6] or [38]), together with the existence of the near-critical scaling limit [29, Theorem 1.4].

Lemma 2.5. *Let $A \subset D$ be a piecewise smooth topological annulus (with finitely many non-smooth boundary points). Then the 1-arm, the alternating 4-arm and any polychromatic 6-arm event in A , denoted by \mathcal{A}_1 , \mathcal{A}_4 and \mathcal{A}_6 , respectively, are measurable w.r.t. the scaling limit of critical percolation in D , and one has*

$$\lim_{\eta \rightarrow 0} \mathbf{P}_\eta^\lambda[\mathcal{A}_i] = \mathbf{P}_\infty^\lambda[\mathcal{A}_i].$$

Moreover, in any coupling of the measures $\{\mathbf{P}_\eta^\lambda\}$ and $\mathbf{P}_\infty^\lambda$ on $(\mathcal{H}_D, \mathcal{T}_D)$ in which $\omega_\eta^\lambda \rightarrow \omega_\infty^\lambda$ a.s. as $\eta \rightarrow 0$, we have

$$\mathbf{P}[\{\omega_\eta^\lambda \in \mathcal{A}_i\} \Delta \{\omega_\infty^\lambda \in \mathcal{A}_i\}] \rightarrow 0 \quad (\text{as } \eta \rightarrow 0). \quad (2.4)$$

2.3 Pivotal and pivotal measures

In [29], we managed to describe the changes of macroscopic connectivities in a percolation configuration under the stationary or the asymmetric near-critical dynamics using just the pivotal measures of [28], without making explicit use of notions like clusters or the set of pivotal sites in continuum percolation. Unfortunately, the situation is slightly more complicated for the models in the present paper, hence we need some foundational work in addition to what was done in [28, Section 2.4].

Let x be a point surrounded (with a positive distance) by a piecewise smooth Jordan curve $\gamma \subset D$, where “surrounded” means that $D \setminus \gamma$ has two connected components, with the one containing x being homeomorphic to a disk. For any $\epsilon > 0$, fix a lattice $\epsilon\mathbb{Z}^2$ in D , and let $B_\epsilon(x)$ be the ϵ -square $[i, i+1)\epsilon \times [j, j+1)\epsilon$ in the lattice that contains x . We say that x is **pivotal for** γ in ω_∞^λ if, for any $\epsilon > 0$ such that $B_{2\epsilon}(x)$ is surrounded by γ , the alternating 4-arm event occurs in the annulus with boundary pieces $\partial B_{2\epsilon}(x)$ and γ , as defined in Subsection 2.2. We let \mathcal{P}^γ denote the set of pivotal points for γ in D . Furthermore, we can identify the **color of a pivotal point** $x \in \mathcal{P}^\gamma$ as **open** (black) versus **closed** (white, empty), as follows. We let $\mathcal{P}_{\text{open}}^{\gamma, \epsilon}$ denote the set of points x for which γ surrounds x without intersecting or touching $B_{2\epsilon}(x)$, and there exist quads $Q_{\epsilon, i}$, $i = 1, 2, 3, 4$, exhibiting the 4-arm event from $\partial B_{2\epsilon}(x)$ to γ such that the quad \bar{U}_ϵ , given by taking the union of $U_\epsilon := Q_{\epsilon, 1} \cup Q_{\epsilon, 3} \cup B_{2\epsilon}(x)$ and the bounded components of $\mathbb{C} \setminus U_\epsilon$, is crossed between the boundary pieces $Q_{\epsilon, 1}(\{1\} \times [0, 1])$ and $Q_{\epsilon, 3}(\{1\} \times [0, 1])$; see the right side of Figure 2.2 in the previous subsection. Then, we let the set of open pivotals for γ be

$$\mathcal{P}_{\text{open}}^\gamma := \{x \in D : x \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon} \text{ for all } \epsilon > 0 \text{ s.t. } B_{2\epsilon}(x) \text{ is surrounded by } \gamma\}.$$

Clearly, the event $x \in \mathcal{P}_{\text{open}}^\gamma$ is measurable w.r.t. the quad-crossing topology. We will use the notation $x \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon, \delta}$ for the event that all the crossing events in $\mathcal{P}_{\text{open}}^{\gamma, \epsilon}$ are satisfied even with a δ margin of safety. Finally, we set $x \in \mathcal{P}_{\text{closed}}^\gamma$ if the analogous dual crossing holds in the quad given by $Q_{\epsilon, 2} \cup Q_{\epsilon, 4} \cup B_{2\epsilon}(x)$, for each small enough $\epsilon > 0$.

For a discrete percolation configuration ω_η^λ , the above definitions do not work: instead of taking all small enough $\epsilon > 0$, we just need to take the annulus between γ and the hexagon of the point $x \in D$. In particular, the previous definition would render every η -hexagon that intersects γ to be pivotal automatically, something that we will not do. Also, here it is clear what the sets $\mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda)$ and $\mathcal{P}_{\text{closed}}^\gamma(\omega_\eta^\lambda)$ are: their disjoint union is the set of pivotal hexagons $\mathcal{P}^\gamma(\omega_\eta^\lambda)$, and the color is determined by the color of the hexagon itself. We will nevertheless use notation like $x \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon}(\omega_\eta^\lambda)$: it has the meaning given above, using quad-crossings, and of course it cannot hold unless η is small enough, say $\epsilon > \eta$, so that $\partial B_{2\epsilon}(x)$ already intersects at least four η -hexagons.

Proposition 2.6 (The set of pivotals, with colors). *In any coupling of the measures $\{\mathbf{P}_\eta^\lambda\}$ and $\mathbf{P}_\infty^\lambda$ on $(\mathcal{H}_D, \mathcal{T}_D)$ in which $\omega_\eta^\lambda \xrightarrow{a.s.} \omega_\infty^\lambda$ as $\eta \rightarrow 0$, for any piecewise smooth null-homotopic Jordan curve $\gamma \subset D$ we have the following statements:*

- (i) $\mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda)$ converges in probability to $\mathcal{P}_{\text{open}}^\gamma(\omega_\infty^\lambda)$ in the Hausdorff metric of closed sets. Same for $\mathcal{P}_{\text{closed}}^\gamma$ and \mathcal{P}^γ .
- (ii) Almost surely, $\mathcal{P}_{\text{open}}^\gamma(\omega_\infty^\lambda) \cup \mathcal{P}_{\text{closed}}^\gamma(\omega_\infty^\lambda) = \mathcal{P}^\gamma(\omega_\infty^\lambda)$, a disjoint union.
- (iii) Almost surely, whenever $x \in \mathcal{P}^\gamma(\omega_\infty^\lambda)$ for some γ , the color of x is the same for all such γ .

In (i) it is not obvious that the sets $\mathcal{P}_{\text{open}}^\gamma(\omega_\infty^\lambda)$, $\mathcal{P}_{\text{closed}}^\gamma(\omega_\infty^\lambda)$, $\mathcal{P}^\gamma(\omega_\infty^\lambda)$ are topologically closed, but this will follow from the proof. Note also that (ii) is not a tautology (neither that the two colored sets are disjoint, nor that their union is the set of all the pivotals), since in ω_∞^λ we did not define the set of closed pivotals as the complement of open pivotals.

The main difficulty in proving (i) is that the event $x \in \mathcal{P}_{\text{open}}^\gamma$ is not an open set in the quad-crossing topology $(\mathcal{H}_D, \mathcal{T}_D)$: perturbing a configuration even by an arbitrary small amount may destroy a pivotal for γ , making the 4-arm event happen only from a strictly positive distance $\epsilon > 0$ to γ . In terms of discrete percolation configurations, if there is an open pivotal connecting two halves of a cluster, then making the connection between the two halves a bit thicker is a small change w.r.t. the quad-crossing topology, but it kills the pivotal. In particular, the harder direction in (i) will be to prove that there are “enough” pivotals in ω_∞^λ , since this requires controlling all scales simultaneously.

Proof. For (i), we need to prove that for any $\epsilon > 0$, if $\eta > 0$ is small enough, then with probability at least $1 - \epsilon$, for every $x_\eta \in \mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda)$ there exists some $x \in \mathcal{P}_{\text{open}}^\gamma(\omega_\infty^\lambda)$ within distance ϵ from x_η , and vice versa, for every $x \in \mathcal{P}_{\text{open}}^\gamma(\omega_\infty^\lambda)$ there exists $x_\eta \in \mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda)$.

There will be three key ingredients. Firstly, we have the following relationship between the discrete and the quad-crossing definitions of open pivotality: for every $\alpha > 0$ if $\epsilon > 0$ is small enough, then for all $\eta > 0$ small enough,

$$\mathbf{P}[\mathcal{P}_{\text{open}}^{\gamma, \epsilon}(\omega_\eta^\lambda) \supseteq \mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda)] > 1 - \alpha. \quad (2.5)$$

To see this, note that, since γ is piecewise smooth, the 3-arm half-plane exponent being 2 (larger than 1) implies that with high probability there are no pivotals close to γ . (See, e.g., the end of the proof of Theorem 1.1 in [28, Section 4.6].) That is, for ϵ small enough, $B_{2\epsilon}(x)$ is surrounded by γ for every $x \in \mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda)$. It is also easy to check that the condition $x \in \mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda)$ implies the existence of the necessary quads $Q_{\epsilon,1}$ and $Q_{\epsilon,3}$ in the definition of $\mathcal{P}_{\text{open}}^{\gamma, \epsilon}(\omega_\eta^\lambda)$: taking either open arm from x to γ , and the arc

on $\partial B_{2\epsilon}(x)$ between its first and last exit from $B_{2\epsilon}(x)$, the quad $Q_{\epsilon,i}$ should contain this arc.

Secondly, for any small $\alpha, \epsilon > 0$ there exists $\delta, \bar{\eta} > 0$ such that for all $0 < \eta < \bar{\eta}$,

$$\mathbf{P}[\mathcal{P}_{\text{open}}^{\gamma,\epsilon}(\omega_\eta^\lambda) = \mathcal{P}_{\text{open}}^{\gamma,\epsilon,\delta}(\omega_\eta^\lambda)] > 1 - \alpha. \quad (2.6)$$

The existence of a δ that still depends on $x \in \mathcal{P}_{\text{open}}^{\gamma,\epsilon}(\omega_\eta^\lambda)$, or rather on its lattice square $B_{2\epsilon}(x)$, is just a special case of [28, Corollary 2.10]. Then, taking the probability α of the error much smaller than ϵ^2 , we can find a $\delta > 0$ that, with large probability, works for all points in $\mathcal{P}_{\text{open}}^{\gamma,\epsilon}(\omega_\eta^\lambda)$ simultaneously, proving (2.6).

The point of introducing the δ margin of safety is that now (2.6) immediately implies that there exists some monotone function $f = f_{\alpha,\epsilon} : [0, \infty) \rightarrow [0, \infty)$ that could be described using the dyadic uniformity structures of [28, Lemma 2.5] and [29, Proposition 3.9]) such that

$$\begin{aligned} \mathbf{P}[\forall x \in \mathcal{P}_{\text{open}}^{\gamma,\epsilon}(\omega_\eta^\lambda) \text{ and } \forall \tilde{\omega} \in \mathcal{H}_D \text{ with } d_{\mathcal{H}}(\tilde{\omega}, \omega_\eta^\lambda) < f(\delta), \\ \text{we have } x \in \mathcal{P}_{\text{open}}^{\gamma,\epsilon,\delta/2}(\tilde{\omega})] > 1 - \alpha, \end{aligned} \quad (2.7)$$

for some $\delta > 0$ and any $0 < \eta < \bar{\eta}$, as given by (2.6).

The third key ingredient is that for any small $\alpha, \beta > 0$, if $\epsilon, \hat{\eta} > 0$ are small enough, then

$$\mathbf{P}[\forall x \in \mathcal{P}_{\text{open}}^{\gamma,\epsilon}(\omega_\eta^\lambda) \exists \tilde{x} \in \mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda) \text{ with } d(\tilde{x}, x) < \beta] > 1 - \alpha \quad (2.8)$$

for all $0 < \eta < \hat{\eta}$. Before proving this, let us see how (2.5), (2.7) and (2.8) imply item (i). We start with the first direction.

Fix $\alpha, \beta > 0$ small. Corresponding to them, (2.8) gives some $\epsilon_0, \hat{\eta}_0 > 0$. We also make sure that they are small enough for (2.5) to hold. Now, corresponding to α and this ϵ_0 , there are $\delta_0, \bar{\eta}_0 > 0$ given by (2.7). Take $0 < \eta_0 < \bar{\eta}_0 \wedge \hat{\eta}_0$ so small that

$$\mathbf{P}[d_{\mathcal{H}}(\omega_\eta^\lambda, \omega_\infty^\lambda) < f(\delta_0)/2 \text{ holds for all } \eta \leq \eta_0] > 1 - \alpha \quad (2.9)$$

in the coupling $\omega_\eta^\lambda \xrightarrow{a.s.} \omega_\infty^\lambda$ that we have. If the event of (2.9) holds, then we also have $d_{\mathcal{H}}(\omega_\eta^\lambda, \omega_{\eta_0}^\lambda) < f(\delta_0)$, and hence, together with (2.7) and (2.8), we get that, for all $\eta \leq \eta_0$,

$$\mathbf{P}[\forall x \in \mathcal{P}_{\text{open}}^{\gamma,\epsilon_0}(\omega_{\eta_0}^\lambda) \exists \tilde{x} \in \mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda) \text{ with } d(\tilde{x}, x) < \beta] > 1 - 3\alpha.$$

Similarly, for $k \geq 1$, corresponding to $\alpha/2^k$ and $\beta/2^k$, there are $\epsilon_k, \hat{\eta}_k > 0$ given by (2.8); we can make sure that $\epsilon_k < \epsilon_{k-1}/2$. Then, corresponding to

$\alpha/2^k$ and ϵ_k , there are $\delta_k, \bar{\eta}_k > 0$ given by (2.7). Take $0 < \eta_k < \eta_{k-1}/2 \wedge \bar{\eta}_k \wedge \hat{\eta}_k$ so small that $d_{\mathcal{H}}(\omega_\eta^\lambda, \omega_\infty^\lambda) < f(\delta_k)/2$ is satisfied for all $\eta \leq \eta_k$ with probability at least $1 - \alpha/2^k$. Then, for all $\eta < \eta_k$, (2.7) and (2.8) together give that

$$\mathbf{P}[\forall x \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon_k}(\omega_{\eta_k}^\lambda) \exists \tilde{x} \in \mathcal{P}_{\text{open}}^\gamma(\omega_\eta^\lambda) \text{ with } d(\tilde{x}, x) < \beta/2^k] > 1 - 3\alpha/2^k.$$

Iterating this procedure, we get that there exist sequences $\eta_k \rightarrow 0$ and $\epsilon_k \rightarrow 0$ such that with probability at least $1 - 3\alpha \sum_{k \geq 0} 2^{-k} = 1 - 6\alpha$, for any $x_0 \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon_0}(\omega_{\eta_0}^\lambda)$ there exist

$$x_k \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon_k}(\omega_{\eta_k}^\lambda) \text{ for } k = 0, 1, 2, \dots, \text{ satisfying } d(x_{k+1}, x_k) < \beta/2^k. \quad (2.10)$$

These points have a limit $x_k \rightarrow \tilde{x}_0$, which satisfies $d(x_0, \tilde{x}_0) < 2\beta$. Unsurprisingly, we claim that $\tilde{x}_0 \in \mathcal{P}_{\text{open}}^\gamma(\omega_\infty^\lambda)$. Indeed, otherwise there would exist some $\tilde{\epsilon} > 0$ such that $\tilde{x}_0 \notin \mathcal{P}_{\text{open}}^{\gamma, \tilde{\epsilon}}(\omega_\infty^\lambda)$, but for some small enough ϵ , this would clearly contradict the existence of an $\omega_{\eta_k}^\lambda$ satisfying $d_{\mathcal{H}}(\omega_{\eta_k}^\lambda, \omega_\infty^\lambda) < \epsilon$ and having an almost-pivotal $x_k \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon}(\omega_{\eta_k}^\lambda)$ at distance $d(x_k, \tilde{x}_0) < \epsilon$, which we have from (2.10). Since we can take α and β arbitrarily small, this argument and (2.5) together finish the proof of the first direction of item (i).

For the other direction, if $x \in \mathcal{P}_{\text{open}}^\gamma(\omega_\infty^\lambda)$, then, by definition, for all $\epsilon > 0$ with $B_{2\epsilon}(x)$ surrounded by γ , there is some $\delta > 0$ such that $x \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon, \delta}(\omega_\infty^\lambda)$. Now, if ω_η^λ is close enough to ω_∞^λ (again quantifiable in the sense of dyadic uniformity structures), then $x \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon, \delta/2}(\omega_\eta^\lambda)$ also occurs. By (2.8), if $\epsilon > 0$ is small enough, this means with large probability that there is an actual pivotal of ω_η^λ close to x , as required.

We still owe the proof of (2.8). Assume that $x \in \mathcal{P}_{\text{open}}^{\gamma, \epsilon}(\omega_\eta^\lambda)$ but there are no open pivotal sites in $B_\beta(x)$. This implies that there is a 6-arm event from $\partial B_{2\epsilon}(x)$ to $\partial B_\beta(x)$: the interfaces between the open and closed arms cannot touch each other within $B_\beta(x)$, hence their open sides form two disjoint open paths, creating four open arms besides the two closed ones; see the left side of Figure 2.3. Since the 6-arm exponent is strictly larger than 2 at any fixed near-critical level λ (see [50, Corollary A.8] for $\lambda = 0$, and [29, Proposition 11.6] or Proposition 1.3 in the present paper for general λ), we can take $\beta := \epsilon^\zeta$ with $\zeta > 0$ small enough for the polychromatic 6-arm probability satisfy $\alpha_6(2\epsilon, \beta) = o(\epsilon^2)$, and then the probability that such a 6-arm event occurs anywhere in the domain tends to zero as $\epsilon \rightarrow 0$, and we are done.

For $\mathcal{P}_{\text{closed}}^\gamma$, the proof is of course identical. For \mathcal{P}^γ , the harder direction (that there exist points of $\mathcal{P}^\gamma(\omega_\infty^\lambda)$ close to the points of $\mathcal{P}^\gamma(\omega_\eta^\lambda)$) is automatic from the same results for $\mathcal{P}_{\text{open}}^\gamma$ and $\mathcal{P}_{\text{closed}}^\gamma$, since in the discrete configuration ω_η every pivotal has a color. For the easier direction, the proof we gave for

$\mathcal{P}_{\text{open}}^\gamma$ via (2.8) did not in fact use that the point in $\mathcal{P}^\gamma(\omega_\infty^\lambda)$ had a definite color, only the 4-arm events on all scales, hence it also applies here.




Figure 2.3: Left: an open ϵ -almost pivotal event without actual pivots in the β -square implies a 6-arm event (four open and two closed arms) between radii ϵ and β . Right: having both an open and a closed γ -pivotal in an ϵ -square implies a 6-arm event from 2ϵ to γ .

In item (ii), the fact that the union of the two colored sets gives all the pivots follows immediately from the discrete analogue and item (i). To prove the disjointness claim, by part (i) it is enough to prove that the probability of having a closed and an open pivotal for γ within distance ϵ from each other goes to 0 as $\epsilon \rightarrow 0$. But this event implies the existence of 6 disjoint arms from ϵ to γ (see the right side of Figure 2.3), and hence, as usual, the 6-arm exponent being larger than 2 implies the claim.

For item (iii), if γ_1 and γ_2 both surround x , with $x \in \mathcal{P}_{\text{open}}^{\gamma_1}(\omega_\infty^\lambda) \cap \mathcal{P}_{\text{closed}}^{\gamma_2}(\omega_\infty^\lambda)$, then we would also have $x \in \mathcal{P}_{\text{open}}^\gamma(\omega_\infty^\lambda) \cap \mathcal{P}_{\text{closed}}^\gamma(\omega_\infty^\lambda)$, where γ is the part of $\gamma_1 \cup \gamma_2$ that is visible from x . However, this is impossible by item (ii). \square

Beyond the set of pivots, we are also interested in the normalized counting measure on them. In [29, Subsection 2.6], for any fixed $\epsilon > 0$, we defined the set of ϵ -**important points** $\mathcal{P}^\epsilon(\omega_\eta)$ of any discrete percolation configuration in a bounded domain $D \subset \hat{\mathbb{C}}$, relative to the $(\epsilon, 3\epsilon)$ -annuli given by a fixed lattice $\epsilon\mathbb{Z}^2$. Namely, for $x \in D$, we let $B_\epsilon(x)$ be the lattice square as before, let $\tilde{B}_\epsilon(x)$ be the 3ϵ -square centered at $B_\epsilon(x)$, and let $x \in \mathcal{P}^\epsilon$ iff $x \in \mathcal{P}^{\partial\tilde{B}_\epsilon(x)}$. Then we considered the normalized counting measure $\mu^\epsilon(\omega_\eta)$ on this set \mathcal{P}^ϵ . Of course, the same discrete definition works for near-critical percolation configurations ω_η^λ . Then, the main result of [28] is the following

convergence of μ^ϵ for $\lambda = 0$, extended to general $\lambda \in \mathbb{R}$ by [29, Theorem 11.5]:

Theorem 2.7. *For any $\epsilon > 0$, there exists a measurable map $\mu^\epsilon : \mathcal{H}_D \rightarrow \mathfrak{M}_D$, into the space of finite Borel measures on D , such that, for $\lambda \in \mathbb{R}$, as $\eta \rightarrow 0$,*

$$(\omega_\eta^\lambda, \mu^\epsilon(\omega_\eta^\lambda)) \xrightarrow{d} (\omega_\infty^\lambda, \mu^\epsilon(\omega_\infty^\lambda))$$

in the quad-crossing topology $(\mathcal{H}_D, \mathcal{T}_D)$ in the first coordinate and in the Lévy-Prokhorov distance of measures in the second one. Furthermore, the above Proposition 2.6 implies immediately the convergence

$$(\mathcal{P}_{\text{open}}^\epsilon(\omega_\eta^\lambda), \mathcal{P}_{\text{closed}}^\epsilon(\omega_\eta^\lambda)) \xrightarrow{d} (\mathcal{P}_{\text{open}}^\epsilon(\omega_\infty^\lambda), \mathcal{P}_{\text{closed}}^\epsilon(\omega_\infty^\lambda))$$

in the Hausdorff metric of closed sets.

3 Enhanced networks and cut-off forests built from the near-critical ensemble

The pivotal measures of [28] that we recalled in Theorem 2.7 were used in [29] as the intensity measures for the Poisson point processes of pivotal sites that switch as the near-critical parameter $\lambda \in \mathbb{R}$ changes. Here is the exact notation that we will use:

Definition 3.1. *Let $\bar{\lambda} = (\lambda, \lambda') \in \mathbb{R}^2$ be any pair of near-critical parameters with $\lambda < \lambda'$, and let $\epsilon > 0$ be fixed. Let ω^λ be a near-critical configuration ω_η^λ or ω_∞^λ in \mathbb{T}_M^2 . We will denote by $\text{PPP}_{\bar{\lambda}}^\epsilon = \text{PPP}_{\bar{\lambda}}(\mu^\epsilon(\omega^\lambda))$ the Poisson point process*

$$\text{PPP}_{\bar{\lambda}}^\epsilon = \{(x_i, t_i), 1 \leq i \leq p\} \subset \mathcal{P}^\epsilon(\omega^\lambda) \times [\lambda, \lambda']$$

of intensity measure $\mu^\epsilon(\omega^\lambda)(dx) \times \mathbf{1}_{[\lambda, \lambda']}(t) dt$. The set $\{x_1, \dots, x_p\}$ of pivots will usually be denoted by X . For the case of ω_η^λ , the process $\text{PPP}_{\bar{\lambda}}^\epsilon$ can clearly be constructed measurably from $\omega_\eta^{[\lambda, \lambda']}$, and we will always work in this natural coupling.

In Section 6 and Subsection 11.2 of [29], for any quad $Q \subset \mathbb{C}$, any $\epsilon > 0$, any discrete or continuum near-critical percolation configuration ω^λ and the associated Poisson point process $\text{PPP}_{\bar{\lambda}}^\epsilon(\omega^\lambda)$, we constructed an edge-colored graph $N_Q(\omega^\lambda, \text{PPP}_{\bar{\lambda}}^\epsilon)$, called an ϵ -network, whose vertex set was the Poisson point set $X = \{x_1, \dots, x_p\}$ of pivots together with the four boundary arcs of Q , and whose edge set was given by the primal and dual connections in

ω^λ between the vertices. Since in this paper we are primarily interested in spanning trees, not in quad-crossings, it will be useful to change the boundary conditions in the definition slightly (but still using the quad-crossing topology). We will also need to add a bit more structure to these networks: roughly, we will need to know which pivotals in X are connected together by an open cluster of $\omega^\lambda \setminus X$, and will need to know the colors of these pivotals in ω^λ . The resulting structures will be called enhanced networks. Just as in [29], we start with the following simple definition:

Definition 3.2 (A nested family of dyadic coverings). *For any $b > 0$ in $2^{-\mathbb{N}} = \{2^{-k} : k = 0, 1, 2, \dots\}$, let G_b be a disjoint covering of \mathbb{T}_M^2 using the lattice b -squares $\{[k, k+1)b \times [\ell, \ell+1)b : (k, \ell) \in \mathbb{Z}^2\}$. Now, for any $r \in 2^{-\mathbb{N}}$ and any finite subset $X = \{x_1, \dots, x_p\} \subset \mathbb{T}_M^2$, one can associate uniquely r -squares $B_{x_1}^r, \dots, B_{x_p}^r$ in the following manner: for all $1 \leq i \leq p$, there is a unique square $\tilde{B}_{x_i} \in G_{r/2}$ which contains x_i and we define $B_{x_i}^r$ to be the r -square in the grid $r\mathbb{Z}^2 - (r/4, r/4)$ centered around the $r/2$ -square \tilde{B}_{x_i} . We will denote by $B^r(X)$ this family of r -squares. This family has the following two properties:*

- (i) *Each point x_i is at distance at least $r/4$ from $\partial B_{x_i}^r$.*
- (ii) *For any set X , $\{B^r(X)\}_{r \in 2^{-\mathbb{N}}}$ forms a nested family of squares in the sense that for any $r_1 < r_2$ in $2^{-\mathbb{N}}$, and any $x \in X$, we have $B_x^{r_1} \subset B_x^{r_2}$.*

For a finite set of points $X \subset \mathbb{T}_M^2$, let $r^*(X) > 0$ denote one-tenth of the smallest distance between any pair $x_i \neq x_j \in X$. With minor changes from the case of a domain with a boundary to the case of a torus, it is proved in [29, Proposition 5.2] that for X being the pivotals in $\text{PPP}_\lambda^\epsilon$, the random variable $r^*(\text{PPP}_\lambda^\epsilon)$ is almost surely positive (with a small abuse of notation, since $\text{PPP}_\lambda^\epsilon$ is a subset of $\mathbb{T}_M^2 \times [\lambda, \lambda']$).

Definition 3.3. *For $0 < r < r^*(\text{PPP}_\lambda^\epsilon)$, the r -mesoscopic ϵ -network $\text{N}_M^{r\text{-meso}}(\omega^\lambda, \text{PPP}_\lambda^\epsilon)$ associated to a near-critical percolation configuration ω^λ in the torus \mathbb{T}_M^2 and the Poisson point process $\text{PPP}_\lambda^\epsilon$ of Definition 3.1 is the graph with vertex set $\{x_1, \dots, x_p\}$ and two types of edges, labelled primal or dual, with a primal edge connecting x_i and x_j if there exists a quad R such that $\partial_1 R$ and $\partial_3 R$ remain strictly inside $B_{x_i}^r$ and $B_{x_j}^r$, and R remains strictly away from the squares $B_{x_k}^r, k \notin \{i, j\}$, and for which $\omega^\lambda \in \Xi_R$. Dual edges are defined analogously (still w.r.t. ω^λ).*

We consider two r -mesoscopic networks to be the same if the r -squares for the vertices (as embedded in \mathbb{T}_M^2) and the labelled graph structures coincide. For $r_1 < r_2$, we can compare an r_1 -mesoscopic network with an

r_2 -mesoscopic network by considering the unique r_2 -squares containing the r_1 -squares of the first network.

We will now take $r \rightarrow 0$, get a network $N_M(\omega^\lambda, \text{PPP}_\lambda^\epsilon)$, and then compare these networks for ω_η^λ and ω_∞^λ . The following results were proved in [29, Theorem 6.14] and [29, Subsection 7.4] for $\lambda = 0$, extended to general λ in [29, Subsection 11.2], for networks defined using slightly different boundary conditions than here, but with the same proofs working fine:

Proposition 3.4 (r -stabilization and η -convergence of networks).

- (i) *There exists a measurable $r_M = r_M(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon) < r^*(\text{PPP}_\lambda^\epsilon(\omega_\infty^\lambda)) > 0$ such that for all $r \in (0, r_M)$ we get the same r -mesoscopic ϵ -network $N_M^{r\text{-meso}}(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$. This stabilized network will be called the ϵ -network $N_\infty^{\bar{\lambda}, \epsilon} = N_M(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$. For discrete percolation configurations, the definition of $N_\eta^{\bar{\lambda}, \epsilon} = N_M(\omega_\eta^\lambda, \text{PPP}_\lambda^\epsilon)$ is the obvious one.*
- (ii) *For any $\alpha > 0$ there is a scale $r_\alpha = r_\alpha(M, \bar{\lambda}, \epsilon)$ such that in any coupling with $\omega_\eta^\lambda \xrightarrow{a.s.} \omega_\infty^\lambda$ in \mathbb{T}_M^2 , for all sufficiently small $\eta > 0$ there is a coupling of $\text{PPP}_\lambda^\epsilon(\omega_\eta^\lambda)$ and $\text{PPP}_\lambda^\epsilon(\omega_\infty^\lambda)$ such that with probability at least $1 - \alpha$ the following holds: r_α is less than both $r_M(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon) < r^*(\text{PPP}_\lambda^\epsilon(\omega_\infty^\lambda))$ and $r^*(\text{PPP}_\lambda^\epsilon(\omega_\eta^\lambda))$, and for all $r < r_\alpha$ we have*

$$N_M^{r\text{-meso}}(\omega_\eta^\lambda, \text{PPP}_\lambda^\epsilon(\omega_\eta^\lambda)) = N_M^{r\text{-meso}}(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon(\omega_\infty^\lambda));$$

i.e., $N_\eta^{\bar{\lambda}, \epsilon} = N_M(\omega_\eta^\lambda, \text{PPP}_\lambda^\epsilon)$ coincides with $N_\infty^{\bar{\lambda}, \epsilon} = N_M(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$ in this sense. (Only in this sense, not exactly, since the vertex sets $\text{PPP}_\lambda^\epsilon(\omega_\infty^\lambda)$ and $\text{PPP}_\lambda^\epsilon(\omega_\eta^\lambda)$ are only close to each other, but do not coincide.)

Note that a network in itself may completely fail to describe the structure of clusters: see Figure 3.1. This is a bit of a problem for the purposes of the present paper, hence we are going to add some extra structure to our networks that will be measurable w.r.t. the quad-crossing topology (in particular, it makes sense for ω_∞^λ), while it describes how the pivotals of $\text{PPP}_\lambda^\epsilon$ are connected to each other in ω^λ .

Definition 3.5 (Mesoscopic sub-routers). *Fix $0 < r < \rho < \infty$. Utilizing the notation introduced in Definition 3.2, let $B^r(\mathbb{T}_M^2)$ be the finite covering of \mathbb{T}_M^2 by overlapping r -squares. Given a subset Y of the set $X = \{x_1, \dots, x_p\}$ of the pivotals in $\text{PPP}_\lambda^\epsilon$, with $|Y| \geq 2$, an (r, ρ) -mesoscopic sub-router for Y is an r -square $\bar{B} \in B^r(\mathbb{T}_M^2)$ with the following properties:*




Figure 3.1: The same graph structure in a network (the middle picture) may correspond to very different cluster structures (on the two sides).

- *it is at distance at least 2ρ from each $x_i \in X$;*
- *there is an open circuit (i.e., no dual arm) in the square annulus with inner face B and outer radius ρ ; the largest s -square with some $s \in 2^{-\mathbb{N}}$ that is concentric with B , contains it, and is surrounded by the open circuit will be denoted by \widehat{B} ;*
- *for each $x_i \in Y$, there exists a quad R with $\partial_1 R$ contained in \widehat{B} , $\partial_3 R$ contained in $B_{x_i}^r$, remaining strictly away from all the squares $B_{x_k}^r$ with $x_k \in X \setminus \{x_i\}$, and for which $\omega^\lambda \in \Xi_R$.*

Let $\mathcal{R}_Y(B)$ denote the event that an r -square B is an (r, ρ) -mesoscopic sub-router for some $Y \subseteq X$. This is measurable w.r.t. ω^λ , and using Lemmas 2.3 and 2.5, in the coupling of Proposition 3.4 (ii), the set of r -squares B for which $\mathcal{R}_Y(B)$ holds in ω_η^λ is the same with probability tending to 1 (as $\eta \rightarrow 0$) as in ω_∞^λ . Furthermore, by choosing (r, ρ) appropriately, this set will turn out to be non-empty with high probability, for all possible Y . For this, a key proposition, interesting in its own right, is the following:

Proposition 3.6 (The volume of clusters). *For any $\lambda \in \mathbb{R}$, $M > \rho > 0$ and $\zeta > 0$ fixed, for percolation ω_η^λ in \mathbb{T}_M^2 , with probability tending to 1 as $\eta \rightarrow 0$, all clusters of diameter at least ρ have at least $(\rho/\eta)^{91/48-\zeta}$ sites. (Note that $91/48$ equals 2 minus the one-arm exponent $5/48$ [40].)*

Similarly, with probability tending to 1 as $r \rightarrow 0$, uniformly in the mesh η , all these clusters have a “large r -volume” in the following sense: the number of r -squares in $B^r(\mathbb{T}_M^2)$ that intersect the cluster is at least $(\rho/r)^{91/48-\zeta}$.

After the first version of this paper was posted, Rob van den Berg pointed out that this proposition follows from (3.15) of [35]. However, since the proof there is hard to read, we decided to keep our proof for the sake of completeness. Earlier, similar but weaker results were proved in [37, Lemma 3.20] and [14, Theorem 3.3]. Finally, [13, Lemma 2.7] gives a bit more elegant version of our argument, but proving a little less; in particular, it

is not proved there that all the radial crossings of a $(\rho/3, \rho)$ -annulus are everywhere well-separated from each other (see our proof below).

Proof. The proof will rely only on multi-arm exponents, hence, in view of Proposition 1.3, the reader may just think of $\lambda = 0$. We will do the case of the standard volume (number of sites in the η -mesh); the proof works the same way for the case of the r -volume.

Take the lattice $(\rho/3)\mathbb{Z}^2$, and centered around each $\rho/3$ -square, consider the square of side-length ρ and the annulus between these two square boundaries. It is easy to check that any cluster of diameter at least ρ produces a radial crossing of such a $(\rho/3, \rho)$ -annulus. The number of such annuli is $\asymp (M/\rho)^2$.

Whether a given $(\rho/3, \rho)$ -annulus A_ρ is radially crossed can be decided using the radial exploration process started at any point along the boundary at radius $\rho/3$, with open hexagons on the right side, closed hexagons on the left, stopped when reaching the boundary at radius ρ . (See around Figure 2.6 of [28] or [57, Section 4.3] for the definition of this exploration process.) If the annulus is crossed, there are two cases: either (a) there is also an open circuit, or (b) there is also at least one radial dual crossing.

(a) Condition on having an open circuit; this is slightly more general than the first of the two above cases, since we do not condition on having also a radial crossing. Condition on the smallest open circuit, Γ . The radial exploration process finds it from inside, hence the configuration in the annulus between Γ and $\partial_2 A_\rho$, denoted by A_Γ , is undisturbed percolation. Moreover, by the half-plane 3-arm exponent being 2, the probability that the distance between Γ and $\partial_2 A_\rho$ is smaller than $\delta\rho$ is $O(\delta)$. Let this distance be the random variable $\delta_\Gamma\rho$, take any $0 < \delta < \delta_\Gamma$, and take the set of points of A_Γ whose distance from Γ is less than $\delta\rho$. It is clear that this set, denoted by $\tilde{A}_{\Gamma, \delta}$, contains a collection of $K = K(\delta) \geq c/\delta$ disjoint balls of diameter $\delta\rho$, denoted by \tilde{A}_i , $i = 1, \dots, K$, such that all their pairwise distances are at least $\delta\rho$; for instance, take a family of vertical parallel lines with mesh $\delta\rho$, and in every other slab, take the uppermost ball of diameter $\delta\rho$ that touches Γ . See the first picture in Figure 3.2. We will still fine-tune the value of δ later.

If a site in some \tilde{A}_i has an open arm to distance at least $c\delta\rho$, then with a uniformly positive probability it is connected to Γ , within the $\delta\rho/2$ -neighborhood of \tilde{A}_i that will be denoted by \tilde{B}_i . Vice versa, most sites in \tilde{A}_i need to have an arm of length $c\delta\rho$ in order to be connected to Γ . Thus, letting $X_i(\delta)$ be the number of sites in \tilde{A}_i that are connected to Γ within \tilde{B}_i , and using quasi-multiplicativity of the one-arm probability $\alpha_1(\cdot, \cdot)$ (see




Figure 3.2: If the annulus A_ρ has an open circuit or is crossed radially, then the radial exploration process gives an open path Γ that has macroscopically wide unexplored space on one side, collecting large enough volume connected to Γ with high probability.

Remark 1.4), we have

$$\mathbf{E}_\eta^\lambda[X_i(\delta)] \asymp (\delta\rho/\eta)^2 \alpha_1(\eta, \delta\rho) = (\delta\rho/\eta)^{91/48+o(1)}.$$

It is a standard argument using quasi-multiplicativity and a summation over dyadic scales that the second moment of X_i is comparable to the square of the first moment (see, e.g., [26, Lemma 3.1] for the second moment of the number of pivotals). Thus, by the Paley-Zygmund second moment inequality (a simple consequence of Cauchy-Schwarz; see, e.g., [42, Section 5.5]), there exists a uniform constant $c = c_\lambda > 0$ such that

$$\mathbf{P}_\eta^\lambda[X_i(\delta) > c \mathbf{E}_\eta^\lambda X_i(\delta)] > c.$$

Using the independence of the variables X_i (conditionally on Γ) that follows from the disjointness of the neighborhoods \tilde{B}_i , we get that

$$\begin{aligned} & \mathbf{P}_\eta^\lambda \left[\text{cluster of } \Gamma \text{ has volume} \leq c \mathbf{E}_\eta^\lambda[X_i(\delta)] \mid \Gamma \right] \\ & \leq \mathbf{P}_\eta^\lambda \left[X_i < c \mathbf{E}_\eta^\lambda[X_i(\delta)] \text{ for all } i = 1, \dots, K(\delta) \mid \Gamma \right] \quad (3.1) \\ & \leq (1 - c)^{K(\delta)} = \exp(-c'/\delta). \end{aligned}$$

Now we want to choose δ so that the bound $c \mathbf{E}_\eta^\lambda[X_i(\delta)] = (\delta\rho/\eta)^{91/48+o(1)}$ on the cluster size becomes at least $(\rho/\eta)^{91/48-\zeta}$. This means that δ should be at least $(\rho/\eta)^{-48\zeta/91+o(1)}$, but this choice is allowed only if this value is less than δ_Γ . As mentioned above, this fails with probability $(\rho/\eta)^{-48\zeta/91+o(1)}$,

which, for η small enough, is much smaller than $(\rho/M)^2$. Therefore, with probability tending to 1 as $\eta \rightarrow 0$, in all the at most $O((M/\rho)^2)$ annuli where case (a) occurs, δ_Γ is large enough and the event of (3.1) fails to hold, hence the cluster of Γ has volume at least $(\rho/\eta)^{91/48-\zeta}$.

(b) Condition on the second case, and let Γ be the clockwisemost radial open crossing that the exploration process has found. We claim that, similarly to case (a), there is a random variable δ_Γ , uniformly positive in η , such that no hexagons have been explored in the clockwise $\delta_\Gamma\rho$ -neighborhood of Γ . Indeed, this was already used in [28, Lemma 2.9] in the proof of the quad-measurability of the 1-arm event, and the reason is simply that this maximal distance δ_Γ can be less than some $\delta > 0$ only if the radial exploration path comes to distance $\delta\rho$ to itself without touching, which would imply a full plane 6-arm event from distance $\delta\rho$ to distance of order ρ (or a half-plane 3-arm event, if it happens close to one of the boundary components of A_ρ). See the second and third pictures in Figure 3.2. Now, we can repeat the rest of the proof of case (a) within this unexplored space of width $\delta_\Gamma\rho$, and we are almost done: we have just proved that, with very high probability as $\eta \rightarrow 0$, the cluster found by the radial exploration process started at some arbitrary (say, uniform random) point at radius $\rho/3$ has large volume. However, we want this for *all* clusters that cross A_ρ , while the above procedure finds larger clusters with larger probability.




Figure 3.3: Consecutive radial exploration processes.

To this end, once we have found one crossing cluster, we start a new radial exploration from radius $\rho/3$, at the first point on $\partial B_{\rho/3}$ to the right

of the last boundary touching point of the first exploration path that has an open site on the right and a closed site on the left side. We stop the process either when it reaches an open site explored by the previous exploration path and hence turns inside, towards $\partial B_{\rho/3}$, or when it reaches ∂B_ρ (which we may call a “success”). Then we take the next point on $\partial B_{\rho/3}$ that has an open site on the right and a closed site on the left side, and so on, until the entire boundary $\partial B_{\rho/3}$ has been explored and hence all radially crossing clusters have been found. Now, before each success, the right boundary of what has been built by the sequence of unsuccessful explorations is an open arm from $\partial B_{\rho/3}$ to ∂B_ρ , and from each point of this open arm, there is also a closed arm to $\partial B_{\rho/3}$. Therefore, if the next successful exploration path comes $\delta\rho$ -close to this right boundary, then it creates a full plane 6-arm or a half-plane 3-arm event (the third picture of Figure 3.2 applies locally), which do not happen anywhere in \mathcal{A}_ρ if δ is small enough. Therefore, all these right boundaries have the open unexplored space to their right that is required for our argument to work. Since each radially crossing cluster has, as a subset, such a right boundary (not necessarily the right boundary of the entire cluster), the proof of Proposition 3.6 is complete. \square

We can now prove that the (r, ρ) -mesoscopic sub-routers of Definition 3.5 exist:

Lemma 3.7. *With probability tending to 1 as $\eta \rightarrow 0$ and then $r \rightarrow 0$, for any $0 < \rho < r^*(X)$, for all $Y \subseteq X$ with $|Y| \geq 2$ whose points are connected together by a single open cluster of $\omega_\eta^\lambda \setminus X$ (more precisely, there is a cluster of $\omega_\eta^\lambda \setminus X$ that neighbours each hexagon in Y), the set of (r, ρ) -mesoscopic sub-routers for Y is non-empty.*

Proof. Assume that in a configuration ω_η^λ , some $Y \subseteq X$ satisfies the above conditions. Let ρ be less than $r^*(X)$, take $r \ll \rho$, and consider any r -square B that intersects the cluster and whose distance from $B_r(X)$ is at least ρ . By the definition of $r^*(X)$ and by $|Y| \geq 2$, such a B certainly exists. We are going to examine when such a B could be a mesoscopic sub-router. Simply aiming at $\widehat{B} = B$, the required quad connecting B with an $x_i \in Y$ can fail to exist only if all the connections from B to $B_{x_i}^r$ are r -close to some $x_k \in X \setminus \{x_i\}$; however, this would imply a 6-arm event from radius r to $\rho \wedge \epsilon$ (see Figure 3.4), which does not occur anywhere in \mathbb{T}_M^2 if r is small enough.

We still need to show that, among the r -squares B as above, there is at least one that also has the open circuit in the (r, ρ) -annulus around it.

If $\rho < r^*(X)$, then any cluster \mathcal{C} connecting the points of Y has a connected subset \mathcal{C}' of diameter at least ρ that has a distance at least ρ from




Figure 3.4: Connections from possible sub-routers B can avoid other r -squares $B_{x_k}^r$ unless a 6-arm event happens.

all points of X . (We used here the definition of $r^*(X)$ and that $|Y| \geq 2$.) For the maximal such \mathcal{C}' , the proof of Proposition 3.6 clearly applies, and for $r \ll \rho$, the number of r -squares in $B^r(\mathbb{T}_M^2)$ intersected by \mathcal{C}' is at least $(\rho/r)^{91/48-\zeta}$ with probability tending to 1 as $r \rightarrow 0$. On the other hand, any of these r -squares fails to be an (r, ρ) -mesoscopic sub-router only if there is no open circuit in the (r, ρ) -annulus around B . In such a case, we have both a primal and a dual arm in the (r, ρ) -annulus, which event has probability $(r/\rho)^{1/4+o(1)}$, uniformly in $\eta > 0$, by the 2-arm exponent [52]. Thus the number of such r -squares is $(\rho/r)^{7/4+o(1)}$ in expectation, and by Markov's inequality, it is unlikely to be much larger, for any of the possible subsets $Y \subseteq X$ (whose number is independent of r). Since $(\rho/r)^{7/4+o(1)}$ is negligible compared to the r -volume $(\rho/r)^{91/48-\zeta}$ if $\zeta > 0$ is small enough, with probability going to 1 as $r \rightarrow 0$, we do have (r, ρ) sub-routers in every cluster spanned by some $Y \subseteq X$. \square

If B_1, B_2 are (r, ρ) sub-routers for $Y_1, Y_2 \subseteq X$, respectively, we will call them connected if there exists a quad R with $\partial_1 R$ contained in \widehat{B}_1 , $\partial_3 R$ contained in \widehat{B}_2 , remaining strictly away from all the squares B_X^r , and for which $\omega^\lambda \in \Xi_R$. As before, in the coupling of Proposition 3.4 (ii), for $\rho < r_M$, the relation of being connected converges in probability as $\omega_\eta \rightarrow \omega_\infty$, which also implies that it is an equivalence relation even in ω_∞ . If B_i is an (r, ρ) sub-router for $Y_i \subseteq X$, $i = 1, 2$, and B_1 and B_2 are connected, then both B_i 's are (r, ρ) sub-routers for $Y_1 \cup Y_2$, since we can glue the path between B_1 and B_2 , the circuit around B_2 , and the path from B_2 to any of the r -squares $B \in B^r(Y_2)$ to get a path from B_1 to B . Therefore, for each equivalence class of (r, ρ) sub-routers there exists a maximal subset $Y \subseteq X$ for which all elements of the equivalence class are sub-routers. Such a maximal subset Y will sometimes be called a **cluster of pivots**, and a corresponding

equivalence class is said to be **spanned by** Y . For instance, in Figure 3.1, the left configuration has two clusters, spanned by the same three pivots, while the right configuration has three clusters, each with a maximal Y of two elements.

In each equivalence class of sub-routers, we want to single out one of them. In order to do this in a way that is typically continuous w.r.t. ω^λ , we need to restrict ourselves to the case when every open cluster of ω^λ in \mathbb{T}_M^2 has diameter less than $M/3$; this will turn out to be typically the case when λ is very negative. (For continuum percolation configurations, the diameter is the $\limsup_{r \rightarrow 0}$ of distances between r -boxes that are connected in the usual sense that there is a crossed quad with its opposite sides contained in the r -boxes. It is clear from Lemma 2.3 that, in any coupling with $\omega_\eta^\lambda \xrightarrow{a.s.} \omega_\infty^\lambda$, the event that this diameter is at most $M/3$ converges almost surely.) Then, for $r < M/10$, the set of (r, ρ) sub-routers for any Y has an isometric embedding into $[-M/2, M/2]^2 \subset \mathbb{C}$. The leftmost sub-router of the lowermost ones in such an embedding will be the same in any of these embeddings; moreover, its location in \mathbb{T}_M^2 can change only a little if we move each sub-router a little, to at most distance $M/10$. The set of these “leftmost of lowermost” sub-routers will be the (r, ρ) -**mesoscopic routers of X** , or, after fixing $\rho = r_M/2$ (from Proposition 3.4), the set of r -**mesoscopic routers**. Note that by restricting ourselves to subsets $|Y| \geq 2$, clusters containing only one pivotal from X will not have routers.

Although we will not really need them, for the sake of symmetry in our presentation, analogously to the above routers that used primal (open) connections, we also define **dual clusters of pivots** and **dual r -mesoscopic routers**.

We can now define the enhanced networks we promised.

Definition 3.8. *Assuming that the diameter of every cluster in ω^λ is at most $M/3$, the r -mesoscopic enhanced ϵ -network $\text{EN}_M^{r\text{-meso}}(\omega^\lambda, \text{PPP}_\lambda^\epsilon)$ is the following vertex- and edge-labeled bipartite graph. One part of the vertex set is the set X of the pivots of $\text{PPP}_\lambda^\epsilon$, the other part is the r -mesoscopic routers of X (both the primal and dual ones). The vertices in $\text{PPP}_\lambda^\epsilon$ are colored open or closed, according to the definitions before Proposition 2.6; the routers are colored in the obvious way. The edge set consists of the connections between the routers and the elements of their maximal $Y \subset X$, labelled primal or dual according to the color of the router. The edges are drawn on the torus so that they are homotopic (with fixed endpoints) to the connections they represent; clearly, one can also achieve that they do not intersect each other. See Figure 3.5.*

If the assumption on the diameters is not satisfied, let $\text{EN}_M^{r\text{-meso}}(\omega^\lambda, \text{PPP}_\lambda^\epsilon)$

be empty.




Figure 3.5: A schematic picture of a percolation configuration ω_η^λ with the pivots of $\text{PPP}_\lambda^\epsilon$ on a torus, and the corresponding enhanced network. Pivots are represented by circles, routers are represented by squares. Primal connections are shown using red solid lines, dual connections are shown using blue dashed lines.

Note that the networks of Definition 3.3 are measurable functions of these enhanced networks in a very simple way: there exists a primal (or dual) router with edges to $x_i, x_j \in X$ in $\text{EN}_M^{r\text{-meso}}(\omega^\lambda, \text{PPP}_\lambda^\epsilon)$ if and only if there is a primal (dual, resp.) edge between x_i and x_j in $\text{N}_M^{r\text{-meso}}(\omega^\lambda, \text{PPP}_\lambda^\epsilon)$. Moreover, the same proof as for Proposition 3.4, together with Theorem 2.7, implies the following:

Proposition 3.9 (r -stabilization and η -convergence of enhanced networks).

- (i) *There is a measurable scale $\tilde{r}_M = \tilde{r}_M(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon) \in (0, r^*)$ such that for all $r \in (0, \tilde{r}_M)$ we get the same r -mesoscopic enhanced ϵ -network $\text{EN}_M^{r\text{-meso}}(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$ in the sense that the networks are the same, plus the colors in X and the collections of primal and dual clusters of pivots are also the same. (The corresponding routers do not exactly stabilize, since for a smaller r new $(r, r_M/2)$ sub-routers can appear; but they cannot disappear, and hence each router does converge to a point in \mathbb{T}_M^2 as $r \rightarrow 0$.) This stabilized network will be called the **enhanced ϵ -network** $\text{EN}_\infty^{\tilde{\lambda}, \epsilon} = \text{EN}_M(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$. For discrete percolation configurations, the definition of $\text{EN}_\eta^{\tilde{\lambda}, \epsilon} = \text{EN}_M(\omega_\eta^\lambda, \text{PPP}_\lambda^\epsilon)$ is the obvious one.*

- (ii) In any coupling with $\omega_\eta^\lambda \xrightarrow{a.s.} \omega_\infty^\lambda$ in \mathbb{T}_M^2 , there is a coupling of $\text{PPP}_\lambda^\epsilon(\omega_\eta^\lambda)$ and $\text{PPP}_\lambda^\epsilon(\omega_\infty^\lambda)$ such that with probability tending to 1 as $\eta \rightarrow 0$, $\text{EN}_\eta^{\bar{\lambda}, \epsilon} = \text{EN}_M(\omega_\eta^\lambda, \text{PPP}_\lambda^\epsilon)$ is the same as $\text{EN}_\infty^{\bar{\lambda}, \epsilon} = \text{EN}_M(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$ in the sense that the vertex sets for η and ∞ (consisting of the pivotals in $\text{PPP}_\lambda^\epsilon$ and the routers) are arbitrarily close to each other, and the labelled graph structures coincide.

Remark 3.10. These enhanced networks are very useful planar (more precisely, toroidal) representations of the discrete and continuous percolation configurations, which was not a priori obvious how to achieve, since the quad-crossing space allows for non-planar configurations and hence is not ideal to express planarity.

Using the enhanced networks, we are now going to define a spanning forest $\text{MSF}^{\bar{\lambda}, \epsilon}$ with vertices being the primal routers in $\text{EN}^{\bar{\lambda}, \epsilon}$. We will show in Section 4 that, for $\lambda < 0$ very negative, $\lambda' > 0$ very large, and $\epsilon > 0$ small, this forest has a unique giant tree component, which will be the cut-off tree $\text{MST}^{\bar{\lambda}, \epsilon}$ that approximates well the macroscopic structure of MST in \mathbb{T}_M^2 . This forest $\text{MSF}^{\bar{\lambda}, \epsilon}$ may have edges that intersect each other besides their endpoints; nevertheless, it consists of trees immersed into the torus \mathbb{T}_M^2 in the sense of Subsection 2.1, and the fact it turns out to approximate MST_η implies in particular that these possible intersections vanish in the limit.




Figure 3.6: Building the cut-off forest $\text{MSF}^{\bar{\lambda}, \epsilon}$ from the enhanced network of Figure 3.5. On level λ there was a cycle that had to be broken. The numbers from -6 to 6 on the closed pivotals of ω_η^λ represent their levels $t_i \in (\lambda, \lambda')$ in $\text{PPP}_\lambda^\epsilon$ at which they become open. The resulting spanning forest has two components.

Definition 3.11 (Constructing the cut-off spanning forest $\text{MSF}^{\bar{\lambda},\epsilon}$ on \mathbb{T}_M^2).

1. The vertices are the primal routers in $\text{EN}^{\bar{\lambda},\epsilon}$. Connect two routers by an edge if they are both connected in $\text{EN}^{\bar{\lambda},\epsilon}$ to the same open pivotal of ω^λ . The resulting graph usually has several components (e.g., six of them on the left-hand picture of Figure 3.6), which more-or-less represent the λ -clusters in ω^λ (this will be made more precise in the next section).
2. In each component of this graph, choose a spanning tree in an arbitrary deterministic way, and label each edge of this tree by λ .
3. For each pivotal x_i of $\text{PPP}_{\bar{\lambda}}^\epsilon$ that is closed in ω^λ , add an edge between the corresponding routers, and label it by its $t_i \in (\lambda, \lambda')$ value. Note that these edges may be loops, as the one labelled by -5 on the left-hand picture of Figure 3.6, for instance.
4. As in the so-called reversed Kruskal algorithm, from each cycle delete the edge with the largest label, and get a minimal spanning tree in each component of the above graph.
5. Draw all the edges of the thus constructed forest as straight line segments, respecting the torus topology (i.e., choosing the line segment on the torus that is homotopic to the concatenation of the embedded edges of $\text{EN}_M^{\text{r-meso}}(\omega^\lambda, \text{PPP}_{\bar{\lambda}}^\epsilon)$ that gave rise to this edge of the forest). See the right-hand picture of Figure 3.6. Note that the edges may intersect each other besides their endpoints (even if this does not happen on this particular picture).

4 Approximation of MST_η by the cut-off trees $\text{MST}_\eta^{\bar{\lambda},\epsilon}$

4.1 Preparatory lemmas and the definition of $\text{MST}_\eta^{\bar{\lambda},\epsilon}$

Our first lemma is a RSW-type result that is interesting even in the critical case. Nevertheless, the simplest proof we have found uses our dynamical and near-critical stability results from [29, Section 8].

Lemma 4.1 (Local Ring Lemma). *There exists $\delta > 0$ such that for any $\lambda \leq 0$ and any radius $R < |\lambda|^{-4/3}$, for all small enough mesh $0 < \eta < \eta_0(\lambda, R)$, one has*

$$\mathbf{P}_\eta[\mathcal{A}_{R,\lambda,\delta}] > 1 - \frac{1}{100},$$

where $\mathcal{A}_{R,\lambda,\delta}$ stands for the event that there exist λ -clusters for the restriction of ω_η^λ to the annulus $\mathcal{A}_{R,2R}$: $\mathcal{C}_1, \dots, \mathcal{C}_N, \mathcal{C}_{N+1} = \mathcal{C}_1$ which satisfy the following conditions:

1. for each $i \in [1, N]$, $\text{diam}(\mathcal{C}_i) \geq \delta R$; in particular, the clusters of the percolation configuration non-restricted to $\mathcal{A}_{R,2R}$ that contain these \mathcal{C}_i 's also have diameter $\geq \delta R$;
2. for each $i \in [1, N]$, there exists at least one closed site y_i neighboring both \mathcal{C}_i and \mathcal{C}_{i+1} ;
3. the circuit $\{\mathcal{C}_1, \dots, \mathcal{C}_N\}$ disconnects the annulus $\mathcal{A}_{R,2R}$ in the sense that the two boundaries of the annulus are not connected in the graph $\mathcal{A}_{R,2R} \setminus \bigcup_{i=1}^N (\mathcal{C}_i \cup \{y_i\})$.

Moreover, we can choose the clusters \mathcal{C}_i and the points y_i such that all the y_i 's are elements of the Poisson point set $\text{PPP}_\lambda^\epsilon$, with $\epsilon = \delta R$ and λ' large enough (depending on R).

Proof. Consider the near-critical coupling $(\omega_\eta^t)_{t \in \mathbb{R}}$. For $\lambda' \geq \lambda$ large enough (on the order of $R^{-3/4}$), there is a probability at least $995/1000$ that $\omega_\eta^{\lambda'}$ has an open circuit even in the smaller annulus $\mathcal{A}_{5R/4, 7R/4}$; this follows from known results on the correlation length, e.g., [29, Theorem 10.7]. Now sample ω_η^λ , consider some small $\epsilon > 0$ to be fixed in a second, and let $\tilde{\omega}_\eta^{\lambda'}$ be the configuration where we open only those vertices in the coupling while getting from λ to λ' that are given in $\text{PPP}_\lambda^\epsilon$. The assumption $R < |\lambda|^{-4/3}$, below the correlation length given by [29, Theorem 10.7], implies that we can choose $\eta > 0$ small enough compared to R so that ω_η^λ has 4-arm probabilities inside the domain $\mathcal{A}_{R,2R}$ that are comparable to the critical ones. Therefore, the critical case computations of [29, Section 8] apply uniformly in $\lambda \leq 0$ and $R < |\lambda|^{-4/3}$, and by a straightforward modification of [29, Proposition 8.6] from quad-crossings to annulus circuits, for $\epsilon = \delta R > 0$ with $\delta > 0$ small enough (uniformly in λ and R), the probability that $\omega_\eta^{\lambda'}$ has an open circuit in $\mathcal{A}_{5R/4, 7R/4}$ but $\tilde{\omega}_\eta^{\lambda'}$ does not have one in $\mathcal{A}_{R,2R}$ is less than $5/1000$. Altogether, the probability that $\tilde{\omega}_\eta^{\lambda'}$ has an open circuit in $\mathcal{A}_{R,2R}$ is at least $99/100$. But such a circuit must be composed of λ -clusters and ϵ -important points that have become open, which implies that all these λ -clusters must have diameter at least ϵ , and the lemma is proved. \square

Lemma 4.2 (Global Ring Lemma). *Fix $\delta > 0$ as in Lemma 4.1. For any $\lambda \leq 0$ and $\alpha > 0$, there is a radius $r = r(\lambda, \alpha) < \frac{\delta}{2} |\lambda|^{-4/3}$ such that, for any small enough η , with probability at least $1 - \alpha$, one can find around all points*

$x \in \mathbb{T}_M^2$ an annulus $A_{R,2R}$ surrounding x with $\bar{r} = r/\delta \leq R \leq |\lambda|^{-4/3}$ that satisfies the event $\mathcal{A}_{R,\lambda,\delta}$. (The choice $\bar{r} = r/\delta$ is made so that the clusters we find are at least of diameter r .)

Proof. Consider the covering of \mathbb{T}_M^2 by the squares given by $\bar{r}\mathbb{Z}^2$, and around each such \bar{r} -square, consider the dyadic annuli up to scale $|\lambda|^{-4/3}$. By Lemma 4.1, the probability that there is an \bar{r} -square for which all the dyadic annuli fail to have the required ring of clusters is at most

$$O(1)(M/\bar{r})^2(1/100)^{\log_2 \frac{|\lambda|^{-4/3}}{\bar{r}}} = O(M^2) |\lambda|^{4/3 \log_2 100 \bar{r}^{-2} + \log_2 100},$$

which can be made arbitrarily small as $\bar{r} \rightarrow 0$. \square

Part (ii) of the next lemma again has a RSW feeling to it, and is again proved using [29, Section 8].

Lemma 4.3 (Subcritical lakes joining the supercritical ocean). *Consider percolation ω_η^λ on \mathbb{T}_M^2 with $\lambda < -1$, and fix an arbitrarily small $\alpha > 0$.*

- (i) *For any $s > 0$, if $\lambda < -1$ is small enough, then for all $\eta > 0$ small enough, with probability at least $1 - \alpha$, all clusters in ω_η^λ have diameter less than s .*
- (ii) *Using part (i), take $\lambda < -1$ small enough so that with probability more than $1 - \alpha/3$, all clusters in ω_η^λ have diameter less than $M/3$. Then, for any $r > 0$ there is a $\lambda'_0(\lambda, r, \alpha) > 0$ and an $\epsilon_0(\lambda, r, \alpha) > 0$ such that for all $\lambda' \geq \lambda'_0$ and $\epsilon \leq \epsilon_0$, for all small enough $\eta > 0$, with probability at least $1 - \alpha$, all the clusters in ω_η^λ of diameter at least r are connected via primal paths in the enhanced network $\text{EN}_\eta^{\lambda, \epsilon} = \text{EN}_M(\omega_\eta^\lambda, \text{PPP}_\lambda^\epsilon)$ with $\bar{\lambda} = (\lambda, \lambda')$, defined in Proposition 3.9, in the sense that each such cluster contains a primal router and these routers are all connected by primal edges (through closed or open pivotals, as in Definition 3.11) in the enhanced network.*

Proof. It is proved in [29, Theorem 10.7] that, for any fixed $s > 0$, as $\lambda' \rightarrow \infty$, the probability of having an open circuit in a given annulus $A_{s/3, s}$ in $\omega_\eta^{\lambda'}$ converges to 1. Consider a tiling of \mathbb{T}_M^2 by $s/3$ -squares, and the annuli of side-length s centered around them. By a union bound, the probability of having open circuits in all of them converges to 1. When all these circuits are present, their union is a single component, and any subset of \mathbb{T}_M^2 of diameter at least s intersects this cluster.

Running the above argument for dual circuits in ω_η^λ with $\lambda \rightarrow -\infty$ gives that, with probability tending to 1, the diameter of the largest open cluster must be less than s , which proves item (i).

For item (ii), we will use the first paragraph with $s = r/10$. Note that any cluster \mathcal{C} of ω_η^λ with diameter at least r will radially cross two such $(s/3, s)$ -annuli at distance at least $r/2$ from each other, A_1 and A_2 . Moreover, assuming that \mathcal{C} has diameter at most $M/3$ (which is satisfied for all clusters \mathcal{C} with probability more than $1 - \alpha/3$), we can choose A_1 and A_2 with the additional property that, for each of them, not all of the eight neighboring inner squares are intersected by \mathcal{C} . On the other hand, the first paragraph says that if $\lambda' \geq \lambda'_0$ is large enough, then with probability more than $1 - \alpha/3$, all the $(s/3, s)$ -annuli will have open circuits in $\omega_\eta^{\lambda'}$. Now we use [29, Proposition 8.1], which implies that in the configuration $\tilde{\omega}_\eta^{\lambda'}$ that we get by starting from the configuration ω_η^λ and opening only the pivotal points of $\text{PPP}_{\bar{\lambda}}^\epsilon$, with $\bar{\lambda}_0 = (\lambda, \lambda')$, all these open circuits in the $(s/3, s)$ -annuli will already be there with probability more than $1 - 2\alpha/3$, provided that $0 < \epsilon \leq \epsilon_0(\lambda', r, \alpha)$ is small enough. This can happen in our two above annuli A_1, A_2 only if \mathcal{C} neighbours points of $\text{PPP}_{\bar{\lambda}}^\epsilon$ in both annuli that are closed in ω_η^λ but open in $\tilde{\omega}_\eta^{\lambda'}$. These ϵ -pivotal points appear in the enhanced network $\text{EN}_{\bar{\lambda}}^{\lambda, \epsilon} = \text{EN}_M(\omega_\eta^\lambda, \text{PPP}_{\bar{\lambda}}^\epsilon)$, and \mathcal{C} has a primal router connecting these pivotals with high probability (as $\eta \rightarrow 0$).

Furthermore, the open circuits of $\tilde{\omega}_\eta^{\lambda'}$ in the $(s/3, s)$ -annuli must be composed of clusters of ω_η^λ with diameter at least $\min\{s/3, \epsilon\}$, joined by ϵ -pivotals. The expected number of disjoint such clusters is bounded from above by some $K = K(M, \lambda, r, \epsilon) < \infty$, uniformly in η (as follows from [3]), hence in all of them we have primal routers with probability tending to 1, as $\eta \rightarrow 0$. That is, altogether, with probability at least $1 - \alpha$, all the λ -clusters of diameter at least r have diameter at most $M/3$, and they are all connected in the enhanced network $\text{EN}_{\bar{\lambda}}^{\lambda, \epsilon}$, for λ' chosen large enough, then $\epsilon(\lambda', r, \alpha)$ and then η chosen small enough. \square

Using the above lemmas, we can now see why there is typically a unique giant component in the cut-off forests $\text{MSF}_\eta^{\bar{\lambda}, \epsilon}$ and $\text{MSF}_\infty^{\bar{\lambda}, \epsilon}$ of Definition 3.11:

Lemma 4.4 (Defining the cut-off trees $\text{MST}_\eta^{\bar{\lambda}, \epsilon}$ and $\text{MST}_\infty^{\bar{\lambda}, \epsilon}$). *For any $M > 0$, any small $s > 0$ and $\alpha > 0$, if $\lambda < -1$ is very negative, $\epsilon > 0$ is small, and $\lambda' > 1$ is large enough, then with probability at least $1 - \alpha$ for any small enough mesh $\eta > 0$, there is a **unique giant component** in the cut-off forest $\text{MSF}_\eta^{\bar{\lambda}, \epsilon}$ (and hence, by Proposition 3.9, in $\text{MSF}_\infty^{\bar{\lambda}, \epsilon}$), with the properties that it comes to distance at most s from any point of \mathbb{T}_M^2 , while all other components of $\text{MSF}_\eta^{\bar{\lambda}, \epsilon}$ have diameter at most s . This giant tree component will be our **approximating cut-off tree**, denoted by $\text{MST}_\eta^{\bar{\lambda}, \epsilon}$ and $\text{MST}_\infty^{\bar{\lambda}, \epsilon}$; whenever the above large probability event fails to occur, we set*

$\text{MST}^{\bar{\lambda},\epsilon}$ to be a single point in \mathbb{T}_M^2 , and call this tree degenerate.

Proof. Take $\lambda < -1$ such that $2|\lambda|^{-4/3} < s$ holds and the diameter bound of Lemma 4.3 (i) applies. By Lemma 4.2, with probability at least $1 - \alpha/2$, every point of \mathbb{T}_M^2 has in its s -neighborhood a ring of λ -clusters of diameter at least $r(\lambda, \alpha)$ each (possibly a single cluster, but still of diameter in \mathbb{T}_M^2 less than s). Now, if we take $\epsilon > 0$ small and $\lambda' > 0$ large, then Lemma 4.3 (ii) says that with probability at least $1 - \alpha/2$ all λ -clusters of diameter at least $r(\lambda, \alpha)$ get connected in the enhanced network $\text{EN}_\eta^{\bar{\lambda},\epsilon}$. Therefore, with probability altogether at least $1 - \alpha$, there is a component of the graph of Definition 3.11 that has distance at most s from any point of \mathbb{T}_M^2 , while all other components have diameter at most $r(\lambda, \alpha) \leq s$. The spanning trees of these components inherit these properties, hence we are done. \square

So, we have finally defined the trees $\text{MST}_\eta^{\bar{\lambda},\epsilon}$ and $\text{MST}_\infty^{\bar{\lambda},\epsilon}$. Naturally, they are also immersed spanning trees in the sense of Definition 2.2. Now, using Proposition 3.9, we can easily show that they are close to each other in the space of essential spanning forests, introduced in that definition:

Corollary 4.5. *For any $M > 0$ and $s, \alpha > 0$, if $\lambda < -1$ is very negative, $\epsilon > 0$ is small, and $\lambda' > 1$ is large enough, then, in the coupling of Proposition 3.9 (ii) between $(\omega_\eta^\lambda, \text{PPP}_\lambda^\epsilon)$ and $(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$, we have*

$$\mathbf{P}[d_{\Omega_M}(\text{MST}_\eta^{\bar{\lambda},\epsilon}, \text{MST}_\infty^{\bar{\lambda},\epsilon}) < s] > 1 - \alpha,$$

for all $\eta > 0$ small enough.

Proof. The parameters $\lambda, \lambda', \epsilon$ can be set so, by Lemma 4.4, that both $\text{MST}_\eta^{\bar{\lambda},\epsilon}$ and $\text{MST}_\infty^{\bar{\lambda},\epsilon}$ are non-degenerate with probability at least $1 - \alpha/2$, for any $\eta > 0$ small enough. Now, by Proposition 3.9 (ii), we can take $\eta > 0$ so small that, with probability at least $1 - \alpha/2$, the enhanced networks $\text{EN}_\eta^{\bar{\lambda},\epsilon}$ and $\text{EN}_\infty^{\bar{\lambda},\epsilon}$ agree as graphs and the Hausdorff distance between their vertex sets is less than s . On the event that both trees exist, the networks agree, and the vertex sets are closer than s to each other, which occurs with probability at least $1 - \alpha$, the uniform distance between the corresponding ℓ -trees is always less than s , and hence the sum with the weights $2^{-\ell}$ is also less than s , and we are done. \square

4.2 Approximation as $\epsilon \rightarrow 0$ and $(\lambda, \lambda') \rightarrow (-\infty, \infty)$.

After these preparations, we can turn to approximating MST_η on the lattice $\eta\mathbb{T} \cap \mathbb{T}_M^2$ by the cut-off trees $\text{MST}_\eta^{\bar{\lambda},\epsilon}$. First of all, note that if two vertices

of $\eta\mathbb{T} \cap \mathbb{T}_M^2$ are in the same λ -cluster, then the path in MST_η that connects them remains in that λ -cluster. This also means that for any two λ -clusters of $\eta\mathbb{T} \cap \mathbb{T}_M^2$, there is a unique path in MST_η that connects them.

The next lemma is a key step in approximating MST_η by cut-off trees:

Lemma 4.6 (Paths through macroscopic clusters). *Fix any $\alpha > 0$. If $\lambda < -1$ is small enough, $\rho > 0$ is arbitrary, $\epsilon > 0$ is small and $\lambda' > 0$ is large enough, then, for any $\eta > 0$ small enough, with probability at least $1 - \alpha$, for any two clusters of diameter at least ρ in ω_η^λ , there is a unique path in $\text{MSF}_\eta^{\bar{\lambda}, \epsilon}$ that connects the two clusters, and the unique path in MST_η doing the same goes through the same closed pivotals of $\text{PPP}_\lambda^\epsilon$, and hence the distance of these two paths in the uniform metric is at most the maximal diameter of all λ -clusters.*

Proof. Choose $\lambda < -1$ then $\epsilon_1 > 0$ and $\lambda'_1 > 1$ so that the event of Lemma 4.3 (ii), with $r := \rho$ there, occurs with probability at least $1 - \alpha/3$. Condition on this event, set $\bar{\lambda}_1 = (\lambda, \lambda'_1)$, and consider the path in $\text{MSF}_\eta^{\bar{\lambda}_1, \epsilon_1}$ that connects two of the clusters. There is a corresponding path in $\eta\mathbb{T}$, going through the same finitely many ϵ_1 -important points of ω_η^λ and some λ -clusters, using labels at most λ'_1 . Therefore, the true path in MST_η also uses labels at most λ'_1 . Assume now that this path goes through some λ -cluster \mathcal{C} of diameter at most $r \ll \rho$. This path must go through a vertex x of $\eta\mathbb{T}$, neighboring \mathcal{C} , with the following properties (see Figure 4.1):

- it is closed in ω_η^λ but open in $\omega_\eta^{\lambda'_1}$;
- it has two λ -closed arms emanating from it, which together separate the two clusters of diameter at least ρ that we started with;
- on the side of these two closed arms that contains \mathcal{C} , there is a λ -open arm from x only to distance at most r . (Note here that x might be neighbouring λ -clusters of diameter more than r , but we can choose the closed arms so that they separate those clusters from \mathcal{C} .)

If x had the alternating 4-arm event to a distance more than r in ω_η^λ , that could happen only if the two open arms out of these four were on the side of the two long closed arms that does not contain \mathcal{C} . This would altogether yield a 5-arm event to distance r , of type closed-open-closed-open-closed, where the extra closed arm in the middle is needed to get an *alternating* 4-arm event. Moreover, since the labels along the path in MST_η are all at most λ'_1 , we would get a (λ, λ'_1) -near-critical six-arm event from x to distance r , as defined in Proposition 1.3 (the sixth arm is of length at least




Figure 4.1: The path in MST_η connecting two large λ -clusters does not go through very small λ -clusters, basically because of the near-critical stability of 4-arm probabilities.

$\rho/2$, because of being λ'_1 -connected to the λ -cluster of diameter at least ρ). By that proposition and by the 6-arm exponent being larger than two (see [50, Corollary A.8]) this happens with very small probability if η is small enough. So, we can basically assume that x is not r -pivotal in ω_η^λ . On the other hand, if we now change all the labels above λ along the path in MST_η to λ , then, in the new configuration, x will have the alternating 4-arm event to distance at least ϵ_1 . Since the labels we have changed are all in $[\lambda, \lambda'_1]$, we can apply another version of near-critical stability, Lemma 8.5 of [29], saying that the probability that the importance of a given vertex x can be changed from r to ϵ_1 by such label changes is at most

$$O_{\lambda, \lambda'_1}(1) r^2 \alpha_4(\eta, r) \alpha_4(\epsilon_1, 1)^{-1}.$$

Additionally, the probability that the status of this vertex x is different in $\omega_\eta^{\lambda'_1}$ than in ω_η^λ is $O_{\lambda, \lambda'_1}(1) \eta^2 / \alpha_4(\eta, 1)$, independently of everything. The number of possible vertices $x \in \mathbb{T}_M^2$ is $O_M(\eta^{-2})$. Altogether, the expected number of vertices x satisfying this scenario is

$$O_{M, \lambda, \lambda'_1}(1) r^2 \alpha_4(r, 1)^{-1} \alpha_4(\epsilon_1, 1)^{-1}.$$

This is arbitrarily small if r is small, hence the probability that this scenario happens is also small. Summarizing, there exists $r > 0$ depending on M, α, λ and ρ , such that for all small enough $\eta > 0$, with probability at least $1 - \alpha/3$,

the path in MST_η connecting any two λ -clusters of diameter at least ρ does not go through λ -clusters smaller than r .

Now choosing $\epsilon > 0$ small and $\lambda' > 1$ large, again by Lemma 4.3 (ii), the enhanced network $\text{EN}_\eta^{\bar{\lambda}, \epsilon}$ will connect all the λ -clusters of diameter at least r with probability at least $1 - \alpha/3$. Altogether, with probability at least $1 - \alpha$, for any two λ -clusters of diameter at least ρ , the unique paths in MST_η and $\text{MSF}_\eta^{\bar{\lambda}, \epsilon}$ both go through the same λ -clusters, connected by λ -closed pivotals of importance at least r . The last half sentence of the lemma follows immediately from the way Definition 3.11 is done. \square

We can now easily prove the main result of this section:

Proposition 4.7. *For any $M > 0$ and $s, \alpha > 0$, if $\lambda < -1$ is very negative, $\epsilon > 0$ is small, and $\lambda' > 1$ is large enough, then we have*

$$\mathbf{P}[d_{\Omega_M}(\text{MST}_\eta, \text{MST}_\eta^{\bar{\lambda}, \epsilon}) < s] > 1 - \alpha,$$

for all $\eta > 0$ small enough.

Proof. As in the proof of Lemma 4.4, take $\lambda < -1$ such that with probability at least $1 - \alpha/2$, all λ -clusters in \mathbb{T}_M^2 have diameter less than s , and every point of \mathbb{T}_M^2 has in its $s/2$ -neighborhood a ring of λ -clusters of diameter at least r each, for some $0 < r < s/2$, uniformly in η , as provided by Lemma 4.2. Now, if we take $\epsilon > 0$ small and $\lambda' > 0$ large, then, with probability at least $1 - \alpha/2$, all λ -clusters of diameter at least r are connected in $\text{MST}_\eta^{\bar{\lambda}, \epsilon}$, and, by Lemma 4.6, the paths connecting them are at a uniform distance at most s from the corresponding paths of MST_η . We will assume that both events of probability at least $1 - \alpha/2$ hold.

Consider any path γ of MST_η connecting some $x, y \in \mathbb{T}_M^2$. Both x and y have the above-mentioned ring of macroscopic λ -clusters around them, and γ must intersect at least one member of each ring. See Figure 4.2. But then, the part of γ connecting the intersected members closest to x and y , denoted by \mathcal{C}_x and \mathcal{C}_y , respectively, by the previous paragraph, is uniformly s -close to a path in $\text{MST}_\eta^{\bar{\lambda}, \epsilon}$, denoted by $\gamma^{\bar{\lambda}, \epsilon}$. And this $\gamma^{\bar{\lambda}, \epsilon}$ is of course s -close to the entire γ , since the parts of γ going from x to \mathcal{C}_x and from y to \mathcal{C}_y are contained in the s -neighborhoods of \mathcal{C}_x and \mathcal{C}_y .

In the other direction, consider any path $\gamma^{\bar{\lambda}, \epsilon}$ in $\text{MST}_\eta^{\bar{\lambda}, \epsilon}$, connecting two routers. The clusters of pivotals corresponding to these routers have diameter at most s , but could be rather small. Nevertheless, fixing one point in each cluster, there is a ring of macroscopic λ -clusters around each, which certainly contains a cluster of pivotals that $\gamma^{\bar{\lambda}, \epsilon}$ goes through. The rest of the proof is just as above.




Figure 4.2: Paths in MST_η can be approximated by paths through macroscopic λ -clusters.

Now that we have good approximations for paths in the two trees connecting any two vertices, the extension to trees with $\ell > 2$ leaves is straightforward. \square

5 Proof of the main result

5.1 Putting the pieces together for MST on tori \mathbb{T}_M^2

In this subsection, we prove convergence in any fixed torus \mathbb{T}_M^2 .

Theorem 5.1 (Limit of MST_η and $\text{MST}_\infty^{\bar{\lambda}, \epsilon}$ in \mathbb{T}_M^2). *In the metric space Ω_M of spanning trees in the torus \mathbb{T}_M^2 , as defined in Definition 2.2, the spanning tree MST_η on the lattice $\eta\mathbb{T} \cap \mathbb{T}_M^2$ converges in law to a translation invariant MST_∞ , which is also the distributional limit of the cut-off trees $\text{MST}_\infty^{\bar{\lambda}, \epsilon}$, as $\bar{\lambda} \rightarrow (-\infty, \infty)$ and $\epsilon \rightarrow 0$.*

Proof. Using the results of the previous section, the proof is classical; e.g., the exact same strategy was used in [29, Section 9]. By Proposition 4.7, for any $k \in \mathbb{N}$ there exists $\bar{\lambda}_k = (\lambda_k, \lambda'_k)$ and $\epsilon_k > 0$, such that, for all $0 < \eta < \eta_k$ sufficiently small,

$$\mathbf{P} [d_{\Omega_M}(\text{MST}_\eta, \text{MST}_\infty^{\bar{\lambda}_k, \epsilon_k}) < 2^{-k}] > 1 - 2^{-k}. \quad (5.1)$$

Now, by Corollary 4.5, there is a coupling between $(\omega_\eta^\lambda, \text{PPP}_\lambda^\epsilon)$ and $(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$, and by the same token, between $\omega_\eta^{[\lambda, \lambda']}$ and $(\omega_\infty^\lambda, \text{PPP}_\lambda^\epsilon)$, such that, for all $0 < \eta < \eta'_k$ sufficiently small,

$$\mathbf{P}[d_{\Omega_M}(\text{MST}_\eta^{\bar{\lambda}_k, \epsilon_k}, \text{MST}_\infty^{\bar{\lambda}_k, \epsilon_k}) < 2^{-k}] > 1 - 2^{-k}. \quad (5.2)$$

Combining (5.1) and (5.2) using the triangle inequality, in the same coupling,

$$\mathbf{P}[d_{\Omega_M}(\text{MST}_\eta, \text{MST}_\infty^{\bar{\lambda}_k, \epsilon_k}) < 2^{-k+1}] > 1 - 2^{-k+1}.$$

We can now couple all the trees $\text{MST}_\infty^{\bar{\lambda}_k, \epsilon_k}$ to MST_η one-by-one, and given MST_η , conditionally independently to each other, such that, again using the triangle inequality,

$$\mathbf{P}[d_{\Omega_M}(\text{MST}_\infty^{\bar{\lambda}_k, \epsilon_k}, \text{MST}_\infty^{\bar{\lambda}_\ell, \epsilon_\ell}) < 2^{-k+2} \text{ for all } \ell \geq k] > 1 - 2^{-k+2}. \quad (5.3)$$

Using the Borel-Cantelli lemma for the events appearing on the left hand side of (5.3) shows that the sequence $\text{MST}_\infty^{\bar{\lambda}_k, \epsilon_k}$ is almost surely a Cauchy sequence in Ω_M . The space is complete, hence there is an almost sure limit MST_∞ . Of course, this limit may a priori depend on the sequences $\{\bar{\lambda}_k\}$, $\{\epsilon_k\}$ and on the coupling. However, using the triangle inequality again, going through $\text{MST}_\infty^{\bar{\lambda}_k, \epsilon_k}$, we have that for any $\delta > 0$, if $\eta > 0$ small enough, then

$$\mathbf{P}[d_{\Omega_M}(\text{MST}_\eta, \text{MST}_\infty) < \delta] > 1 - \delta.$$

Therefore, in this coupling, MST_η converges in probability, and hence in law, to MST_∞ , in the metric space Ω_M . Since MST_η and the metric d_{Ω_M} are translation invariant, the limit MST_∞ is also invariant.

To prove the convergence of $\text{MST}_\infty^{\bar{\lambda}_k, \epsilon_k}$, note that the bounds (5.1) and (5.2) hold not just for $\bar{\lambda}_k$ and ϵ_k , but for all $\epsilon < \epsilon_k$ and $\lambda < \lambda_k$ and $\lambda' > \lambda'_k$, thus we have that $\text{MST}_\infty^{\bar{\lambda}_k, \epsilon}$ is close in distribution in the d_{Ω_M} -metric to MST_η and hence to MST_∞ . \square

5.2 Extension to the full plane; invariance under translations, scalings and rotations

We are now ready to prove the main result of this paper.

Proof of Theorem 1.1. We will use the notation MST_η^M and MST_∞^M for MST_η and its scaling limit on the torus \mathbb{T}_M^2 . We will also use the approximations $\text{MST}^{\bar{\lambda}, \epsilon, M}$.

It was proved in [4, equation (8.1)] that MST_η is **uniformly quasi-local** in the sense that for any $\delta > 0$ and compact $\Lambda \subset \mathbb{C}$ there exists a $\bar{\Lambda}_\delta \subset \mathbb{C}$ such that for any small enough $\eta > 0$, with probability at least $1 - \delta$, all trees with leaves in Λ are contained in $\bar{\Lambda}_\delta$. Since this event is measurable w.r.t. the percolation ensemble inside $\bar{\Lambda}_\delta$, by taking $\delta > 0$ small and $M > 0$ so large that $\bar{\Lambda}_\delta \subset [-M, M]^2$, we get that the law of MST_η restricted to Λ is close in total variation distance to the law of MST_η^M restricted to Λ . By the $(\bar{\lambda}, \epsilon)$ -approximation result Proposition 4.7, the same holds for $\text{MST}_\eta^{\bar{\lambda}, \epsilon, M}$, and by the uniformity in $\eta > 0$, also for $\text{MST}_\infty^{\bar{\lambda}, \epsilon, M}$. In the proof of Theorem 5.1, we have constructed MST_∞^M as a limit of $\text{MST}_\infty^{\bar{\lambda}, \epsilon, M}$, thus we also have that the law of MST_∞^M restricted to Λ converges as $M \rightarrow \infty$, in the metric d_{Ω_M} that is based on the flat Euclidean metric on \mathbb{T}_M^2 .

Now we take $\Lambda = [-L, L]^2$, with $L \rightarrow \infty$. As pointed out at the beginning of Subsection 2.1, the metric defined in (2.1) for $\hat{\mathbb{C}}$ is equivalent to the Euclidean metric in bounded domains, while the distance between any two points in $\hat{\mathbb{C}} \setminus [-L, L]^2$ is at most $O(1/L)$. Thus the uniform distance between any two trees embedded in $\hat{\mathbb{C}} \setminus [-L, L]^2$ is at most $O(1/L)$, and if two essential spanning forests are δ -close in the metric d_{Ω_M} restricted to $[-L, L]^2$, then their distance in d_Ω is $O_L(\delta) + O(1/L)$. Therefore, the convergence in d_{Ω_M} for any given $\Lambda \subset \mathbb{C}$, established in the previous paragraph, implies convergence in d_Ω .

Translation invariance of the limit measure MST_∞ follows from a standard trick: for any compact $\Lambda \subset \mathbb{C}$, quasi-locality implies that the limit of $\text{MST}_\infty^{[-M, M]^2}$ restricted to Λ , as $M \rightarrow \infty$, is the same as the limit of $\text{MST}_\infty^{[-M+x, M+x]^2}$ restricted to the same Λ , for any $x \in \mathbb{R}$, and hence MST_∞ restricted to Λ has the same distribution as restricted to $\Lambda - x$.

To prove scale-invariance, consider the scaling $f_\alpha(z) := \alpha z$. The conformal covariance of the pivotal measures, proved in [28, Theorem 6.1], says that

$$(f_\alpha)_*(\mu^\epsilon(\omega_\infty^\lambda)) = \alpha^{-3/4} \mu^{\alpha\epsilon}(f_\alpha(\omega_\infty^\lambda)). \quad (5.4)$$

Also, by the conformal covariance of ω_∞^λ , proved in [29, Theorem 10.3], we have

$$f_\alpha(\omega_\infty^\lambda) \stackrel{d}{=} \omega_\infty^{\alpha^{-3/4}\lambda}. \quad (5.5)$$

Scaling the spatial intensity measure of a Poisson point process by $\alpha^{-3/4}$ as in (5.4) is the same as scaling the time duration by the same factor, in the sense that there is a natural coupling in which the spatial coordinates of the arrivals are the same, and there is a simple scaling between the time

coordinates. Thus, combining (5.4) and (5.5), and denoting the notion of “same” in the previous sentence by \approx , we have

$$f_\alpha(\text{PPP}_{\bar{\lambda}}^\epsilon(\omega_\infty^\lambda)) \stackrel{d}{\approx} \text{PPP}_{\alpha^{-3/4}\bar{\lambda}}^{\alpha\epsilon}(\omega_\infty^{\alpha^{-3/4}\lambda}). \quad (5.6)$$

Our constructions of $\text{MSF}_\infty^{\bar{\lambda},\epsilon}$ and $\text{MST}_\infty^{\bar{\lambda},\epsilon}$ in Definition 3.11 and Lemma 4.4 are equivariant under spatial and time scalings, hence the identities (5.5) and (5.6) imply that

$$f_\alpha(\text{MST}_\infty^{\bar{\lambda},\epsilon,M}) \stackrel{d}{=} \text{MST}_\infty^{\alpha^{-3/4}\bar{\lambda},\alpha\epsilon,\alpha M}.$$

Since we obtained MST_∞ as a limit of $\text{MST}_\infty^{\bar{\lambda},\epsilon,M}$ with $\bar{\lambda} \rightarrow (-\infty, \infty)$, $\epsilon \rightarrow 0$, $M \rightarrow \infty$, the last identity gives that $f_\alpha(\text{MST}_\infty) \stackrel{d}{=} \text{MST}_\infty$.

Next, let $f_\theta : \mathbb{C} \rightarrow \mathbb{C}$ be the rotation by angle θ . Now [28, Theorem 6.1] and [29, Theorem 10.3] give for full plane configurations that

$$f_\theta(\omega_\infty^\lambda, \text{PPP}_{\bar{\lambda}}^\epsilon(\omega_\infty^\lambda)) \stackrel{d}{=} (\omega_\infty^\lambda, \text{PPP}_{\bar{\lambda}}^{\epsilon,\theta}(\omega_\infty^\lambda)), \quad (5.7)$$

where $\text{PPP}_{\bar{\lambda}}^{\epsilon,\theta}$ is constructed using a rotated grid to define ϵ -importance. (As pointed out in [28, Remark 6.3], this rotational equivariance of the ϵ -importance measure and hence the Poisson point process is not a tautology, since the normalization factor in the definition of the measure is not changed with the rotation.) Now, if we want to consider $\text{MST}_\infty^{\bar{\lambda},\epsilon}$ on the torus \mathbb{T}_M^2 , the rotated ϵ - and r -grids cannot be exactly defined; nevertheless, we can consider the squares in the grid fully contained in $f_\theta([-M, M]^2)$, and make some arbitrary definition close to the boundary — due to quasi-locality, this will not matter. Hence, from (5.7) we get that for large $M > 0$, the distribution of $f_\theta(\text{MST}_\infty^{\bar{\lambda},\epsilon,M})$ restricted to some fixed domain Λ , which is close to $f_\theta(\text{MST}_\infty)$ restricted to Λ , is close to $\text{MST}_\infty^{\bar{\lambda},\epsilon,M,\theta}$ restricted to Λ . On the other hand, Corollary 4.5 and Proposition 4.7 work fine with the rotated grids, giving that $\text{MST}_\infty^{\bar{\lambda},\epsilon,M,\theta}$ is close to $\text{MST}_\eta^{\bar{\lambda},\epsilon,M,\theta}$, and the latter is close to MST_η . Finally, since MST_η is close to MST_∞ , after taking all the limits we get that $f_\theta(\text{MST}_\infty)$ agrees with MST_∞ in distribution. \square

6 Geometry of the limit tree MST_∞

6.1 Degree types and pinching

The **degree** of a point $x \in \hat{\mathbb{C}}$ in an immersed tree $f : \tau \rightarrow \hat{\mathbb{C}}$, where we assume that f is locally injective from each edge of τ , is

$$\text{deg}_f(x) := \sum_{f(v)=x} \text{deg}_\tau(v), \quad (6.1)$$

where the sum is over all points v of τ , meaning a vertex in $V(\tau)$ or a point on an edge in $E(\tau)$, and $\deg_\tau(v)$ is the combinatorial degree in the first case, while equals 2 in the second case. (Note that for any immersed tree $f : \tau \rightarrow \hat{\mathbb{C}}$ there exists a minimal $\tau^* \prec \tau$ on which f can be defined, and here f is equivalent to a locally injective immersion $f^* : \tau^* \rightarrow \hat{\mathbb{C}}$. Hence the local injectivity assumption is not a real restriction.) For an essential spanning forest \mathcal{F} ,

$$\deg_{\mathcal{F}}(x) := \sup_{\ell \geq 1} \sup_{f \in \mathcal{F}^{(\ell)}} \deg_f(x). \quad (6.2)$$

The **degree type** of a point x in an immersed tree $f : \tau \rightarrow \hat{\mathbb{C}}$ is the vector of summands in (6.1), ordered in decreasing order, and the degree type in an essential spanning forest \mathcal{F} is the supremum as in (6.2), now w.r.t. a natural partial order on the vectors of degree types: after padding vectors with zeros at the end, use the lexicographic ordering. The supremum in this partial order exists because of condition (iii) of Definition 2.2. (This supremum vector may start with a few ∞ entries if $\deg_{\mathcal{F}}(x) = \infty$, but this is fine.) See Figure 6.1 for a few examples (ignoring at this point the dual trees on the pictures).




Figure 6.1: Degree type (5) and two examples of (2, 1, 1, 1) in a spanning tree of the plane, giving degree types (1, 1, 1, 1, 1), (4, 1) and (3, 2) in a dual spanning tree.

For instance, saying that $x \in \mathbb{C}$ is a **pinching point** for \mathcal{F} if $\mathcal{F}^{(2)}$ includes a path which passes through x twice without terminating there can be expressed as saying that x has degree type at least (2, 2). If one of the two branches terminates at x , the other does not, i.e., degree type at least (2, 1), then we talk about a **figure of 6**, while degree type at least (1, 1) is called a point of non-uniqueness, or a **loop** at x . Points of degree type at least (2) constitute the **trunk** of \mathcal{F} : the union of curves in $\mathcal{F}^{(2)}$ excluding the endpoints. A **branching point** is a point with degree type at least (3).

Lemma 6.1 (Dual spanning tree). *There is a spanning tree $\text{MST}_\infty^\dagger$ of \mathbb{C} coupled with MST_∞ that is dual in the sense that none of its paths cross any of the paths of MST_∞ , and whose distribution is again that of MST_∞ .*

Note that we are not claiming that $\text{MST}_\infty^\dagger$ is measurable w.r.t. MST_∞ , nor that there is a unique such spanning tree. These claims should be possible to prove, but we will not need them. For all subsequential scaling limits of the Uniform Spanning Tree on \mathbb{Z}^2 , they were proved in [48] via first establishing that the trunk is a topological tree that is everywhere dense in \mathbb{C} , and then defining the dual tree in the complement of the trunk.

Proof of Lemma 6.1. The planar dual of the triangular lattice \mathbb{T} is the hexagonal lattice \mathbb{T}^* , and, as usual, MST_η^M on $\eta\mathbb{T} \cap \mathbb{T}_M^2$ has a dual graph on $\eta\mathbb{T}^* \cap \mathbb{T}_M^2$, denoted by $\text{MST}_\eta^{M\dagger}$. Because of the torus geometry, this dual has some cycles, but it is easy to check that for any null-homotopic cycle in $\eta\mathbb{T}^* \cap \mathbb{T}_M^2$, the edge whose dual in $\eta\mathbb{T} \cap \mathbb{T}_M^2$ has the minimal weight must be present in MST_η^M , hence must be missing from $\text{MST}_\eta^{M\dagger}$, and thus we are almost talking about the *Maximal Spanning Tree* on $\eta\mathbb{T}^* \cap \mathbb{T}_M^2$, denoted by MaxST_η^{M*} , which is defined from the $\text{Unif}[0,1]$ vertex labels $\{V(x)\}$ on \mathbb{T} analogously to MST_η^M : each edge $e^* \in E(\eta\mathbb{T}^* \cap \mathbb{T}_M^2)$ is the dual edge to some $e = (x, y) \in E(\eta\mathbb{T} \cap \mathbb{T}_M^2)$, then we let

$$U(e^*) := (V(x) \vee V(y), V(x) \wedge V(y)),$$

and get MST_η^{M*} (or MaxST_η^{M*}) by removing the maximal (or minimal) edge from each cycle of $E(\eta\mathbb{T}^* \cap \mathbb{T}_M^2)$ w.r.t. the lexicographic ordering. See Figure 6.2.

To make the connection between $\text{MST}_\eta^{M\dagger}$ and MaxST_η^{M*} stronger, note that the two edges in which they differ form a translation invariant random finite subset in the torus, hence, as $M \rightarrow \infty$, the probability that this set intersects any given finite set of edges in \mathbb{Z}^2 goes to 0. Therefore, inside any bounded window, the distributions of $\text{MST}_\eta^{M\dagger}$ and MaxST_η^{M*} are very close to each other, for large M . On the other hand, MaxST_η^{M*} has the same distribution as MST_η^{M*} on $\eta\mathbb{T}^* \cap \mathbb{T}_M^2$, hence if we can understand the geometry of MST_η^{M*} , and prove, for instance, uniform quasi-locality, then we can use that, just as in Subsection 5.2, to understand the geometry of $\text{MST}_\eta^{M\dagger}$.

Now, we claim that, in the spirit of the remark after Figure 1.2, the macroscopic structure of MST_η^{M*} on $\eta\mathbb{T}^* \cap \mathbb{T}_M^2$ can be described using the near-critical ensemble on $\eta\mathbb{T}$. Indeed, if x^* and y^* are vertices in $\eta\mathbb{T}^* \cap \mathbb{T}_M^2$




Figure 6.2: (a) MST_η^M and its dual almost-tree $\text{MST}_\eta^{M\dagger}$. (b) The tree MaxST_η^{M*} associated with the same vertex labels, which is just $\text{MST}_\eta^{M\dagger}$ minus two edges. (c) The tree MST_η^{M*} .

such that the triangular faces they represent have vertices x and y in the same percolation p -cluster of $\eta\mathbb{T} \cap \mathbb{T}_M^2$, then it is immediate to see that the path in MST_η^{M*} between x^* and y^* cannot cross any cycle of $\eta\mathbb{T} \cap \mathbb{T}_M^2$ whose vertices all have labels above p . In this sense, the path does not leave the p -cluster. It follows that two distinct p -clusters cannot be connected by two different paths of MST_η^{M*} (we would otherwise get a cycle in MST_η^{M*}), and hence our entire paper applies to this version of the Minimal Spanning Tree. Thus we get that MST_η^{M*} and $\text{MST}_\eta^{M\dagger}$ have the same unique scaling limit as $\eta \rightarrow 0$ then $M \rightarrow \infty$, denoted by $\text{MST}_\infty^\dagger$, with the same distribution as MST_∞ .

The fact that the paths of $\text{MST}_\infty^\dagger$ do not cross the paths of MST_∞ is clear from obtaining them as scaling limits of discrete dual graphs. \square

It was proved in [4] that any subsequential limit of MST_η in $\hat{\mathbb{C}}$ is a spanning tree of $\hat{\mathbb{C}}$, and hence, using Lemma 6.1, it has one end (a single route to infinity). Furthermore, regularity properties of MST paths proved in that paper implied that the degrees in MST_∞ are almost surely bounded from above by some absolute deterministic constant $k_0 \in \mathbb{N}$, and that the set of points with loops has Hausdorff dimension strictly between 1 and 2. It was also shown, using a Burton-Keane-type argument with trifurcation points and the amenability of the graph \mathbb{Z}^2 (see [17] or [42, Section 7.3]) that the set of branching points is at most countable. It was conjectured in [4] that there are no branching points of degree 4 or larger, and that there are no pinching points. We are now able to establish the latter conjecture, and get close to the former:

Theorem 6.2 (Degree types in MST). *Almost surely in MST_∞ on \mathbb{C} :*

- (i) *there are no points of degree type at least $(2,2)$; in other words, for any two points $x, y \in \mathbb{C}$, none of the paths connecting the two vertices has a pinching point;*
- (ii) *there are no points of degree at least 5 (with any degree type);*
- (iii) *the set of points of degree 4 (with any degree type) is at most countable.*

These hold not only for the scaling limit on $\eta\mathbb{T}$ but also for any subsequential limit on $\eta\mathbb{Z}^2$.

Proof. (i) We want to show that, in any given unit square in \mathbb{C} and for any $0 < \rho < 1$, the probability of MST_∞ having a point of degree type $(2,2)$ with the four strands not being connected within the radius ρ ball of the pinching point is zero. (Note that if there was no positive radius in which the four strands are not connected, then this would in fact be a point of degree type at least 4, not $(2,2)$. On the other hand, the strands must be connected somewhere in MST_∞ , hence we can just assume that they are parts of one path connecting two vertices, which explains the second part of the statement.) For this, it is enough to show that for any $M > 1$, the probability in $\text{MST}_\eta = \text{MST}_\eta^M$ that there is an r -square $B \in B^r([0,1]^2)$ (as in Definition 3.2, with $r < \rho$) such that there is a path γ with two disjoint subpaths, γ_1 and γ_2 , that both enter B but are not connected to each other in MST_η within the ρ -neighbourhood of B tends to 0 as $r \rightarrow 0$, uniformly in $\eta > 0$.

Fix $\alpha > 0$ arbitrarily small. As in the proof of Proposition 4.7, we can take $\lambda < -1$, $\epsilon > 0$, and $\lambda' > 0$ such that with probability at least $1 - \alpha/2$, all λ -clusters in \mathbb{T}_M^2 have diameter less than $\rho/10$, every point of \mathbb{T}_M^2 has in its $\rho/20$ -neighborhood a ring of λ -clusters of diameter at least δ each, for some $0 < \delta < \rho/20$ (uniformly in η), and all λ -clusters of diameter at least δ are connected in $\text{MST}_\eta^{\lambda, \epsilon}$, with these paths going through the same closed pivotals of $\text{PPP}_\lambda^\epsilon$ as the corresponding paths of MST_η . We will assume that this event of probability at least $1 - \alpha/2$ holds, and also that the above r -square B exists, with some $r \ll \rho$ to be determined later.

Any path in MST_η that connects two points in the same λ -cluster must stay in that cluster. Thus, the paths γ_1 and γ_2 that are connected in MST_η but not inside the ρ -neighbourhood of B (denoted by B_ρ), must go through disjoint λ -clusters inside B_ρ . These λ -clusters all have diameter at most $\rho/10$, connected by λ -closed pivotals. Close to each end of each γ_i , there is such a λ -closed pivotal, at distance at least $\rho - \rho/5$ from B . Thus there

must exist two λ -closed paths, separating the λ -clusters of γ_1 from those of γ_2 , going through B , of radius at least $4\rho/5$. See Figure 6.3.




Figure 6.3: Pinching would imply a near-critical 6-arm event.

We would like to bound now the labels from above on the MST_η paths. To this end, let x_i be the point where γ_i leaves the ρ -neighborhood of B , at the end of γ_i that is opposite from γ_{3-i} along γ , for $i = 1, 2$. Around each x_i , there is a ring of macroscopic λ -clusters, the MST_η path from x_1 to x_2 must intersect at least one λ -cluster from each ring, and the part of the path connecting the two rings must go through λ -clusters connected by pivotals with labels at most λ' . Thus, besides the two λ -closed arms between radii r and $4\rho/5$ we also have four λ' -open arms between the same radii. By the near-critical stability of 6-arm probabilities, Proposition 1.3, the probability of this happening anywhere in \mathbb{T}_M^2 is smaller than $\alpha/2$ if r/ρ is chosen small enough. Therefore, the probability of the existence of B is less than α if $r > 0$ is chosen small enough, uniformly in the mesh $\eta > 0$, and we are done.

(ii) It is proved in [12] that the critical monochromatic 5-arm exponent is strictly larger than the polychromatic one, which is 2 (see [50, Corollary A.8]). Therefore, near-critical stability for the monochromatic 5-arm exponent (again, Proposition 1.3) tells us that no near-critical monochromatic 5-arm event between radii r and ρ happens anywhere in $[0, 1]^2$ if r/ρ is small enough. Based on this, as before, we will exclude the existence of an r -square

$B \in B^r([0, 1]^2)$ with degree 5 to distance at least ρ .




Figure 6.4: Degree 5 would imply a near-critical monochromatic 5-arm or a polychromatic 6-arm event.

We look at the λ -clusters traversed by the five branches, for some small $\lambda < -1$. As in part (i), the branches contributed by components at least 2 in the vector of the degree type traverse macroscopic λ -clusters, and hence the labels of their λ -closed pivotals are all at most some uniform λ' . That is, a degree component $k \geq 2$ implies k λ' -open arms from r to ρ (e.g., the open arms labelled 1 to 4 on the left side of Figure 6.4) and if there are more than one such components, we also have λ -closed arms separating them (e.g., the closed arms labelled 6 to 9 on the right side of Figure 6.4). On the other hand, if we have $\ell \geq 1$ branches contributed by components of size 1 in the vector of the degree type, they are necessarily separated from the other branches by λ -closed paths. Thus, they

- either contribute $\ell + 1$ closed arms to the existing 4 open arms,
- or if we have at least two degree components with at least 2 branches, then they raise the number of closed arms by ℓ ,
- or if we do not have any degree components with at least 2 branches, then $\ell \geq 5$, the degree type is all 1's, and we again get ℓ closed arms.

Altogether, we either have at least 5 λ' -open arms, or at least 5 λ -closed arms, or a (λ, λ') near-critical polychromatic event with at least 6 arms. None of these happens if r/ρ is small enough, and we are done.

(iii) Degree 4 points can have five different degree types: (4), (3, 1), (2, 2), (2, 1, 1), (1, 1, 1, 1). The countability of the first two types follows from the countability of branching points proved in [4]. Points of the third type do not exist, by part (i) above. At a point of the fourth type, the dual MST_∞ tree defined in Lemma 6.1 would either have a branching of degree 3, for which we already know countability, or a degree type (2, 2), which does not exist by part (i). (See Figure 6.1 for examples of dual degree types.) Finally, if a point has degree type (1, 1, 1, 1), then the dual tree has a branching point of degree 4 there, so we have countability again.

Since the well-known 5- and 6-arm bounds and Proposition 1.3 hold also for \mathbb{Z}^2 , all the above arguments work fine for subsequential limits of MST_η on $\eta\mathbb{Z}^2$, as well. \square

It is tempting to try and argue that a figure of 6 should imply 5 arms with labels bounded suitably by λ and λ' , and hence by the near-critical stability of the 5-arm exponent (which is 2), the set of points with degree type (2, 1) should be at most countable, but we did not manage to make this argument work.

6.2 A dimension bound for the trunk

Our present techniques reveal very little about the dimension of different subsets of interest in MST_∞ . It was proved in [4] that all the curves connecting any two points almost surely have Hausdorff dimension at least some unspecified deterministic $d_{\min} > 1$ and at most another constant $d_{\max} < 2$. Note that, having a countable number of branching points, the trunk is a countable union of such curves, hence we can equivalently talk about the dimension of the trunk. We will now slightly improve the upper bound to $d_{\max} = 2 - \alpha'_2 < 7/4$, where α'_2 is the monochromatic two-arm (or backbone) exponent of critical percolation, shown to be strictly larger than the polychromatic two-arm exponent $\alpha_2 = 1/4$ in [12]. According to simulations, the true value of the Hausdorff dimension is close to 1.22^- [58, 54], while $2 - \alpha'_2$ is close to $79/48 = 1.646^-$ [11].

Theorem 6.3. *The Hausdorff dimension of the trunk of MST_∞ is almost surely at most $2 - \alpha'_2 < 7/4$, where α'_2 is the monochromatic two-arm exponent of critical percolation.*

Proof. For any $\rho > 0$, let Trunk_η^ρ (resp. Trunk_∞^ρ) be the set of points in $[0, 1]^2$ that have a path of MST_η (resp. MST_∞) passing through them, going to distance at least ρ in both directions. Since the trunk of MST_∞ is a countable union of sets of the form Trunk_∞^ρ , it is enough to prove the

dimension bound on each Trunk_∞^ρ . Consider our usual grid $B^r([0, 1]^2)$ of r -squares, with $r \ll \rho$. The subset of those r -squares that are intersected by Trunk_η^ρ (resp. Trunk_∞^ρ) will be denoted by $\text{Trunk}_\eta^{\rho, r}$ (resp. $\text{Trunk}_\infty^{\rho, r}$), and it is clear that in any coupling where MST_η converges to MST_∞ almost surely, for small enough $\eta > 0$ we have $|\text{Trunk}_\eta^{\rho, r}|/9 \leq |\text{Trunk}_\infty^{\rho, r}| \leq 9|\text{Trunk}_\eta^{\rho, r}|$ with probability close to 1, where the factors of 9 accommodate the possibility of the points of Trunk^ρ moving across the boundaries of r -squares. Therefore, it suffices to prove that for any $\beta > 0$ there is a sequence $r_k \rightarrow 0$ such that

$$\mathbf{P}[|\text{Trunk}_\eta^{\rho, r_k}| > r_k^{-2+\alpha'_2-\beta}] < 2^{-k} \quad (6.3)$$

for all small enough $\eta = \eta_k > 0$, because then Borel-Cantelli gives that the Hausdorff (even the lower Minkowski) dimension of Trunk_∞^ρ is almost surely at most $2 - \alpha'_2 + \beta$.

To prove (6.3), take $\bar{\lambda}_k$ and ϵ_k such that with probability at least $1 - 3^{-k}$ all λ_k -clusters have diameter at most $\rho/10$, all points have a ring of λ_k -clusters of diameter at least $\delta > 0$ in their $\rho/20$ -neighborhood, and all λ_k -clusters of diameter at least δ are connected in $\text{MST}_\eta^{\bar{\lambda}_k, \epsilon_k}$. Condition on this event, denoted by \mathcal{G}_k . Then, just as in the proof of Theorem 6.2, every element of $\text{Trunk}_\eta^{\rho, r}$ has a $\bar{\lambda}_k$ -near-critical monochromatic 2-arm event from radius r to $\delta/2$. From near-critical stability, we know that, for any $B \in B^r([0, 1]^2)$, denoting this 2-arm event by $\mathcal{A}'_2(B, r, \delta/2, \bar{\lambda}_k)$, we have

$$\mathbf{P}[\mathcal{A}'_2(B, r, \delta/2, \bar{\lambda}_k)] < C_{\delta, k} r^{\alpha'_2}.$$

Since $1/(1 - 3^{-k}) < 2$, the previous line gives

$$\mathbf{P}[\mathcal{A}'_2(B, r, \delta/2, \bar{\lambda}_k) \mid \mathcal{G}_k] < 2C_{\delta, k} r^{\alpha'_2},$$

and, summing up over B ,

$$\mathbf{E}[|\text{Trunk}_\eta^{\rho, r}| \mid \mathcal{G}_k] < 2C_{\delta, k} r^{-2+\alpha'_2}.$$

Then, by Markov's inequality, for any $\beta > 0$,

$$\mathbf{P}[|\text{Trunk}_\eta^{\rho, r}| > 2C_{\delta, k} r^{-2+\alpha'_2-\beta/2} \mid \mathcal{G}_k] < r^{\beta/2}.$$

By taking $r_k > 0$ so small that $2C_{\delta, k} < r_k^{-\beta/2}$ and $r_k^{\beta/2} < 3^{-k}$, we get that

$$\mathbf{P}[|\text{Trunk}_\eta^{\rho, r_k}| > r_k^{-2+\alpha'_2-\beta} \mid \mathcal{G}_k] < 3^{-k}.$$

Since we have

$$\begin{aligned} \mathbf{P}[|\text{Trunk}_\eta^{\rho, r_k}| > r_k^{-2+\alpha'_2-\beta}] &< \mathbf{P}[|\text{Trunk}_\eta^{\rho, r_k}| > r_k^{-2+\alpha'_2-\beta} \mid \mathcal{G}_k] + \mathbf{P}[\mathcal{G}_k^c] \\ &< 3^{-k} + 3^{-k}, \end{aligned}$$

we have verified (6.3) and completed the proof. \square

7 Invasion percolation

The Invasion Tree in a finite graph is simply the MST itself, hence it cannot provide us with a good finite approximation to InvPerc in the infinite plane. Instead, we will consider the following finite versions:

- $\text{InvPerc}_\eta^{M,\partial}$ will be the tree built by the invasion process started from the origin, stopped at the first time that it reaches $\partial[-M, M]^2$.
- For a fixed vertex $x \in V(\eta\mathbb{T})$ and M large enough so that $x \in [-M, M]^2$, we will denote by $\text{InvPerc}_\eta^{M,x}$ the invasion process in the torus \mathbb{T}_M^2 , started from the origin, stopped at the first time when it reaches x .

When $M \rightarrow \infty$, the weak limits of the above measures are $\text{InvPerc}_\eta = \text{InvPerc}_\eta(0)$ and $\text{InvPerc}_\eta(0) \cup \text{InvPerc}_\eta(x)$, respectively. Of course, the latter coincides with $\text{InvPerc}_\eta(0)$ with positive probability, and the symmetric difference $\text{InvPerc}_\eta(0) \triangle \text{InvPerc}_\eta(x)$ is almost surely finite. These results are classical [20, 7, 6, 43].

Given the enhanced ϵ -networks $\text{EN}_\eta^{\bar{\lambda},\epsilon}$ and $\text{EN}_\infty^{\bar{\lambda},\epsilon}$ defined in Proposition 3.9, the cut-off versions of the above invasion trees, both in the discrete case and in the continuum, can be defined quite similarly to $\text{MSF}^{\bar{\lambda},\epsilon}$ (done in Definition 3.11) and $\text{MST}^{\bar{\lambda},\epsilon}$ (in Lemma 4.4):

Definition 7.1 (The cut-off invasion trees $\text{InvPerc}^{\bar{\lambda},\epsilon,s}$ in \mathbb{T}_M^2 , with target set ∂ or x).

1. Consider the edge-labelled graph defined in steps 1-3 of Definition 3.11 on the set of the primal routers of $\text{EN}^{\bar{\lambda},\epsilon}$ as vertices.
2. Take the giant component of this graphs, which exists with large probability by Lemma 4.4. On the bad event that this giant cannot be defined, our cut-off invasion trees will be just degenerate one-point trees.
3. Take the router closest to the origin 0; in case of a tie, decide in some arbitrary but fixed manner. This will be called the origin router. Furthermore, consider all routers that are at most distance $s > 0$ from the target set $\partial[-M, M]^2$ or x . By Lemma 4.4, for any $s > 0$, if λ is very negative, λ' is very positive, and ϵ is small, then with high probability the set of these target routers is not empty. When it is empty, the invasion tree will consist of just the origin router.
4. Take the invasion tree process in the above graph, started from the origin router, stopped when reaching any of the target routers. There

may be steps in the invasion process when more than one edge with label λ lead out of the invaded set; in such a case, all these edges get invaded simultaneously.

It was proved already in [4] that the set of points with degree larger than 1 (i.e., points in the trunk or having a loop) in any subsequential scaling limit of MST_η is of zero measure. Therefore, almost surely there is a unique path of MST_∞ that goes to the origin, and hence we did not lose any information in the above definition by taking the router closest to the origin instead of considering all routers that are s -close to it.

Given this definition, we immediately have the following analogues of Corollary 4.5 and Proposition 4.7. Note the double meaning of the parameter s : if we want to reach precision $s > 0$ in d_{Ω_M} , it is enough to get s -close to the target sets.

Lemma 7.2. *For any $M > 0$, target set ∂ or $x \in \mathbb{T}_M^2$, and any $s, \alpha > 0$, if $\lambda < -1$ is very negative, $\epsilon > 0$ is small, and $\lambda' > 1$ is large enough, then, in the coupling of Proposition 3.9 (ii) between $(\omega_\eta^\lambda, \text{PPP}_\eta^\epsilon)$ and $(\omega_\infty^\lambda, \text{PPP}_\infty^\epsilon)$, for all $\eta > 0$ small enough,*

$$\mathbf{P}[d_{\Omega_M}(\text{InvPerc}_\eta^{\bar{\lambda}, \epsilon, s}, \text{InvPerc}_\infty^{\bar{\lambda}, \epsilon, s}) < s] > 1 - \alpha.$$

Lemma 7.3. *For any $M > 0$, target set ∂ or $x \in \mathbb{T}_M^2$, and $s, \alpha > 0$, if $\lambda < -1$ is very negative, $\epsilon > 0$ is small, and $\lambda' > 1$ is large enough, then, for all $\eta > 0$ small enough,*

$$\mathbf{P}[d_{\Omega_M}(\text{InvPerc}_\eta, \text{InvPerc}_\eta^{\bar{\lambda}, \epsilon, s}) < s] > 1 - \alpha,$$

where InvPerc_η is only a shorthand now for $\text{InvPerc}_\eta^{M, \partial}$ or $\text{InvPerc}_\eta^{M, x}$.

Using these lemmas, the proof of the following theorem follows exactly the proofs of Theorem 5.1 and Theorem 1.1.

Theorem 7.4. *For any $M > 0$, the invasion trees $\text{InvPerc}_\eta^{M, \partial}$ and $\text{InvPerc}_\eta^{M, x}$ started at the origin of $\eta\mathbb{T} \cap \mathbb{T}_M^2$ converge in distribution as $\eta \rightarrow 0$, in the metric d_{Ω_M} of Definition 2.2, to the unique scaling limits $\text{InvPerc}_\infty^{M, \partial}$ and $\text{InvPerc}_\infty^{M, x}$, respectively.*

The invasion tree InvPerc_η started at the origin of $\eta\mathbb{T}$ converges in distribution to a unique scaling limit InvPerc_∞ that is invariant under scalings and rotations.

As $M \rightarrow \infty$, the weak limit of $\text{InvPerc}_\infty^{M, \partial}$ is InvPerc_∞ and the weak limit of $\text{InvPerc}_\infty^{M, x}$ is $\text{InvPerc}_\infty(0) \cup \text{InvPerc}_\infty(x)$.

8 Questions and conjectures

We start with a very natural and interesting open problem:

Conjecture 8.1.

- (i) *Show that MST_∞ is not conformally invariant. In particular, show that it is different from the scaling limit of the Uniform Spanning Tree, described in [41].*
- (ii) *Show that InvPerc_∞ is not conformally invariant.*

This is of course supported by simulation results [59]. Moreover, it was explained in [27] why our description of these scaling limits using the near-critical ensemble gives serious support to this conjecture, and why it is nevertheless not at all an easy issue. The case of InvPerc_∞ might be simpler, using the results of [23].

Probably the simplest open problem in this section is the following one, left open by Lemma 6.1:

Conjecture 8.2. *Show that there is a unique dual tree $\text{MST}_\infty^\dagger$, measurable w.r.t. MST_∞ .*

The following questions are left open by Theorem 6.2:

Question 8.3 (Topology of MST_∞).

- (i) *Are there non-simple paths giving figures of 6, i.e., points with degree type $(2, 1)$?*
- (ii) *Show that almost surely there are no points of degree 4.*

Finally, sharpening the bound of Theorem 6.3 would probably require new techniques:

Question 8.4. *Find the Hausdorff dimension of the paths of MST_∞ .*

References

- [1] Louigi Addario-Berry, Nicolas Broutin and Christina Goldschmidt. The continuum limit of critical random graphs. *Probab. Theory Related Fields* **152** (2012), 367–406. [arXiv:arXiv:0903.4730](https://arxiv.org/abs/0903.4730) [math.PR]

- [2] Louigi Addario-Berry, Nicolas Broutin, Christina Goldschmidt and Grégory Miermont. The scaling limit of the minimum spanning tree of the complete graph. *Ann. Probab.*, to appear. [arXiv:1301.1664](https://arxiv.org/abs/1301.1664) [math.PR].
- [3] Michael Aizenman. On the number of incipient spanning clusters. *Nuclear Phys. B* **485** (1997), no. 3, 551–582. [[arXiv:cond-mat/9609240](https://arxiv.org/abs/cond-mat/9609240)]
- [4] Michael Aizenman, Almut Burchard, Charles M. Newman and David B. Wilson. Scaling limits for minimal and random spanning trees in two dimensions. *Random Structures Algorithms* **15** (1999), 319–365. [[arXiv:math.PR/9809145](https://arxiv.org/abs/math.PR/9809145)]
- [5] David Aldous and J. Michael Steele. The objective method: probabilistic combinatorial optimization and local weak convergence. *Probability on discrete structures*, vol. 110 of Encyclopaedia Math. Sci., Springer, Berlin, 2004, pp. 1–72. <http://www.stat.berkeley.edu/~aldous/Papers/me101.pdf>
- [6] Kenneth S. Alexander. Percolation and minimal spanning forests in infinite graphs. *Ann. Probab.* **23** (1995), 87–104.
- [7] K. S. Alexander and S. Molchanov. Percolation of level sets for two-dimensional random fields with lattice symmetry. *J. Stat. Phys.* **77** (1994), 627–643.
- [8] Omer Angel, Jesse Goodman, Frank den Hollander, and Gordon Slade. Invasion percolation on regular trees. *Ann. Probab.* **36** (2008), 420–466. [[arXiv:math.PR/0608132](https://arxiv.org/abs/math.PR/0608132)]
- [9] Omer Angel, Jesse Goodman and Mathieu Merle. Scaling limit of the invasion percolation cluster on a regular tree. *Ann. Probab.* **41** (2013), 229–261. [arXiv:0910.4205](https://arxiv.org/abs/0910.4205) [math.PR]
- [10] Simon Aumann. Singularity of full scaling limits of planar near-critical percolation. *Stochastic Processes and their Applications* **124** (2014), no. 11, 3807–3818. [arXiv:1301.5175](https://arxiv.org/abs/1301.5175) [math.PR]
- [11] Vincent Beffara and Pierre Nolin. Numerical estimates for monochromatic percolation exponents. Unpublished, 2009. <http://perso.ens-lyon.fr/vincent.beffara/pdfs/Monochromatic-simu.pdf>
- [12] Vincent Beffara and Pierre Nolin. On monochromatic arm exponents for 2D critical percolation. *Ann. Probab.* **39** (2011), 1286–1304. [arXiv:0906.3570](https://arxiv.org/abs/0906.3570) [math.PR]

- [13] Jacob van den Berg and Rene Conijn. The gaps between the sizes of large clusters in 2D critical percolation. *Electron. Commun. Probab.* **18** (2013), no. 92, 1–9. [arXiv:1310.2019](#) [[math.PR](#)]
- [14] Christian Borgs, Jennifer T. Chayes, Harry Kesten, and Joel Spencer. The birth of the infinite cluster: finite-size scaling in percolation. *Comm. Math. Phys.* **224** (2001), no. 1, 153–204. <http://research.microsoft.com/en-us/um/people/borgs/papers/iicbirth.pdf>
- [15] Otokar Borůvka. O jistém problému minimálním. *Práce Moravské Přírodovědecké Společnosti Brno* **3** (1926), 37–58.
- [16] Simon R. Broadbent and John M. Hammersley. Percolation processes I. Crystals and mazes. *Proc. Cambridge Phil. Soc.* **53** (1957), 629–641.
- [17] Robert M. Burton and Michael Keane. Density and uniqueness in percolation. *Comm. Math. Phys.* **121** (1989), 501–505.
- [18] Federico Camia, Luiz Renato Fontes and Charles M. Newman. Two-dimensional scaling limits via marked non-simple loops. *Bull. Braz. Math. Soc. (N.S.)* **37** (2006), 537–559. [[arXiv:cond-mat/0604390](#)]
- [19] Federico Camia and Charles M. Newman. Two-dimensional critical percolation: the full scaling limit. *Comm. Math. Phys.* **268** (2006), no. 1, 1–38. [[arXiv:math.PR/0605035](#)]
- [20] Jennifer T. Chayes, Lincoln Chayes, and Charles M. Newman. The stochastic geometry of invasion percolation. *Comm. Math. Phys.* **101** (1985), 383–407.
- [21] Jennifer T. Chayes, Lincoln Chayes, and Charles M. Newman. Bernoulli percolation above threshold: an invasion percolation analysis. *Ann. Probab.* **15** (1987), 1272–1287.
- [22] Sourav Chatterjee and Sanchayan Sen. Minimal spanning trees and Stein’s method. *Ann. App. Probab.*, to appear. [arXiv:1307.1661](#) [[math.PR](#)]
- [23] Michael Damron, Artëm Sapozhnikov and Bálint Vágvolgyi. Relations between invasion percolation and critical percolation in two dimensions. *Ann. Probab.* **37** (2009), 2297–2331. [arXiv:arXiv:0806.2425](#) [[math.PR](#)]

- [24] Michael Damron and Artëm Sapozhnikov. Limit theorems for 2D invasion percolation. *Ann. Probab.* **40** (2012), 893–920. [arXiv:1005.5696](#) [math.PR]
- [25] Michael Damron. Recent work on chemical distance in critical percolation. *Preprint*, [arXiv:1602.00775](#) [math.PR]
- [26] Christophe Garban, Gábor Pete, and Oded Schramm. The Fourier spectrum of critical percolation. *Acta Math.* **205** (2010), no. 1, 19–104. [arXiv:0803.3750](#) [math.PR]
- [27] Christophe Garban, Gábor Pete, and Oded Schramm. The scaling limit of the minimal spanning tree — a preliminary report. *XVIth International Congress on Mathematical Physics*, ed. P. Exner, pages 475–480, World Sci. Publ., Hackensack, NJ, 2010. [arXiv:0909.3138](#) [math.PR]
- [28] Christophe Garban, Gábor Pete, and Oded Schramm. Pivotal, cluster and interface measures for critical planar percolation. *J. Amer. Math. Soc.* **26** (2013), 939–1024. [arXiv:1008.1378](#) [math.PR]
- [29] Christophe Garban, Gábor Pete, and Oded Schramm. The scaling limits of near-critical and dynamical percolation. *J. Europ. Math. Soc.*, to appear, [arXiv:1305.5526](#) [math.PR]
- [30] Christophe Garban and Jeffrey E. Steif. Lectures notes on noise sensitivity and percolation. *Probability and Statistical Physics in Two and more Dimensions (Buzios, Brazil)*, ed. D. Ellwood, C. Newman, V. Sidoravicius and W. Werner. Clay Math. Proc. 15, 49–154, 2012. [arXiv:1102.5761](#) [math.PR]
- [31] Ronald L. Graham and Pavol Hell. On the history of the minimum spanning tree problem. *Ann. Hist. Comput.* **7** (1985), 43–57. https://www.researchgate.net/publication/3331240_On_the_History_of_the_Minimum_Spanning_Tree_Problem
- [32] Geoffrey Grimmett. *Percolation. Second edition.* Grundlehren der Mathematischen Wissenschaften, 321. Springer-Verlag, Berlin, 1999.
- [33] Olle Häggström, Yuval Peres and Roberto H. Schonmann. Percolation on transitive graphs as a coalescent process: Relentless merging followed by simultaneous uniqueness. *Perplexing Problems in Probability, Festschrift in honor of Harry Kesten*, ed. M. Bramson and R. Durrett, pp. 69–90. Birkhäuser, 1999.

- [34] Alan Hammond, Gábor Pete, and Oded Schramm. Local time on the exceptional set of dynamical percolation, and the Incipient Infinite Cluster. *Ann. Probab.* **43** (2015), 2949–3005. [arXiv:1208.3826](https://arxiv.org/abs/1208.3826) [math.PR].
- [35] Antal A. Járai. Incipient infinite percolation clusters in 2D. *Ann. Probab.* **31** (2003), 444–485. <http://projecteuclid.org/euclid.aop/1046294317>
- [36] Harry Kesten. *Percolation Theory for Mathematicians*. Birkhäuser, Boston, 1982.
- [37] Harry Kesten. Subdiffusive behavior of random walk on a random cluster. *Ann. Inst. H. Poincaré Probab. Statist.* **22** (1986), no. 4, 425–487.
- [38] Harry Kesten. Scaling relations for 2D-percolation. *Comm. Math. Phys.* **109** (1987), no. 1, 109–156.
- [39] Joseph B. Kruskal. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc. Amer. Math. Soc.* **7** (1956), 48–50.
- [40] Gregory F. Lawler, Oded Schramm, and Wendelin Werner. One-arm exponent for critical 2D percolation. *Electron. J. Probab.* **7** (2002), paper no. 2, 1–13. [arXiv:math.PR/0108211]
- [41] Gregory F. Lawler, Oded Schramm and Wendelin Werner. Conformal invariance of planar loop-erased random walks and uniform spanning trees. *Ann. Probab.* **32** (2004), 939–995. [arXiv:math.PR/0112234]
- [42] Russell Lyons and Yuval Peres. *Probability on trees and networks*. Cambridge University Press, 2016. <http://mypage.iu.edu/~rdlyons>
- [43] Russell Lyons, Yuval Peres and Oded Schramm. Minimal spanning forests. *Ann. Probab.* **34** (2006), 1665–1692. [arXiv:math.PR/0412263v5]
- [44] Charles M. Newman, Vincent Tassion, and Wei Wu. Critical percolation and the minimal spanning tree in slabs. *Preprint*, arXiv:1512.09107 [math.PR]
- [45] Pierre Nolin. Near-critical percolation in two dimensions. *Electron. J. Probab.* **13** (2008), paper no. 55, 1562–1623. <https://projecteuclid.org/euclid.ejp/1464819128>

- [58] Benjamin Wieland and David B. Wilson. Winding angle variance of Fortuin-Kasteleyn contours. *Phys. Rev. E* **68** (2003), 056101. [arXiv:1002.3220](https://arxiv.org/abs/1002.3220) [cond-mat.stat-mech]
- [59] David B. Wilson. On the Red-Green-Blue model. *Phys. Rev. E* **69** (2004), 037105. [[arXiv:cond-mat/0212042](https://arxiv.org/abs/cond-mat/0212042)]
- [60] Joseph E. Yukich. *Probability theory of classical Euclidean optimization problems*. LNM 1675, Springer, Berlin, 1998.

Christophe Garban

Institut Camille Jordan, Université Lyon 1
<http://math.univ-lyon1.fr/~garban>

Gábor Pete

Rényi Institute, Hungarian Academy of Sciences, Budapest, and
Institute of Mathematics, Budapest University of Technology and Economics
<http://www.math.bme.hu/~gabor>

Oded Schramm (December 10, 1961 – September 1, 2008)

Microsoft Research
<http://research.microsoft.com/en-us/um/people/schramm/>