

ON THE SCALING LIMITS OF PLANAR PERCOLATION

BY ODED SCHRAMM¹ AND STANISLAV SMIRNOV²

(WITH AN APPENDIX BY CHRISTOPHE GARBAN)

*Microsoft Research, Université de Genève and
CNRS, ENS Lyon*

We prove Tsirelson's conjecture that any scaling limit of the critical planar percolation is a black noise. Our theorems apply to a number of percolation models, including site percolation on the triangular grid and any subsequential scaling limit of bond percolation on the square grid. We also suggest a natural construction for the scaling limit of planar percolation, and more generally of any discrete planar model describing connectivity properties.

1. Introduction.

1.1. *Motivation.* This paper has a two-fold motivation: to propose a new construction for the (subsequential) scaling limits of the critical and near-critical percolation in the plane, and to show that such limits are two-dimensional *black noises* as suggested by Boris Tsirelson.

Percolation is perhaps the simplest statistical physics model exhibiting a phase transition. We will be interested in planar percolation models, the archetypical examples being site percolation on the triangular lattice and bond percolation on the square lattice. More generally, our theorems apply to a wide class of models on planar graphs, satisfying the assumptions outlined below.

Consider a planar graph and fix a *percolation probability* $p \in [0, 1]$. In *site percolation*, each site is declared *open* or *closed* with probabilities p and $1 - p$, independently of others. In the pictures, we represent open and closed sites by blue and yellow colors correspondingly. One then studies open *clusters*, which are maximal connected subgraphs of open sites. In *bond percolation*, each bond is similarly declared open or closed independently of others, and one studies clusters of bonds.

Percolation as a mathematical model was originally introduced by Broadbent and Hammersley to model the flow of liquid through a random porous medium. It

Received February 2010; revised February 2011.

¹December 10, 1961–September 1, 2008.

²Supported by the European Research Council AG CONFRA, the Swiss National Science Foundation and by the Chebyshev Laboratory (Faculty of Mathematics and Mechanics, St. Petersburg State University) under the grant of the government of the Russian Federation.

MSC2010 subject classifications. Primary 60K35; secondary 28C20, 82B43, 60G60.

Key words and phrases. Percolation, noise, scaling limit.

is assumed that the liquid can move freely between the neighboring open sites, and so it can flow between two regions if there is an open path between them. Fixing a quad Q (= a topological quadrilateral), we say that there is an *open crossing* (between two distinguished opposite sides) for a given percolation configuration, if for its restriction to the quad, there is an open cluster intersecting both opposite sides. One then defines the crossing probability for Q as the probability that such a crossing exists.

We will be interested in the (subsequential) scaling limits of percolation, as the lattice mesh tends to zero. It is expected that for a large class of models the *sharp threshold* phenomenon occurs: there is a critical value p_c such that for a fixed quad Q the crossing probability tends to 1 for $p > p_c$ and to 0 for $p < p_c$ in the scaling limit. For bond percolation on the square lattice and site percolation on the triangular lattice, it is known that $p_c = 1/2$ by the classical theorem of Kesten [20].

In the critical case $p = p_c$ the crossing probabilities are expected to have nontrivial scaling limits. Moreover, those are believed to be universal, conformally invariant, and given by Cardy's formula [8, 23]. For the models we are concerned with, this was established only for site percolation on the triangular lattice [34, 35]. However, the Russo–Seymour–Welsh theory provides existence of nontrivial subsequential limits, which is enough for our purposes. Our theory also applies to near-critical scaling limits like in [29], when the mesh tends to zero and the percolation probability to its critical value in a coordinated way, so as the crossing probabilities have nontrivial limits, different from the critical one. The Russo–Seymour–Welsh theory still applies below any given scale.

Having a (subsequential) scaling limit of the (near) critical percolation, we can ask how to describe it. The discrete models have finite σ -fields, but in the scaling limit the cardinality blows up, so we have to specify how we pass to a limit. First, we remark that we restrict ourselves to the *crossing events* (which form the original physical motivation for the percolation model), disregarding events of the sort “number of open sites in a given region is even” or “number of open sites in a given region is bigger than the number of closed ones.” While such events could be studied in another context, they are of little interest in the framework motivated by physics. Furthermore, we restrict ourselves to the *macroscopic* crossing events, that is, crossings of quads of fixed size. Otherwise, one could add to the σ -field crossing events for microscopic quads (of the lattice mesh scale or some intermediate scale), which have infinitesimal size in the scaling limit. The resulting construction would be an extension of ours, similar to a nonstandard extension of the real numbers. But because of the locality of the percolation, different scales are independent, so it won't yield new nontrivial information (as our results in fact show). However, we would like to remark that in dependent models such extensions could provide new information. For example, Kenyon has calculated probabilities of local configurations appearing in domino tilings and loop erased random walks [19], while such information is lost in describing the latter by an SLE curve.

Summing it up, we restrict ourselves only to the scaling limit of the *macroscopic crossing events* for percolation—it describes the principal physical properties (though adding microscopic events can be of interest). The question then arises, how to describe such a scaling limit, and several approaches were proposed. We list some of them along with short remarks:

(1) *Random coloring of a plane* into yellow and blue colors would be a logical candidate, since in the discrete setting we deal with random colorings of graphs. However, it is difficult to axiomatize the allowed class of “percolation-induced colorings,” and connectivity properties are hard to keep track of. In particular, in the scaling limit almost every point does not belong to a cluster, but is rather surrounded by an infinite number of nested open and closed clusters, and so has no well-defined color.

(2) *Collection of open clusters* as random compact subsets of the plane would satisfy obvious precompactness (for Borel probability measures on collections of compacts endowed with some version of the Hausdorff distance), guaranteeing existence of the subsequential scaling limits. However, crossing probabilities cannot be extracted: as a toy model imagine two circles, O and C , intersecting at two points, x and y . Consider two different configurations: in the first one, $O \setminus \{x\}$ is an open crossing and $C \setminus \{y\}$ is a closed one; in the second one, $O \setminus \{y\}$ is an open crossing and $C \setminus \{x\}$ is a closed one. In both configurations, the cluster structure is the same (with O being the open and C being the closed cluster), while some crossing events differ. Thus one has to add the state of the “pivotal” points (like x and y) to the description, making it more complicated. Constructing such a limit, proving its uniqueness, and studying it would require considerable technical work.

(3) *Aizenman’s web* represents percolation configuration by a collection of all curves (= crossings) running inside open clusters. Here pre-compactness follows from the Aizenman–Burchard work [1], based on the Russo–Seymour–Welsh theory, and crossing events are easy to extract. Moreover, it turns out that the curves involved are almost surely Hölder-continuous, so this provides a nice geometric description of the connectivity structure. However, there is a lot of redundancy in taking *all* the curves inside clusters, and there are deficiencies similar to those of the previous approach.

(4) *Loop ensemble*—instead of clusters themselves, it is enough to look at the interfaces between open and closed clusters, which form a collection of nested loops. Such a limit was constructed by Camia and Newman in [7] from Schramm’s SLE loops. However, it was not shown that such a limit is full (i.e., contains all the crossing events), and in general, extracting geometric information is far from being straightforward.

(5) *Branching exploration tree*—there is a canonical way to explore all the discrete interfaces: we follow one, and when it makes a loop, branch into two created components. The continuum analogue was studied (in relation with the loop ensembles) by Sheffield in [33], and this is perhaps the most natural approach for

general random cluster models with dependence; see [18]. Construction is rather technical, however the scaling limit can be described by branching SLE curves and so is well suited for calculations.

(6) *Height function*—in the discrete setting one can randomly orient interfaces, constructing an integer-valued height function, which changes by ± 1 whenever an interface is crossed. This stores all the connectivity information (at least in our setup of interfaces along the edges of trivalent graphs), but adds additional randomness of the height change. It is expected that after appropriate manipulations (e.g., coarse-graining and compactifying—i.e., projecting mod 1 to a unit interval) the height function would converge in the scaling limit to a random distribution, which is a version of the free field. This is the cornerstone of the “Coulomb gas” method of Nienhuis [28], which led to many a prediction. It is not immediate how to connect this approach to the more geometric ones, retrieving the interfaces (= level curves) from the random distribution, but SLE theory suggests some possibilities. Also it is not clear whether we can implement this approach so that the extra randomness disappears in the limit, while the connectivity properties remain.

(7) *Correlation functions*—similarly to the usual Conformal Field Theory approach, instead of dealing with a random field as a random distribution, we can restrict our attention to its n -point correlations. For example, that would be the usual physics framework for fields arising in approach (6), but to implement it mathematically we would need to reconstruct height lines of a field, based on its correlations, which is perhaps more difficult than developing (6) by itself, and possibly requires some extra assumptions on the field. A more geometric approach would be to study for every finite collection of points $\{z_j^k\}_{j,k}$ the probability $F_r(z_j^k)_{j,k}$ that for every k all the balls $\{B(z_j^k, r)\}_j$ are connected by an open cluster. To obtain the correlation function, one takes the double limit, passing to the scaling limit first and sending $r \rightarrow 0$ afterwards. Since the probability of a radius r ball being touched by a macroscopic cluster (conjecturally) decays as $r^{5/48}$, one has to normalize accordingly, considering F as a $\prod_{j,k} |dz_j^k|^{5/48}$ -form. The geometric approach seems more feasible, but along the way one would have to further develop a version of CFT corresponding to such “connectivity” correlations.

(8) *Collection of quads crossed*—in the discrete setting one can describe a percolation configuration by listing the finite collection of all discrete quads (i.e., topological quadrilaterals) having an open crossing. We propose a continuous analogue, described in detail below. The percolation configuration space thus becomes a space of quad collections, satisfying certain additional properties.

(9) *Continuous product of probability spaces, or noise*—this approach was advocated by Tsirelson [38–40]. In the discrete setting, the site percolation (on a set V of vertices) is given by the *product probability space* ($\Omega_V = \prod_V \{\text{open, closed}\}, \mathcal{F}_V, \mu_V$). Note that when V is decomposed into two disjoint subsets, we obviously have

$$(1.1) \quad \mathcal{F}_V = \mathcal{F}_{V_1} \times \mathcal{F}_{V_2}.$$

Straightforward generalization of the product space to the continuous case does not work: the product σ -field for the space $\prod_{\mathcal{C}}\{\text{open, closed}\}$ won't contain crossing events. Instead Tsirelson proposed to send the percolation measure space to a scaling limit, described by a *noise* or a *homogeneous continuous product of probability spaces*. Namely, for every smooth domain D the percolation scaling limit inside it would be given by a probability space $(\Omega_D, \mathcal{F}_D, \mu_D)$, which is translation invariant, continuous in D , and satisfies the property (1.1) for a smooth domain V cut into two smooth domains V_1, V_2 by a curve. There is a well-developed theory of such spaces, but establishing the latter property was problematic [40].

To most constructions, one can also add the information about the closed clusters, though in principle it should be retrievable from the open ones by duality. This list is not exhaustive, but it shows the variety of possible descriptions of the same object. Indeed, in the discrete settings the approaches above contain all the information about macroscopic open clusters, and are expected to keep it in the scaling limit; on the other hand with an appropriate setup [one has to be careful, especially with (6)], we do not expect to pick up any extra information on the way. However, showing that most of these approaches lead to equivalent results, and in particular to isomorphic σ -fields, is *far from easy*.

We decided to proceed with the approach (8), since it follows the original physical motivation for the percolation model, and also provides us with enough precompactness to establish existence of subsequential scaling limits *before* we apply any percolation techniques.

It also serves well our second purpose: to provide description of the percolation (subsequential) scaling limits, following Tsirelson's approach (9), and show that it is a *noise*, that is, a homogeneous continuous (parameterized by an appropriate Boolean algebra of piecewise-smooth planar domains) product of probability spaces.

The question whether critical percolation scaling limit is a noise was posed by Tsirelson in [38–40], where he has also noted that such noise must be nonclassical, or *black* (compared to the classical white noise). The theory of black noises was started by Tsirelson and Vershik in [41]. In particular, they constructed the first examples of zero and one-dimensional black noises (and more of those were found since); we provide the first genuinely two-dimensional example.

To establish that percolation leads to noise, it is enough to prove that its scaling limit, restricted to two adjacent rectangles, determines the scaling limit in their union. Following is a version of this result, perhaps more tangible (though stated slightly informally here). It is proved in Proposition 4.1, which is the most technical result of our paper. *Consider a rectangle, and a smooth path α cutting it. We show that for every $\epsilon > 0$ there is a finite number of percolation crossing events in quads disjoint from α , from which one can reliably predict whether or not the rectangle is crossed, in the sense that the probability for a mistake is less than ϵ .* The essential point is that the set of crossing events that one looks at *does not depend on the mesh of the lattice*, which allows us to deduce the needed result, thus

showing that any subsequential scaling limit of the critical percolation is indeed a *black noise*.

1.2. *Percolation basics and notation.* For completeness of exposition, we present most of the needed background below; interested readers can also consult the books [4, 16, 21, 42].

We start with a rather general setup, which in particular includes all planar site and bond models (as well as percolation on planar hypergraphs; cf. [5]). Since we will be interested in scaling limits, it is helpful from the very beginning to work with subsets of the plane, rather than graphs.

We consider a locally finite *tiling* H of the plane (or its subdomain) by topologically closed polygonal *tiles* P . Furthermore, we ask it to be trivalent, meaning that at most three tiles meet at every vertex and so any two tiles are either disjoint or share a few edges. A *percolation model* fixes a percolation probability $p(P) \in [0, 1]$ for every tile P , and declares it open (or closed) with probability $p(P)$ [or $1 - p(P)$] independently of the others. In principle, we can work with a random tiling H , with percolation probabilities $p(P, H)$ satisfying appropriate measurability conditions (so that crossing events are measurable).

More generally, we can consider a *random coloring* of tiles into open and closed ones, given by some measure μ on the space $\prod_{P \in H} \{\text{open, closed}\}$ with a product σ -field (which contains all events concerned with a finite number of tiles). Percolation models correspond to product measures μ .

We define open (closed) *clusters* as connected topologically closed subsets of the plane—components of connectivity of the union of open (closed) tiles taken with boundary.

All site and bond percolation models on planar graphs can be represented in this way. Indeed, given a planar site percolation model, we replace sites by tiles, so that two of them share a side if and only if two corresponding sites are connected by a bond. It is easy to see that such a collection always exists (e.g., tiles being the faces of the dual graph; see Figure 1 for the triangular lattice case). Each tile is declared open (or closed) with the same probability p (or $1 - p$) as the corresponding site.




FIG. 1. *Critical site percolation on the triangular lattice. Each site is represented by a hexagonal tile, which is open (blue) or closed (yellow) with probabilities 1/2 independently of others. Interfaces between open and closed clusters are pictured in bold.*

If our original graph is a triangulation, then the dual one is trivalent and only three tiles can meet at a point. However, when our graph has a face with four or more sides, two nonadjacent sites can correspond to tiles sharing a vertex (but not an edge). To resolve the resulting ambiguity we insert a small tile P around every tiling vertex, where four or more tiles meet, and set it to be *always closed*, taking $p(P) = 0$. After the modification, at most three tiles can meet at a point, as required.

Bond percolation can also be represented by partitioning the plane into tiles. There will be a tile for every bond, site, and face of the original bond graph, chosen so that two different face-tiles or two different site-tiles are disjoint. Moreover, a bond-tile shares edges with two site-tiles and two face-tiles, corresponding to adjacent sites and faces. The bond-tiles are open (or closed) independently, with the same probability p (or $1 - p$) as the corresponding bonds. The site-tiles P are then always open [$p(P) = 1$], while the face-tiles P are always closed [$p(P) = 0$]. With such construction, no ambiguities arise, and at most three tiles can meet at a point. A possible construction for the bond percolation (with octagons being the bond-tiles) is illustrated in Figure 2.

Following the original motivation of Broadbent and Hammersley, we represent holes in some random porous medium by open tiles. A liquid is assumed to move freely through the holes, thus it can flow between two regions if there is a path of open tiles between them. More precisely, we say that there is an *open (closed) crossing* between two sets K_1 and K_2 inside the set U , if for the percolation configuration restricted to U , there is an open (closed) cluster intersecting both K_1 and K_2 . We will be mostly interested in crossings of quads (= topological quadrilaterals).

Note the following duality property: *there is an open crossing between two opposite sides of a quad if and only if there is no dual (i.e., closed) crossing between two other sides*. Exchanging the notions of open and closed, we conclude that, on a given tiling, percolation models with probabilities $p(P)$ and $p^*(P) = 1 - p(P)$




FIG. 2. *Left: critical bond percolation on the square lattice, every bond is open (blue) or closed (yellow) with probabilities $1/2$ independently of others. Right: representation by critical percolation on the bathroom tiling, with direction of original bonds marked on the corresponding octagonal tiles. Rhombic tiles alternate in color, with blue (always open) corresponding to the sites of the original square lattice, while yellow (always closed)—to the faces. Octagonal tiles correspond to the bonds, and are open or closed independently, with probabilities $1/2$.*

are dual to each other. In particular, the model with $p(P) \equiv 1/2$ is self-dual, which can often be used to show that it is critical, like for the hexagonal lattice. The bathroom-tile model of the square lattice bond percolation in Figure 2 is dual to itself translated by one lattice step (so that always open rhombi are shifted to always closed ones), which also leads to its criticality, [20]. There are, however, dualities (e.g., arising from the star-triangle transformations) with more complicated relations between p and p^* , and so there are critical models with p away from $1/2$; see, for example, [5, 43].

Though we have in mind critical models, we do not use criticality in the proof, rather some crossing estimates, which are a part of the Russo–Seymour–Welsh theory. Those also hold (below any given scale) for the near-critical models, like those discussed in [29]. In a sense the estimates we need should hold whenever the crossing probabilities do not tend to a trivial limit; see Remark 1.3.

We work with a collection of percolation models μ_η (and sometimes more general random colorings) on tilings H_η , indexed by a set $\{\eta\}$ (e.g., square lattices of different mesh), and denote by μ_η also the corresponding measure on percolation configurations (= random coloring of tiles). By $|\eta| \in \mathbb{R}_+$, we will denote a “scale parameter” of the model μ_η . We assume that all tiles have diameter at most $|\eta|$ (for a percolation on random tiling, our proof would work under the relaxed assumption that for any positive r , the probability to find in a bounded region a tile of diameter bigger than r tends to zero as $|\eta| \rightarrow 0$). This parameter also appears in the Russo–Seymour–Welsh-type estimates below, which have to hold on scales larger than $|\eta|$. Note that in most models $|\eta|$ can be taken to be the supremum of tile diameters (or the lattice mesh for the percolation on a lattice).

We will be interested in scaling limits as the scale parameter tends to zero, looking at sequences of percolation models with $|\eta| \rightarrow 0$, which do not degenerate in the limit. We will show in Remark 1.17, that in our setting working with sequences and nets yields the same results. As discussed above, all critical percolation models are conjectured to have a universal and conformally invariant scaling limit. Moreover, its SLE description allows to calculate exactly many scaling exponents and dimensions and we will use this conjectured picture in informal discussions below. However, to make our paper applicable to a wider class of models, we will work under much weaker assumptions.

The famous Russo–Seymour–Welsh theory provides universal bounds for crossing probabilities of the same shape, superimposed over lattices with different mesh. Originally established for the critical bond percolation on the square lattice, it has been since generalized to many critical (and near-critical, when the percolation probabilities tend to their critical values at appropriate speeds as the mesh tends to zero) models with several different approaches; see [4, 16, 21, 42]. We will use one of the recurrent techniques: conditioning on the lowest crossing. If a quad is crossed horizontally, one can select the lowest possible crossing (the closest to the bottom side), which will then depend only on the configuration below. If we condition on the lowest crossing, the configuration above it is unbiased, allowing for

easy estimates of events there. Besides using this neat trick, we will assume certain bounds on the crossing probabilities for the annuli.

A prominent role is played by the so-called k -arm crossing events in annuli, since they control dimensions of several important sets. In particular, for a given percolation model μ_η we denote by $P_\eta^1(z, r, R)$ the probability of a *one-arm* event, that *there is an open crossing connecting two opposite circles of the annulus* $A(z, r, R)$. Then this is roughly the probability that the disc $B(z, r)$ intersects an open cluster of size $\approx R$, which is expected for $|\eta| < r$ to have a power law like

$$P_\eta^1(z, r, R) = (r/R)^{a_1+o(1)} \quad \text{as } r/R \rightarrow 0,$$

morally meaning (modulo some correlation estimates) that percolation clusters in the scaling limit have dimension $2 - a_1$.

Similarly, denote by $P_\eta^4(z, r, R)$ the probability of a *four-arm* event, that *there are alternating open–closed–open–closed crossings connecting two opposite circles of the annulus* $A(z, r, R)$. Then this is roughly the probability that changing the percolation configuration on the disc $B(z, r)$ changes the crossing events on the scale $\approx R$. Indeed, if all the tiles in $B(z, r)$ are made open, that connects two open arms; while making them closed connects two closed arms, destroying the open connection. For $|\eta| < r$ a power law

$$P_\eta^4(z, r, R) = (r/R)^{a_4+o(1)} \quad \text{as } r/R \rightarrow 0$$

is expected, roughly meaning that the *pivotal* tiles (such that altering their state changes crossing events on large scale) have dimension $2 - a_4$ in the scaling limit.

Whenever the RSW theory applies, the probabilities for k arm events have power law bounds from above and below as $r \rightarrow 0$. In fact, the probabilities are conjectured to satisfy universal (independent of a particular percolation model) power laws with rational powers, predicted by the Conformal Field Theory; see discussion in [36]. Those predictions have been so far proved for the critical site percolation on triangular lattice only [26, 36], giving $a_1 = 5/48$ and $a_4 = 5/4$.

Other k -arm probabilities are also of interest, but estimating just the mentioned two allows to apply our methods. Roughly speaking, we need to know that the percolation clusters have dimension smaller than 2 and that pivotal points have dimension smaller than 1, or, in other words, that $a_1 > 0$ and $a_4 > 1$. Note that we do not actually need the power laws, but rather weaker versions of the corresponding upper estimates.

Namely, we need the following two estimates to hold *uniformly* for percolation models under consideration.

ASSUMPTIONS 1.1. There exist positive functions $\Delta_1(r, R)$ and $\Delta_4(r, R)$, such that

$$\lim_{r \rightarrow 0} \Delta_j(r, R) = 0 \quad \text{for any fixed } R < R_0,$$

and the following estimates hold whenever $z \in \mathbb{C}$ and $0 < |\eta| < r < R < R_0$. The probability of one open arm event and the probability of a similar event for one

closed arm satisfy

$$(1.2) \quad P_{\eta}^1(z, r, R) \leq \Delta_1(r, R),$$

while the probability of a four arm event satisfies

$$(1.3) \quad P_{\eta}^4(z, r, R) \leq \left(\frac{r}{R}\right) \cdot \Delta_4(r, R).$$

For $r \geq R$, when the annulus is empty, we set $\Delta_j(r, R) := 1$, so that the function is defined for all positive arguments. Without loss of generality, we can assume that functions $\Delta_j(r, R)$ are increasing in r and decreasing in R .

REMARK 1.2 (A stronger RSW estimate). For most percolation models where the RSW techniques work, their application would prove a stronger estimate, with $\Delta_1(r, R)$ replaced in (1.2) by $C(r/R)^{a_1}$ with $a_1 > 0$. However, we can imagine situations where only our weaker version can a priori be established.

REMARK 1.3 (RSW estimates and scaling limits). For most percolation models where the RSW techniques work, the estimate (1.2) would be equivalent to saying that crossing probability for any given quad (or for all quads) is uniformly bounded away from 0 and 1. So morally our noise characterization would apply whenever one can speak of nontrivial subsequential scaling limits of crossing probabilities.

The first assumption (1.2) is used several times throughout the paper. In a stronger form, it was one of the original Russo–Seymour–Welsh results, and has since been shown to hold for a wide range of percolation models; see [4, 5, 16, 21]. For site percolation on the triangular lattice, one can take $\Delta_1(r, R) \approx (r/R)^{5/48}$; see [26].

The second assumption (1.3) is used only once, namely in the proof of Theorem 1.5. It is known [in a sharp form with $\Delta_4(r, R) \approx (r/R)^{1/4}$] for site percolation on the triangular lattice; see [36]. Christophe Garban kindly allowed us to include as Appendix B his proof for bond percolation on the square lattice, which is partly inspired by the noise-sensitivity arguments in [3]. It seems possible to extend the proof to a wider range of percolation models (by changing the values of the C_j random variables so they become centered, rather than using symmetry arguments).

QUESTION 1.4 (Estimating pivots). For which models can (1.3) be proved? Can it be deduced from (1.2)? Is there a geometric argument?

Summing up the discussion above, we study collections of percolation models satisfying Assumptions 1.1. In particular, *our results apply to critical and near-critical site percolation on the triangular lattice and bond percolation on the square lattice.*

1.3. *Definition of the scaling limits.* Every discrete percolation configuration (or a coloring of a tiling) is clearly described by the collection of quads (= topological quadrilaterals) which it crosses (and for a fixed lattice mesh in a bounded domain a finite collection of quads would suffice). We introduce a setup, which allows to work independently of the lattice mesh and pass to the scaling limit, but the idea remains the same: a percolation configuration is encoded by a collection of quads—those which are crossed by clusters. Our setup is inspired by the Dedekind’s sections and uses the metric and ordering on the space of quads. The ordering comes from the obvious observation that crossing of a quad automatically contains crossings of “shorter” sub-quads, and so possible percolation configurations (= quad collections) should satisfy certain monotonicity properties.

We will work in a domain $D \subset \mathbb{C} = \mathbb{R}^2$. A *quad* in D is a topological quadrilateral, that is, a homeomorphism $Q : [0, 1]^2 \rightarrow Q([0, 1]^2) \subset D$. We introduce some redundancy when considering quads as parameterized quadrilaterals, but we avoid certain technicalities. The space of all such quads will be denoted by $\mathcal{Q} = \mathcal{Q}_D$. It is a metric space under the uniform metric

$$d(Q_1, Q_2) = \sup_{z \in [0, 1]^2} |Q_1(z) - Q_2(z)|.$$

A *crossing* of Q is a connected compact subset of $[Q] := Q([0, 1]^2)$ that intersects both opposite sides $\partial_0 Q := Q(\{0\} \times [0, 1])$ and $\partial_2 Q := Q(\{1\} \times [0, 1])$. We also denote the remaining two sides by $\partial_1 Q := Q([0, 1] \times \{0\})$ and $\partial_3 Q := Q([0, 1] \times \{1\})$, see Figure 3. The whole boundary of Q we denote by ∂Q . In the discrete setting there is no difference between connected and path-connected crossings, but one can imagine lattice models where the probabilities of the two would be different in the scaling limit. However, for the models under consideration, every crossing in the scaling limit almost surely can be realized by a Hölder continuous curve [1], so the two notions are essentially the same.

In addition to the structure of \mathcal{Q}_D as a metric space, we will also use the following partial order on \mathcal{Q}_D . If $Q_1, Q_2 \in \mathcal{Q}_D$, we write $Q_1 \leq Q_2$ if every crossing of Q_2 contains a crossing of Q_1 . The simplest example is seen in Figure 4. Also, we write $Q_1 < Q_2$ if there are open (in the uniform metric) neighborhoods U_1 of Q_1 and U_2 of Q_2 in \mathcal{Q}_D , such that for every $Q \in U_1$ and $Q' \in U_2$ we have




FIG. 3. A quad with a crossing.

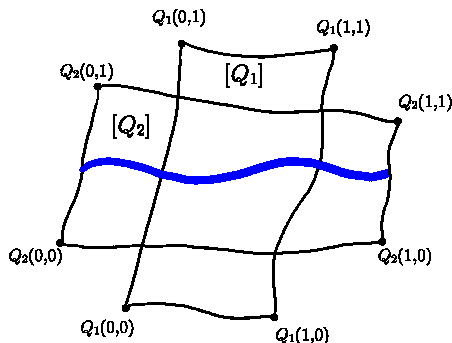


FIG. 4. Any crossing of Q_2 contains a crossing of Q_1 , so we write $Q_1 \leq Q_2$.

$Q \leq Q'$. Thus, the set of pairs $(Q, Q') \in \mathcal{Q}^2$ satisfying $Q < Q'$ is the interior of $\{(Q, Q') \in \mathcal{Q}^2 : Q \leq Q'\}$ in the product topology.

A subset $S \subset \mathcal{Q}_D$ is a *lower set* 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 lower subsets of \mathcal{Q}_D will be denoted \mathcal{H}_D . Any discrete percolation configuration ω in \mathbb{C} is naturally associated with an element S_ω of \mathcal{H}_D : the set of all quads for which ω contains a crossing formed by an open cluster. Thus, a percolation model on the tiling H_η induces a probability measure μ_η on $\mathcal{H}_\mathbb{C}$ (and more generally on \mathcal{H}_D for any domain D contained in its domain of definition).

A topology is defined on \mathcal{H}_D by specifying a subbase. If $U \subset \mathcal{Q}_D$ is topologically open and $Q \in \mathcal{Q}_D$, let

$$V_U := \{S \in \mathcal{H}_D : S \cap U \neq \emptyset\}$$

and

$$V^Q := \{S \in \mathcal{H}_D : Q \notin S\}.$$

If we regard S as a collection of crossed quads, then

$$V_U := \{S : \text{some quads from } U \text{ are crossed}\}$$

and

$$V^Q := \{S : \text{quad } Q \text{ is not crossed}\}.$$

Our topology of choice on \mathcal{H}_D will be the minimal topology \mathcal{T}_D which contains every such V_U and V^Q . As we will see, $(\mathcal{H}_D, \mathcal{T}_D)$ is a metrizable compact Hausdorff space. The percolation scaling limit configuration will be defined as an element of $\mathcal{H}_\mathbb{C}$, and the scaling limit measure is a Borel probability measure on $\mathcal{H}_\mathbb{C}$.

Then the event of a percolation configuration crossing a quad Q corresponds to

$$\square_Q := \neg V^Q = \{S \in \mathcal{H} : Q \in S\} \subset \mathcal{H}.$$

We will prove that if $Q^0 \subset Q$ is dense in Q , then the events \square_Q for $Q \in \mathcal{Q}^0$, generate the Borel σ -field of \mathcal{H} .

1.4. *Statement of main results.* Construction of the (subsequential) scaling limits is accomplished for all random colorings of trivalent tilings without additional assumptions. Other results of this paper apply to all *percolation models on trivalent tilings, satisfying Assumptions 1.1*, that is, the RSW estimate (1.2) and the pivotal estimate (1.3).

As discussed above, we fix a collection of such percolation models indexed by a set $\{\eta\}$, and denote by μ_η the (locally) discrete probability measure on percolation configurations for the model corresponding to parameter η . By $|\eta| \in \mathbb{R}_+$, we denote the scale parameter of the model (which enters in the RSW estimates). Later we will show that μ_η has subsequential scaling limits as $|\eta| \rightarrow 0$, and will work with one of those, denoted by μ_0 .

First, we present a simplified discrete version of the gluing theorem. Informally, the statement is that if you fix a quad Q_0 and a finite length path α cutting it into two pieces, then the discrete percolation configuration outside a small neighborhood of α reliably predicts the existence or nonexistence of a crossing of Q_0 with probability close to one, uniformly in the mesh size.

Following is the precise version, estimating the random variable which is the conditional probability of observing an open crossing given the configuration s -away from α . As a regularity assumption, we need a constant $C(\alpha)$ such that for every $s > 0$ the set α can be covered by at most $C(\alpha)/s$ balls of radius s . This is the case if and only if α has finite one-dimensional upper Minkowski content $\mathcal{M}^{*1}(\alpha) = \limsup_{s \rightarrow 0^+} (\text{area}\{z : \text{dist}(z, \alpha) < s\})$.

THEOREM 1.5 (Discrete gluing). *Consider a collection of percolation models satisfying Assumptions 1.1. Let $Q_0 \in \mathcal{Q}$ be some quad, and let $\alpha \subset [Q_0]$ be a finite union of finite length paths, or more generally a set of finite one-dimensional upper Minkowski content. Let \square_{Q_0} denote the event that Q_0 is crossed by ω , and for each $s > 0$ let \mathcal{F}_s denote the σ -field generated by the restriction of ω to the complement of the s -neighborhood of α . Then for every $\epsilon > 0$,*

$$\lim_{s \searrow 0} \sup_{|\eta| \in (0, s)} \mu_\eta(\epsilon < \mu_\eta(\square_{Q_0} \mid \mathcal{F}_s) < 1 - \epsilon) = 0.$$

REMARK 1.6 (Gluing is nonconstructive). Theorem 1.5 essentially states that event \square_{Q_0} for any given η can be reconstructed with high accuracy by sampling crossing events away from α . This does not lead to the noise characterization of the scaling limit, since a priori the number of events we need to sample can grow out of control as $|\eta| \rightarrow 0$. We address this by providing a uniform bound on the number of events to be sampled in Proposition 4.1. However, we do not give a recipe, and in general the gluing procedure may still depend on η . We deal with this possibility by assuming that there is a (subsequential) scaling limit, then, since the number of possible outcomes of sampling a finite number of crossing events is finite, we can choose a sub-subsequence along which the same procedure is used. The downside is that, in the current form, our results alone cannot be used to show the uniqueness of the percolation scaling limit.

The previous remark leads to the following:

QUESTION 1.7 (Constructive gluing). Can one show that there is an independent of η and preferably constructive procedure to predict \square_{Q_0} by sampling crossing events away from α ?

A positive answer would provide an elegant proof of the uniqueness of the scaling limit (provided Cardy's formula is known), make the universality phenomenon more accessible, and lay foundation for the renormalization theory approach to percolation, similar to one discussed by Langlands; see [24] and the references therein. A particularly optimistic scenario is suggested in Conjecture 1.11.

Theorem 1.5 is a toy version of our main gluing results, and the logic behind it was folklore for a long time (but strangely enough appeared few times in print, as was pointed out in [37]). The proof proceeds by consequently resampling radius $2s$ balls covering the s -neighborhood of α and using the pivotal estimate (1.3). For this to work, the bounds have to sum up to something small, and so with better pivotal estimates we can relax assumptions on α (and the number of balls needed to cover it).

REMARK 1.8 (Regularity assumptions on α). When applying the pivotal estimate (1.3), we cannot relax the regularity assumption that α has finite one-dimensional upper Minkowski content. However, as discussed above, most percolation models are expected to satisfy stronger estimates, and then the regularity assumption can be considerably weakened (though it cannot be dropped). The set of pivotal points conjecturally has Hausdorff dimension $3/4$ in the scaling limit (proved on the triangular lattice in [36]), so if the Hausdorff dimension of the set α is greater than $5/4$, the two should intersect with positive probability (needed correlation estimates are expected to hold). In such a situation, there is little hope to reliably determine the crossing without knowledge of α . To be more rigorous, instead of pivots one has to work with the percolation spectrum, but it has the same dimension (see [13] for the definition and discussion), so α has to be of dimension at most $5/4$. On the other hand, our proof works for sets α of dimension smaller than 2 minus dimension of the pivotal points, provided also that $\alpha \cap \partial Q_0$ is finite. So $5/4$ seems to be the correct-dimensional threshold, and this can be established for the critical site percolation on the triangular lattice.

REMARK 1.9 (Crossings with finite intersection number). Suppose that a smooth curve α cuts the quad Q_0 into two quads Q_+ and Q_- . The following informal discussion concerns the critical percolation scaling limit, assuming its existence and conformal invariance (but it can be easily made more precise by working with the limits of discrete configurations and events). First observe that open clusters inside Q_{\pm} intersect the smooth boundary α on a set of Hausdorff dimension $2/3$ (and we have a control for the correlation of points inside it). Being

independent, two large open clusters on opposite sides of α have a positive probability of touching. But if they touch at some point $x \in \alpha$, we cannot have a large closed cluster passing through x and separating them. Indeed, that would mean a four-arm event occurring at the *pivotal* point x on the smooth curve α , and since $a_4 = 5/4$, the set of pivotal points has Hausdorff dimension $3/4$ and so almost surely does not intersect the smooth α . Thus, the two open clusters in Q_{\pm} which touch on α are in fact parts of the same large open cluster in Q_0 . We conclude, that given a smooth curve α , with positive probability there is a crossing of Q_0 which intersects α only *once*.

The last remark says that for a smooth curve α , if we condition on the quad being crossed, the probability of finding a crossing intersecting α only finitely many times is positive. We can ask whether this conditional probability is one, leading to the following:

CONJECTURE 1.10 (Finite intersection property). *Let Q be a quad and α a smooth curve. Then in the scaling limit, given existence of a crossing, there is almost surely a crossing intersecting α only a finite number of times.*

Alternatively, one can formulate a discrete version, which is a priori stronger (for an equivalent one we would have to count small scale packets of intersections):

CONJECTURE 1.11 (Tightness of intersections). *Let Q be a quad and α a smooth curve. For a discrete percolation measure μ_{η} , let the random variable $N_{\eta} = N_{\eta}(Q, \alpha)$ be the minimal possible number of intersections with α for a crossing of Q (and 0 if Q is not crossed). Then N_{η} is tight as $|\eta| \rightarrow 0$.*

Proving either of these conjectures would not only simplify our proof of the main theorem, but also provide an easy and constructive “gluing procedure”: we take all “excursions” (= crossings in the complements of α) and check whether a crossing of Q can be assembled from finitely many of them (as discussed in Remark 1.9, the finite number of excursions will necessarily connect into one crossing). Such a procedure would give a very short proof of the uniqueness of the full percolation scaling limit, provide an elegant setup for the renormalization theory of percolation crossings, and show that the “finite model” discussed by Langlands and Lafortune in [25] indeed approximates the critical percolation. Note that if conjectures above fail, the constructive procedure for gluing crossings will have to deal with countably many excursions intersecting α on a Cantor set C . Then the criterion for possibility of gluing the excursions into one crossing is likely to be in terms of some capacity (or dimensional) characteristics of C .

The critical percolation scaling limit is conjectured to be conformally invariant, which was established for site percolation on the triangular lattice [34, 35]. In Remark 1.9, we mentioned that, for a given quad Q_0 , a smooth curve α has the

property that two large clusters inside two components of $[Q_0] \setminus \alpha$ touch each other (on α) with positive probability. One can ask, whether an arbitrary curve α would always enjoy this property, and the answer is negative. Moreover, under conformal invariance assumption, this property (equivalent to positive probability of finding a crossing intersecting α only finitely many times) can be reformulated in terms of multifractal properties of harmonic measure on two sides of α . Below we present an informal argument to this effect, note also similar discussions in the SLE context in [15].

To simplify the matters, we discuss a related (and probably equivalent) property that the expected number of such touching points for two large clusters inside two components of $[Q_0] \setminus \alpha$ is positive (or even infinite). Let $f(a_+, a_-)$ be the two-sided dimension spectrum of harmonic measure on α ; see [2]. Roughly speaking, $f(a_+, a_-)$ is the dimension of the set of $z \in \alpha$ where harmonic measures ω_{\pm} on two sides of α have power laws a_{\pm} :

$$\omega_+(B(z, r)) \approx r^{a_+}, \quad \omega_-(B(z, r)) \approx r^{a_-}.$$

Consider a set $E \subset \alpha$ such that harmonic measures ω_{\pm} have power laws a_{\pm} on E . Cover most of E by a collection of balls B_j of small radii r_j . By conformal invariance, the probability of a ball B_j to be touched by a large open cluster on the “+ side” of α is governed by its harmonic measure and Cardy’s formula [35] that gives the exponent $1/3$:

$$P_+(B_j) \approx \omega_+(B_j)^{1/3} \approx r_j^{a_+/3}.$$

Combining with the similar estimate on the “– side,” we write an estimate for the expected number of balls B_j touched by large clusters on both sides:

$$\sum_j P_+(B_j)P_-(B_j) \approx \sum_j r_j^{(a_++a_-)/3}.$$

Here the right-hand side is small (large) when $(a_+ + a_-)/3$ is bigger (smaller) than $\dim(E)$, since $\sum_j r_j^{\dim(E)} \approx 1$. We conclude that the expected number of points touched by large clusters on both sides is positive if and only if we can find E such that $(a_+ + a_-)/3 < \dim(E)$.

REMARK 1.12 (Sharpness of the finite intersection property). It seems that positive probability of having a crossing intersecting α a finite number of times is roughly equivalent to the existence of positive exponents a_+ and a_- such that the two-sided multifractal spectrum of harmonic measure on α satisfies $3f(a_+, a_-) > a_+ + a_-$; see [2] for a discussion of such spectra. The latter property seems to fail even for some nonsmooth α of dimension close to one, so we do not expect the factorization property to be equivalent to the finite intersection property, but rather to be a weaker one.

Before proving the main theorem, we discuss a new framework for the percolation scaling limit and show that collection of discrete random colorings is precompact, so in this framework the subsequential scaling limits exist for any random model.

Recall that we work with \mathcal{H}_D , which is the collection of all closed lower (i.e., monotonicity of the crossing events is obeyed) subsets of \mathcal{Q}_D , which one should think of as sets of quads crossed. The topology \mathcal{T}_D is the minimal one containing all sets V_U (of S such that some quads in U are crossed) and V^Q (of S such that Q is not crossed).

We now list some important properties of \mathcal{H}_D .

THEOREM 1.13 (The space of percolation configurations). *Let $D \subset \hat{\mathbb{C}}$ be topologically open and nonempty.*

- (1) $(\mathcal{H}_D, \mathcal{T}_D)$ is a compact metrizable Hausdorff space.
- (2) Let A be a dense subset of \mathcal{Q}_D . If $S_1, S_2 \in \mathcal{H}_D$ satisfy $S_1 \cap A = S_2 \cap A$, then $S_1 = S_2$. Moreover, the σ -field generated by V^Q , $Q \in A$ is the Borel σ -field of $(\mathcal{H}_D, \mathcal{T}_D)$.
- (3) If S_1, S_2, \dots is a sequence in \mathcal{H}_D and $S \in \mathcal{H}_D$, then $S_j \rightarrow S$ in \mathcal{T}_D is equivalent to $S = \limsup_j S_j = \liminf_j S_j$.
- (4) Let $D' \subset D$ be a subdomain, and for $S \in \mathcal{H}_D$ let $S' := S \cap \mathcal{Q}_{D'}$. Then $S \mapsto S'$ is a continuous map from \mathcal{H}_D to $\mathcal{H}_{D'}$.

In the above, $\liminf_j S_j$ is the set of $Q \in \mathcal{Q}_D$ such that every neighborhood of Q intersects all but finitely many of the sets S_j , and $\limsup_j S_j$ is the set of $Q \in \mathcal{Q}_D$ such that every neighborhood of Q intersects infinitely many of the sets S_j .

REMARK 1.14 (Applicability of Theorem 1.13). The result concerns the topological properties of the space of percolation configurations (colorings of tilings) and involves no probability measures. Thus, it can be applied to *any random coloring* of a trivalent tiling (into open and closed tiles), yielding results for dependent models, like the Fortuin–Kasteleyn percolation.

It follows from the theorem above that the space of continuous functions on $(\mathcal{H}_D, \mathcal{T}_D)$ is separable, and hence the unit ball in its dual space $M(\mathcal{H}_D)$ of Borel measures is compact (by the Banach–Alaoglu theorem, Theorem 3.15 in [31]), metrizable (by Theorem 3.16 in [31]) and obviously Hausdorff in the weak-* topology. The space of probability measures being a closed convex subset of the unit ball, we arrive at the following.

COROLLARY 1.15 (The space of percolation measures). *The space $\text{Prob}(\mathcal{H}_D)$ of Borel probability measures on $(\mathcal{H}_D, \mathcal{T}_D)$ with weak-* topology is a compact metrizable Hausdorff space.*

Discrete percolation measures μ_η clearly belong to $\text{Prob}(\mathcal{H}_D)$ and hence the existence of percolation subsequential scaling limits (or, more generally, for any random colorings) is an immediate corollary:

COROLLARY 1.16 (Precompactness). *The collection of all discrete measures $\{\mu_\eta\}$ (corresponding to random colorings of trivalent tilings) is precompact in the topology of weak-* convergence on $\text{Prob}(\mathcal{H}_D)$.*

REMARK 1.17 (Nets vs. sequences). We can make $\{\eta\}$ a directed set by writing $\eta_1 \leq \eta_2$ whenever $|\eta_1| \geq |\eta_2|$. Then the scaling limit $\lim_{|\eta| \rightarrow 0} \mu_\eta$ is just the limit of the net μ_η . Since the target space $\text{Prob}(\mathcal{H}_D)$ is metrizable and hence first-countable, it is enough to work with the notion of convergence along sequences, rather than nets. In particular, if $\lim_{j \rightarrow \infty} \mu_{\eta_j} = \mu_0$ for every sequence η_j with $\lim_{j \rightarrow \infty} |\eta_j| = 0$, then $\lim_{|\eta| \rightarrow 0} \mu_\eta = \mu_0$.

REMARK 1.18. Our construction keeps information about all macroscopic percolation events, and so gives the *full percolation scaling limit* at or near criticality. For the critical site percolation on triangular lattice, Garban, Pete and Schramm explain in Section 2 of [14] that the scaling limit is unique, appealing to the work [7] of Camia and Newman. Basically, we know the probabilities of individual crossing events, and it allows to reconstruct the whole picture, albeit in a difficult way. It would be interesting to establish the uniqueness using an appropriate modification of Proposition 4.1.

So that our theorems apply even to models where the uniqueness of the scaling limit has not been proved yet, we work with one of the subsequential limits μ_0 —a Borel probability measure on \mathcal{H}_C , provided by Corollary 1.16. By Theorem 1.13, the Borel σ -field of \mathcal{H}_D is $\sigma(\boxplus_Q : Q \in \mathcal{Q}_D)$, that is, generated by the crossing events inside D . Let $\mathcal{F}_D := \sigma(\boxplus_Q : Q \in \mathcal{Q}_D)$ also denote the corresponding sub-field of the Borel σ -field of \mathcal{H}_C .

THEOREM 1.19 (Factorization). *Consider a collection of percolation models, satisfying Assumptions 1.1. Let D be a domain, and let $\alpha \subset \mathbb{C}$ be a finite union of finite length paths with finitely many double points. Denote the components of $D \setminus \alpha$ by D_j . Then for any subsequential scaling limit*

$$(1.4) \quad \mathcal{F}_D = \mathcal{F}_{D \setminus \alpha} = \bigvee_j \mathcal{F}_{D_j},$$

up to sets of measure zero.

Again, we note that this does not hold when α is a sufficiently wild path.

Fix a subsequential scaling limit of the critical percolation for the bond model on the square lattice or for the site model on the triangular lattice. By the results

above for a domain D , it is described by a probability space $(\mathcal{H}_D, \mathcal{F}_D, \mu_D)$, which is invariant under translations of the plane (in fact, under all conformal transformations, but it is so far established for the triangular lattice only), depends continuously on D and satisfies (1.4). Then in the language of Tsirelson (see [38], Definition 3d1) it is a *noise* or a *homogeneous continuous product of probability spaces* with a Boolean base given by an appropriate algebra of piecewise-smooth planar domains (e.g., generated by rectangles).

There is a well developed and beautiful theory of noises, and we refer the reader to two expositions [38, 39] by Tsirelson. In particular, every noise can be decomposed into a “classical,” or “stable” part and a “nonclassical,” or “sensitive” part. White noise is purely classical, and purely nonclassical noises are called “black” by Tsirelson. The black noises are difficult to construct, with the first example given by Tsirelson and Vershik in [41]. As pointed out by Tsirelson in [40] and in Remark 8a2 of [39], percolation noise has to be nonclassical, and so our paper provides the first genuinely two-dimensional example of a black noise.

COROLLARY 1.20 (Percolation is a noise). *Thus, we conclude that, in the language of Tsirelson [38], any subsequential scaling limit as above is a noise with a Boolean base given by an appropriate algebra of piecewise-smooth planar domains (e.g., generated by rectangles). Therefore, it has to be a black noise, as explained in [39], Remark 8a2.*

By the conjectured universality, all the critical percolation models should have the same scaling limit and so correspond to the same black noise. The situation with near-critical noises is less clear, so the following question was proposed by Christophe Garban (the answer is expected to be negative):

QUESTION 1.21. Are the noises arising from the critical and near-critical models isomorphic, or do we get a family of different noises?

The notion of isomorphism for noises is discussed by Tsirelson in Section 4a of [38]. Roughly speaking, we ask *whether one can get a near-critical noise from the critical one by a local deterministic procedure?* Note that the current construction of the near-critical percolation scaling limit would use the critical percolation along with some extra randomness, as in [6, 12].

Our construction of the scaling limit applies to general random colorings, and much of what we do afterwards mostly uses “arms estimates,” which are known for a wider class of models, than just percolation. Therefore, it is logical to ask:

QUESTION 1.22. To what extent our results apply to the models with dependence? One has to be more careful in formulating the results, since now the configuration inside a domain depends on the boundary conditions, but this can be addressed by working with a full-plane model and stripe domains.

2. The discrete gluing theorem. In this section, we prove Theorem 1.5.

Applying several times in succession Lemma A.1, we can find a smooth (i.e., given by a diffeomorphism) quad Q , such that the crossing events for Q_0 and Q are arbitrarily close. By the same lemma, as $\varepsilon \rightarrow 0$, the crossing event for Q is well approximated by the crossing event of its perturbation $Q^\varepsilon := Q([\varepsilon, 1 - \varepsilon]^2)$. On the other hand, α has finite length and so for almost every ε the intersection $\alpha \cap \partial Q^\varepsilon$ is finite. This can be deduced, for example, from the area theorem in the geometric measure theory [10, 27], which implies that an orthogonal projection of a rectifiable curve on a given line covers almost every point at most finitely many times.

Thus, approximating if necessary, we can assume that α intersects ∂Q_0 at finitely many points $x_i, i = 1, \dots, k$. Fix $\epsilon_0 > 0$, then by the RSW estimate (1.2), there is an $r > 0$ such that, if $0 < |\eta| < r$, then the probability under μ_η that there is a crossing of Q_0 which intersects any of the disks $B(x_i, r)$ is less than ϵ_0 . Fix such an r , remove from the curve α the disks, the rest will be away from the boundary, simplifying future estimates:

$$\alpha' := \alpha \setminus \bigcup_i B(x_i, r)$$

and

$$d := \inf\{|x - y| : x \in \alpha', y \in \partial Q_0\} > 0.$$

Fix $s \in (0, d/4)$. By the regularity assumption, we can choose $n \leq C(\alpha)/s$ balls of radius s , entirely covering α . Denote their centers by w_1, w_2, \dots, w_n , and set $B_j := B(w_j, 2s)$; see Figure 5.

Let M denote the union of the set of tiles of H_η that intersect the s -neighborhood of α' minus $\bigcup_i B(x_i, r)$. Assume $|\eta| < s$, so that M would be contained in the $2s$ -neighborhood of α' . Let ω and ω' be samples from μ_η that agree on the com-

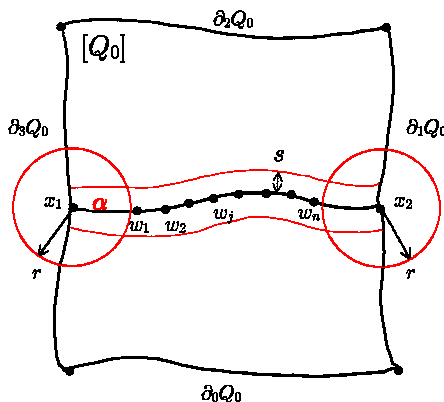


FIG. 5. Setup for the proof of the discrete gluing theorem in the simplest case.

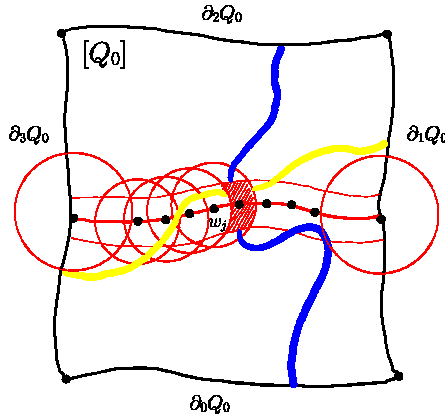


FIG. 6. If resampling a part of the ball B_j changes the crossing event, then there are two open crossings from B_j to $\partial_0 Q_0$ and $\partial_2 Q_0$, and two dual closed crossings from B_j to $\partial_0 Q_0$ and $\partial_2 Q_0$, a four arm event.

plement of M and are independent in M . For $j = 0, 1, \dots, n$, let ω_j be the configuration that agrees with ω' on the tiles of H_η lying inside $B_1 \cup B_2 \cup \dots \cup B_j$ and agrees with ω elsewhere. Then each ω_j is a fair sample from μ_η . Also observe that $\omega_0 = \omega$ and ω_n differs from ω' only inside $\cup_i B(x_i, r)$.

Our goal is to estimate

$$\mathbf{P}[\omega \in \square_{Q_0}, \omega' \notin \square_{Q_0}],$$

which we will do by successively comparing ω_{j-1} to ω_j . Fix $j \in \{1, 2, \dots, n\}$. In order for $\{\omega_{j-1} \in \square_{Q_0}\} \cap \{\omega_j \notin \square_{Q_0}\}$ to hold, there must be a crossing in ω_{j-1} from $\partial_0 Q_0$ to $\partial_2 Q_0$ that goes through B_j and there must be a dual closed crossing in ω_j from $\partial_1 Q_0$ to $\partial_3 Q_0$ that goes through B_j ; see Figure 6. Thus, in ω_j we have the four arm event from ∂B_j to ∂Q_0 . Since the radius of B_j is $2s$ and the distance from B_j to ∂Q_0 is at least $d - 2s > d/2$, by assumption (1.3) we have

$$\mathbf{P}[\omega_{j-1} \in \square_{Q_0}, \omega_j \notin \square_{Q_0}] \leq P_\eta^4(z, 2s, d/2) < \left(\frac{2s}{d/2}\right) \cdot \Delta_4(2s, d/2).$$

Summing over j , we conclude that

$$\begin{aligned} \mathbf{P}[\omega_0 \in \square_{Q_0}, \omega_n \notin \square_{Q_0}] &\leq \sum_{j=1}^n \mathbf{P}[\omega_{j-1} \in \square_{Q_0}, \omega_j \notin \square_{Q_0}] \\ &\leq n \cdot \left(\frac{2s}{d/2}\right) \cdot \Delta_4(2s, d/2) \\ &\leq \frac{C(\alpha)}{s} \cdot \frac{4s}{d} \cdot \Delta_4(2s, d/2) \\ &\leq \frac{4C(\alpha)}{d} \cdot \Delta_4(2s, d/2). \end{aligned}$$

By Assumption 1.1 the right-hand side tends to 0 with s , and so we can make sure that $\mathbf{P}[\omega_0 \in \square_{Q_0}, \omega_n \notin \square_{Q_0}] < \epsilon_0$ by choosing s sufficiently small. This is the only place in the paper, where we use the pivotal estimate (1.3).

Now, recall that $\omega_0 = \omega$, and ω_n differs from ω' only near endpoints of curves in α , that is, inside $\bigcup_i B(x_i, r)$. By our choice of r , we have the estimate $\mathbf{P}[\omega_n \in \square_{Q_0}, \omega' \notin \square_{Q_0}] < \epsilon_0$. Summing the above, we get the bound

$$\mathbf{P}[\omega \in \square_{Q_0}, \omega' \notin \square_{Q_0}] < 2\epsilon_0.$$

Now note that by first conditioning on \mathcal{F}_s and then using the fact that ω and ω' are conditionally independent given \mathcal{F}_s , we obtain

$$2\epsilon_0 > \mathbf{P}[\omega \in \square_{Q_0}, \omega' \notin \square_{Q_0}] = \mathbf{E}[\mu_\eta(\square_{Q_0} \mid \mathcal{F}_s)(1 - \mu_\eta(\square_{Q_0} \mid \mathcal{F}_s))].$$

In particular, by Chebyshev inequality,

$$\mu_\eta(\epsilon < \mu_\eta(\square_{Q_0} \mid \mathcal{F}_s) < 1 - \epsilon) < 2\epsilon_0 / (\epsilon(1 - \epsilon)).$$

Since ϵ_0 was an arbitrary positive number, this completes the proof.

3. The space of lower sets. The main goal of this section is to provide an abstract setup for the percolation scaling limit and to prove Theorem 1.13. We discuss topology on the space \mathcal{H}_X of closed lower subsets of an abstract partially ordered topological space X . In the percolation case, X corresponds to the space \mathcal{Q}_D of quads, and the space $(\mathcal{H}_X, \mathcal{T}_\mathcal{H})$ to the space $(\mathcal{H}_D, \mathcal{T}_D)$ of percolation configurations in a domain $D \subset \mathbb{C}$.

With these conventions, Theorem 1.13 directly follows from Theorem 3.10 below. The latter applies to the space of percolation configurations (colorings of tilings), since \mathcal{Q}_D is clearly second-countable, (3.1) holds by definition and (3.2) holds since a quad Q_0 can be easily approximated by smaller quads [e.g., by quads Q^q with $q = (-\epsilon, \epsilon, 1 + \epsilon, 1 - \epsilon)$ from the proof of Lemma 5.1].

We now start the (mostly classical) abstract construction, inspired by the Dedekind’s sections. Let (X, τ) be a topological space, let \mathfrak{F}_X denote the collection of topologically closed subsets of X . For a topologically open $U \subset X$ and a compact $K \subset X$, let

$$W_U := \{F \in \mathfrak{F}_X : F \cap U \neq \emptyset\},$$

$$W^K := \{F \in \mathfrak{F}_X : F \cap K = \emptyset\}.$$

Let $\hat{\mathcal{T}}$ be the topology on \mathfrak{F}_X generated by all such sets W^K and W_U . It is a more convenient version of the Vietoris topology, called *Fell topology* or *topology of closed convergence*.

The following result is Lemma 1 in [11].

PROPOSITION 3.1. $(\mathfrak{F}_X, \hat{\mathcal{T}})$ is compact.

For completeness, we include a short proof, based on Alexander’s subbase theorem (see, e.g., Theorem 5.6 in [17]), which states that if \mathcal{B} is a subbase for the topology on a space X and every cover of X by elements of \mathcal{B} contains a finite subcover, then X is compact.

PROOF OF PROPOSITION 3.1. By Alexander’s subbase theorem, it suffices to show that every cover of \mathfrak{F}_X by sets of the form W_U (where $U \in \tau$) and W^K (where $K \subset X$ is compact) has a finite subcover. Let

$$\{W_U : U \in I\} \cup \{W^K : K \in J\} \supset \mathfrak{F}_X$$

be such a cover. Observe that $\bigcup_{U \in I} W_U = W_{\tilde{U}}$, where $\tilde{U} = \bigcup_{U \in I} U$. Let $F := X \setminus \tilde{U}$, and note that F is topologically closed. Since $F \in \mathfrak{F}_X \subset W_{\tilde{U}} \cup \bigcup_{K \in J} W^K$, we know that $F \in \bigcup_{K \in J} W^K$, namely, there is some $K_0 \in J$ with $F \in W^{K_0}$. In other words, $F \cap K_0 = \emptyset$; that is, $K_0 \subset \tilde{U}$. Since K_0 is compact and $\{U : U \in I\}$ covers K_0 , there is some finite $I' \subset I$ such that $\{U : U \in I'\}$ covers K_0 . Then $\{W_U : U \in I'\} \cup \{W^{K_0}\}$ is a finite cover of \mathfrak{F}_X . \square

Now suppose that “ $<$ ” denotes a partial order on X such that

$$(3.1) \quad \mathcal{R} := \{(x, y) \in X^2 : x < y\} \text{ is a topologically open subset of } X^2.$$

Define $\mathcal{R}^x := \{y \in X : y < x\}$ and $\mathcal{R}_x := \{y \in X : x < y\}$. Then for every $x \in X$ these sets are topologically open. A set $H \subset X$ is said to be *lower* if $x \in H$ implies $\mathcal{R}^x \subset H$. Let \mathcal{H}_X denote the collection of all topologically closed lower sets.

PROPOSITION 3.2. *Assuming (3.1), $\mathcal{H}_X \subset \mathfrak{F}_X$ is closed with respect to the topology \hat{T} ; that is, $\mathfrak{F}_X \setminus \mathcal{H}_X \in \hat{T}$.*

PROOF. Set $U_x := W^{\{x\}} \cap W_{\mathcal{R}_x}$ and

$$U := \bigcup_{x \in X} U_x = \{F : \exists x, y \text{ with } x < y, y \in F, x \notin F\}.$$

Then U is topologically open and $\mathcal{H}_X = \mathfrak{F}_X \setminus U$. \square

Let $\mathcal{T}_{\mathcal{H}}$ be the topology of \mathcal{H}_X as a subspace of \mathfrak{F}_X . It is the topology generated by the sets $V_U := W_U \cap \mathcal{H}_X$ and $V^K := W^K \cap \mathcal{H}_X$, for topologically open U and compact K subsets of X . As we will later see, in our setting it is sufficient to restrict ourselves to one-point compact sets, working with $V^{\{x\}}$, which we will often abbreviate V^x .

For simplicity, we will also abbreviate $\mathcal{H} := \mathcal{H}_X$ and $\mathcal{T} := \mathcal{T}_{\mathcal{H}}$. Propositions 3.1 and 3.2 imply that $(\mathcal{H}, \mathcal{T})$ is compact.

In general, \mathcal{H} does not have to be a Hausdorff space, even if X is Hausdorff. Indeed, the sufficient condition would be for X to be Hausdorff and compact, but for the application we have in mind, X is not compact. However, the following result gives another sufficient condition for \mathcal{H} to be Hausdorff.

PROPOSITION 3.3. *Suppose, in addition to (3.1), that*

$$(3.2) \quad \forall x \in X \quad x \in \overline{\mathcal{R}^x}.$$

Then $(\mathcal{H}, \mathcal{T})$ is a Hausdorff space.

PROOF. Let $H_1, H_2 \in \mathcal{H}$ be different, then one of the two differences $H_1 \setminus H_2$ and $H_2 \setminus H_1$ is nonempty, without loss of generality the first one. Take some $x \in H_1 \setminus H_2$. Since $X \setminus H_2$ is topologically open in X and contains $x \in \overline{\mathcal{R}^x}$, we can find some $y \in \mathcal{R}^x \cap (X \setminus H_2) = \mathcal{R}^x \setminus H_2$. Then $V_{\mathcal{R}_y}$ is disjoint from V^y . Moreover, $H_1 \in V_{\mathcal{R}_y}$ and $H_2 \in V^y$. So there are disjoint topologically open sets in \mathcal{T} containing H_1 and H_2 , respectively, and the space is Hausdorff. \square

REMARK 3.4. Let \mathcal{T}' be the topology on \mathfrak{F}_X generated by the sets W_U , $U \in \mathcal{T}_X$ and $W^{\{x\}}$, $x \in X$. Since this topology is coarser than $\hat{\mathcal{T}}$, Proposition 3.1 implies that $(\mathfrak{F}_X, \mathcal{T}')$ is compact. The proof of Proposition 3.2 shows that \mathcal{H} is a topologically closed subset of \mathfrak{F}_X , also with respect to \mathcal{T}' . Moreover, assuming (3.2), the proof of Proposition 3.3 shows that \mathcal{H} is Hausdorff with respect to \mathcal{T}' . It is well known (and not hard to verify; see Theorem 5.8 in [17]) that if a compact topology on a space is finer than a Hausdorff topology on the same space, then they must be equal. We conclude that the topology on \mathcal{H} induced by $\hat{\mathcal{T}}$ is the same as that induced by \mathcal{T}' whenever (3.2) and (3.1) hold.

LEMMA 3.5. *Suppose that (3.1) and (3.2) hold, and that X_0 is a dense subset of X . If $H_1, H_2 \in \mathcal{H}$ and $H_1 \neq H_2$, then $H_1 \cap X_0 \neq H_2 \cap X_0$.*

PROOF. Take $x \in H_1 \setminus H_2$, if the latter is nonempty. Then (3.2) implies that $\mathcal{R}^x \setminus H_2 \neq \emptyset$. Since X_0 is dense and $\mathcal{R}^x \setminus H_2$ is topologically open and nonempty, we have $X_0 \cap \mathcal{R}^x \setminus H_2 \neq \emptyset$. Because $\mathcal{R}^x \subset H_1$, this implies that $H_1 \cap X_0 \neq H_2 \cap X_0$. A symmetric argument applies if $H_2 \setminus H_1 \neq \emptyset$. \square

In the next lemma, we characterize convergence of nets in \mathcal{H}_X . As we shall see, $(\mathcal{H}_X, \mathcal{T}_{\mathcal{H}})$ is first countable, thus convergence of sequences actually leads to the same properties, but we decided to work with the more general notion for now, since it does not lead to additional difficulties.

Suppose that \mathcal{D} is a directed set (i.e., partially ordered by “ \succeq ” so that every two elements have a common upper bound), and $\{Y_n\}_{n \in \mathcal{D}}$ is a net (i.e., a sequence indexed by \mathcal{D}) of subsets of X . We write $\limsup_n Y_n$ for the set of all $x \in X$ with the property that for every topologically open $U \subset X$ containing x and for every $m \in \mathcal{D}$ there is some $n \succeq m$ in \mathcal{D} satisfying $Y_n \cap U \neq \emptyset$. Similarly, $\liminf_n Y_n$ is the set of all $x \in X$ with the property that for every topologically open $U \subset X$ containing x there is some $m \in \mathcal{D}$ such that $Y_n \cap U \neq \emptyset$ for every $n \succeq m$. We write $Y = \lim_n Y_n$ if $Y = \limsup_n Y_n = \liminf_n Y_n$.

LEMMA 3.6. *Assume that (3.1) and (3.2) hold, and that $S_n, n \in \mathcal{D}$, is a net with $S_n \in \mathcal{H}$ for every n .*

- (1) *If $\lim_n S_n \subset X$ exists, then $\lim_n S_n \in \mathcal{H}$.*
- (2) *If the net S_n converges to S with respect to \mathcal{T} , then $S = \lim_n S_n$.*
- (3) *Conversely, if $S = \lim_n S_n$, then the net S_n converges to S with respect to \mathcal{T} .*

PROOF. To prove (1), suppose that $S = \lim_n S_n$, $x \in S$ and $y < x$. Then $x \in \mathcal{R}_y$. Therefore, there is some $m \in \mathcal{D}$ such that $S_n \cap \mathcal{R}_y \neq \emptyset$ for every $n \geq m$. Since $S_n \in \mathcal{H}$, this implies that $y \in S_n$ for every $n \geq m$. Thus, $y \in S$, which proves that S is a lower set. The definition of $\lim_n S_n$ makes it clear that S is topologically closed. This proves (1).

To prove (2), suppose that S_n converges to S with respect to \mathcal{T} . Let $x \in S$. If $U \subset X$ is topologically open and contains x , then $S \in V_U$. Therefore, there is some $m \in \mathcal{D}$ such that $S_n \in V_U$ for every $n \geq m$. But $S_n \in V_U$ is equivalent to $S_n \cap U \neq \emptyset$. Therefore, $x \in \liminf_n S_n$; that is, $S \subset \liminf_n S_n$. Now suppose that $y \notin S$. By (3.2) there is some $z \in \mathcal{R}^y \setminus S$. Then $S \in V^z$. Consequently, there is some $m \in \mathcal{D}$ such that $S_n \in V^z$ for all $n \geq m$. Hence, $S_n \cap \mathcal{R}_z = \emptyset$ for all $n \geq m$. Since $y \in \mathcal{R}_z$ and \mathcal{R}_z is topologically open, this implies that $y \notin \limsup_n S_n$. Thus, $S \supset \limsup_n S_n$. Summing it up, we have $\liminf_n S_n \supset S \supset \limsup_n S_n$. Since $\liminf_n S_n \subset \limsup_n S_n$, this proves (2).

To prove (3), suppose that $S = \lim_n S_n$. Then we know from (1) that $S \in \mathcal{H}$. Suppose that S_{n_j} is a subnet of S_n that converges to some $S' \in \mathcal{H}$. By (2), we have $S' = \lim_j S_{n_j}$. Therefore, $S' = S$. Thus, every convergent subnet of S_n converges to S . Since \mathcal{H} is compact, this proves that S_n converges to S , and completes the proof of the lemma. \square

LEMMA 3.7. *Assuming (3.1) and (3.2), if X is second-countable (i.e., with a countable base), then \mathcal{H} is second-countable and hence metrizable.*

PROOF. Assume that X is second-countable. In particular, there is a countable dense set $X_0 \subset X$. Let \mathcal{B} be a countable basis for \mathcal{T}_X . We claim that $\mathcal{B}' := \{V_U : U \in \mathcal{B}\} \cup \{V^x : x \in X_0\}$ is a subbase for \mathcal{T} .

Indeed, if $U' \in \mathcal{T}_X$, then there is a subset $J \subset \mathcal{B}$ such that $U' = \bigcup_{U \in J} U$. Then $V_{U'} = \bigcup_{U \in J} V_U$; thus, $V_{U'}$ is in the minimal topology containing \mathcal{B}' .

Now suppose that $x \in X$, then

$$(3.3) \quad V^x = \bigcup_{y \in X_0 \cap \mathcal{R}^x} V^y.$$

Indeed, if $S \in V^x$, then $x \notin S$ and by (3.2) the topologically open set $\mathcal{R}^x \setminus S$ is nonempty. Take some $y \in X_0 \cap \mathcal{R}^x \setminus S$, then $S \in V^y$. This proves that $V^x \subset$

$\bigcup_{y \in X_0 \cap \mathcal{R}^x} V^y$. Since $y < x$, we also have $V^y \subset V^x$ for $y \in \mathcal{R}^x$, concluding that $V^x = \bigcup_{y \in X_0 \cap \mathcal{R}^x} V^y$. Thus, V^x is also in the minimal topology containing \mathcal{B}' .

Now Remark 3.4 shows that \mathcal{B}' is a subbase for \mathcal{T} . Since \mathcal{B}' is countable, \mathcal{T} also has a countable basis (of all finite intersections of elements of \mathcal{B}').

By Proposition 3.3, topology \mathcal{T} is Hausdorff. Since a second-countable compact Hausdorff space is metrizable (which follows from Urysohn's Theorem 4.16 in [17]), this completes the proof. \square

LEMMA 3.8. *Suppose that (3.1) and (3.2) hold, and that X_0 is a dense subset of X . Then the Borel σ -field of $(\mathcal{H}, \mathcal{T})$ is generated by $V^y, y \in X_0$.*

PROOF. Observe that the Borel σ -field of $(\mathcal{H}, \mathcal{T})$ is generated by topologically open sets, which in turn are generated by the sets of the form V_U and V^x for topologically open $U \subset X$ and $x \in X$. So we must prove that those can be obtained from $V^y, y \in X_0$, by the operations of countable union, intersection and taking complements.

For $x \in X$, the set V^x can be obtained as $V^x = \bigcup_{y \in X_0 \cap \mathcal{R}^x} V^y$, by (3.3).

For a topologically open $U \subset X$, note that by the definitions $V_U = \bigcup_{x \in U} \neg V^x$. Therefore, $V_U \supset \bigcup_{y \in U \cap X_0} \neg V^y$. On the other hand, take $x \in U$. By (3.2) we can find $y \in X_0$ inside the topologically open set $U \cap \mathcal{R}^x$, then $y < x$ and $\neg V^x \subset \neg V^y$. Thus $V_U \subset \bigcup_{y \in U \cap X_0} \neg V^y$. We conclude that $V_U = \bigcup_{y \in U \cap X_0} \neg V^y$, and so the sets V_U can also be obtained. \square

Finally, observe that a topologically open $X' \subset X$ inherits the properties (3.1) and (3.2) from X , and the following continuity under inclusions holds.

LEMMA 3.9. *Let $X' \subset X$ be topologically open, and for $S \in \mathcal{H}_X$ let $\Phi(S) := S \cap X'$. Then Φ is a continuous map from \mathcal{H}_X to $\mathcal{H}_{X'}$.*

PROOF. We have to check that preimages of topologically open sets are also topologically open. This follows easily since a topologically open subset $U \subset X'$ is topologically open inside X and $\Phi^{-1}(V^U) = V^U$, while for $x \in X'$ we have $\Phi^{-1}(V^x) = V^x$. \square

Summing it up, we arrive at the following.

THEOREM 3.10. *Let (X, τ) be a partially ordered second-countable topological space, such that the ordering satisfies (3.1) and (3.2). Then:*

- (1) $(\mathcal{H}_X, \mathcal{T}_\mathcal{H})$ is a compact metrizable Hausdorff space.
- (2) Let X_0 be a dense subset of X . If $S_1, S_2 \in \mathcal{H}_X$ satisfy $S_1 \cap X_0 = S_2 \cap X_0$, then $S_1 = S_2$. Moreover, the σ -field generated by $V^y, y \in X_0$, is the Borel σ -field of $(\mathcal{H}_X, \mathcal{T}_\mathcal{H})$.

(3) If $\{S_n\}$ is a sequence in \mathcal{H}_X and $S \in \mathcal{H}_X$, then $S_n \rightarrow S$ in $\mathcal{T}_{\mathcal{H}}$ is equivalent to $S = \limsup_n S_n = \liminf_n S_n$.

(4) Let $X' \subset X$ be a topologically open subset, and for $S \in \mathcal{H}_X$ let $\Phi(S) := S \cap X'$. Then Φ is a continuous map from \mathcal{H}_X to $\mathcal{H}_{X'}$.

PROOF. Part 1 follows from Propositions 3.1, 3.2 and Lemma 3.7; part (2) follows from Lemmas 3.5 and 3.8; part (3) follows from Lemma 3.6; and part (4) follows from Lemma 3.9. \square

4. Uniformity in the mesh size. In this section, we prove a mesh-independent version of the gluing theorem.

PROPOSITION 4.1 (Mesh-independent gluing). *Let $D \subset \mathbb{C}$ be some domain, and let $Q_0 \in \mathcal{Q}_D$ be a piecewise smooth quad in D . Suppose that $\alpha \subset \mathbb{C}$ is a finite union of finite length paths, with finitely many double points. Consider a collection of percolation models indexed by a set $\{\eta\}$. Assume that μ_η converges to (a subsequential) scaling limit μ_0 as mesh $|\eta| \rightarrow 0$. Then for every $\epsilon > 0$ there is a finite set of piecewise smooth quads $\mathcal{Q}_\epsilon \subset \mathcal{Q}_{D \setminus \alpha}$ and a subset $\mathcal{W}_\epsilon \subset \mathcal{H}$, which is measurable with respect to the finite σ -field $\sigma(\boxminus_Q : Q \in \mathcal{Q}_\epsilon)$, such that*

$$(4.1) \quad \lim_{|\eta| \rightarrow 0} \mu_\eta(\mathcal{W}_\epsilon \Delta \boxminus_{Q_0}) < \epsilon.$$

PROOF. This is the most technically difficult part of our paper. The rough plan is as follows: we set up a thin strip around α , and cut it into “bays,” forming a neighborhood of α , and a part away from α , bounded by “beaches.” Partition depends on the percolation configuration, and has the properties that (i) percolation in the neighborhood of α is decoupled from the rest, and (ii) the crossings outside this neighborhood can be encoded by a finite information, stable under small perturbations of the percolation configuration. We use a complicated construction of the neighborhood for the sake of the property (ii). To effectively bound the needed crossing information away from α (whose amount grows as we pass to the scaling limit), we further coarse-grain this procedure to a fixed small scale, when it is described by a finite σ -field. This will allow to write the required estimates.

We now start the proof by fixing a countable subsequence $\{\eta_j\}$ from our set, such that $|\eta_j| \rightarrow 0$ (and μ_{η_j} converges to μ_0). Recall that by Remark 1.17 in our setting it is sufficient to work with sequences, rather than nets.

Setup for the neighborhood of α . Reasoning like in the proof of Theorem 1.5 (the first paragraph of Section 2), we can approximate the quad Q_0 , by quads whose boundaries intersect α on a finite set. Thus, we can assume that α has finitely many double points, intersects ∂Q_0 on a finite set. Prolonging some of the curves in α if necessary (to cut the non-simply-connected components into smaller pieces), we can further assume that each component of $[Q_0] \setminus \alpha$ is simply connected.

Fix small $\epsilon_0 > 0$. Define two random variables: one given by the indicator function of the crossing event $Y_0 := 1_{\square_{Q_0}}$; and another by the conditional probability of a crossing, given the percolation configuration s -away from α :

$$Y_s := \mu_\eta(\square_{Q_0} \mid \mathcal{F}_s).$$

Here $\mathcal{F}_s, s > 0$, denotes the σ -field generated by the restriction of ω to the complement of the s -neighborhood of α , as in Theorem 1.5. By that theorem, for all sufficiently small $s > 0$,

$$(4.2) \quad \sup_{|\eta| \in (0, s)} \mu_\eta(\epsilon_0 < Y_s < 1 - \epsilon_0) < \epsilon_0.$$

Denote by $x_i, i = 1, \dots, n$, the finitely many points where α intersects ∂Q_0 , as well as the endpoints and the double points of curves in α . Let d be the minimum over i of the distances from x_i to other x_j 's and the farthest of $\partial_0 Q_0$ and $\partial_2 Q_0$. We now fix $s > 0$ sufficiently small so that (4.2) holds, and sufficiently small so that for $i = 1, \dots, n$, the balls $B(x_i, 2s)$ are disjoint, and

$$(4.3) \quad \sup_{|\eta| \in (0, s)} \mu_\eta(\exists \text{ a crossing from } B(x_i, 2s) \text{ to } \partial B(x_i, d/3)) < \epsilon_0/n.$$

The latter follows for small s from the RSW estimate (1.2). We also require that the intersection of α with $\bigcup_i \partial B(x_i, s)$ is finite (which holds for almost every value of s , since the total length of α is finite).

For the following construction, the reader is advised to refer to Figure 7. Let β and β' be finite unions of disjoint smooth simple paths in $[Q_0] \setminus \bigcup_i B(x_i, s)$, and




FIG. 7. Curves in α are between β' which are in turn between β . We modify configuration on the small discs so that long quads bounded by β and β' on long sides, have open tiles on one short side and closed on another.

let K be the closure of the union of connected components of

$$[Q_0] \setminus \left(\bigcup_i B(x_i, s) \cup \beta \cup \beta' \right),$$

whose boundary intersects both β and β' . We can choose these unions of paths so that:

- all their endpoints belong to $\bigcup_i \partial B(x_i, s)$,
- β' separates β from α ,
- the connected component of $[Q_0] \setminus \beta$ that contains α is contained in the $(s/2)$ -neighborhood of α ,
- each component of connectivity K_j of K is a quad with two “long” sides on β and β' and two “short” sides on two of the circles $\partial B(x_i, s)$.

It is immediate to verify that there exist such unions of paths.

Bays and beaches. The following construction is illustrated in Figure 8. Let ω be a sample from μ_η for some small $\eta > 0$. For somewhat technical reasons, it will be convenient to consider in place of ω the configuration $\tilde{\omega}$ which is modified on the tiles intersecting disks $B(x_j, s)$. We alter ω so that for every component K_i one of two short sides is covered by open tiles and another by closed. Clearly, (4.3) implies that if $\eta < s$,

$$(4.4) \quad \mu_\eta(\{\omega \in \Xi_{Q_0}\} \Delta \{\tilde{\omega} \in \Xi_{Q_0}\}) \leq \epsilon_0.$$

For this reason, studying $\tilde{\omega}$ will still be useful.

In what follows, we start on the curve β and explore all the potential crossings of the strip between β and β' . To be more precise, denote by $\partial\tilde{\omega}$ the set of percolation interfaces—curves, separating open tiles in $\tilde{\omega}$ from the closed ones. Let I be the set of connected components of the intersection of $\partial\tilde{\omega}$ with the interior of K (i.e., percolation interface in the interior of K), and let \tilde{I} denote the set of elements of




FIG. 8. Interfaces between open and closed clusters, starting from β , cut the strips K_i into components. Bays are the components touching β' , and are alternatively open and closed. Beaches are the inner boundaries of the bays, when entered through β' .

I that have at least one endpoint on β . The definition of $\tilde{\omega}$ guarantees that in each component K_i there is at least one interface of $\tilde{\omega}$ with one endpoint on β and the other on β' .

Let Γ denote the union of the elements of \tilde{I} ; that is, $\Gamma = \bigcup \tilde{I} = \bigcup_{\gamma \in \tilde{I}} \gamma$. Each connected component of $K \setminus \Gamma$ which meets β' will be called a *bay*. Clearly, each point in β' is in the closure of some bay, and for each bay its intersection with β' is connected. Let B be some bay. The union of the tiles of H_η that intersect B as well as $\partial B \setminus \beta'$ will be called the *beach* of B . It is easy to see that the beach is connected, and all the tiles in the beach are either open (i.e., in $\tilde{\omega}$) or closed (not in $\tilde{\omega}$). If they are all open (resp., closed), then B itself will be referred to as an open (resp., closed) bay.

Let $M' = M'(\omega)$ be the union of the component of $[Q_0] \setminus \beta'$ containing α , and of all the bays without beaches. Denote by $M = M(\omega)$ the complement in $[Q_0]$ of M' , as well as the corresponding set of tiles. The following property of M is essential: given M and the restriction of ω to it, the conditional law of the restriction of ω to M' (i.e., the complement of M) is unbiased. In other words, if we take ω_1 independent from ω and of the same law, and we define ω_2 to agree with ω on M and to agree with ω_1 elsewhere, then ω_2 also has the law μ_η . This follows directly from the fact that $M(\omega_2) = M(\omega)$. Below we will denote the restriction of a percolation configuration ω to M by $\omega \llcorner M$.

Fix some small $s' \in (0, s)$, and let $S = S_{s'}$ denote the event that there is a small bay B , that is, with $\text{diam}(B \cap \beta') < s'$. Observe that for every bay B the endpoints of $B \cap \beta'$ are endpoints of interfaces of ω which meet both β and β' . Thus, three crossings of K (which are alternating, e.g., closed, open and closed) land on the arc $B \cap \beta'$ and Lemma A.2 applied to all components of K implies that

$$\lim_{s' \rightarrow 0} \sup_{|\eta| \in (0, s')} \mu_\eta(S) = 0.$$

By choosing s' sufficiently small, we therefore assume, with no loss of generality, that $\mu_\eta(S) < \epsilon_0$ holds for every $\eta \in (0, s')$ and that $s' < s$. Observe that on $\neg S$ the number of bays is bounded by a constant depending only on s', β, β' and Q_0 .

We will find the structure of the bays to be a convenient component in the estimation of the conditional probability (given the bays and some additional information) of $\tilde{\omega} \in \Xi_{Q_0}$. However, as $|\eta| \rightarrow 0$, the number of possible bay configurations increases without bound, and therefore their structure cannot be completely captured by a finite σ -field. For this purpose, we need to introduce a discretization, and any sort of coarse-graining, for example, by the square lattice, would do.

Discretization to a finite σ -field by coarse-graining. We need some tessellation of K , and the following one is somewhat ad hoc, but suitable. Let $\delta \in (0, s')$ be small, to be chosen later. Fix some finite set of points $W \subset K$ such that the distance between any two points in W is at least $\delta/4$ and every point in K is within distance $\delta/2$ of some point in W . Let T be the collection of Voronoi tiles in K with

respect to W ; that is, each T_w , $w \in W$, is the set of points $x \in K$ whose distance to W is equal to $|x - w|$. Then the diameters of the tiles in the tessellation T are all less than δ (but the lattice H_η tiles are still much smaller). Let \hat{T} denote the tessellation obtained by adjoining to T the closures of the (finite in number) connected components of $[Q_0] \setminus K$ that do not contain α . We think of tessellations T and \hat{T} as cell complexes. The vertices of T are defined as the vertices of the Voronoi tiles together with the points in $(\beta \cup \beta') \cap \partial Q_0$. The vertices of \hat{T} are defined as those of T plus the four corners of Q_0 . Note that some of the edges of \hat{T} are not necessarily straight. Denote by $[\hat{T}]$ the union of the tiles in \hat{T} .

We say that a quad $Q \in Q_D$ is *compatible* with \hat{T} if $[Q]$ is a union of tiles of \hat{T} and each of its four corners $Q(0, 0)$, $Q(0, 1)$, $Q(1, 0)$, $Q(1, 1)$ is a vertex of \hat{T} . Since \hat{T} is finite, it is clear that up to reparameterization preserving corners there are only finitely many compatible Q . Let $Q_{\hat{T}}$ denote a set of representatives of equivalence classes of compatible quads up to reparameterization preserving corners. Then $Q_{\hat{T}}$ is a finite collection of piecewise smooth quads in $D \setminus \alpha$. Let \mathcal{F}_T denote the σ -field generated by the events \Box_Q , $Q \in Q_{\hat{T}}$.

Without loss of generality, we assume that along our chosen sequence $\{\eta_j\}$, for each edge e of the tessellation T that does not touch ∂Q_0 :

- each edge of H_η intersects the tessellation edge e in at most finitely many points,
- the tessellation edge e does not contain any vertices of H_η ,
- the endpoints of the tessellation edge e are not on edges of H_η .

There is no loss of generality in this assumption, because we may slightly perturb β , β' and T (and we do not even need T to be a Voronoi tessellation).

Let B be some bay. Define $T(B)$ to be its T -discretization, namely the union of all edges of T inside $B \cap \beta'$ and all tiles of T inside B , which can be connected to β' by a path of tiles of T also contained inside B . By construction, each point in $\partial T(B)$ is within distance δ from ∂B . Note that there is a finite number of possibilities for the choices of $T(B)$, and also observe that whenever S does not hold, all T -discretizations are nonempty and so $T(B) = T(B')$ implies $B = B'$.

Connectivity in the coarse-grained model. The next objective is to show that on the complement of S the collections of discretized bays

$$Z_o := \{T(B) : B \text{ an open bay}\}$$

and

$$Z_c := \{T(B) : B \text{ a closed bay}\}$$

can be determined from \mathcal{F}_T . That is, there are \mathcal{F}_T -measurable random variables Z'_o and Z'_c such that $Z_o = Z'_o$ and $Z_c = Z'_c$ on the complement of S . Here and in the following, we ignore events that have zero μ_η -measure for each $\eta \in \{\eta_j\}$.

Note that on the event $\neg S$ the cardinality of Z_o and Z_c is bounded by a constant depending on Q_0 , β' and s' only.

Let e be an edge of T that lies on the component $\beta'_j = \beta' \cap K_j$, denote also $\beta_j := \beta \cap K_j$. Let Q_e and Q'_e be the quads in $\mathcal{Q}_{\hat{T}}$ that satisfy $\partial_0 Q_e = \partial_1 Q'_e = e$, $\partial_2 Q_e = \partial_3 Q'_e = \beta_j$ and $[Q_e] = [Q'_e] = K_j$. Then the event that there is an interface $\gamma \in \tilde{I}$ that meets e is the same as the event $\square_{Q_e} \setminus \square_{Q'_e}$.

If E is the set of all edges of T on β'_j which meet interfaces in \tilde{I} , then on the event $\neg S$, the connected components of $\beta'_j \setminus \cup E$ are the intersections of the elements of $Z_o \cup Z_c$ with β'_j , and since the bays alternate in color, it follows that the sets $\{\tau \cap \beta'_j : \tau \in Z_o\}$ and $\{\tau \cap \beta'_j : \tau \in Z_c\}$ can each be determined from \mathcal{F}_T , up to an event contained in S .

Next, suppose that $\gamma \subset \beta'_j$ is of the form $\tau_0 \cap \beta'_j$ for some $\tau_0 \in Z_c$ (assuming $\neg S$, such intersection is always nonempty). Let T_0 be some tile of T . Then $T_0 \subset \tau_0$ if and only if there is a simple path γ' of edges of T with the endpoints of γ' on γ such that γ' separates T_0 from β_j in K_j and there is no open crossing in $K_j \cap \tilde{\omega}$ from γ' to β_j . This, along with the dual argument, implies the existence of Z'_o and Z'_c , as claimed.

Describing the crossing structure by a finite graph. We now introduce a graph that describes the connectivity (by open crossings) structure of the various open beaches and the two boundary edges $\partial_0 Q_0$ and $\partial_2 Q_0$.

We start by defining a random graph G which would describe the connectivity away from α . The vertices of G are $Z_o \cup \{\partial_0 Q_0, \partial_2 Q_0\}$, where an edge is placed between vertices $v, v' \in Z_o$ if the beaches of the corresponding bays are connected by an open crossing in the restriction of $\tilde{\omega}$ to M , denoted by $\tilde{\omega}_\perp M$; while if v and/or v' are in $\{\partial_0 Q_0, \partial_2 Q_0\}$, the connectivity to the beach is simply replaced by the connectivity to the corresponding boundary edge itself. Note that there is a finite number of possibilities for the choice of G .

Suppose that $\tau_0, \tau_1 \in Z_o$. Let $e_j, j = 0, 1$, be an edge of T that is contained in β' and has an endpoint in τ_j , but e_j itself is not in τ_j . Then e_j meets an interface on the boundary of the corresponding bay. On $\neg S$, the two edges e_j are connected by an open crossing in $\tilde{\omega}_\perp[\hat{T}]$ if and only if $[\tau_0, \tau_1]$ is an edge of G . Thus, on the event $\neg S$ the subgraph of G induced by Z_o is determined from \mathcal{F}_T . A similar argument applies also to the edges with endpoints in $\{\partial_0 Q_0, \partial_2 Q_0\}$. Thus, on the complement of S the graph G can be determined from \mathcal{F}_T .

Now we define a random graph G^* so that it describes the connectivity near α . Consider $L = (\partial_0 Q_0 \cup \partial_2 Q_0) \setminus [\hat{T}]$, which contains finitely many arcs on ∂Q_0 . Denote by G^* the graph whose vertices are Z_o and the connected components of L , where $v, v' \in Z_o$ are connected by an edge in G^* if the two corresponding beaches of $\tilde{\omega}$ are connected by an open crossing in $\tilde{\omega}_\perp([Q_0] \setminus M)$, and if v and/or v' are components of L , then the connectivity to the beaches is replaced by connectivity to v and/or v' .

Clearly, $\tilde{\omega} \in \square_{Q_0}$ if and only if there is a path from $\partial_0 Q_0$ (or a subarc thereof) to $\partial_2 Q_0$ in $G \cup G^*$. Thus, on the event $\neg S$, the graphs G and G^* determine whether $\tilde{\omega} \in \square_{Q_0}$.

Denoting by $\omega_{\perp}M$ the restriction of ω to M , we set

$$\tilde{Y} = \tilde{Y}(\omega) := \mu_{\eta}(\tilde{\omega} \in \Xi_{Q_0} \mid \omega_{\perp}M).$$

We now show that on $\neg S$, the knowledge of G , Z_c and Z_o can be used to approximate \tilde{Y} .

Let

$$T(M) := [\hat{T}] \setminus \left(\bigcup Z_c \cup \bigcup Z_o \right),$$

that is, the complement in $[\hat{T}]$ of the union of $T(B)$, where B runs over all the bays.

Let G^{\dagger} be the graph on the same vertex set as G^* , where an edge appears between $v, v' \in Z_o$ if the common boundary of v and $T(M)$ is connected in $\tilde{\omega}_{\perp}([Q_0] \setminus T(M))$ with the common boundary of v' and $T(M)$, while if v and/or v' are components of L , then the connectivity is instead to v and/or v' , but still within $\tilde{\omega}_{\perp}([Q_0] \setminus T(M))$. We now show that if δ is sufficiently small, $\eta < \delta$ and $\eta \in \{\eta_n\}$, then

$$(4.5) \quad \mu_{\eta}(G^{\dagger} \neq G^*, \neg S \mid \omega_{\perp}M) < \epsilon_0.$$

As we have mentioned, on the event $\neg S$ there is a finite upper bound on the size of $Z_c \cup Z_o$, depending on Q_0 , β' and s' only. Therefore, it is enough to have a good estimate for the probability that a particular pair of vertices v and v' of G^* are connected by an edge in one of the two graphs but not in the other, and we are free to impose additional conditions on δ . Connectivity of any pair of vertices is reduced to δ -perturbations of quads of size at least s' , and so is easily obtained by several applications of Lemma A.1 (some of them applied with the corners of the quads appropriately permuted) proving (4.5).

Final estimates. Define $\tilde{Y}_0 = \tilde{Y}_0(\omega) := 1_{\tilde{\omega} \in \Xi_{Q_0}}$ and

$$\tilde{Y}_T := \mu_{\eta}(\partial_0 Q_0 \text{ and } \partial_2 Q_0 \text{ are connected by a path in } G^{\dagger} \cup G \mid \omega_{\perp}M).$$

Then \tilde{Y}_T is a function of the triple (Z_c, Z_o, G) , and is therefore \mathcal{F}_T -measurable. By the above observation that on $\neg S$ the event Ξ_{Q_0} is described by connectivity in $G \cup G^*$, we have

$$|\tilde{Y}_T - \tilde{Y}|_{1-\neg S} \leq \mu_{\eta}(G^{\dagger} \neq G^*, \neg S \mid \omega_{\perp}M).$$

By taking expectations, applying (4.5) and recalling that s' was chosen to guarantee $\mu(S) < \epsilon_0$ one gets

$$(4.6) \quad \|\tilde{Y}_T - \tilde{Y}\|_2 < \sqrt{2\epsilon_0},$$

where the norm refers to the measure μ_{η} .

By (4.2) and the definition of Y_s , we can write

$$\|Y_0 - Y_s\|_2^2 < 2\epsilon_0 + \epsilon_0^2 < 3\epsilon_0,$$

and by (4.4), we have

$$(4.7) \quad \|Y_0 - \tilde{Y}_0\|_2 < \sqrt{\epsilon_0}.$$

Hence, $\|\tilde{Y}_0 - Y_s\|_2 < \sqrt{3\epsilon_0} + \sqrt{\epsilon_0} < 3\sqrt{\epsilon_0}$. Since Y_s is \mathcal{F}_s -measurable, it is also $\omega_\perp M$ -measurable. By its definition, \tilde{Y} minimizes $\|\tilde{Y}_0 - X\|_2^2$ among $\omega_\perp M$ -measurable random variables X . Therefore, comparison to Y_s (recall that the set M includes the complement of the s -neighborhood of α) yields

$$\|\tilde{Y}_0 - \tilde{Y}\|_2 \leq \|\tilde{Y}_0 - Y_s\|_2 < 3\sqrt{\epsilon_0}.$$

Combining this with (4.6) and (4.7), we conclude that

$$\|Y_0 - \tilde{Y}_T\|_2 < 6\sqrt{\epsilon_0}.$$

Because \tilde{Y}_T is \mathcal{F}_T -measurable and ϵ_0 may be chosen arbitrarily small, this proves that there is an \mathcal{F}_T -measurable event \mathcal{W} (which may depend on η), such that $\mu_\eta(\mathcal{W} \Delta \square_{Q_0}) < \epsilon/2$. However, since \mathcal{F}_T is finite, there are finitely many possibilities for the event \mathcal{W} , and one of those works for a sub-subsequence η_{j_k} . Since we work with a (subsequential) scaling limit, the limit in (4.1) exists along the original subsequence and is equal to $\mu_0(\mathcal{W} \Delta \square_{Q_0})$. Since along a sub-subsequence the quantity in question is bounded by ϵ , so is the limit and we deduce (4.1). \square

5. Factorization. In this section, we prove the Factorization Theorem 1.19.

LEMMA 5.1. *Let μ_0 be some subsequential scaling limit, then the boundary (in topology \mathcal{T} on \mathcal{H}) of a crossing event has probability zero. Namely,*

$$\mu_0(\partial \square_{Q_0}) = 0$$

holds for every $Q_0 \in \mathcal{Q}_D$.

PROOF. Fix $Q_0 \in \mathcal{Q}_D$ and let $\epsilon > 0$. It is easy to see (e.g., using the Riemann map onto $\hat{\mathbb{C}} \setminus [Q_0]$) that there is a continuous injective $\hat{Q}_0 : [-1, 2]^2 \rightarrow \mathbb{C}$, whose restriction to $[0, 1]^2$ is Q_0 . Let $M := [-1, 1/3]^2 \times [2/3, 2]^2$. For every quadruple $q = (x_0, y_0, x_1, y_1) \in M$, define the quad $Q^q : [0, 1]^2 \rightarrow \mathbb{C}$ by

$$Q^q(x, y) := \hat{Q}_0(x_0 + (x_1 - x_0)x, y_0 + (y_1 - y_0)y).$$

It is a perturbation of Q_0 , obtained as an image of $[x_0, x_1] \times [y_0, y_1]$ by \hat{Q}_0 . Lemma A.1 implies that there is some positive $\delta_0 > 0$ (depending on \hat{Q}_0), such that if q and q' are in M and differ in exactly one coordinate, and the difference in that coordinate is at most δ_0 , then $\limsup_{|\eta| \rightarrow 0} \mu_\eta(\square_{Q^q} \Delta \square_{Q^{q'}}) < \epsilon$. Consequently, we have

$$\limsup_{|\eta| \rightarrow 0} \mu_\eta(\square_{Q^q} \Delta \square_{Q^{q'}}) < 4\epsilon \quad \text{provided } \|q - q'\|_\infty \leq \delta_0.$$

Set $q_s := (-s, s, 1 + s, 1 - s)$, $Q' := Q^{q-\delta_0/2}$ and $Q'' := Q^{q\delta_0/2}$. Then $Q' < Q_0 < Q''$ and

$$(5.1) \quad \limsup_{|\eta| \rightarrow 0} \mu_\eta(\Box_{Q'} \Delta \Box_{Q''}) < 4\epsilon.$$

Also, $Q', Q'' \in \mathcal{Q}_D$, if δ_0 is chosen sufficiently small.

Recall that $V^{Q_0} = \mathcal{H}_D \setminus \Box_{Q_0}$ is open in \mathcal{H}_D . Hence \Box_{Q_0} is closed. Let U' be the set of quads $Q \in \mathcal{Q}_D$ satisfying $Q_0 < Q < Q''$. Then

$$\Box_{Q_0} \supset V_{U'} \supset \Box_{Q''}.$$

By passing to the complements, we see that there is a closed subset (the complement of $V_{U'}$) that contains $V^{Q_0} = \neg \Box_{Q_0}$ and is contained in $V^{Q''}$. Hence, the closure of V^{Q_0} is contained in $V^{Q''}$, which gives

$$\partial \Box_{Q_0} = \partial V^{Q_0} \subset \overline{V^{Q_0}} \subset V^{Q''}.$$

On the other hand, let U be the set of quads $Q \in \mathcal{Q}_D$ satisfying $Q' < Q < Q''$, then

$$\partial \Box_{Q_0} \subset \Box_{Q_0} \subset V_U.$$

Now observe that $V_U \subset \Box_{Q'}$ while $V^{Q''} = \neg \Box_{Q''}$. We have shown that the open set $V_U \cap V^{Q''}$ contains $\partial \Box_{Q_0}$. The portmanteau theorem (see, e.g., [9], Theorem 11.1.1) therefore gives

$$\begin{aligned} \mu_0(\partial \Box_{Q_0}) &\leq \mu_0(V_U \cap V^{Q''}) \leq \liminf_{|\eta| \rightarrow 0} \mu_\eta(V_U \cap V^{Q''}) \\ &\leq \liminf_{|\eta| \rightarrow 0} \mu_\eta(\Box_{Q'} \Delta \Box_{Q''}). \end{aligned}$$

Since ϵ was arbitrary, the result now follows from (5.1) \square

COROLLARY 5.2. *If $\mathcal{M} \subset \mathcal{H}_D$ is in the Boolean algebra generated by finitely many of the events \Box_Q , $Q \in \mathcal{Q}_D$ (i.e., it can be expressed using finitely many \Box_Q and the operations of union, intersection and taking complements), then $\mu_0(\mathcal{M}) = \lim_{|\eta| \rightarrow 0} \mu_\eta(\mathcal{M})$.*

PROOF. Lemma 5.1 implies that $\mu_0(\partial \mathcal{M}) = 0$. Thus, the corollary follows from the portmanteau theorem [9], Theorem 11.1.1, and the weak convergence of μ_η to μ_0 . \square

PROOF OF THEOREM 1.19. Clearly, $\mathcal{F}_{D \setminus \alpha} = \bigvee_j \mathcal{F}_{D_j}$. It therefore remains to prove the left-hand equality stated in the theorem. Note that we work up to sets of μ_0 -measure zero.

Take a smooth (given by a diffeomorphism) quad $Q_0 \in \mathcal{Q}_D$. By Proposition 4.1 and Corollary 5.2, we have

$$(5.2) \quad \square_{Q_0} \in \mathcal{F}_{D \setminus \alpha}.$$

Since such collection of quads is dense in \mathcal{Q}_D , Theorem 1.19 follows from Proposition 1.13(2). \square

APPENDIX A: CONTINUITY OF CROSSING EVENTS

In this section, we apply Russo–Seymour–Welsh techniques to deduce estimates of various crossing probabilities. We start by showing that crossing events are stable under small perturbations of quads.

LEMMA A.1. *There exist a positive function $\Delta_c(\delta, d)$, such that*

$$\lim_{\delta \rightarrow 0} \Delta_c(\delta, d) = 0 \quad \text{for any fixed } d,$$

and the following estimates hold. Let Q be a quad. Let $d_j, j = 0, 1$, be the infimum diameter of any path in $[Q]$ connecting $\partial_j Q$ and $\partial_{j+2} Q$, and define $d := \max\{d_0, d_1\}$. Fix some $\delta < d/2$. Let Q' be another quad, satisfying at least one of the following conditions; see Figures 9 and 10 for a graphical interpretation.

- (1) $[Q'] = [Q]$, $\partial_0 Q' = \partial_0 Q$, $\partial_1 Q' = \partial_1 Q$, and there is a path $\gamma \subset [Q]$ of diameter at most δ that separates $\{Q(1, 1), Q'(1, 1)\}$ from $\partial_0 Q \cup \partial_1 Q$ inside $[Q]$.
- (2) $[Q'] \subset [Q]$, $\partial_0 Q' = \partial_0 Q$, $\partial_1 Q' \subset \partial_1 Q$, $\partial_3 Q' \subset \partial_3 Q$, and each point on $\partial_2 Q'$ can be connected to $\partial_2 Q$ by a path $\alpha \subset [Q]$ with $\text{diam}(\alpha) \leq \delta$.
- (3) $[Q'] \subset [Q]$, $\partial_0 Q' \subset \partial_0 Q$, $\partial_1 Q' = \partial_1 Q$, $\partial_2 Q' \subset \partial_2 Q$, and each point on $\partial_3 Q'$ can be connected to $\partial_3 Q$ by a path $\alpha \subset [Q]$ with $\text{diam}(\alpha) \leq \delta$.

Then for every $|\eta| < \delta$ we have

$$\mu_\eta(\square_Q \Delta \square_{Q'}) < \Delta_c(\delta, d).$$




FIG. 9. Case (1): quads Q and Q' differ by only one vertex, separated from the other vertices by a short cut γ . Configurations with only one of two quads crossed have an open crossing from γ to $\partial_0 Q$ and a dual closed crossing from γ to $p_1 Q$, making their probability small.




FIG. 10. Case (2): side $\partial_2 Q'$ is close to the side $\partial_2 Q$. We take the lowest crossing γ of Q' , landing at x . A dual closed crossing prevents it from landing on $\partial_2 Q$, making the probability small. Case (3) is symmetric.

For site percolation on the triangular lattice, there is a short proof based on Cardy’s formula. The proof below is a simple application of the RSW estimate (1.2) and the “lowest crossing” concept.

PROOF OF LEMMA A.1. First, we deal with case (1) in the statement and prove the estimate with Δ_c equal to Δ_1 from assumption (1.2). Suppose first that $d = d_0$. The event $\Xi_Q \Delta \Xi_{Q'}$ is contained in the event that there is a percolation cluster meeting γ and $\partial_0 Q$. The latter has an open crossing of the annulus $A(Q(1, 1), \delta, d_0)$, which by the RSW estimate (1.2) has probability at most $\Delta_1(\delta, d)$, and we are done with case (1). If $d = d_1$, a similar argument shows the symmetric difference between the event of a closed crossing from $\partial_1 Q$ to $\partial_3 Q$ in $[Q]$ and the corresponding event in Q' is bounded by $\Delta_1(\delta, d_1)$. Duality shows that the latter symmetric difference is the same as $\Xi_Q \Delta \Xi_{Q'}$, which proves case (1).

We now deal with case (2), when clearly $\Xi_Q \Delta \Xi_{Q'} = \Xi_{Q'} \setminus \Xi_Q$. We start by estimating the crossing probability for Q' . Recall our convention that $\Delta(r, R) = 1$ for $r \geq R$ (i.e., when annulus in question is empty). Let x be some point on the path τ of diameter d_1 , connecting $\partial_1 Q$ to $\partial_3 Q$. Any crossing of Q' has diameter at least $d_0 - \delta$ and passes within distance δ of τ , so in particular it crosses an annulus $A(x, d_1 + \delta, (d_0 - \delta)/2)$ if well defined, and by the RSW estimate (1.2) we have

$$(A.1) \quad \mu_\eta(\Xi_{Q'}) \leq \Delta_1\left(d_1 + \delta, \frac{d_0 - \delta}{2}\right).$$

Next, we cut $\partial_2 Q'$ by a point z into two parts, σ_1 and σ_3 , so that σ_j cannot be connected to $\partial_j Q$ by a path of diameter less than $d_1/2$ inside Q . If $\Xi_{Q'}$ occurs, there is a crossing to at least one of σ_1 and σ_3 . We will work with configurations with a crossing landing on σ_3 , the other case being symmetric, with the lowest crossing replaced by the uppermost, so the total estimate will be double of what we obtain.

Consider a percolation configuration ω . On the event $\Xi_{Q'}$, let γ be the “lowest” ω -crossing of Q' (in the sense that it is the closest to $\partial_1 Q'$, i.e., separates $\partial_1 Q'$ from all other ω -crossings of Q' within $[Q']$), and let x be the endpoint of γ on $\partial_2 Q'$. Our assumption implies that $x \in \sigma_3$, or there would be a lower crossing.

We now estimate $\mu_\eta(\neg \Xi_Q \mid \Xi_{Q'}, \gamma, x \in \sigma_3)$. Let M be the connected component of $[Q'] \setminus \gamma$ which has $\partial_1 Q'$ on its boundary. The event $\neg \Xi_Q$ would imply that γ cannot be connected to $\partial_2 Q$ by a crossing inside $[Q] \setminus M$. Hence, there is a dual closed crossing from $\partial_3 Q$ to $(\sigma_3 \cap \bar{M}) \cup \partial_1 Q$, in particular crossing the annulus $A(x, \delta, d_1/2)$ inside $[Q] \setminus M$.

The lowest crossing γ depends only on the configuration inside M , so the restriction of ω to $[Q] \setminus M$ is unbiased. Therefore, the conditional probability of the mentioned annulus crossing is can be bound by the RSW estimate (1.2) and we conclude that

$$(A.2) \quad \mu_\eta(\neg \Xi_Q \mid \Xi_{Q'}, \gamma, x \in \sigma_3) \leq \Delta_1(\delta, d_1/2).$$

Now we are ready to prove the estimate, working out separately three possibilities:

- (i) $d = d_1$,
- (ii) $d = d_0 > \delta \geq d_1$,
- (iii) $d = d_0 > d_1 > \delta$.

When (i) occurs, we majorate the probability of the event in question by the maximum of its conditional probability, estimated by (A.2), to arrive at

$$(A.3) \quad \mu_\eta(\Xi_{Q'} \setminus \Xi_Q) \leq 2\Delta_1\left(\delta, \frac{d_1}{2}\right) = 2\Delta_1\left(\delta, \frac{d}{2}\right).$$

When (ii) occurs, we use the estimate (A.1):

$$(A.4) \quad \mu_\eta(\Xi_{Q'} \setminus \Xi_Q) \leq \mu_\eta(\Xi_{Q'}) \leq \Delta_1\left(d_1 + \delta, \frac{d_0 - \delta}{2}\right) \leq \Delta_1\left(2\delta, \frac{d}{4}\right).$$

Finally, when (iii) occurs, we apply the total expectation law, using both estimates (A.2) and (A.1) along the way:

$$(A.5) \quad \begin{aligned} \mu_\eta(\Xi_{Q'} \setminus \Xi_Q) &\leq \mathbf{E}[\mu_\eta(\neg \Xi_Q \mid \Xi_{Q'}, \gamma, x \in \sigma_3)] \\ &\quad + \mathbf{E}[\mu_\eta(\neg \Xi_Q \mid \Xi_{Q'}, \gamma, x \in \sigma_1)] \\ &\leq 2\mu_\eta(\Xi_{Q'})\Delta_1\left(\delta, \frac{d_1}{2}\right) \\ &\leq 2\Delta_1\left(d_1 + \delta, \frac{d_0 - \delta}{2}\right)\Delta_1\left(\delta, \frac{d_1}{2}\right) \\ &\leq 2\Delta_1\left(2d_1, \frac{d}{4}\right)\Delta_1\left(\delta, \frac{d_1}{2}\right). \end{aligned}$$

To prove lemma in case (2), we have to show that for fixed d and any $\epsilon > 0$, the right-hand side of the estimates (A.3), (A.4) and (A.5) is smaller than ϵ , if δ is small enough. For estimates (A.3) and (A.4), this follows directly from Assumption 1.2. For the remaining (A.5), let $\rho > 0$ be such that for $r < \rho$ we have $\Delta_1(2r, \frac{d}{4}) < \epsilon/2$. Then if $d_1 \leq \rho$, then the right-hand side of (A.5) can be bounded by

$$2\Delta_1\left(2d_1, \frac{d}{4}\right)\Delta_1\left(\delta, \frac{d_1}{2}\right) \leq 2\Delta_1\left(2d_1, \frac{d}{4}\right) < \epsilon,$$

we are done. Otherwise, $d_1 > \rho$ and the right-hand side of (A.5) can be bounded by

$$2\Delta_1\left(2d_1, \frac{d}{4}\right)\Delta_1\left(\delta, \frac{d_1}{2}\right) \leq 2\Delta_1\left(\delta, \frac{\rho}{2}\right),$$

which tends to zero with δ by Assumption 1.2, and we finish the proof in case (2).

The proof for case (3) is easily obtained by considering the dual closed crossing from ∂_1 to ∂_3 and applying case (2). The details are left to the reader. \square

The following lemma shows that it is unlikely to see three crossings approaching the same boundary point. This would make the crossing event unstable under boundary perturbation, so in principle the lemma can be deduced from the previous one, but instead we give a self-contained proof. The lemma would follow from considering the three arm event in *half-annuli*, for whose probability Aizenman proposed an argument to be comparable to $(r/R)^2$. We will use a version of his reasoning along with RSW techniques to show a weaker estimate, sufficient for our purposes.

LEMMA A.2. *Let Q be a quad with smooth sides, two opposite being labeled β and β' , and with the tiles on each of the two other sides ν and ν' possibly declared all open or all closed. Then there is a function $\Delta_Q(\delta)$, tending to zero with δ and such that the following estimate holds.*

Denote by $S(\theta)$ the event that there are three crossings with some prescribed alternating order (say closed, open, and closed) inside $[Q]$ between the set $\theta \subset \beta'$ and β . Then for $|\eta| < \delta$, we have

$$\mu_\eta\{\exists\theta \subset \beta' : S(\theta), \text{diam}(\theta) < \delta\} \leq \Delta_Q(\delta).$$

PROOF. Fix $\epsilon > 0$. We shall prove that for sufficiently small δ the probability of the event in question is less than ϵ .

Since crossings can in principle use tiles on the sides ν and ν' (which are declared open or closed), we have to treat θ 's near the corners of Q separately. Let z_1 and z_2 be the endpoints of β' . Since tiles on each of the arcs ν and ν' can be used by at most one of the three crossings, event $S(\theta)$ implies existence of at least one crossing from θ to β inside $[Q]$, not touching the sides. By assumption (1.2), there is $r \in (0, \delta)$ such that probability to have a crossing between $\cup_i B(z_i, 2r)$ and β is

smaller than $\epsilon/2$:

$$\mu_\eta \left\{ \exists \theta \subset \bigcup_i B(z_i, 2r) : S(\theta) \right\} \leq \frac{\epsilon}{2}.$$

Set $\beta'' := \beta' \setminus \bigcup_i B(z_i, r)$, then to prove lemma it remains to show that

$$(A.6) \quad \mu_\eta \{ \exists \theta \subset \beta'' : S(\theta), \text{diam}(\theta) < \delta \} \leq \frac{\epsilon}{2}.$$

Using assumption (1.2), choose $\rho \in (0, r/4)$ such that $\Delta_1(2\rho, r/2) < 1/2$.

Cover β'' by a finite number $n = n(\rho, \beta'')$ of overlapping arcs θ_j of diameter at most ρ . Let $\delta < \rho$ be so small that $\Delta_1(4\delta, r) < \epsilon/(12n)$. Since β' is smooth, and decreasing δ if necessary, we can cover each arc θ_j by arcs θ_j^l of length 3δ , overlapping at most thrice, and such that any set $\theta \subset \beta''$ of diameter less than δ is entirely contained within one of the arcs θ_j^l .

To deduce (A.6) it is sufficient to show that, for a fixed j ,

$$(A.7) \quad \sum_l \mu_\eta [S(\theta_j^l)] \leq \frac{\epsilon}{2n}.$$

Fix j and let $B(z, \rho)$ be a radius ρ ball, containing θ_j . The point z_1 splits $\partial Q \setminus \theta_j$ into two arcs: $\lambda' \subset \beta'$ and a $\lambda \supset \beta$. Both start at endpoints of θ_j inside $B(z, \rho)$ and end at least r -away.

Observe that event $S(\theta_j^l)$ is contained in the event $S'(\theta_j^l)$ that *there are three alternating crossings inside $[Q]$ between θ and $\beta \cup v \cup v'$* .

Condition on the event $S'(\theta_j^l)$, and, starting from z_2 , denote the three such crossings, closest to z_2 , by $\gamma_1, \gamma_2, \gamma_3$. Namely, let γ_1 be the closed crossing between θ and $\beta \cup v \cup v'$ inside $[Q]$, closest to z_2 . Let M_1 be the component of connectivity of $[Q] \setminus \gamma_1$, not containing z_2 . Then γ_1 depends only on the percolation configuration in $[Q] \setminus M_1$.

Now, take γ_2 be the open crossing between θ and $\beta \cup v \cup v'$ inside M_1 , closest to γ_1 . Let M_2 be the component of connectivity of $M_2 \setminus \gamma_2$, not containing γ_1 . Then γ_1 and γ_2 depend only on the percolation configuration in $[Q] \setminus M_2$.

We define γ_3 and M_3 similarly, and observe that $\gamma_1, \gamma_2, \gamma_3$ depend only on the percolation configuration in $[Q] \setminus M_3$. Therefore, given M_3 and the restriction $\omega_\perp([Q] \setminus M_3)$ (i.e., the restriction of percolation configuration ω to $[Q] \setminus M_3$), the conditional law of the restriction of ω to M_3 is unbiased.

Let $S^* = S^*(M_3)$ be the event that *there is a closed crossing γ_4 between λ and γ_3 inside M_3* . Using that the restriction of ω to M_3 is unbiased, we can then write

$$(A.8) \quad \mu_\eta [S^* \mid M_3, \omega_\perp([Q] \setminus M_3)] = \mu_\eta [S^*] > 1 - P_\eta^1(z, \rho, r) > \frac{1}{2}.$$

Above we use that (in an unbiased percolation configuration), if there is no open crossing of the annulus $A(z, \rho, r)$, then by duality there is a closed circuit (a path going around the annulus), whose intersection with M_3 would then contain a crossing required for S^* .

Note that when $S'(\theta_j^l)$ and S^* occur, the crossing γ_3 from $[S(\theta_j^l)]$ and the crossing γ_4 (from S^*) together give a closed crossing between θ and λ , and so the following event $S''(\theta_j^l)$ occurs: *there is a closed crossing γ_3' from θ to λ , an open crossing from γ_2' from θ to $\beta \cup \nu \cup \nu'$, and a closed crossing from γ_1' from θ to $\beta \cup \nu \cup \nu'$.* Then the estimate (A.8) can be rephrased as

$$\mu_\eta[S''(\theta_j^l) \mid S'(\theta_j^l)] \geq \frac{1}{2},$$

and therefore

$$(A.9) \quad \mu_\eta[S(\theta_j^l)] \leq \mu_\eta[S'(\theta_j^l)] \leq 2\mu_\eta[S''(\theta_j^l)].$$

Consider some percolation configuration ω in $[Q]$, and assume that $S''(\theta_j^l)$ holds for some $\theta_j^l \subset \theta_j$.

Choose inside Q a closed crossing $\gamma_3'(\omega)$ between λ' and θ_j ; and an open crossing $\gamma_2'(\omega)$ between λ which are the closest to each other (and the point z_1 separating λ' from λ). Denote by $L = L(\omega)$ the union of tiles in γ_2', γ_3' , and in the part of $[Q]$ between them.

Then γ_2' and γ_3' depend only on the restriction of ω to L , and the rest is unbiased. Moreover, if γ_2' ends at a tile on θ_j , then γ_3' ends at a neighboring tile on θ_j —otherwise using duality reasoning we can choose two closer crossings. Denote by x the only common vertex of these two tiles lying on the boundary of N . Note that each configuration ω has a unique such point $x = x(\omega)$ and it depends only on the restriction of ω to N .

Now for $S(\theta_j^l)$ to occur, besides γ_2' and γ_3' we must also have a closed crossing γ_1' from θ to $\beta \cup \nu \cup \nu'$. The crossing γ_1' necessarily lies in $[Q] \setminus N$, which is unbiased on $\omega \llcorner N$. Note also that γ_1' contains a crossing of the annulus $A(x, 2\delta, r)$, whose center x is a random point depending on $\omega \llcorner L$ only. Summing it up, we can estimate

$$\mu_\eta\{\omega : S''(\theta_j^l) \mid x(\omega) = x\} \leq P_\eta^1(x, 4\delta, r) \leq \Delta_1(4\delta, r) < \frac{\epsilon}{12n}.$$

Since every point x is covered by at most three arcs θ_j^l , we recall (A.9) and conclude that

$$\begin{aligned} \sum_l \mu_\eta[S(\theta_j^l)] &\leq \sum_l 2\mu_\eta[S''(\theta_j^l)] \\ &= \sum_x \sum_l 2\mu_\eta\{\omega : S(\theta_j^l) \mid x(\omega) = x\} \mu_\eta\{\omega : x(\omega) = x\} \\ &\leq \sum_x \left(2 \cdot 3 \cdot \frac{\epsilon}{12n}\right) \mu_\eta\{\omega : x(\omega) = x\} \\ &\leq \frac{\epsilon}{2n} \sum_x \mu_\eta\{\omega : x(\omega) = x\} \leq \frac{\epsilon}{2n}, \end{aligned}$$

so estimate (A.7) and the lemma follow. \square

APPENDIX B: MULTI-SCALE BOUND ON THE FOUR-ARM EVENT ON \mathbb{Z}^2 (BY CHRISTOPHE GARBAN)

In this Appendix, we give a proof of assumption (1.3) for critical bond percolation on the square lattice. Namely:

LEMMA B.1. *For critical bond percolation on the square lattice with mesh $|\eta|$, there is a positive ϵ such that the probability of a four arm event satisfies*

$$(B.1) \quad P_{\eta}^4(z, r, R) \leq \text{const} \left(\frac{r}{R} \right)^{1+\epsilon},$$

whenever $|\eta| < r$.

The case $r = |\eta|$ can be extracted from [22] as well as [3] or [32]. Note that the above multi-scale generalization is not a consequence of the so-called “quasi-multiplicativity” property. If one wanted to use this property, the problem would more or less boil down to the *existence* of the *four-arm* critical exponent. Since its existence is still open, we believe that Lemma B.1 cannot be extracted directly from the above mentioned papers.

In the papers [3, 32], the general idea which enables one to prove that points are unlikely to be pivotal is the observation that a macroscopic crossing event for critical percolation is asymptotically uncorrelated with the Majority Boolean function defined on the same set of bits. This asymptotic uncorrelation can be understood using an exploration path which will compute the percolation event while giving little information on the Majority event. A careful analysis then shows that this decorrelation is possible only if points are unlikely to be pivotal for the percolation event. This heuristical program has been carried out in a very convenient manner in [30].

The present proof has the same flavor except that one now looks at the correlation of a percolation event in a domain of diameter $O(R)$ with a *two-layers majority function*: namely the bits of this Majority function are now indexed by the $O(R^2/r^2)$ r -squares included in the domain and for each such r -square, its corresponding bit ± 1 depends on how “connected” the percolation configuration is within this square. This setup allows to “interpolate” the above program from the macroscopic scale R to the mesoscopic scale r .

PROOF OF LEMMA B.1. We leave out a few details, which are easy to fill in for readers familiar with the applications of the RSW theory, similar to [13, 22, 36].

After rescaling and, if needed, changing the radii by bounded factors, we can assume without loss of generality that the lattice mesh is 1 and r, R are positive integers.

Denote by Q an $R \times R$ square, and cut the concentric $\frac{1}{3}Q$ square into $r \times r$ squares, denoted by Q_j , with $j = 1, \dots, (\frac{R}{3r})^2$. Denote by $X = 2 \cdot 1_{\square_Q} - 1$ the “crossing random variable” equal to 1 when Q is crossed and -1 otherwise. Given a percolation configuration ω , we say that Q_j is pivotal for X , if altering ω so that all the bonds in Q_j are open, and so that all the bonds in Q_j are closed yields two different values of X . Note that Q_j is pivotal for X if and only if then there are two open arms connecting ∂Q_j to $\partial_0 Q$ and $\partial_2 Q$, and two dual closed arms connecting ∂Q_j to $\partial_1 Q$ and $\partial_3 Q$. By arm separation properties similar to ones used in [36],

$$(B.2) \quad P^4(r, R) \asymp \mathbf{P}[Q_j \text{ pivotal for } X],$$

and so it is sufficient to estimate the latter probability, or its sum over all Q_j 's.

Let S_j be an annulus $(1 - \delta)Q_j \setminus (1 - 2\delta)Q_j$ where αQ_j denotes a copy shrank of Q_j by a factor of α (if r is not exactly divisible by α , the copy shrank is understood modulo integer parts). The fixed parameter $\delta > 0$ will be chosen later. Let S_j^* denote the same annulus but on the dual lattice: that is, $S_j^* = S_j + (1/2, 1/2)$.

Define a random variable C_j equal to:

- 1, if there is an open circuit in S_j and no dual-closed circuit in S_j^* ,
- -1 , if there is a dual-closed circuit in S_j^* and no open circuit in S_j ,
- 0, otherwise.

Note that by symmetry and by the RSW theory, up to the constant depending on δ ,

$$(B.3) \quad \mathbf{E}[C_j] = 0, \quad \mathbf{E}[C_j^2] \asymp 1.$$

Now, let us argue that if $\delta > 0$ is chosen small enough, one has

$$(B.4) \quad \mathbf{E}[XC_j] \asymp \mathbf{P}[Q_j \text{ pivotal for } X] \quad [\asymp P^4(r, R)].$$

Indeed, one has

$$(B.5) \quad \mathbf{E}[XC_j] = \mathbf{P}[Q_j \text{ pivotal for } X] \mathbf{E}[XC_j \mid Q_j \text{ pivotal for } X],$$

since, conditioned on the event that Q_j is not pivotal for X , C_j is independent of X and is such that its (conditional) expectation is still 0. Therefore, it remains to bound from below $\mathbf{E}[XC_j \mid Q_j \text{ pivotal for } X]$. Let $\rho = \rho(\delta) > 0$ be such that $\mathbf{P}[C_j = 1] = \mathbf{P}[C_j = -1] > \rho$.

$$\begin{aligned} & \mathbf{E}[XC_j \mid Q_j \text{ pivotal for } X] \\ &= \mathbf{P}[C_j = 1] \mathbf{E}[X \mid C_j = 1 \text{ and } Q_j \text{ pivotal for } X] \\ & \quad - \mathbf{P}[C_j = -1] \mathbf{E}[X \mid C_j = -1 \text{ and } Q_j \text{ pivotal for } X]. \end{aligned}$$

Now, again by arm separation properties similar to ones used in [32, 36], and using the important fact that the event $\{C_j = 1\}$ is increasing which enables to use FKG for the given conditional law inside Q_j , one can see that, if δ is chosen small enough,

$$\mathbf{E}[X \mid C_j = 1 \text{ and } Q_j \text{ pivotal for } X] > \frac{1}{4}.$$

One has the opposite bound for the term conditioned on $\{C_j = -1\}$. All together this gives

$$\mathbf{E}[XC_j \mid Q_j \text{ pivotal for } X] > \frac{\rho}{2},$$

which implies the desired estimate (B.4) if δ is chosen to be small enough.

Consider Q with open boundary conditions on the side $\partial_0 Q$ and dual closed on the complementary three sides. Let γ be the interface running between the two ends of the side $\partial_0 Q$ and separating the open cluster rooted on it from the dual closed cluster rooted on the three other sides.

Denote by Y_j the event that γ intersects Q_j , as well as its indicator function. Note that, by the RSW theory,

$$(B.6) \quad \mathbf{E}[Y_j] \lesssim \left(\frac{r}{R}\right)^{2\epsilon}$$

for some positive ϵ . Note that the interface γ drawn from $Q(0, 0)$ until some stopping time depends only on the percolation configuration in the immediate neighborhood of the drawn part. Drawing γ until the first time it hits ∂Q_j , we conclude that Y_j is independent from the inside of Q_j and hence from C_j .

Note also that, if Q_j is pivotal for X , then there are four alternating arms connecting ∂Q_j to four sides of Q which forces the interface γ to intersect Q_j , and Y_j to occur. Therefore, similarly as for the estimate (B.5) above, one can check that

$$(B.7) \quad \mathbf{E}[XC_j] = \mathbf{E}[XC_j Y_j],$$

which combined with (B.4) gives

$$(B.8) \quad \mathbf{P}[Q_j \text{ pivotal for } X] \asymp \mathbf{E}[XC_j Y_j].$$

Summing over all Q_j 's, we use the Cauchy–Schwarz inequality to write

$$(B.9) \quad \begin{aligned} \sum_j \mathbf{E}[XC_j Y_j] &= \mathbf{E}\left[X \sum_j C_j Y_j\right] \leq \sqrt{\mathbf{E}[X^2] \cdot \mathbf{E}\left[\left(\sum_j C_j Y_j\right)^2\right]} \\ &= \sqrt{\mathbf{E}[1] \cdot \mathbf{E}\left[\sum_{i,j} C_i Y_i C_j Y_j\right]} = \sqrt{\sum_{i,j} \mathbf{E}[C_i Y_i C_j Y_j]}. \end{aligned}$$

In the last sum, nondiagonal terms vanish. Indeed, let $i \neq j$ and stop the interface γ the first time it touched both squares Q_i and Q_j , or when it ends.

Denote by Q_k the first square of Q_i, Q_j to be hit (or Q_i if none was) and by Q_l the other one, so that (k, l) is a superposition of (i, j) . Let W be the percolation configuration in the immediate neighborhood of the interface so far, and in Q_k .

Then W determines Y_i, Y_j, C_k , while being independent of C_l . Thus by the identity in (B.3),

$$\mathbf{E}[C_i Y_i C_j Y_j | W] = \mathbf{E}[Y_i Y_j C_k | W] \cdot \mathbf{E}[C_l] = 0.$$

By the total expectation formula we conclude that, for $i \neq j$,

$$\mathbf{E}[C_i Y_i C_j Y_j] = 0.$$

Thus, we can continue (B.9), leaving only the diagonal terms, and use (B.6) and (B.3) to write

$$\begin{aligned} \dots &= \sqrt{\sum_j \mathbf{E}[C_j^2 Y_j^2]} = \sqrt{\sum_j \mathbf{E}[C_j^2] \cdot \mathbf{E}[Y_j^2]} \\ &\lesssim \sqrt{\sum_j 1 \cdot \left(\frac{r}{R}\right)^{2\epsilon}} = \left(\frac{r}{R}\right)^{\epsilon-1}. \end{aligned}$$

Combining this with (B.2), (B.8), (B.9), we conclude

$$\begin{aligned} P^4(r, R) &\lesssim \left(\frac{3r}{R}\right)^2 \sum_j \mathbf{P}[Q_j \text{ pivotal for } X] \\ &\lesssim \left(\frac{r}{R}\right)^2 \left(\frac{r}{R}\right)^{\epsilon-1} = \left(\frac{r}{R}\right)^{1+\epsilon}, \end{aligned}$$

proving the lemma. \square

Acknowledgments. We are thankful to Christophe Garban and Wendelin Werner for useful conversations. The second author is grateful to Hugo Duminil-Copin, Clément Hongler, Gábor Pete, Jeff Steif, the referee, and especially Christophe Garban and Boris Tsirelson for their comments on the manuscript.

REFERENCES

- [1] AIZENMAN, M. and BURCHARD, A. (1999). Hölder regularity and dimension bounds for random curves. *Duke Math. J.* **99** 419–453. [MR1712629](#)
- [2] ASTALA, K., PRAUSE, I. and SMIRNOV, S. (2010). Holomorphic motions and harmonic measure. Preprint.
- [3] BENJAMINI, I., KALAI, G. and SCHRAMM, O. (1999). Noise sensitivity of Boolean functions and applications to percolation. *Publ. Math. Inst. Hautes Études Sci.* **90** 5–43. [MR1813223](#)
- [4] BOLLOBÁS, B. and RIORDAN, O. (2006). *Percolation*. Cambridge Univ. Press, New York. [MR2283880](#)
- [5] BOLLOBÁS, B. and RIORDAN, O. (2010). Percolation on self-dual polygon configurations. Preprint. Available at [arXiv:1001.4674](#).
- [6] CAMIA, F., FONTES, L. R. G. and NEWMAN, C. M. (2006). Two-dimensional scaling limits via marked nonsimple loops. *Bull. Braz. Math. Soc. (N.S.)* **37** 537–559. [MR2284886](#)

- [7] CAMIA, F. and NEWMAN, C. M. (2006). Two-dimensional critical percolation: The full scaling limit. *Comm. Math. Phys.* **268** 1–38. [MR2249794](#)
- [8] CARDY, J. L. (1992). Critical percolation in finite geometries. *J. Phys. A* **25** L201–L206. [MR1151081](#)
- [9] DUDLEY, R. M. (1989). *Real Analysis and Probability*. Wadsworth & Brooks/Cole Advanced Books & Software, Pacific Grove, CA. [MR0982264](#)
- [10] FEDERER, H. (1969). *Geometric Measure Theory. Die Grundlehren der Mathematischen Wissenschaften* **153**. Springer, New York. [MR0257325](#)
- [11] FELL, J. M. G. (1962). A Hausdorff topology for the closed subsets of a locally compact non-Hausdorff space. *Proc. Amer. Math. Soc.* **13** 472–476. [MR0139135](#)
- [12] GARBAN, C., PETE, G. and SCHRAMM, O. (2010). The scaling limit of the Minimal Spanning Tree—A preliminary report. In *XVIIth International Congress on Mathematical Physics, Prague, 2009* (P. Exner, ed.) 475–480. World Scientific, Singapore.
- [13] GARBAN, C., PETE, G. and SCHRAMM, O. (2010). The Fourier spectrum of critical percolation. *Acta Math.* **205** 19–104. [MR2736153](#)
- [14] GARBAN, C., PETE, G. and SCHRAMM, O. (2010). Pivotal, cluster and interface measures for critical planar percolation. Preprint. Available at [arXiv:1008.1378](#).
- [15] GARBAN, C., ROHDE, S. and SCHRAMM, O. (2011). Continuity of the SLE trace in simply connected domains. *Israel J. Math.* To appear. Available at [arXiv:0810.4327](#).
- [16] GRIMMETT, G. (1999). *Percolation*, 2nd ed. *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]* **321**. Springer, Berlin. [MR1707339](#)
- [17] KELLEY, J. L. (1975). *General Topology. Graduate Texts in Mathematics* **27**. Springer, New York. [MR0370454](#)
- [18] KEMPPAINEN, A. and SMIRNOV, S. (2010). Conformal invariance in random cluster models. III. Full scaling limit. Unpublished manuscript.
- [19] KENYON, R. (2009). Lectures on dimers. In *Statistical Mechanics. IAS/Park City Mathematics Series* **16** 191–230. Amer. Math. Soc., Providence, RI. [MR2523460](#)
- [20] KESTEN, H. (1980). The critical probability of bond percolation on the square lattice equals $\frac{1}{2}$. *Comm. Math. Phys.* **74** 41–59. [MR0575895](#)
- [21] KESTEN, H. (1982). *Percolation Theory for Mathematicians. Progress in Probability and Statistics* **2**. Birkhäuser, Boston, MA. [MR0692943](#)
- [22] KESTEN, H. (1987). Scaling relations for 2D-percolation. *Comm. Math. Phys.* **109** 109–156. [MR0879034](#)
- [23] LANGLANDS, R., POULIOT, P. and SAINT-AUBIN, Y. (1994). Conformal invariance in two-dimensional percolation. *Bull. Amer. Math. Soc. (N.S.)* **30** 1–61. [MR1230963](#)
- [24] LANGLANDS, R. P. (2005). The renormalization fixed point as a mathematical object. In *Twenty Years of Białowieża: A Mathematical Anthology. World Sci. Monogr. Ser. Math.* **8** 185–216. World Sci. Publ., Hackensack, NJ. [MR2181552](#)
- [25] LANGLANDS, R. P. and LAFORTUNE, M. A. (1994). Finite models for percolation. In *Representation Theory and Analysis on Homogeneous Spaces (New Brunswick, NJ, 1993). Contemporary Mathematics* **177** 227–246. Amer. Math. Soc., Providence, RI. [MR1303608](#)
- [26] LAWLER, G. F., SCHRAMM, O. and WERNER, W. (2002). One-arm exponent for critical 2D percolation. *Electron. J. Probab.* **7** 13 pp. (electronic). [MR1887622](#)
- [27] MATTILA, P. (1995). *Geometry of Sets and Measures in Euclidean Spaces: Fractals and Rectifiability. Cambridge Studies in Advanced Mathematics* **44**. Cambridge Univ. Press, Cambridge. [MR1333890](#)
- [28] NIENHUIS, B. (1984). Critical behavior of two-dimensional spin models and charge asymmetry in the Coulomb gas. *J. Stat. Phys.* **34** 731–761. [MR0751711](#)
- [29] NOLIN, P. (2008). Near-critical percolation in two dimensions. *Electron. J. Probab.* **13** 1562–1623. [MR2438816](#)

- [30] O'DONNELL, R. and SERVEDIO, R. A. (2007). Learning monotone decision trees in polynomial time. *SIAM J. Comput.* **37** 827–844 (electronic). [MR2341918](#)
- [31] RUDIN, W. (1973). *Functional Analysis*. McGraw-Hill, New York. [MR0365062](#)
- [32] SCHRAMM, O. and STEIF, J. E. (2010). Quantitative noise sensitivity and exceptional times for percolation. *Ann. of Math. (2)* **171** 619–672. [MR2630053](#)
- [33] SHEFFIELD, S. (2009). Exploration trees and conformal loop ensembles. *Duke Math. J.* **147** 79–129. [MR2494457](#)
- [34] SMIRNOV, S. (2001). Critical percolation in the plane. Preprint. Available at [arXiv:0909.4499](#).
- [35] SMIRNOV, S. (2001). Critical percolation in the plane: Conformal invariance, Cardy's formula, scaling limits. *C. R. Math. Acad. Sci. Paris Sér. I* **333** 239–244. [MR1851632](#)
- [36] SMIRNOV, S. and WERNER, W. (2001). Critical exponents for two-dimensional percolation. *Math. Res. Lett.* **8** 729–744. [MR1879816](#)
- [37] STEIF, J. E. (2009). A survey of dynamical percolation. In *Fractal Geometry and Stochastics IV. Proceedings of the 4th Conference, Greifswald, Germany, September 8–12, 2008. Progress in Probability* **61** 145–174. Birkhäuser, Basel.
- [38] TSIRELSON, B. (2004). Nonclassical stochastic flows and continuous products. *Probab. Surv.* **1** 173–298 (electronic). [MR2068474](#)
- [39] TSIRELSON, B. (2004). Scaling limit, noise, stability. In *Lectures on Probability Theory and Statistics. Lecture Notes in Math.* **1840** 1–106. Springer, Berlin. [MR2079671](#)
- [40] TSIRELSON, B. (2005). Percolation, boundary, noise: An experiment. Preprint. Available at [Arxiv:math/0506269](#).
- [41] TSIRELSON, B. S. and VERSHIK, A. M. (1998). Examples of nonlinear continuous tensor products of measure spaces and non-Fock factorizations. *Rev. Math. Phys.* **10** 81–145. [MR1606855](#)
- [42] WERNER, W. (2009). Lectures on two-dimensional critical percolation. In *Statistical Mechanics. IAS/Park City Mathematics Series* **16** 297–360. Amer. Math. Soc., Providence, RI. [MR2523462](#)
- [43] WIERMAN, J. C. (1981). Bond percolation on honeycomb and triangular lattices. *Adv. in Appl. Probab.* **13** 298–313. [MR0612205](#)

S. SMIRNOV
 SECTION DE MATHÉMATIQUES
 UNIVERSITÉ DE GENÈVE
 2-4 RUE DU LIÈVRE, CP 64
 1211 GENÈVE 4
 SWITZERLAND
 E-MAIL: stanislav.smirnov@unige.ch
 URL: <http://www.unige.ch/~smirnov>

C. GARBAN
 CNRS, UMPA
 ENS DE LYON
 46, ALLÉE D'ITALIE
 69364 LYON CEDEX 07
 FRANCE
 E-MAIL: christophe.garban@umpa.ens-lyon.fr
 URL: <http://www.umpa.ens-lyon.fr/~cgarban/>