

An invariance principle for ergodic scale-free random environments

by

EWAIN GWYNNE

*University of Chicago
Chicago, IL, U.S.A.*

JASON MILLER

*University of Cambridge
Cambridge, U.K.*

SCOTT SHEFFIELD

*Massachusetts Institute of Technology
Cambridge, MA, U.S.A.*

Contents

1. Introduction	304
1.1. Basic definitions	304
1.2. Overview of results	310
1.3. Random walk on the faces of an embedded lattice	312
1.4. Variants of Theorem 1.5	316
1.5. Why $d=2$ is special	323
1.6. Outline	324
Basic notation	327
2. Existence and sublinearity of the corrector	327
2.1. Dyadic systems	327
2.2. Ergodic averages using dyadic systems	331
2.3. Existence of the corrector	335
2.4. Sublinearity of the corrector	340
2.5. Random walk on \mathcal{H} is recurrent	346
3. Convergence of random walk to Brownian motion	349
3.1. Ergodic averages for simple random walk	349
3.2. Proof of Theorem 1.16	355
3.3. Convergence is uniform in the starting point	359
3.4. Proof of Theorems 1.5, 1.11, and 1.13	362
4. Applications to planar maps and Liouville quantum gravity	365
Appendix A. Specific Dirichlet energy	367
Appendix B. Equivalence of translation invariance modulo scaling definitions	369
Appendix C. No macroscopic edges in Theorem 1.13	374
References	382

1. Introduction

1.1. Basic definitions

The goal of this paper is to prove that simple random walks on certain random planar lattices have Brownian motion as a scaling limit. We will consider random planar lattices whose laws are “translation invariant modulo scaling” in a sense we will define, but not necessarily translation invariant in the usual sense. Along the way, we will also explain how certain familiar tools (such as the ergodic theorem) can be adapted to this setting. In this subsection, we define and discuss three terms: embedded lattice (with conductances), translation invariant modulo scaling, and ergodic modulo scaling. The impatient reader can skim this subsection quickly on a first read (referring back later as needed) and proceed to the results overview in §1.2.

Definition 1.1. An *embedded lattice* \mathcal{M} is an infinite, locally finite, planar, undirected graph (multiple edges and self-loops allowed) embedded into \mathbb{C} in such a way that each edge is a simple (unparameterized) curve with zero Lebesgue measure, the edges intersect only at their endpoints, each compact subset of \mathbb{C} intersects at most finitely many edges of \mathcal{M} , and each connected component of $\mathbb{C} \setminus \mathcal{M}$ is compact. For an embedded lattice \mathcal{M} , we write \mathcal{EM} for the set of edges of \mathcal{M} and \mathcal{FM} for the set of faces of \mathcal{M} , i.e., the set of closures of connected components of $\mathbb{C} \setminus \mathcal{M}$. An *embedded lattice with conductances* is an embedded lattice \mathcal{M} together with a function $\mathbf{c} = \mathbf{c}_{\mathcal{M}}: \mathcal{EM} \rightarrow (0, \infty)$. By a slight abuse of notation, we write \mathcal{M} instead of $(\mathcal{M}, \mathbf{c})$. We write \mathbb{M} for the space of all embedded lattices with conductances.

A *finite embedded lattice* is defined in the same manner as above, except that the graph is required to be finite and \mathbb{C} is replaced by \bar{U} , where U is a bounded sub-domain of \mathbb{C} .

For $\mathcal{M} \in \mathbb{M}$ and a set $A \subset \mathbb{C}$, we write $\mathcal{M}(A)$ for the embedded lattice consisting of all of the vertices and edges of \mathcal{M} which lie on the boundaries of faces of \mathcal{M} which intersect A (the conductances are the same). We endow \mathbb{M} with a metric in which \mathcal{M}' and \mathcal{M} are “close” when they look very similar on a large finite ball. There are many ways to do this, but we fix the following for concreteness: for embedded lattices \mathcal{M} and \mathcal{M}' , with conductances \mathbf{c} and \mathbf{c}' , we define their distance by

$$\mathbf{d}^{\text{EL}}(\mathcal{M}, \mathcal{M}') := \int_0^\infty e^{-r} \wedge \inf_{f_r} \left\{ \sup_{z \in \mathbb{C}} |z - f_r(z)| + \max_{e \in \mathcal{EM}(B_r(0))} |\mathbf{c}(e) - \mathbf{c}'(f_r(e))| \right\} dr \quad (1.1)$$

where the infimum is over all homeomorphisms $f_r: \mathbb{C} \rightarrow \mathbb{C}$ such that f_r takes each vertex (resp. edge) of $\mathcal{M}(B_r(0))$ to a vertex (resp. edge) of $\mathcal{M}'(B_r(0))$, and f_r^{-1} does the same with \mathcal{M} and \mathcal{M}' reversed. Note that the integrand is equal to e^{-r} for each $r > 0$ such

that no such homeomorphism exists. When we speak of random embedded lattices, we implicitly assume that their laws are defined w.r.t. the Borel σ -algebra associated with the topology defined by (1.1).

It is a classical problem to determine whether a random walk on a random embedded lattice converges to Brownian motion (or to some other limiting process). There is a vast literature on this question, falling under the heading of *random walk in random environment*. See, e.g., [KLO], [BF], [Bi], [Ku] for recent surveys. Within this literature, the term *random conductance model* is used to describe walks which choose an edge to traverse (starting from the current vertex) with probability proportional to the conductance assigned to that edge. These are the types of walks we consider in this paper.

Many existing results show the convergence (in the *quenched* sense) of random walk to Brownian motion for random embedded lattices that are *translation invariant* (stationary in law with respect to spatial translations), i.e., $\mathcal{M} - z \stackrel{d}{=} \mathcal{M}$ for each $z \in \mathbb{C}$ or for each $z \in \mathbb{Z}^2$, in the discrete case. Examples of such lattices include periodic lattices such as \mathbb{Z}^2 with conductances assigned in a stationary way; the infinite cluster of supercritical percolation on such a lattice [BB]; and various graph structures on stationary point processes in \mathbb{C} . Much of this work is based on the “environment seen from the particle” approach initiated in [KV], [DFGW]. Brownian motion convergence results for random walks in random environments are often called *invariance principles*.

In this paper, we will be interested in random embedded lattices whose laws are in some sense only stationary with respect to translations *modulo spatial scaling*, so that there is no universal length scale between different points. To get some intuition about what this means, suppose the embedded lattice \mathcal{M} is “fractal” in nature (like the embedded lattice illustrated in Figure 1, right), so that there are some regions in the plane where the vertices of \mathcal{M} are much more concentrated than others. To quantify this effect, we can consider a point z sampled uniformly from Lebesgue measure on a large ball $B_r(0)$, independently from \mathcal{M} , and look at the law of the diameter of the face of \mathcal{M} which contains z . We are interested in the case when the law of this diameter does not have a weak limit as $r \rightarrow \infty$, so that the possible face sizes become more and more spread out at larger scales. In other words, there is no “typical” size for a face.

If \mathcal{M} has no typical face size in the above sense, then \mathcal{M} cannot be stationary with respect to spatial translations. However, it is still possible for such a lattice \mathcal{M} to satisfy a weaker analog of translation invariance, which can be described as “translation invariance modulo scaling”.

Before explaining what this means, let us first define embedded lattices \mathcal{M}_1 and \mathcal{M}_2 to be *equivalent modulo scaling* if $\mathcal{M}_1 = C\mathcal{M}_2$ for some real $C > 0$, where $C\mathcal{M}$ denotes the image of \mathcal{M} under the map $z \mapsto Cz$ (the conductances are unchanged, i.e.,




Figure 1. *Left*: The subgraph of an embedded lattice which intersects a square in \mathbb{C} . *Right*: A numerical simulation of random planar map (in particular, the $\sqrt{2}$ -mated-CRT map studied in [GMS3]) embedded into \mathbb{C} via the Tutte embedding, which means that the two coordinates of the embedding function are discrete harmonic. Note that the faces have very different sizes.

$\mathfrak{c}_{\mathcal{C}\mathcal{M}}(Ce) = \mathfrak{c}_{\mathcal{M}}(e)$). Note that if \mathcal{M}_1 and \mathcal{M}_2 are equivalent modulo scaling, then the constant C is determined by \mathcal{M}_1 and \mathcal{M}_2 : indeed, C is the ratio of the diameters of the origin-containing faces in the two lattices.

If \mathcal{M}_1 and \mathcal{M}_2 are both random, we say their *laws agree modulo scaling* if they assign the same probabilities to any scale-invariant Borel measurable subset of \mathbb{M} (or, equivalently, if there is a random $C > 0$ such that $C\mathcal{M}_1$ agrees in law with \mathcal{M}_2). We will mostly prove statements about \mathcal{M} that would obviously remain true if its law were replaced by a different-but-equivalent-modulo-scaling law. Note that the law of a random embedded lattice \mathcal{M} is equivalent modulo scaling to a unique law w.r.t. which the origin-containing face a.s. has area one; thus, for most of this paper, there is no lost generality if one assumes that the law of \mathcal{M} is such that the origin-containing face a.s. has area one.⁽¹⁾

Definition 1.2 below describes what it means for a law to be translation invariant modulo scaling. On a first read, it is enough to internalize *any one* of the equivalent conditions in Definition 1.2. Most of the proofs in this paper will make use of a fifth

⁽¹⁾ There are two (subtly different) ways to formulate “modulo scaling” analogs of probability measures on the space \mathbb{M} of embedded lattices. The first approach is to define the probability measure only for the σ -algebra of Borel measurable (w.r.t. (1.1)) events that are *scale invariant* (i.e., that are unions of modulo scaling equivalence classes). Such a σ -algebra contains no information about scale, so a “sample” from such a measure could be interpreted as a random modulo-scaling equivalence class. The second approach is to require that the probability measure be defined for all Borel measurable subsets of \mathbb{M} , but to prove statements that *only depend* on the restriction of the measure to the scale-invariant σ -algebra. We take the latter approach (which has the minor cosmetic advantage of allowing us to speak of “random embedded lattices” rather than “random equivalence classes of embedded lattices”), but the former would work as well.

condition, defined in terms of a so-called *dyadic system* (see §2.1), which is implied by any of the four conditions listed below. Conceptual motivation for the conditions in Definition 1.2 is given *just after* the full statement of Definition 1.2.

Definition 1.2. A random embedded lattice \mathcal{M} , endowed with conductance function \mathfrak{c} , is *translation invariant modulo scaling* if a.s. all of its faces are compact and it satisfies one of the following equivalent conditions.

(1) *Attainability as Lebesgue-centered infinite volume limit (modulo scaling)*: There is a sequence $\{\mathcal{M}_n\}_{n \in \mathbb{N}}$ of random finite embedded lattices such that the following is true.

- (a) The union of the closures of the bounded faces of each \mathcal{M}_n is either a closed disk or a closed square (as in Figure 1).
- (b) For each n , conditional on \mathcal{M}_n let z_n be sampled uniformly from Lebesgue measure on the union of faces in \mathcal{M}_n . There is a random $C_n > 0$ (depending on z_n and \mathcal{M}_n) such that $C_n(\mathcal{M}_n - z_n)$ converges in law to \mathcal{M} as $n \rightarrow \infty$ with respect to the metric (1.1).

(2) *Attainability as spatial average of sample instance (modulo scaling)*: There is a way of associating with \mathcal{M} a (possibly random) sequence U_n of origin-containing sets, each of which is a square or a disk, such that $\bigcup_{n=1}^{\infty} U_n = \mathbb{C}$ a.s., and the following is true: suppose that, given \mathcal{M} and U_n , a point z_n is sampled from Lebesgue measure on U_n . Then, there is a random $C_n > 0$ (depending on z_n , U_n , and \mathcal{M}) such that the random embedded lattices $C_n(\mathcal{M} - z_n)$ converge in law to \mathcal{M} as $n \rightarrow \infty$.

(3) *Invariance under repositioning origin within an origin-containing “block”*: Suppose we associate with each embedded lattice \mathcal{M} a random partition $\mathcal{B}(\mathcal{M})$ of the plane into a countable collection of *blocks*, where each block is a closed set with finite diameter and positive area, and each $z \in \mathbb{C}$ a.s. belongs to exactly one block⁽²⁾ (we allow $\mathcal{B}(\mathcal{M})$ to depend on \mathcal{M} and possibly on additional randomness; see examples after this definition below). Suppose further that the procedure for constructing \mathcal{B} from \mathcal{M} commutes with translations and dilations in the sense that

$$(C(\mathcal{M} - z), C(\mathcal{B}(\mathcal{M}) - z)) \stackrel{d}{=} (C(\mathcal{M} - z), \mathcal{B}(C(\mathcal{M} - z)))$$

for each $C > 0$ and $z \in \mathbb{C}$. For each possible choice of $\mathcal{B}(\mathcal{M})$, the following is true. Conditional on \mathcal{M} and $\mathcal{B}(\mathcal{M})$, let z be sampled uniformly from Lebesgue measure on the

⁽²⁾ One can define a topology, hence a σ -algebra, on block decompositions in a similar manner to (1.1), but with the f_r 's required to preserve blocks instead of vertices and edges. See also (1.15). Also note that a block decomposition differs from a cell configuration (Definition 1.15) in that blocks are not required to be connected.

origin-containing block. Then, for some random $C > 0$ (depending on \mathcal{M} and z), the embedded lattice $C(\mathcal{M} - z)$ agrees in law with \mathcal{M} .

(4) *Mass transport*: Consider the set of non-negative measurable functions $F: \mathbb{M} \times \mathbb{C}^2 \rightarrow [0, \infty)$ on the space of embedded lattices with two marked points with the property that F is covariant with respect to dilations and translations of the plane in the sense that

$$F(C(\mathcal{M} - z), C(w_0 - z), C(w_1 - z)) = C^{-2} F(\mathcal{M}, w_0, w_1) \quad \text{for all } C > 0 \text{ and } z \in \mathbb{C}. \quad (1.2)$$

For every such function F ,

$$\mathbb{E} \left[\int_{\mathbb{C}} F(\mathcal{M}, 0, w) dw \right] = \mathbb{E} \left[\int_{\mathbb{C}} F(\mathcal{M}, w, 0) dw \right]. \quad (1.3)$$

It is not hard to show that the four conditions listed in Definition 1.2 are equivalent for embedded lattices with compact faces. We will prove that this is the case in Appendix B (see also Lemma 2.3).

The infinite volume limit condition (1) of Definition 1.2 is natural from the point of view of random planar maps: if $\{\mathcal{M}_n\}_{n \in \mathbb{N}}$ is any sequence of random planar maps embedded in the unit disk in some way (e.g., via a circle packing procedure), then we may choose z uniformly on the unit disk and consider the random embedded lattices $\mathcal{M}_n - z$. If $\mathcal{M}_n - z$ (or some subsequence) converges in law modulo scaling to an unbounded embedded lattice \mathcal{M} , then \mathcal{M} satisfies translation invariance modulo scaling.

The spatial averaging condition (2) says that although sizes of the faces of \mathcal{M} near two different points of \mathbb{C} can differ dramatically, the shapes of these faces and their adjacency structure are such that the local behavior near zero (modulo scaling) looks like the average behavior over a sufficiently large set. Note that when we speak of the law of $C_n(\mathcal{M} \setminus z_n)$, we mean the overall law not the conditional law given \mathcal{M} . (In other settings, analogous statements involving conditional laws are sometimes used to define ergodicity; but for the moment we are only defining translation invariance modulo scaling, not ergodicity modulo scaling.)

For a simple example of a block decomposition satisfying the condition (3) given in Definition 1.2, one may let the blocks of \mathcal{M} be the faces of \mathcal{M} . Another option is to “mark” each vertex of \mathcal{M} independently with some small probability p , and let the blocks be the cells of the Voronoi tessellation corresponding to the marked vertices. §2.1 will introduce a particularly convenient block decomposition defined in terms of a so-called *dyadic system* of squares, which we will use throughout most of our proofs.

The mass transport condition (4) is a continuum analog of the discrete mass transport principle used to define the concept of *unimodularity* for a random rooted graph [AL],

which says that the root vertex is *some* sense “typical” w.r.t. counting measure on vertices. The analogous statement in our story is that the origin-containing face is “typical” w.r.t. the measure that assigns each face its Lebesgue measure. We interpret F itself as a “mass transport” rule in which

$$\int_A \int_B F(\mathcal{M}, w, z) dz dw$$

is the amount of mass “transported” from the set A to the set B . The C^{-2} in (1.2) ensures that neither integral in (1.3) changes if we replace \mathcal{M} by $C\mathcal{M}$. It also implies that, for $C > 0$, we have

$$\int_{CA} \int_{CB} F(C\mathcal{M}, w, z) dz dw = C^2 \int_A \int_B F(\mathcal{M}, w, z) dz dw,$$

so that in a sense the total mass transported scales like area. For a concrete example, suppose that every face H of \mathcal{M} is assigned a *target* face H' , and let

$$F(\mathcal{M}, w, z) = \frac{1_{z \in H'_w}}{\text{Area}(H'_w)},$$

where H_w is the face containing w . Then, F describes a transport rule that sends an $\text{Area}(H)$ amount of mass from H to H' . Informally, the transport rule takes all of the Lebesgue measure contained in H and spreads it uniformly over H' . And (1.3) states that the expected amount of mass flowing into (an infinitesimal neighborhood of) the origin is equal to the expected amount flowing out.

The following is one of two main conditions we need to impose to show that the simple random walk on \mathcal{M} converges to Brownian motion with a deterministic covariance matrix.

Definition 1.3. We say the law of \mathcal{M} is *ergodic modulo scaling* if it is translation invariant modulo scaling and it assigns probability zero or one to every Borel measurable event which is invariant under translations/dilations of the form $\mathcal{M} \mapsto C(\mathcal{M} - z)$ for $C > 0$ and $z \in \mathbb{C}$.

We remark that if \mathcal{I} is the σ -algebra of events invariant under translations/dilations of the form $\mathcal{M} \mapsto C(\mathcal{M} - z)$, and the law of \mathcal{M} is translation invariant modulo scaling, then the regular conditional law of \mathcal{M} given \mathcal{I} (which exists because \mathbb{M} is a separable metric space) is a.s. ergodic modulo scaling. In other words, the law of \mathcal{M} is a weighted average of laws that are ergodic modulo scaling.

The other main condition involved in our theorem statements is a bound on some notion of “expected Dirichlet energy per unit area”. Let us now briefly explain what this means.

Definition 1.4. For a graph G with a conductance function $\mathfrak{c}: \mathcal{E}(G) \rightarrow (0, \infty)$ and a function $f: \mathcal{V}(G) \rightarrow \mathbb{C}$, we define its *Dirichlet energy* to be the sum over unoriented edges

$$\text{Energy}(f; G) := \sum_{\{x,y\} \in \mathcal{E}(G)} \mathfrak{c}(x,y) |f(x) - f(y)|^2.$$

If the law of a random embedded lattice \mathcal{M} is translation invariant modulo scaling, then the *specific Dirichlet energy* is defined to be (roughly speaking) the expected amount of Dirichlet energy per area for the function which takes each vertex of \mathcal{M} to its position in \mathbb{C} . Informally, each embedded edge has an energy (equal to its conductance times the square of its diameter) and the specific Dirichlet energy is the expected energy per area of \mathcal{M} in a region containing the origin.

In fact, there are various ways to make this definition precise (depending how one chooses the origin-containing region and how each edge’s energy is “localized”). Appendix A formalizes these definitions and uses mass transport to show that they are all equivalent. One of these definitions (the first precise definition of specific Dirichlet energy we present) appears as (1.12) within the statement of Theorem 1.13 (where the main theorem hypothesis is the finiteness of the specific Dirichlet energy of both \mathcal{M} and its dual). Several other theorems in this paper require the finiteness of the expectation of different but closely related quantities (which provide upper bounds on the “Dirichlet energy per unit area”; see, e.g., condition (2) in §1.3). These alternative *finite expectation* conditions will be explained within the theorem statements themselves.

1.2. Overview of results

§1.3 and §1.4 present our main theorem statements, all of which are variants of the following: if a random embedded lattice \mathcal{M} is ergodic modulo scaling and satisfies an appropriate “finite specific Dirichlet energy” condition, then simple random walk on \mathcal{M} scales to Brownian motion modulo time parametrization in the *quenched* sense. That is, given \mathcal{M} , it is a.s. the case that simple random walk on $\varepsilon\mathcal{M}$ (appropriately time changed) converges as $\varepsilon \rightarrow 0$ to Brownian motion with some deterministic diffusion matrix.

We will prove variants of this statement for random walk on the faces (Theorem 1.5) and vertices (Theorems 1.11, 1.13, and 1.16) of \mathcal{M} with slightly different energy conditions. In Theorem 1.13 the condition is that both \mathcal{M} and a simultaneously embedded dual lattice \mathcal{M}^* satisfy a “finite specific Dirichlet energy” condition. In this case, roughly speaking, the constraint on \mathcal{M} gives an upper bound on the asymptotically homogenized *conductance*, while the constraint on \mathcal{M}^* gives an upper bound on the asymptotically homogenized *resistance* (i.e., the inverse of the conductance). Taken together, these

bounds imply that the average effective conductance or resistance in any direction is strictly between zero and infinity. Theorem 1.16 is a generalization of Theorem 1.5 that allows for some degree of non-planarity. This is the version cited in the companion paper [GMS3], as well as in [GMS2]. It is not the most general non-planar theorem that we expect to be true, but it is at least reasonably easy to state.

Theorems 1.5 and 1.11 can be derived⁽³⁾ as consequences of either Theorem 1.13 or Theorem 1.16, but it is useful to state them separately since they can be stated without introducing any extra objects besides the embedded lattice itself, so the statements are slightly simpler.

The main novelty of this paper is that we are able to treat environments that are translation invariant modulo scaling, not translation invariant in the usual sense. On the other hand, our main theorems have non-trivial content even for translation-invariant random conductance-weighted subgraphs of \mathbb{Z}^2 . In particular, our results include as special cases many existing theorems for random walk in *stationary* planar random environments (see examples in §1.3 and §1.4 below) which is not too surprising given that we borrow some of our techniques from these earlier papers.

Environments that are only translation invariant modulo scaling (not stationary in the conventional sense) arise in the study of random planar maps. In particular, the results of this paper are a key input in the companion paper [GMS3] which gives the first rigorous proof that certain embedded random planar maps converge to Liouville quantum gravity, as well as the more recent paper [GMS2] which shows that, modulo time parametrization, Brownian motion on the Brownian map is a scaling limit of simple random walks on discretizations of the Brownian map. See §4 for further discussion of these points.

Finally, we remark that there are many other works which study random walks in “fractal” environments: see [Ku], [Ba] for surveys. Recent examples of work include papers by Murugan [M], and work by Biskup, Ding, and Goswami [BDG] on random walk on \mathbb{Z}^2 with conductances determined by the exponential of the discrete Gaussian free field. The setting of the latter work differs from that of the present paper in that the *conductances* (rather than the lattice itself) are what fails to be translation invariant. We do not expect that the random walk in the setting of [BDG] converges to Brownian motion modulo time parametrization.

⁽³⁾ Theorems 1.5 and 1.11 are derived from Theorem 1.16 in §3.4; they can alternatively be derived from Theorem 1.13 by imagining one chooses the “worst possible” choice for \mathcal{M}^* given \mathcal{M} .

Acknowledgements. We thank an anonymous referee for helpful comments on an earlier version of this article. We thank the Mathematical Research Institute of Oberwolfach for its hospitality during a workshop where part of this work was completed. E. G. was partially funded by NSF grant DMS-1209044. S. S. was partially supported by NSF grants DMS-1712862 and DMS-1209044 and a Simons Fellowship with award number 306120. We thank Marek Biskup, Jean-Dominique Deuschel, Jian Ding, and Tom Hutchcroft for helpful conversations.

1.3. Random walk on the faces of an embedded lattice

We will first state and discuss a version of our main result for random walk on the faces of an embedded lattice \mathcal{M} , or equivalently, random walk on the dual planar map \mathcal{M}^* of \mathcal{M} . One convenient aspect of this approach is that Lebesgue-a.e. point of \mathbb{C} corresponds to a unique face (namely, the one containing it) and thus \mathcal{M} comes with a distinguished origin-containing face (which is a natural place to start the walk). Theorems 1.11 and 1.13 below are variants involving random walk on vertices.

For $z \in \mathbb{C}$, we write H_z for the face of \mathcal{M} containing z , chosen in some arbitrary manner if z lies on one of the edges of \mathcal{M} (the union of these edges has zero Lebesgue measure). For two faces $H, H' \in \mathcal{FM}$, we write $H \sim H'$ if H and H' share an edge, in which case we write $\mathfrak{c}(H, H')$ for the conductance of this edge (or the sum of the conductances of the edges, if there is more than one).

For $H \in \mathcal{FM}$, define the primal and dual stationary measures by

$$\pi(H) := \sum_{\substack{H' \in \mathcal{FM} \\ H' \sim H}} \mathfrak{c}(H, H') \quad \text{and} \quad \pi^*(H) := \sum_{\substack{H' \in \mathcal{FM} \\ H' \sim H}} \frac{1}{\mathfrak{c}(H, H')}. \tag{1.4}$$

Note that in the case of unit conductances, $\pi(H) = \pi^*(H)$ is the degree of H .

We will be interested in the simple random walk on the faces of a random embedded lattice \mathcal{M} with conductances \mathfrak{c} satisfying the following hypotheses.

(1) *Ergodicity modulo scaling.* \mathcal{M} is ergodic modulo scaling, i.e., \mathcal{M} satisfies one (equivalently, all) of the conditions of Definition 1.2 as well as the additional condition of Definition 1.3.

(2) *Finite expectation.* With H_0 the face containing zero, and π and π^* as in (1.4), we have

$$\mathbb{E} \left[\frac{\text{diam}(H_0)^2}{\text{Area}(H_0)} \pi(H_0) \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\frac{\text{diam}(H_0)^2}{\text{Area}(H_0)} \pi^*(H_0) \right] < \infty. \tag{1.5}$$

THEOREM 1.5. *Let \mathcal{M} be a random embedded lattice satisfying the above two hypotheses. Let $X: [0, \infty) \rightarrow \mathcal{FM}$ be the continuous-time simple random walk on the dual of \mathcal{M} (viewed as an edge-weighted graph with conductances $\mathfrak{c}(H, H')$) which spends exactly $\text{Area}(H)/\pi(H)$ units of time at each face before jumping to the next. Let*

$$\widehat{X} = \phi \circ X: [0, \infty) \rightarrow \mathbb{C},$$

where ϕ is a function that sends each face H to an (arbitrary) point of H . There is a deterministic covariance matrix Σ with $\det \Sigma \neq 0$, depending on the law of \mathcal{M} , such that a.s. as $\varepsilon \rightarrow 0$ the conditional law of

$$t \mapsto \varepsilon \widehat{X}_{t/\varepsilon^2}$$

given \mathcal{M} converges to the law of Brownian motion on \mathbb{C} started from zero with covariance matrix Σ with respect to the local uniform topology.

In many examples we are interested in, the law of the embedded lattice \mathcal{M} will satisfy some sort of rotational symmetry, which will allow us to conclude that Σ is a positive scalar multiple of the identity matrix.

The choice of time parametrization in Theorem 1.5 is somewhat arbitrary. One can obtain convergence under other parametrizations — e.g., the one where the walk spends $\text{diam}(H)^2$ units of time at each face — using Lemma 3.2 below. In general, however, we do *not* expect that the random walk on \mathcal{M} parameterized according to counting measure on its steps converges to standard Brownian motion. Indeed, if \mathcal{M} looks like the embedded lattice in the right panel of Figure 1, then the random walk on \mathcal{M} will travel much faster through some regions than others. If \mathcal{M} is a discretization of Liouville quantum gravity (such as an embedded random planar map of an appropriate type), then we expect that the random walk on \mathcal{M} parameterized by counting measure on steps converges to a re-parameterized variant of Brownian motion called *Liouville Brownian motion* [Be] [GRV], but this has not been proven.

We will actually prove a generalization of Theorem 1.5 (see Theorem 1.16 below) in which \mathcal{FM} is replaced by a so-called *cell configuration*. A cell configuration is a generalization of the face set of an embedded lattice where the faces are replaced by “cells” which can be arbitrary compact connected sets, so the adjacency graph of cells is not necessarily planar. We will also prove (Theorem 3.10) that the rate of convergence of random walk to Brownian motion in Theorem 1.5 is uniform on compact subsets of \mathbb{C} .

The fact that we are able to replace exact translation invariance with ergodicity modulo scaling is the main novelty of our result. The finite expectation hypothesis (2)

provides an upper bound for the Dirichlet energy per unit area of the given embedding of \mathcal{M} . Indeed, integrating

$$\frac{\text{diam}(H_z)^2}{\text{Area}(H_z)} \pi(H_z)$$

over all z in some domain $D \subset \mathbb{C}$ gives an upper bound on the discrete Dirichlet energy of the function on \mathcal{M}^* which sends each face (dual vertex) which intersects D to a point in the corresponding face (see also Lemmas 2.10 and 2.12). As such hypothesis (2) tells us that the expected Dirichlet energy of this function is locally finite. We need to include the finite expectation condition with π^* in place of π in order to transfer from Dirichlet energy bounds to pointwise bounds (this is done in Lemma 2.17). Since $\pi(H_0) + \pi^*(H_0)$ is at least the degree (number of neighbors) of H_0 , (1.5) implies that this degree has finite expectation.

There is one natural hypothesis which is conspicuously missing from the statement of Theorem 1.5: namely, that \mathcal{M} does not have any macroscopic faces in the sense that the maximal diameter of the faces of \mathcal{M} which intersect $B_r(0)$ grows sublinearly in r . It turns out that this is implied by the given hypotheses; see Lemma 2.9.

A function $f: \mathcal{FM} \rightarrow \mathbb{C}$ is called \mathfrak{c} -discrete harmonic if, for any $H \in \mathcal{FM}$,

$$f(H) = \frac{1}{\pi(H)} \sum_{\substack{H' \in \mathcal{FM} \\ H' \sim H}} \mathfrak{c}(H, H') f(H')$$

As is common in the random walk in random environment literature (see, e.g., [Bi]), Theorem 1.5 will be proven by constructing a discrete harmonic function ϕ_∞ on \mathcal{M}^* which is close to the a-priori embedding (which sends each face to a point in or near the face) at large scales, and then using the martingale central limit theorem to show that the image of random walk on \mathcal{FM} under ϕ_∞ converges to Brownian motion.

THEOREM 1.6. *Recall the restricted embedded lattice $\mathcal{M}(B_r(0))$ for $r > 0$ defined just above (1.1). Almost surely, there exists a \mathfrak{c} -discrete harmonic function $\phi_\infty: \mathcal{FM} \rightarrow \mathbb{C}$ such that*

$$\lim_{r \rightarrow 0} \frac{1}{r} \sup_{H \in \mathcal{FM}(B_r(0))} |\phi_\infty(H) - \phi_0(H)| = 0, \quad (1.6)$$

where

$$\phi_0(H) := \frac{1}{\text{Area}(H)} \int_H z \, dz$$

is the Lebesgue center of mass of H .

In the random walk in random environment (RWRE) literature, $\phi_\infty - \phi_0$ is often referred to as the *corrector* for random walk on \mathcal{FM} . Theorem 1.6 says that the corrector grows sublinearly with respect to the Euclidean metric.

Another important intermediate step in the proof of Theorem 1.5 is the following statement, which is proven in §2.5.

THEOREM 1.7. *Under the same hypotheses as Theorem 1.5, it is a.s. the case that the random walk on the dual map \mathcal{M}^* with conductances \mathfrak{c} is recurrent.*

Gurel–Gurevich and Nachmias [GGN] showed that the simple random walk is a.s. recurrent on any random planar map which is the local limit of finite maps based at a uniform random vertex; and which is such that the law of the degree of the root vertex has an exponential tail. This result was later extended to a wider class of graphs by Lee [L, Theorem 1.6].

The criteria of [GGN], [L] do *not* apply to the graph \mathcal{M} above or to its dual. Indeed, \mathcal{M}^* is not a local limit of finite maps based at a uniform random root vertex, nor is it unimodular (except in trivial cases): rather, by translation invariance modulo scaling, the root face H_0 is a typical face from the perspective of Lebesgue measure, not a typical face from the perspective of the counting measure on faces of \mathcal{M} . Furthermore, our bound on the degree of H_0 from hypothesis (2) is weaker than the exponential tail bound required in [GGN] as well as the degree bound required in [L].

To illustrate the utility of Theorem 1.5, we now give three simple examples of its application to particular random environments. The examples do not use the full force of Theorem 1.5 since environments in these examples are exactly stationary with respect to spatial translations. Applications of Theorem 1.5 to non-stationary random environments are discussed in §4 and worked out in detail in [GMS2], [GMS3].

Example 1.8. (Random conductance model on \mathbb{Z}^2) Consider a random conductance function \mathfrak{c} on the nearest-neighbor edges of \mathbb{Z}^2 which is stationary and ergodic with respect to spatial translations of \mathbb{Z}^2 , and such that the conductances and their reciprocals have finite expectation. Theorem 1.5 implies that the random walk on \mathbb{Z}^2 with these conductances converges in law to Brownian motion in the quenched sense. This convergence was originally proven by Biskup [Bi] (see [SS] for an earlier proof in the case that the conductances are bounded above and below). To deduce the convergence from Theorem 1.5, let \mathcal{M} be the dual lattice $\mathbb{Z}^2 + (\frac{1}{2}, \frac{1}{2})$. If we sample w uniformly from Lebesgue measure on $[0, 1] \times [0, 1]$, then the randomly shifted lattice $\mathcal{M} - w$ is ergodic (hence ergodic modulo scaling) and the finite expectation hypothesis is immediate from our assumption on the conductances.

Example 1.9. (Dual of a random subgraph of \mathbb{Z}^2) Let \mathcal{M} be a random subgraph of \mathbb{Z}^2 which is stationary and ergodic with respect to spatial translations, with unit conductances (e.g., \mathcal{M} could be the infinite cluster of a supercritical bond percolation or Ising model on \mathbb{Z}^2). If the second moment of the number of edges on the boundary

of the face of \mathcal{M} containing zero is finite, then we can apply Theorem 1.5, with unit conductances, to \mathcal{M} (shifted by a uniform element of $[0, 1] \times [0, 1]$ as in Example 1.8) to find that the simple random walk on the faces of \mathcal{M} converges in law to Brownian motion in the quenched sense. If we assume a slightly stronger moment hypothesis, then Theorem 1.11 below shows that random walk on the vertices of \mathcal{M} likewise converges to Brownian motion (Example 1.12).

Example 1.10. (Voronoi tessellation of a homogeneous Poisson point process) Let \mathcal{P} be a homogeneous Poisson point process on \mathbb{C} (so the law of \mathcal{P} is translation invariant). Let \mathcal{M} be the Voronoi tessellation associated with \mathcal{P} , viewed as an embedded lattice with unit conductances whose faces are the Voronoi cells, whose edges are the linear segments on the boundaries of the cells, and whose vertices are the points where these linear segments meet. It is straightforward to verify that \mathcal{M} satisfies the hypotheses of Theorem 1.6, so the simple random walk on the faces of \mathcal{M} converges in law to Brownian motion in the quenched sense.

1.4. Variants of Theorem 1.5

Here we state several variants and extensions of Theorem 1.5 with different setups. The reader who only wants to see the proof of Theorem 1.5 can skip this subsection.

To avoid having to specify time parametrizations of the walks, for some of the statements in this subsection we will consider the topology on curves viewed modulo time parametrization, which we now recall. If $\beta_1: [0, T_{\beta_1}] \rightarrow \mathbb{C}$ and $\beta_2: [0, T_{\beta_2}] \rightarrow \mathbb{C}$ are continuous curves defined on possibly different time intervals, we set

$$\mathbf{d}^{\text{CMP}}(\beta_1, \beta_2) := \inf_{\phi} \sup_{t \in [0, T_{\beta_1}]} |\beta_1(t) - \beta_2(\phi(t))|, \quad (1.7)$$

where the infimum is over all increasing homeomorphisms $\phi: [0, T_{\beta_1}] \rightarrow [0, T_{\beta_2}]$ (the CMP stands for ‘‘curves modulo parametrization’’). It is shown in [AB, Lemma 2.1] that \mathbf{d}^{CMP} induces a complete metric on the set of curves viewed modulo time parametrization.

In the case of curves defined for infinite time, it is convenient to have a local variant of the metric \mathbf{d}^{CMP} . Suppose $\beta_1: [0, \infty) \rightarrow \mathbb{C}$ and $\beta_2: [0, \infty) \rightarrow \mathbb{C}$ are two such curves. For $r > 0$, let $T_{1,r}$ (resp. $T_{2,r}$) be the first exit time of β_1 (resp. β_2) from the ball $B_r(0)$ (or zero if the curve starts outside $B_r(0)$). We define

$$\mathbf{d}_{\text{loc}}^{\text{CMP}}(\beta_1, \beta_2) := \int_1^{\infty} e^{-r} (1 \wedge \mathbf{d}^{\text{CMP}}(\beta_1|_{[0, T_{1,r}]}, \beta_2|_{[0, T_{2,r}]}) dr, \quad (1.8)$$

so that $\mathbf{d}_{\text{loc}}^{\text{CMP}}(\beta^n, \beta) \rightarrow 0$ if and only if for Lebesgue a.e. $r > 0$, β^n stopped at its first exit time from $B_r(0)$ converges to β stopped at its first exit time from $B_r(0)$ with respect

to the metric (1.7). We note that the definition (1.7) of $\mathbf{d}^{\text{CMP}}(\beta_1|_{[0,T_{1,r}]}, \beta_2|_{[0,T_{2,r}]})$ makes sense even if one or both of $T_{1,r}$ and $T_{2,r}$ is infinite, provided we allow

$$\mathbf{d}^{\text{CMP}}(\beta_1|_{[0,T_{1,r}]}, \beta_2|_{[0,T_{2,r}]}) = \infty$$

(this does not pose a problem due to the $1 \wedge$ in (1.8)).

1.4.1. Random walk on vertices

If we assume a slightly stronger variant of hypothesis (2) in Theorem 1.5, we also get quenched convergence of the random walk on the vertices of \mathcal{M} to Brownian motion. To formulate this statement, for an embedded lattice \mathcal{M} and $x \in \mathcal{VM}$, define the *outradius* of x by

$$\text{Outrad}(x) := \text{diam} \left(\bigcup_{\substack{H \in \mathcal{FM} \\ x \in \partial H}} H \right), \tag{1.9}$$

i.e., the diameter of the union of the faces with x on their boundaries. Analogously to (1.4), if \mathcal{M} is equipped with a conductance function \mathbf{c} we also define

$$\pi(x) := \sum_{\substack{y \in \mathcal{VM} \\ y \sim x}} \mathbf{c}(x, y) \quad \text{and} \quad \pi^*(x) := \sum_{\substack{y \in \mathcal{VM} \\ y \sim x}} \frac{1}{\mathbf{c}(x, y)}, \tag{1.10}$$

where $x \sim y$ means that the vertices x and y are joined by an edge.

THEOREM 1.11. *Suppose \mathcal{M} is a random embedded lattice which is ergodic modulo scaling and satisfies the following stronger version of hypothesis (2) of Theorem 1.5:*

(2') *With H_0 the face of \mathcal{M} containing zero,*

$$\mathbb{E} \left[\sum_{x \in \mathcal{VM} \cap \partial H_0} \frac{\text{Outrad}(x)^2 \pi(x)}{\text{Area}(H_0)} \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\sum_{x \in \mathcal{VM} \cap \partial H_0} \frac{\text{Outrad}(x)^2 \pi^*(x)}{\text{Area}(H_0)} \right] < \infty.$$

Almost surely, as $\varepsilon \rightarrow 0$ the conditional law given \mathcal{M} of the simple random walk on $\varepsilon\mathcal{M}$ started from any vertex on the boundary of the origin-containing face and extended from \mathbb{N}_0 to $[0, \infty)$ by linear interpolation converges in law modulo time parametrization to planar Brownian motion started from zero with some deterministic, non-degenerate covariance matrix. Furthermore, the simple random walk on \mathcal{M} is a.s. recurrent and there is a discrete harmonic function $\phi_\infty: \mathcal{VM} \rightarrow \mathbb{C}$ such that a.s.

$$\lim_{r \rightarrow \infty} \frac{1}{r} \max_{x \in \mathcal{VM} \cap B_r(0)} |\phi_\infty(x) - x| = 0.$$

We will deduce Theorem 1.11 from Theorem 1.5 by drawing in the dual map \mathcal{M}^* of \mathcal{M} in such a way that the hypotheses of Theorem 1.5 are satisfied for \mathcal{M}^* . We state the convergence in Theorem 1.11 modulo time parametrization, since (unlike for \mathcal{M}^*) there is not a canonical way of associating an area with each vertex in order to decide how much time the walk spends there.

As in the case of Theorem 1.5, Theorem 1.11 has applications to random walk in random environments on \mathbb{Z}^2 .

Example 1.12. (Random subgraphs of \mathbb{Z}^2) Let \mathcal{M} be a random subgraph of \mathbb{Z}^2 which is stationary and ergodic with respect to spatial translations, as in Example 1.9. Since each vertex of \mathcal{M} has degree at most 4, Theorem 1.11 (applied to $\mathcal{M}-w$ for w a uniform sample from $[0, 1] \times [0, 1]$) implies that the random walk on \mathcal{M} converges to Brownian motion modulo time parametrization in the quenched sense provided

$$\mathbb{E} \left[\left(\sum_{x \in \mathbb{Z}^2 \cap \partial H_0} \text{Outrad}(x) \right)^2 \right] < \infty, \quad (1.11)$$

where H_0 is the face containing $[0, 1] \times [0, 1]$. The bound (1.11) is satisfied, e.g., for supercritical percolation on \mathbb{Z}^2 , since the size of the origin-containing cluster in the dual subcritical percolation has an exponential tail [G, Theorem 5.4]. Other random subgraphs of \mathbb{Z}^2 satisfying this bound arise in subcritical Ising models, FK models, etc. Random walk on the infinite cluster of supercritical bond percolation on \mathbb{Z}^d was shown to converge to Brownian motion uniformly (not just modulo time parametrization) in [BB]; see also [PRS] for a generalization of this result.

1.4.2. Embedded primal/dual lattice pair

In this subsection we formulate an analog of Theorem 1.5 when we have embeddings of both \mathcal{M} and its planar dual \mathcal{M}^* . The hypotheses of this version of the theorem are in some ways more elegant since the conditions are symmetric in \mathcal{M} and \mathcal{M}^* .

If $(\mathcal{M}, \mathcal{M}^*)$ is a pair of embedded lattices, we say that \mathcal{M}^* is *dual* to \mathcal{M} if each face of \mathcal{M} (resp. \mathcal{M}^*) contains exactly one vertex of \mathcal{M}^* (resp. \mathcal{M}) and each edge of \mathcal{M} (resp. \mathcal{M}^*) crosses exactly one edge of \mathcal{M}^* (resp. \mathcal{M}). If \mathcal{M} and \mathcal{M}^* are equipped with edge conductances \mathfrak{c} and \mathfrak{c}^* , we also require that $\mathfrak{c}(e) = 1/\mathfrak{c}^*(e^*)$, for e^* the edge of \mathcal{M}^* which crosses e . See Figure 2, left, for an illustration. There is an obvious generalization of the notions of translation invariance modulo scaling to the case of a pair of dual embedded lattices: indeed, we just replace all of the statements for re-scaled translated versions of \mathcal{M} in Definition 1.2 by joint statements for re-scaled translated versions of $(\mathcal{M}, \mathcal{M}^*)$ (we require that all translation and scaling factors are the same for both \mathcal{M} and \mathcal{M}^*). The

definition of ergodic modulo scaling (Definition 1.3) is similarly well defined for the pair $(\mathcal{M}, \mathcal{M}^*)$.

THEOREM 1.13. *Suppose the pair $(\mathcal{M}, \mathcal{M}^*)$ (where \mathcal{M} is a.s. dual to \mathcal{M}^*) is ergodic modulo scaling and that the following finiteness condition holds. If H_0 (resp. H_0^*) is the face of \mathcal{M} (resp. \mathcal{M}^*) containing the origin and E^* (resp. E) denotes the set of edges of \mathcal{M}^* (resp. \mathcal{M}) which cross edges on the boundary of H_0 (resp. H_0^*), then*

$$\mathbb{E} \left[\sum_{e \in E} \frac{\text{diam}(e)^2 \mathfrak{c}(e)}{\text{Area}(H_0^*)} \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\sum_{e \in E^*} \frac{\text{diam}(e^*)^2 \mathfrak{c}(e^*)}{\text{Area}(H_0)} \right] < \infty. \quad (1.12)$$

Almost surely, as $\varepsilon \rightarrow 0$ the conditional law given $(\mathcal{M}, \mathcal{M}^*)$ of the simple random walk on $\varepsilon\mathcal{M}$ started from any vertex on the boundary of the origin-containing face and extended from \mathbb{N}_0 to $[0, \infty)$ by linear interpolation converges in law modulo time parametrization to planar Brownian motion started from zero with some deterministic, non-degenerate covariance matrix. Furthermore, the simple random walk on \mathcal{M} is a.s. recurrent and there is a discrete harmonic function $\phi_\infty: \mathcal{VM} \rightarrow \mathbb{C}$ such that a.s.

$$\lim_{r \rightarrow \infty} \frac{1}{r} \max_{x \in \mathcal{VM} \cap B_r(0)} |\phi_\infty(x) - x| = 0.$$

The same holds with \mathcal{M}^* in place of \mathcal{M} .

We can take the two quantities in (1.12)—both required to be finite— as the *definition* of the specific Dirichlet energy of \mathcal{M} and \mathcal{M}^* . See Appendix A for further justification of this definition.

As we will see, Theorem 1.13 follows from essentially the same proof as Theorem 1.5 (or Theorem 1.16 below) modulo the following technicality: since the finite expectation hypothesis (1.12) only gives a bound for the expectation of squared edge diameter times conductance, it is much harder to show that there are no macroscopic edges. This is carried out in Appendix C to avoid distracting from the main arguments. We now give an example to illustrate the usage of Theorem 1.13.

Example 1.14. (Random conductance model on \mathbb{Z}^2 without finite expectation) We expect that Theorem 1.13 can be used to prove quenched invariance principles (modulo time parametrization) for some kinds of stationary, ergodic random conductance models on \mathbb{Z}^2 whose conductances and their reciprocals do not have finite expectation. This can be done by “shrinking” edges with large conductance and the dual edges (i.e., edges of $\mathbb{Z}^2 + (\frac{1}{2}, \frac{1}{2})$) corresponding to edges with small conductance to get a new embedded lattice/dual lattice pair satisfying the hypotheses of Theorem 1.13. We do not work out any examples in detail here, but we note that similar ideas are used to deal with i.i.d. random conductance models without a finite expectaton hypothesis, in general dimension, in [ABDH], [BD].




Figure 2. *Left:* An embedded lattice/dual lattice pair. *Right:* A possible arrangement of six cells of a cell configuration. The orange cell consists of two simply connected regions joined by a curve of zero Lebesgue measure. The blue cell is cut in two by the green cell. The pink cell does not share a boundary arc with any cell other than the green one. The purple cell has a “tentacle” joining it to the grey cell. Note that these sorts of situations show that \mathcal{H} , as a graph, need not be planar.

1.4.3. Random walk on cells

We now state a generalization of Theorem 1.5 (Theorem 1.16) where faces are replaced by general compact connected sets. This allows for graph structures which are not planar. This is the variant which will be used in [GMS3] and the proof is identical to that of Theorem 1.5 (which is a special case). Consequently, the version of the theorem stated in this subsection is the first one which we will prove.

Definition 1.15. A *cell configuration* on a domain $D \subset \mathbb{C}$ consists of the following objects.

(1) A locally finite collection \mathcal{H} of compact connected subsets of \bar{D} (“cells”) with non-empty interiors whose union is \bar{D} and such that the intersection of any two elements of \mathcal{H} has zero Lebesgue measure.

(2) A symmetric relation \sim on $\mathcal{H} \times \mathcal{H}$ (“adjacency”) such that if $H \sim H'$, then

$$H \cap H' \neq \emptyset \quad \text{and} \quad H \neq H'.$$

(3) A function $\mathbf{c} = \mathbf{c}_{\mathcal{H}}$ (“conductance”) from the set of pairs (H, H') with $H \sim H'$ to $(0, \infty)$ such that $\mathbf{c}(H, H') = \mathbf{c}(H', H)$.

We typically slightly abuse notation by making the relation \sim and the function \mathbf{c} implicit, so we write \mathcal{H} instead of $(\mathcal{H}, \sim, \mathbf{c})$.

We will almost always consider only the case when $D = \mathbb{C}$. When we refer to “cell configuration” without a specified domain, we mean that $D = \mathbb{C}$. The case when $D \neq \mathbb{C}$ is only used to formulate the cell configuration analog of condition (1) of Definition 1.2. We

view a cell configuration \mathcal{H} (with arbitrary domain) as a weighted graph whose vertices are the cells of \mathcal{H} and whose edge set is

$$\mathcal{E}\mathcal{H} := \{(H, H') \in \mathcal{H} \times \mathcal{H} : H \sim H'\}, \tag{1.13}$$

and such that any edge $\{H, H'\}$ is assigned the weight $\mathbf{c}(H, H')$. Note that any two cells which are joined by an edge intersect but two intersecting cells need not be joined by an edge.

A simple example of a cell configuration is the case that \mathcal{H} is the set of faces of an embedded lattice, with two such faces adjacent if and only if they share an edge. In this case the resulting cell configuration is isomorphic to \mathcal{M}^* . Consequently, one can think of an embedded lattice as a special case of a cell configuration (except that we do not have distinguished vertices and edges for a cell configuration, so we lose some information which does not affect the random walk on the faces). However, the cells in a cell configuration need not be simply connected and they are allowed to overlap (as long as the overlap has zero Lebesgue measure), so a cell configuration need not correspond to a planar graph. An example of a cell configuration is illustrated in Figure 2, right.

Generalizing the above notation for embedded lattices, for a cell configuration \mathcal{H} we define

$$\mathcal{H}(A) := \{H \in \mathcal{F}\mathcal{M} : H \cap A \neq \emptyset\} \quad \text{for all } A \subset \mathbb{C}, \tag{1.14}$$

and we view $\mathcal{H}(A)$ as an edge-weighted graph with edge set consisting of all of the edges in $\mathcal{E}\mathcal{H}$ joining elements of $\mathcal{H}(A)$ and conductances given by the restriction of \mathbf{c} . We define a metric on the space of cell configurations via the obvious extension of (1.1):

$$\begin{aligned} \mathbf{d}^{\mathbb{C}\mathbb{C}}(\mathcal{H}, \mathcal{H}') := & \int_0^\infty e^{-r} \wedge \inf_{f_r} \left\{ \sup_{z \in \mathbb{C}} |z - f_r(z)| \right. \\ & \left. + \max_{\{H_1, H_2\} \in \mathcal{E}\mathcal{H}(B_r(0))} |\mathbf{c}(H_1, H_2) - \mathbf{c}'(f_r(H_1), f_r(H_2))| \right\} dr, \end{aligned} \tag{1.15}$$

where each of the infima is over all homeomorphisms $f_r : \mathbb{C} \rightarrow \mathbb{C}$ such that f_r takes each cell in $\mathcal{H}(B_r(0))$ to a cell in $\mathcal{H}'(B_r(0))$ and preserves the adjacency relation, and f_r^{-1} does the same with \mathcal{H} and \mathcal{H}' reversed. We also say that a cell configuration is *translation invariant modulo scaling* if it satisfies any of the equivalent conditions of Definition 1.2, with \mathcal{H} in place of \mathcal{M} (and cells in place of faces), and *ergodic modulo scaling* if it satisfies the conditions of Definition 1.3, with \mathcal{H} in place of \mathcal{M} .

For $z \in \mathbb{C}$, we write H_z for one of the cells in \mathcal{H} containing z , chosen in some arbitrary manner if there is more than one such cell (the cell is unique for Lebesgue-a.e. z). Exactly as in (1.4), we also define

$$\pi(H) := \sum_{\substack{H' \in \mathcal{H} \\ H' \sim H}} \mathbf{c}(H, H') \quad \text{and} \quad \pi^*(H) := \sum_{\substack{H' \in \mathcal{H} \\ H' \sim H}} \frac{1}{\mathbf{c}(H, H')}. \tag{1.16}$$

The following theorem generalizes Theorems 1.5–1.7. We need an extra hypothesis to ensure that our cell configuration is sufficiently “locally connected” (since we do not assume that cells which intersect are joined by an edge).

THEOREM 1.16. *Let \mathcal{H} be a random cell configuration which is ergodic modulo scaling, satisfies the finite expectation hypothesis (1.5) with \mathcal{H} in place of \mathcal{FM} , and satisfies the following additional condition:*

(3) *Connectedness along lines. Almost surely, for each horizontal or vertical line segment $L \subset \mathbb{C}$, the subgraph of \mathcal{H} consisting of the set of cells which intersect L and the edges joining these cells is connected.*

Let $X: [0, \infty) \rightarrow \mathcal{H}$ be the continuous-time simple random walk on \mathcal{H} (viewed as an edge-weighted graph as above) which spends $\text{Area}(H)/\pi(H)$ units of time at each cell before jumping to the next. Let $\widehat{X}: [0, \infty) \rightarrow \mathbb{C}$ be the process obtained by composing X with a function that sends each cell H to an (arbitrary) point of H . There is a deterministic covariance matrix Σ with $\det \Sigma \neq 0$, depending on the law of \mathcal{H} , such that a.s. as $\varepsilon \rightarrow 0$ the conditional law of $t \mapsto \varepsilon \widehat{X}_{t/\varepsilon^2}$ given \mathcal{H} converges to the law of planar Brownian motion started from zero with covariance matrix Σ with respect to the local uniform topology. Furthermore, the simple random walk on \mathcal{H} is a.s. recurrent and there is a discrete harmonic function $\phi_\infty: \mathcal{VH} \rightarrow \mathbb{C}$ such that a.s.

$$\lim_{r \rightarrow \infty} \frac{1}{r} \max_{H \in \mathcal{H}(B_r(0))} |\phi_\infty(H) - \phi_0(H)| = 0, \tag{1.17}$$

where $\phi_0(H) := \int_H z \, dz$ is the Euclidean center of H .

Theorem 1.16 can be applied to show that random walk on an embedded lattice which is allowed to have finite-range jumps converges to Brownian motion. To illustrate this, we give an example of a random environment on \mathbb{Z}^2 to which Theorem 1.16, but not any of our previous theorems, applies.

Example 1.17. (Finite-range random conductance model on \mathbb{Z}^2) Let

$$\mathbf{c}: \mathbb{Z}^2 \times \mathbb{Z}^2 \longrightarrow (0, \infty)$$

be a conductance function which has finite range in the sense that there exists $N \in \mathbb{N}$ such that $\mathbf{c}(x, y) = 0$ whenever $|x - y| > N$. Assume that \mathbf{c} is stationary and ergodic with respect to spatial translations and satisfies

$$\mathbb{E}[\mathbf{c}(x, y)] < \infty \quad \text{and} \quad \mathbb{E}[\mathbf{c}(x, y)^{-1}] < \infty$$

whenever $|x - y| \leq N$. Then, Theorem 1.16 implies that a.s. the random walk on \mathbb{Z}^2 with conductances \mathbf{c} converges in law to Brownian motion in the quenched sense, under diffusive

scaling. To see this, for $x \in \mathbb{Z}^2$ let H_x be the cell which is the union of the square of side length $\frac{1}{2}$ centered at x and all of the straight line segments from x to points $y \in \mathbb{Z}^2$ with $|x-y| \leq N$. It is easily seen that the hypotheses of Theorem 1.16 are satisfied for the cell configuration with $\mathcal{H} = \{H_x : x \in \mathbb{Z}^2 - w\}$, adjacency defined by $H_x \sim H_y$ whenever $|x-y| \leq N$, and conductances defined by $\mathfrak{c}(H_x, H_y) = \mathfrak{c}(x, y)$ (provided we shift \mathbb{Z}^2 by a uniform random element of $[0, 1] \times [0, 1]$, as in previous examples).

1.5. Why $d=2$ is special

Our arguments are very specific to two dimensions. To give some intuition about why $d=2$ is special, we remark that it is a simple exercise to show that random walk started at the center of Figure 3 approximates an ordinary boundary reflected Brownian motion (up to time change), while the analogous statement in other dimensions is false. For example, the analogous figure in dimension one is a vertical interval, whose upper half is subdivided into twice as many pieces as its lower half; a simple random walk started at the midpoint of the vertical interval would have a $\frac{2}{3}$ chance to hit the bottom before the top.

Similarly, if one sums the squared edge lengths in the upper half of Figure 3, one gets the same value as if one sums the squared edge lengths in the lower half (as both quantities are proportional to area). In other words, the Dirichlet energy per unit area of the embedded lattice is the same in both halves, a statement that would also be false in any dimension $d \neq 2$. Because Dirichlet energy and area both scale as length squared when $d=2$, we will be able to make sense of the expected Dirichlet energy per area of \mathcal{M} as a quantity that does not depend on scaling, and to show that on large regions, the Dirichlet energy per area averages out to be close to a deterministic constant with high probability.

If one wishes to extend the results of this paper to other dimensions, one might try the following: instead of working modulo scalings of the form $(\mathcal{M}, \mathfrak{c}) \rightarrow (C\mathcal{M}, \mathfrak{c})$, work modulo scalings of the form $(\mathcal{M}, \mathfrak{c}) \rightarrow (C\mathcal{M}, C^{d-2}\mathfrak{c})$ (where \mathfrak{c} applied to an edge of $C\mathcal{M}$ is just defined equal \mathfrak{c} applied to the corresponding edge of \mathcal{M}). This way Dirichlet energy per volume would still be a scale-invariant notion. It is an interesting open question to determine the extent to which the results of this paper can be extended to higher dimensions in this context. To start with a specific example, one could try to prove an invariance principle for some discretization of higher-dimensional Gaussian multiplicative chaos. We will not further discuss this problem here.




Figure 3. A unit square with lower half divided into grid of $2^{-k} \times 2^{-k}$ squares and upper half divided into a grid of $2^{-k-1} \times 2^{-k-1}$ squares, where $k=3$. In the $k \rightarrow \infty$ limit, a simple random walk on this figure converges to a time change of (boundary-reflected) Brownian motion, but the analogous statement would be false in any dimension other than 2.

1.6. Outline

In this section we provide a moderately detailed overview of the content of the rest of the paper. See Figure 4 for a schematic illustration.

Most of the paper is devoted to the proof of Theorem 1.16. The proof is essentially identical to that of Theorem 1.5, so the reader who cares only about embedded lattices can think only of the setting of Theorem 1.5, which is a special case of Theorem 1.16. As one might expect from the hypotheses, the theorem will be proven using ergodic theory. Some of our arguments are inspired by those in [BB], [BP], [Bi], but in many places very different techniques are needed due to the lack of exact stationarity with respect to spatial translations.

§2 is devoted to the proof of two parts of Theorem 1.16 (which correspond to Theorems 1.6 and 1.7): the existence of the discrete harmonic function ϕ_∞ satisfying (1.17) and the recurrence of the random walk on \mathcal{H} . We start in §2.1 by defining a particularly convenient collection of squares, called a *uniform dyadic system*, which allows us to formulate another condition which is equivalent to the conditions of Definition 1.2 and which is the only condition which we will use in our proofs. Roughly speaking, a uniform dyadic system \mathcal{D} contains a unique bi-infinite sequence of squares containing any point in \mathbb{C} (except for a Lebesgue measure-zero set of points which lie on boundaries of squares) whose side lengths belong to $\{2^{s+k}\}_{k \in \mathbb{Z}}$, for s a uniform $[0, 1]$ random variable. Throughout the rest of the §2 and §3 we will fix a uniform dyadic system \mathcal{D} which is independent from \mathcal{H} .

In §2.2 we use the backward martingale convergence theorem and ergodicity modulo scaling to prove a.s. limit theorems for the averages of various quantities over the origin-containing squares in \mathcal{D} (Lemma 2.7). Combined with our finite expectation hypothesis,

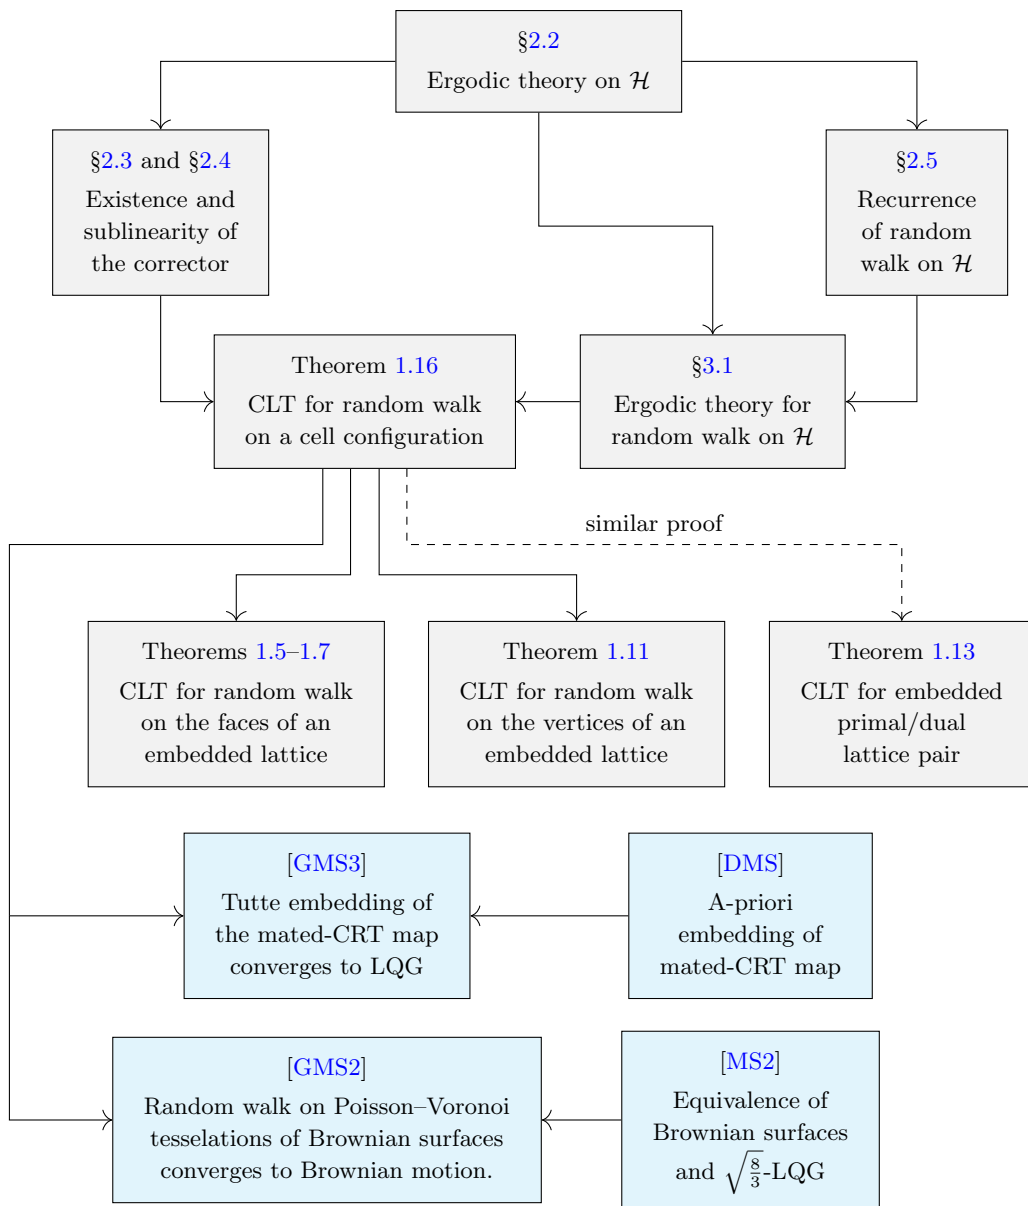


Figure 4. Schematic illustration of how results which are proven in this paper (grey) and related results proven in other papers (cyan) fit together.

this leads to a proof that the maximum diameter of the cells which intersect a large origin-containing square in \mathcal{D} is typically of smaller order than its side length (Lemma 2.9) and a bound for the sum of the squared diameter times degree of the cells which intersect such a square (Lemma 2.10).

In §2.3 we construct the discrete harmonic function ϕ_∞ of Theorem 1.16 as a limit of functions ϕ_m for $m \in \mathbb{N}$ which agree with the a-priori embedding function ϕ_0 on the boundaries of certain large squares in \mathcal{D} and are discrete harmonic on the interiors of these squares. The key input in the proof is a Dirichlet energy bound for these functions (Lemma 2.12), which comes from the diameter² times degree bound of the preceding subsection. This allows us to show that the discrete Dirichlet energy of $\phi_m - \phi_{m'}$ is small when m and m' are large (Lemma 2.14), which in turn implies that the ϕ_m 's are Cauchy in probability and hence the limit ϕ_∞ exists.

In §2.4 we prove that ϕ_∞ approximates ϕ_0 in the sense of (1.17), as follows. We observe that, for fixed $m \in \mathbb{N}$, the maximum of $|\phi_m - \phi_0|$ of a large square in \mathcal{D} is typically of smaller order than the side length of the square since ϕ_m and ϕ_0 agree on the boundaries of many small subsquares of this big square (Lemma 2.18). We then use the Cauchy-Schwarz inequality and the maximum principle to bound the maximum of $|\phi_\infty - \phi_m|$ over a large square in terms of its Dirichlet energy (see in particular Lemmas 2.17), which we know is small when m is large by the results of §2.3.

In §2.5 we prove the recurrence of the simple random walk on \mathcal{H} using the results of §2.2 and the criterion for recurrence in terms of effective resistance (equivalently, Dirichlet energy).

In §3 we will prove our main results, starting with Theorem 1.16. The basic idea of the proof is to apply the martingale functional central limit theorem to show convergence of the image of the random walk on \mathcal{H} under the discrete harmonic function ϕ_∞ , then apply Theorem 3.10. In order to apply the multi-dimensional martingale central limit theorem, we need to know that the quadratic variations and covariation of the two coordinates of the image of the walk differ by (approximately) a constant factor. For this purpose we will need analogs of some of the ergodic theory results of §2.2 when we translate according to the time parametrization of a simple random walk on \mathcal{H} , rather than uniformly from Lebesgue measure on a square. Such results are proven in §3.1 and rely on the recurrence of the walk to obtain tail triviality.

In §3.2, we apply these ergodic theory results and the martingale functional central limit theorem to complete the proof of Theorem 1.16. In §3.3 we prove an extension of Theorem 1.16 (Theorem 3.10) wherein the convergence of the random walk on \mathcal{H} to Brownian motion holds uniformly over all choices of starting point for the walk in a compact subset of \mathbb{C} . In §3.4, we deduce Theorems 1.5 and 1.11 from Theorem 1.16 and

explain the adaptations needed to prove Theorem 1.13.

§4 discusses applications of our work to random planar maps and Liouville quantum gravity (which are explored further in the companion paper [GMS3]).

Appendix B contains a proof of the equivalence of the definitions of translation invariance modulo scaling stated in Definition 1.2. Appendix C contains a proof that the hypotheses of Theorem 1.13 imply that a.s. the maximal size of the edges or dual edges which intersect $B_r(0)$ is $o_r(r)$. This result is not needed for the proofs of any of the other results in the paper or for any of the applications of our results in [GMS3].

Basic notation

We write $\mathbb{N}=\{1, 2, 3, \dots\}$ and $\mathbb{N}_0=\mathbb{N}\cup\{0\}$. For $a < b$, we define the discrete interval $[a, b]_{\mathbb{Z}} := [a, b] \cap \mathbb{Z}$. If $f: (0, \infty) \rightarrow \mathbb{R}$ and $g: (0, \infty) \rightarrow (0, \infty)$, we say that $f(\varepsilon) = O_\varepsilon(g(\varepsilon))$ (resp. $f(\varepsilon) = o_\varepsilon(g(\varepsilon))$) as $\varepsilon \rightarrow 0$ if $f(\varepsilon)/g(\varepsilon)$ remains bounded (resp. tends to zero) as $\varepsilon \rightarrow 0$. We similarly define $O(\cdot)$ and $o(\cdot)$ errors as a parameter goes to infinity. If $f, g: (0, \infty) \rightarrow [0, \infty)$, we say that $f(\varepsilon) \preceq g(\varepsilon)$ if there is a constant $C > 0$ (independent from ε and possibly from other parameters of interest) such that $f(\varepsilon) \leq Cg(\varepsilon)$. We write $f(\varepsilon) \asymp g(\varepsilon)$ if $f(\varepsilon) \preceq g(\varepsilon)$ and $g(\varepsilon) \preceq f(\varepsilon)$.

2. Existence and sublinearity of the corrector

2.1. Dyadic systems

A key technical tool in our proofs is the concept of a *dyadic system*, which leads to a convenient scale/translation invariant way of decomposing space into “blocks” and which we will define in this subsection. Let $S \subset \mathbb{C}$ be a square (not necessarily dyadic). We write $|S|$ for the side length of S . A *dyadic child* of S is one of the four squares $S' \subset S$ (with $|S'| = \frac{1}{2}|S|$) whose corners include one corner of S and the center of S . A *dyadic parent* of S is one of the four squares with S as a dyadic child. Note that each square has four dyadic parents and four dyadic children. A *dyadic descendant* (resp. *dyadic ancestor*) of S is a square which can be obtained from S by iteratively choosing dyadic children (resp. parents) finitely many times.

A *dyadic system* is a collection \mathcal{D} of closed squares (not necessarily dyadic) with the following properties.

- (1) If $S \in \mathcal{D}$, then each of the four dyadic children of S is in \mathcal{D} .
- (2) If $S \in \mathcal{D}$, then exactly one of the dyadic parents of S is in \mathcal{D} .
- (3) Any two squares in \mathcal{D} have a common dyadic ancestor.




Figure 5. The squares S_k which contain the origin (red) in a dyadic system, for $k \leq 0$. Note that the dyadic system also includes the origin-containing squares S_k for $k > 0$, which have side lengths tending to ∞ as $k \rightarrow \infty$, and all of the dyadic descendants of the S_k 's.

See Figure 5 for an illustration of the origin-containing squares of a dyadic system. The set of all side lengths of the squares in a dyadic system is precisely $\{2^{s+k}\}_{k \in \mathbb{Z}}$ for some $s \in [0, 1]$ (determined by the system). If \mathcal{D} is a dyadic system, then for any $z \in \mathbb{C}$ there is a bi-infinite sequence of squares $\{S_k^z\}_{k \in \mathbb{Z}} \subset \mathcal{D}$ which contain z , numbered so that S_k^z is a dyadic parent of S_{k-1}^z for each k . If z does not lie on the boundary of a square in \mathcal{D} , then this sequence is unique up to translation of the indices. If we are given any single square $S \in \mathcal{D}$ and all of its dyadic ancestors, then \mathcal{D} is uniquely determined: \mathcal{D} is the set of all dyadic descendants of S and its dyadic ancestors. This in particular allows us to define a topology on the space of dyadic systems, e.g., by looking at the local Hausdorff distance on the union of the origin-containing squares.

In what follows, there will generally be a dyadic system which is understood from the context. We emphasize that when we refer to a dyadic ancestor of a square S , we mean any of the possible dyadic ancestors of S and not just the one which is contained in the dyadic system.

A *uniform dyadic system* is the random dyadic system defined as follows. Let s be sampled uniformly from $[0, 1]$ and, conditional on s , let w be sampled uniformly from $[0, 2^s] \times [0, 2^s]$. Set $S_0 := ([0, 2^s] \times [0, 2^s]) - w$. For $k \in \mathbb{N}$, inductively let S_k be sampled uniformly from the four dyadic parents of S_{k-1} . Then, let \mathcal{D} be the set of all dyadic descendants of S_k for each $k \in \mathbb{N}$. We view \mathcal{D} as a dyadic system, so in particular \mathcal{D} only determines the sequence of origin-containing squares up to an index shift. The main reason for considering a uniform dyadic system is the following lemma.

LEMMA 2.1. *Let \mathcal{D} be a uniform dyadic system and let $C > 0$ and $z \in \mathbb{C}$ be deterministic. Then, the scaled/translated dyadic system $C(\mathcal{D} - z)$ obtained by translating and scaling the squares agrees in law with \mathcal{D} .*

Proof. We first check translation invariance. Fix $z \in \mathbb{C}$ and set $\tilde{\mathcal{D}} := \mathcal{D} - z$. We seek to show that $\tilde{\mathcal{D}} \stackrel{d}{=} \mathcal{D}$. For $k \in \mathbb{Z}$, let S_k (resp. \tilde{S}_k) be the square in \mathcal{D} (resp. $\tilde{\mathcal{D}}$) containing 0 with side length 2^{k+s} . Since a dyadic system is a.s. determined by any fixed square and all of its dyadic ancestors, it suffices to show that the total variation distance between the laws of $\{S_k\}_{k \geq K}$ and $\{\tilde{S}_k\}_{k \geq K}$ tends to zero as $K \rightarrow \infty$.

We first argue that $\tilde{S}_0 \stackrel{d}{=} S_0$. Indeed, \tilde{S}_0 is obtained by translating the square of \mathcal{D} which has side length 2^s and contains z by $-z$, so $\tilde{S}_0 = [0, 2^s] \times [0, 2^s] - \tilde{w}$, where $\tilde{w} \in [0, 2^s] \times [0, 2^s]$ is chosen so that $(z+w) - \tilde{w} \in 2^s \mathbb{Z}^2$. Since w is uniform on $[0, 2^s] \times [0, 2^s]$ conditional on s , the same is true of \tilde{w} . Thus $\tilde{S}_0 \stackrel{d}{=} S_0$.

We next observe that as $k \rightarrow \infty$, the fraction of the 4^k dyadic ancestors of S_0 with side length 2^{s+k} which are not also dyadic ancestors of the square $\tilde{S}_0 + z \in \mathcal{D}$ decays exponentially in k . Since each S_k is sampled uniformly from the set of dyadic ancestors of S_0 with side length 2^{s+k} , it follows that the total variation distance between the law of $\{S_k - z\}_{k \geq K}$ and the law of a sequence of random squares defined in the same manner as $\{S_k\}_{k \geq K}$ as above but with \tilde{S}_0 used in place of S_0 tends to zero as $K \rightarrow \infty$. On the other hand, a.s. $z \in S_k$ and hence $\tilde{S}_k = S_k - z$ for large enough $k \in \mathbb{N}$. Combining this with the preceding paragraph shows our claim about the total variation distance, and hence that $\mathcal{D} \stackrel{d}{=} \tilde{\mathcal{D}}$.

One checks that $C\mathcal{D} \stackrel{d}{=} \mathcal{D}$ for $C > 0$ using a similar but simpler argument to the one above, based on the fact that the fractional part of $s + \log_2 C$ is uniform on $[0, 1]$. \square

Throughout the rest of this section, we fix a random cell configuration \mathcal{H} and a uniform dyadic system \mathcal{D} independent from \mathcal{H} (in subsequent subsections we will also require that \mathcal{H} satisfies the hypotheses of Theorem 1.16). We let $\{S_k\}_{k \in \mathbb{Z}}$ be the sequence of squares containing the origin in \mathcal{D} , labeled so that the side length of S_0 is 2^s , as above.

It will be convenient to have a natural, scale-invariant way to index the squares in \mathcal{D} . To this end, for $z \in \mathbb{C}$ and $m > 0$, we define

$$\hat{S}_m^z := \left\{ \text{largest square } S \in \mathcal{D} : z \in S \text{ and } \sum_{H \in \mathcal{H}(S)} \frac{\text{Area}(H \cap S)}{\text{Area}(H)} \leq m \right\}. \tag{2.1}$$

Then, \hat{S}_m^z is well defined on the full-probability event that z does not lie on the boundary of any square in \mathcal{D} . We abbreviate $\hat{S}_m^z := \hat{S}_m^0$.

Roughly speaking, \hat{S}_m^z is the largest square containing z which contains at most m cells, but fractional parts of cells which intersect the boundary of the square are also counted. We emphasize that we do *not* restrict to integer values of m .

LEMMA 2.2. *Almost surely, for each square S in \mathcal{D} which contains z in its interior, there is a non-trivial interval $I \subset (0, \infty)$ such that $S = \widehat{S}_m^z$ for all $m \in I$. Furthermore, $\widehat{S}_m^z = \widehat{S}_m^w$ whenever $w \in \widehat{S}_m^z \setminus \partial \widehat{S}_m^z$.*

Proof. If S' is one of the four dyadic parents of S , then

$$\sum_{H \in \mathcal{H}(S')} \frac{\text{Area}(H \cap S')}{\text{Area}(H)} > \sum_{H \in \mathcal{H}(S)} \frac{\text{Area}(H \cap S)}{\text{Area}(H)}.$$

Note that this would not be true if we did not include fractional parts of cells in the sum. The lemma statement is immediate from this observation and the definition (2.1) of \widehat{S}_m^z . \square

Dyadic systems allow us to formulate yet another equivalent definition of translation invariance modulo scaling (which is the only version which we will use in our proofs) and also provide a convenient useful tool for proving the equivalence of the conditions in Definition 1.2.

LEMMA 2.3. *Let \mathcal{H} be a random cell configuration. Each of the conditions of Definition 1.2 (with \mathcal{H} in place of \mathcal{M}) are equivalent to each other and to the following additional condition.*

(5) *Dyadic system re-sampling property. Let \mathcal{D} be a uniform dyadic system independent from \mathcal{H} and define $\widehat{S}_m = \widehat{S}_m^0$ for $m > 0$ as in (2.1). For each $m > 0$, the following is true. Conditional on \mathcal{H} and \mathcal{D} , let z be sampled uniformly from Lebesgue measure on \widehat{S}_m . Then the translated cell configuration/dyadic system pair $(\mathcal{H} - z, \mathcal{D} - z)$ agrees in law with $(\mathcal{H}, \mathcal{D})$ modulo scaling, i.e., there is a random $C > 0$, possibly depending on z , such that*

$$(C(\mathcal{H} - z), C(\mathcal{D} - z)) \stackrel{d}{=} (\mathcal{H}, \mathcal{D}).$$

The proof of Lemma 2.3 proceeds by showing that each of the conditions of Definition 1.2 is equivalent to the dyadic system resampling condition in the lemma statement. The proof is routine in nature, so is deferred to Appendix B to avoiding interrupting the main argument. We emphasize that the condition given in Lemma 2.3 is the only form of translation invariance modulo scaling which is used in the proofs of our main results.

Remark 2.4. The definition of the uniform dyadic system has an obvious d -dimensional generalization for any integer $d \geq 1$, where squares are replaced by d -dimensional cubes, and each cube has $2d$ dyadic children. We will make use of 1-dimensional uniform dyadic systems in §3.1 of this paper when we discuss continuous time random walks indexed by \mathbb{R} .

2.2. Ergodic averages using dyadic systems

Throughout the rest of this section and the next, we assume that the cell configuration \mathcal{H} satisfies the hypotheses of Theorem 1.16. Recall the independent dyadic system \mathcal{D} and the sequence $\{S_k\}_{k \in \mathbb{Z}}$ of origin-containing squares. In this subsection we will establish some basic properties of the pair $(\mathcal{H}, \mathcal{D})$ which we will use frequently in what follows. In particular, we will record an a.s. convergence statement for averages over the squares \widehat{S}_m^0 of (2.1) (Lemma 2.7) which follows from the backward martingale convergence theorem and our ergodicity hypothesis (Definition 1.3). We will then use this to show that a.s. the maximal size of the cells of \mathcal{H} which lie at distance at most r from the origin grows sublinearly in r (Lemma 2.9). We will also show that the sum of

$$\text{diam}(H)^2(\pi(H) + \pi^*(H))$$

over all cells H which intersect one of the origin-containing squares S_k is typically of order $|S_k|^2$ (Lemma 2.10).

Definition 2.5. Recall the squares $\widehat{S}_m = \widehat{S}_m^0$ for $m > 0$ from (2.1). For $m > 0$, let \mathcal{F}_m be the σ -algebra generated by the measurable functions $F = F(\mathcal{H}, \mathcal{D})$ of $(\mathcal{H}, \mathcal{D})$ which satisfy

$$F(C(\mathcal{H} - z), C(\mathcal{D} - z)) = F(\mathcal{H}, \mathcal{D}), \quad \text{for all } z \in \widehat{S}_m \text{ and all } C > 0. \tag{2.2}$$

The σ -algebra \mathcal{F}_m encodes $(\mathcal{H}, \mathcal{D})$, viewed modulo translation within \widehat{S}_m and scaling, but not the location of the origin within the square \widehat{S}_m .

LEMMA 2.6. *Every event in $\bigcap_{m > 0} \mathcal{F}_m$ has probability zero or one.*

Proof. If $F = F(\mathcal{H}, \mathcal{D})$ is a $\bigcap_{m > 0} \mathcal{F}_m$ -measurable function, then since the union of the squares \widehat{S}_m is a.s. equal to all of \mathbb{C} , it follows that

$$F(C(\mathcal{H} - z), C(\mathcal{D} - z)) = F(\mathcal{H}, \mathcal{D})$$

for every $C > 0$ and $z \in \mathbb{C}$. Let \mathfrak{m} be the law of \mathcal{D} and let

$$\widetilde{F} = \widetilde{F}(\mathcal{H}) := \int F(\mathcal{H}, \mathcal{D}) \, d\mathfrak{m}(\mathcal{D}),$$

so that \widetilde{F} depends only on \mathcal{H} .

By Lemma 2.1, if $C > 0$ and $z \in \mathbb{C}$ are random and independent from \mathcal{D} (but allowed to depend on \mathcal{H}), then $C(\mathcal{D} - z)$ is a uniform dyadic system independent from \mathcal{H} , i.e.,

$$(C(\mathcal{H} - z), C(\mathcal{D} - z)) \stackrel{d}{=} (C(\mathcal{H} - z), \mathcal{D}).$$

From this and the scale/translation invariance of F , we infer that for any such choice of C and z , we have $\tilde{F}(C(\mathcal{H}-z))=\tilde{F}(\mathcal{H})$. Our ergodicity modulo scaling hypothesis for \mathcal{H} (Definition 1.3) therefore shows that \tilde{F} is a.s. equal to a deterministic constant, i.e., F is a.s. determined by \mathcal{D} . In fact, by the first paragraph F is a.s. determined by \mathcal{D} viewed modulo translation and scaling.

If we are given any arbitrary *deterministic* dyadic system \mathcal{D}' and $R>0$, then we can find $C>0$ and $z\in\mathbb{C}$ (depending on \mathcal{D} and \mathcal{D}') such that the squares surrounding the origin in $C(\mathcal{D}-z)$ and \mathcal{D}' with side lengths in $[1/R, R]$ are identical. Since $R>0$ is arbitrary, a dyadic system is determined by its origin-containing squares, and the successive origin-containing squares in \mathcal{D} are chosen uniformly at random, it follows that F is equal to a deterministic constant a.s. □

From Lemmas 2.3 and 2.6, we obtain the following ergodicity statement, which will be a key tool in what follows.

LEMMA 2.7. *Let $F=F(\mathcal{H}, \mathcal{D})$ be a measurable function on the space of cell configuration/dyadic system pairs which is scale invariant (i.e., $F(C\mathcal{H}, CD)=F(\mathcal{H}, \mathcal{D})$ for each $C>0$). If either $\mathbb{E}[|F|]<\infty$ or $F\geq 0$ a.s., then a.s.*

$$\lim_{k\rightarrow\infty} \frac{1}{|S_k|^2} \int_{S_k} F(\mathcal{H}-z, \mathcal{D}-z) dz = \mathbb{E}[F]. \tag{2.3}$$

Proof. Suppose first that $\mathbb{E}[|F|]<\infty$. By Lemma 2.3, with \mathcal{F}_m as above, it holds for each $m>0$ that

$$\frac{1}{|\hat{S}_m|^2} \int_{\hat{S}_m} F(\mathcal{H}-z, G-z) dz = \mathbb{E}[F | \mathcal{F}_m].$$

Since \mathcal{F}_m is a decreasing sequence of σ -algebras and each S_k is one of the \hat{S}_m 's, the backward martingale convergence theorem [D, Theorem 5.6.1] implies that the limit on the left-hand side of (2.3) exists (as a random variable) a.s. and in L^1 . The limit is $\bigcap_{m>0} \mathcal{F}_m$ -measurable, so is a.s. equal to $\mathbb{E}[F]$ by Lemma 2.6. If $F\geq 0$ a.s. and $\mathbb{E}[F]=\infty$, then applying the above result with $F\wedge C$ for $C>0$ in place of F gives

$$\liminf_{k\rightarrow\infty} \frac{1}{|S_k|^2} \int_{S_k} F(\mathcal{H}-z, \mathcal{D}-z) dz \geq \lim_{k\rightarrow\infty} \frac{1}{|S_k|^2} \int_{S_k} F(\mathcal{H}-z, \mathcal{D}-z) \wedge C dz = \mathbb{E}[F \wedge C],$$

which tends to ∞ as $C\rightarrow\infty$. □

Lemma 2.7 will allow us to prove a number of properties of the origin-containing squares $\{S_k\}_{k\in\mathbb{Z}}$, but we will also have occasion to consider other squares in \mathcal{D} . The following lemma will enable us to do so. For the statement, we emphasize that in our terminology each square S has 4^l dyadic ancestors of side length $2^l|S|$ for each $l\in\mathbb{N}$.

LEMMA 2.8. *Let $E(S)=E(S, \mathcal{H})$ be an event depending on a square $S \subset \mathbb{C}$ and our cell configuration \mathcal{H} . Suppose that a.s. $E(S_k)$ occurs for each large enough $k \in \mathbb{N}$. Then, for each $l \in \mathbb{N}$, it is a.s. the case that for large enough $k \in \mathbb{N}$, $E(S)$ occurs for each dyadic ancestor S of S_k of side length at most $2^l |S_k|$ (even those which do not belong to \mathcal{D}) and each dyadic descendant S of S_k of side length at least $2^{-l} |S_k|$.*

Proof. We first consider dyadic ancestors. For $k \in \mathbb{N}$, let \mathcal{K}_k be the smallest $k' \geq k$ for which $E(S)^c$ occurs for at least one of the dyadic ancestor S of $S_{k'}$ of side length at most $2^l |S_{k'}|$, or $\mathcal{K}_k = \infty$ if no such k' exists. We want to show that $\mathbb{P}[\mathcal{K}_k = \infty] \rightarrow 1$ as $k \rightarrow \infty$.

We observe that \mathcal{K}_k is a stopping time for the filtration generated by \mathcal{H} and $\{S_j\}_{j \leq k'}$, so if $\mathcal{K}_k < \infty$ and we condition on \mathcal{H} and $\{S_j\}_{j \leq k'}$ then for each $r \in \mathbb{N}$ the square $S_{\mathcal{K}_k+r}$ is sampled uniformly from the 4^r possibilities. On the event $\{\mathcal{K}_k < \infty\}$, let S_* be a dyadic ancestor of $S_{\mathcal{K}_k}$ with side length at most $2^l |S_{\mathcal{K}_k}|$ for which $E(S_*)^c$ occurs, chosen in a manner depending only on \mathcal{H} and $\{S_j\}_{j \leq \mathcal{K}_k}$ and let $l_* \in [1, l]_{\mathbb{Z}}$ be chosen so that $|S_*| = 2^{l_*} |S_{\mathcal{K}_k}|$. Then

$$\mathbb{P}[\mathcal{K}_k < \infty] \leq \frac{\mathbb{P}[E(S_{\mathcal{K}_k+l_*})^c]}{\mathbb{P}[S_{k+l_*} = S_* | \mathcal{K}_k < \infty]} \leq 4^l \mathbb{P}[E(S_{\mathcal{K}_k+l_*})^c],$$

which tends to zero as $k \rightarrow \infty$ since a.s. $E(S_k)$ occurs for each large enough k .

Next, we consider dyadic descendants. We will use the re-sampling property of Lemma 2.3. For $m \in \mathbb{N}$, let M_m be the largest $m' \leq m$ for which $E(S)^c$ occurs for at least one dyadic descendant S of $\widehat{S}_{m'}$ of side length at least $2^{-l} |\widehat{S}_{m'}|$, or $M_m = -\infty$ if there is no such m' . We want to show that a.s. $\limsup_{m \rightarrow \infty} M_m < \infty$, equivalently a.s. $\limsup_{m \rightarrow \infty} K_{M_m} < \infty$.

For $m > 0$, let $K_m \in \mathbb{Z}$ be the largest $k \in \mathbb{Z}$ for which $\widehat{S}_m = S_k$. It suffices to show that $\mathbb{P}[K_{M_m} > k] \rightarrow 0$ as $k \rightarrow \infty$, uniformly over all $m > 0$. Each M_m is a reverse stopping time for the filtration of Definition 2.5. By Lemma 2.3, if we condition on \mathcal{F}_{M_m} then for each $r \in \mathbb{N}$, the square $S_{K_{M_m}-r}$ is sampled uniformly from the 4^r dyadic descendants of $S_{K_{M_m}}$ with side length $2^{-r} |S_{K_{M_m}}|$. As above, on the event $\{M_m \neq -\infty\}$ we let S_* be a dyadic descendant of $S_{K_{M_m}} = \widehat{S}_{M_m}$ with side length at least $2^{-l} |\widehat{S}_{M_m}|$ for which $E(S_*)^c$ occurs, chosen in a \mathcal{F}_{M_m} -measurable manner and we let $l_* \in [1, l]_{\mathbb{Z}}$ be chosen so that $|S_*| = 2^{-l_*} |\widehat{S}_{M_m}|$. Then

$$\mathbb{P}[K_{M_m} > k] \leq \frac{\mathbb{P}[E(S_{K_{M_m}-l_*})^c, K_{M_m} > k]}{\mathbb{P}[S_{K_{M_m}-l_*} = S_* | K_{M_m} > k]} \leq 4^l \mathbb{P}[E(S_{K_{M_m}-l_*})^c, K_{M_m} > k],$$

which tends to zero as $k \rightarrow \infty$, uniformly in m , since a.s. $E(S_k)$ occurs for large enough k . Since K_{M_m} is non-decreasing in m , this implies that a.s. $\limsup_{m \rightarrow \infty} K_{M_m} < \infty$. \square

In what follows, we will apply Lemma 2.7 with

$$F(\mathcal{H}, \mathcal{D}) = \frac{\text{diam}(H_0)^2}{\text{Area}(H_0)} (\pi(H_0) + \pi^*(H_0)). \tag{2.4}$$

The random variable F has finite expectation by hypothesis. We first establish that the maximum cell size of \mathcal{H} grows sublinearly in the distance to the origin.

LEMMA 2.9. *Almost surely, for each $\varepsilon \in (0, 1)$ it holds for large enough $k \in \mathbb{N}$ that*

$$\text{diam}(H) \leq \varepsilon |S_k| \quad \text{for all } H \in \mathcal{H}(S_k). \tag{2.5}$$

Proof. For $\varepsilon \in (0, 1)$, let

$$M_\varepsilon = M_\varepsilon(\mathcal{H}, \mathcal{D}) := \min\{m > 0 : \text{diam}(H_0) \leq \varepsilon |\widehat{S}_m|\}.$$

Since F , defined as in (2.4), has finite expectation there is a deterministic $m_\varepsilon \in \mathbb{N}$ such that $\mathbb{E}[F \mathbf{1}_{(M_\varepsilon > m_\varepsilon)}] \leq \varepsilon^3$. Note that $F \mathbf{1}_{(M_\varepsilon > m_\varepsilon)}$ is a scale-invariant random variable in the sense of Lemma 2.7. By Lemma 2.7 applied with $F \mathbf{1}_{(M_\varepsilon > m_\varepsilon)}$ in place of F , we get that a.s.

$$\limsup_{m \rightarrow \infty} \frac{1}{|\widehat{S}_m|^2} \int_{\widehat{S}_m} \frac{\text{diam}(H_z)^2}{\text{Area}(H_z)} (\pi(H_z) + \pi^*(H_z)) \mathbf{1}_{(\text{diam}(H_z) > \varepsilon |\widehat{S}_m|)} dz \leq \varepsilon^3 \tag{2.6}$$

where here we use that $\text{diam}(H_z) > \varepsilon |\widehat{S}_m|$ implies $M_\varepsilon(\mathcal{H} - z, \mathcal{D} - z) > m_\varepsilon$ whenever $m \geq m_\varepsilon$ and $z \in \widehat{S}_m$. By breaking up the integral over \widehat{S}_m into the sum of the integrals over $H \cap \widehat{S}_m$ for cells $H \in \mathcal{H}(\widehat{S}_m)$, then dropping the boundary terms, and using that

$$\pi(H_z) + \pi^*(H_z) \geq 1,$$

we see that (2.6) implies that

$$\limsup_{m \rightarrow \infty} \sum_{H \in \mathcal{H}(\widehat{S}_m) \setminus \mathcal{H}(\partial \widehat{S}_m)} \frac{\text{diam}(H)^2}{|\widehat{S}_m|^2} \mathbf{1}_{(\text{diam}(H) > \varepsilon |\widehat{S}_m|)} \leq \varepsilon^3. \tag{2.7}$$

Each non-zero summand on the left-hand side of (2.7) is at least ε^2 , so it follows from (2.7) (and the fact that each S_k is one of the \widehat{S}_m 's) that a.s. for large enough $k \in \mathbb{N}$,

$$\text{diam}(H) \leq \varepsilon |S_k| \quad \text{for all } H \in \mathcal{H}(S_k) \setminus \mathcal{H}(\partial S_k). \tag{2.8}$$

It remains to deal with cells which intersect ∂S_k . For this purpose, consider a fixed $k \in \mathbb{N}$ and let k' be the smallest integer $\geq k$ for which each cell in $\mathcal{H}(S_{k'})$ has diameter smaller than $\varepsilon |S_k|$ (such a k' exists since each cell is compact). We claim that a.s. $k' = k$ for large enough k , which will show (2.5).

Indeed, since $\varepsilon < 1$ each $H \in \mathcal{H}(S_k)$ is contained in one of the 16 dyadic ancestors of $S_{k'}$ with side length $4|S_{k'}|$. By (2.8) (applied with $\frac{1}{16}\varepsilon$ in place of ε) together with Lemma 2.8, it is a.s. the case that for large enough $k \in \mathbb{N}$, each cell H which is contained in any one of these 16 dyadic ancestors has diameter at most $\frac{1}{4}\varepsilon|S_{k'}|$. By the definition of k' , if $k' \neq k$ then there is some cell in $\mathcal{H}(S_k) \subset \mathcal{H}(S_{k'})$ with diameter at least $\frac{1}{2}\varepsilon|S_{k'}|$. Therefore a.s. $k' = k$ for large enough k . \square

LEMMA 2.10. *There is a deterministic constant $C > 0$ such that a.s.*

$$\limsup_{k \rightarrow \infty} \frac{1}{|S_k|^2} \sum_{H \in \mathcal{H}(S_k)} \text{diam}(H)^2 (\pi(H) + \pi^*(H)) \leq C. \tag{2.9}$$

Proof. By Lemma 2.7 applied with F as in (2.4), we find that there is a deterministic, finite constant $C_0 > 0$ such that a.s.

$$\lim_{k \rightarrow \infty} \frac{1}{|S_k|^2} \int_{S_k} \frac{\text{diam}(H_z)^2}{\text{Area}(H_0)} (\pi(H_z) + \pi^*(H_z)) dz = C_0.$$

By breaking up the integral over S_k into the sum of the integrals over $H \cap S_k$ for cells $H \in \mathcal{H}(S_k)$, then dropping the boundary terms, we get that a.s.

$$\limsup_{k \rightarrow \infty} \frac{1}{|S_k|^2} \sum_{H \in \mathcal{H}(S_k) \setminus \mathcal{H}(\partial S_k)} \text{diam}(H)^2 (\pi(H) + \pi^*(H)) \leq C_0. \tag{2.10}$$

To deal with the cells which intersect ∂S_k , we observe that Lemma 2.9 (with $\varepsilon = \frac{1}{2}$, say) implies that a.s. for large enough $k \in \mathbb{N}$ each cell in $\mathcal{H}(S_k)$ is contained in the interior of an appropriate dyadic ancestor of S_k of side length $4|S_k|$. By (2.10) and Lemma 2.8 (applied with $l = 2$), we find that (2.3) holds a.s. with $C = 16C_0$. \square

2.3. Existence of the corrector

In this subsection we will construct the discrete harmonic function ϕ_∞ appearing in Theorem 1.6 as a limit of functions $\phi_m: \mathcal{H} \rightarrow \mathbb{C}$ for $m \in \mathbb{N}_0$ which interpolate between the function ϕ_0 of Theorem 1.6 (which sends each cell to its Euclidean center) at $m = 0$ and the discrete harmonic function ϕ_∞ as $m \rightarrow \infty$. Roughly speaking, ϕ_m will agree with the a-priori embedding on the cells which intersect boundary of each square of \mathcal{D} which contains approximately m cells of \mathcal{H} ; and will be discrete harmonic on the cells contained in the interior of each such square.

Throughout this section, we define the square \widehat{S}_m^z for $z \in \mathbb{C}$ and $m > 0$ as in (2.1) and we define the collection of squares

$$\mathcal{S}_m := \{\widehat{S}_m^z : z \in \mathbb{C} \text{ and } z \text{ is not on the boundary of a square in } \mathcal{D}\} \subset \mathcal{D}. \tag{2.11}$$

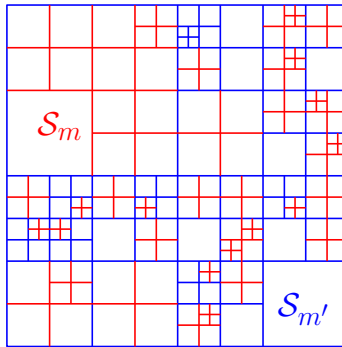


Figure 6. The collections of squares \mathcal{S}_m and $\mathcal{S}_{m'}$ of \mathcal{D} for $m < m'$. The squares of \mathcal{S}_m (resp. $\mathcal{S}_{m'}$) are the minimal squares with red and/or blue (resp. blue) boundary and each contain exactly m (resp. m') cells of \mathcal{H} if we count fractional parts of cells.

If z' is in the interior of \widehat{S}_m^z , then $\widehat{S}_m^{z'} = \widehat{S}_m^z$, so the squares in \mathcal{S}_m intersect only along their boundaries. Clearly, these squares cover all of \mathbb{C} . See Figure 6 for an illustration of these squares.

We now use the set of squares \mathcal{S}_m to define the above-mentioned functions $\phi_m: \mathcal{H} \rightarrow \mathbb{C}$ for $m \in \mathbb{N}_0$. As in Theorem 1.6, let

$$\phi_0(H) := \frac{1}{\text{Area}(H)} \int_H z \, dz.$$

For $m \in \mathbb{N}$, let $\phi_m: \mathcal{H} \rightarrow \mathbb{C}$ be the function which agrees with ϕ_0 on $\mathcal{H}(\partial \widehat{S})$ for each $\widehat{S} \in \mathcal{S}_m$ and is \mathfrak{c} -discrete harmonic on $\mathcal{H}(\widehat{S}) \setminus \mathcal{H}(\partial \widehat{S})$ for each such \widehat{S} .

We observe that if $C > 0$ and $z \in \mathbb{C}$ and we replace $(\mathcal{H}, \mathcal{D})$ by the scaled/translated pair

$$(C(\mathcal{H} - z), C(\mathcal{D} - z)),$$

then the corresponding functions ϕ_m are replaced by $C(\phi_m(\cdot/C + z) - z)$. This will be important when we apply Lemma 2.7.

The main goal of this subsection is to establish the following proposition.

PROPOSITION 2.11. *There exists a \mathfrak{c} -discrete harmonic function $\phi_\infty: \mathcal{H} \rightarrow \mathbb{C}$ such that*

$$\phi_m - \phi_m(H_0) \rightarrow \phi_\infty \quad \text{in probability as } m \rightarrow \infty, \tag{2.12}$$

with respect to the topology of uniform convergence on finite subsets of \mathcal{H} , i.e., for each finite set $\mathcal{A} \subset \mathcal{H}$ (chosen in some measurable manner) and each $\varepsilon > 0$,

$$\lim_{m \rightarrow \infty} \mathbb{P} \left[\max_{H \in \mathcal{A}} |\phi_m(H) - \phi_m(H_0) - \phi_\infty(H)| > \varepsilon \right] = 0.$$

Furthermore, ϕ_∞ is covariant with respect to scaling and translation in the sense that if $C > 0$ and $z \in \mathbb{C}$ and we replace $(\mathcal{H}, \mathcal{D})$ by the scaled/translated pair $(C(\mathcal{H} - z), C(\mathcal{D} - z))$, then ϕ_∞ is replaced by $C(\phi_\infty(\cdot/C + z) - \phi_\infty(H_z))$.

The idea of the proof of Proposition 2.11 is as follows. We first use Lemma 2.10 to argue that ϕ_0 (and hence each ϕ_m) has uniformly bounded Dirichlet energy over S_k when we rescale by $|S_k|^{-2}$ (Lemma 2.12). Once this is established, we will use the orthogonality of the increments $\phi_m - \phi_{m-1}$ with respect to the Dirichlet inner product to show that, for large enough $m \in \mathbb{N}$, the expected discrete “Dirichlet energy per unit area” of $\phi_m - \phi_{m'}$, which is made precise in (2.15), is small uniformly over all $m' \geq m$ (Lemma 2.14). This will allow us to show that the functions $\phi_m - \phi_m(H_0)$ are Cauchy in probability, and thereby establish Proposition 2.11.

LEMMA 2.12. *For each $n, m \in \mathbb{N}_0$ with $n \geq m$, we have (in the notation of Definition 1.4)*

$$\text{Energy}(\phi_m; \mathcal{H}(\widehat{S}_n)) \leq \text{Energy}(\phi_0; \mathcal{H}(\widehat{S}_n)). \tag{2.13}$$

Furthermore, there is a deterministic constant $C > 0$ such that a.s.

$$\limsup_{k \rightarrow \infty} \frac{1}{|S_k|^2} \text{Energy}(\phi_0; \mathcal{H}(S_k)) \leq C. \tag{2.14}$$

Proof. The relation (2.13) is immediate from the fact that discrete harmonic functions minimize Dirichlet energy subject to specified boundary data, so we only need to prove (2.14). For an edge $\{H, H'\} \in \mathcal{EH}(S_k)$, we have $H \cap H' \neq \emptyset$, and so

$$|\phi_0(H) - \phi_0(H')|^2 \leq (\text{diam}(H) + \text{diam}(H'))^2 \leq 2 \text{diam}(H)^2 + 2 \text{diam}(H')^2.$$

Breaking up the sum over edges into a sum over vertices, we get

$$\begin{aligned} \text{Energy}(\phi_0; \mathcal{H}(S_k)) &\leq 2 \sum_{\{H, H'\} \in \mathcal{EH}(S_k)} \mathbf{c}(H, H') (\text{diam}(H)^2 + \text{diam}(H')^2) \\ &\leq 4 \sum_{H \in \mathcal{H}(S_k)} \text{diam}(H)^2 \pi(H). \end{aligned}$$

We now conclude by means of Lemma 2.10. □

For $m_1, m_2 \in \mathbb{N}_0$ with $m_1 < m_2$, let

$$\phi_{m_1, m_2} := \phi_{m_2} - \phi_{m_1}.$$

We define

$$\text{SE}_{m_1, m_2} := \sum_{\substack{H \in \mathcal{H}(\widehat{S}_{m_2}) \\ H \sim H_0}} \frac{\mathbf{c}(H_0, H) |\phi_{m_1, m_2}(H) - \phi_{m_1, m_2}(H_0)|^2}{2 \text{Area}(H_0 \cap \widehat{S}_{m_2})}. \tag{2.15}$$

We can think of SE_{m_1, m_2} as representing the “Dirichlet energy per unit area” of ϕ_{m_1, m_2} , so that $\mathbb{E}[\text{SE}_{m_1, m_2}]$ is the “specific Dirichlet energy” of ϕ_{m_1, m_2} . The following lemma makes this notion precise.

LEMMA 2.13. *For each fixed $m_1, m_2 \in \mathbb{N}_0$ with $m_1 < m_2$, almost surely*

$$\lim_{n \rightarrow \infty} \frac{1}{|\widehat{S}_n|^2} \text{Energy}(\phi_{m_1, m_2}; \mathcal{H}(\widehat{S}_n)) = \mathbb{E}[\text{SE}_{m_1, m_2}], \tag{2.16}$$

and this expectation is bounded above by a finite constant which does not depend on m_1 and m_2 .

Proof. For $z \in \mathbb{C}$, let SE_{m_1, m_2}^z be defined as in (2.15) with the translated cell configuration/dyadic system $(\mathcal{H} - z, \mathcal{D} - z)$ in place of $(\mathcal{H}, \mathcal{D})$. Equivalently,

$$\text{SE}_{m_1, m_2}^z := \sum_{\substack{H \in \mathcal{H}(\widehat{S}_{m_2}^z) \\ H \sim H_z}} \frac{\mathfrak{c}(H_z, H) |\phi_{m_1, m_2}(H) - \phi_{m_1, m_2}(H_z)|^2}{2 \text{Area}(H_z \cap \widehat{S}_{m_2}^z)}. \tag{2.17}$$

Lemma 2.7 implies that a.s.

$$\lim_{n \rightarrow \infty} \frac{1}{|\widehat{S}_n|^2} \int_{\widehat{S}_n} \text{SE}_{m_1, m_2}^z \, dz = \mathbb{E}[\text{SE}_{m_1, m_2}]. \tag{2.18}$$

We will now argue that this integral is exactly equal to $\text{Energy}(\phi_{m_1, m_2}; \mathcal{H}(\widehat{S}_n))$ for $n \geq m_2$.

We first consider the case $n = m_2$. By breaking up the integral over \widehat{S}_{m_2} into integrals over $H \cap \widehat{S}_{m_2}$ for $H \in \mathcal{H}(\widehat{S}_{m_2})$, we find that

$$\begin{aligned} \int_{\widehat{S}_{m_2}} \text{SE}_{m_1, m_2}^z \, dz &= \frac{1}{2} \sum_{H \in \mathcal{H}(\widehat{S}_{m_2})} \sum_{\substack{H' \in \mathcal{H}(\widehat{S}_{m_2}) \\ H' \sim H}} \mathfrak{c}(H, H') |\phi_{m_1, m_2}(H') - \phi_{m_1, m_2}(H)|^2 \\ &= \text{Energy}(\phi_{m_1, m_2}; \mathcal{H}(\widehat{S}_{m_2})). \end{aligned} \tag{2.19}$$

For $n \geq m_2$, we can break up the integral over \widehat{S}_n as a sum of integrals over the squares of the form $\widehat{S}_{m_2}^z$ for $z \in \widehat{S}_n$ to get

$$\int_{\widehat{S}_n} \text{SE}_{m_1, m_2}^z \, dz = \sum_{\substack{\widehat{S} \in \mathcal{S}_{m_2} \\ \widehat{S} \subset \widehat{S}_n}} \text{Energy}(\phi_{m_1, m_2}; \mathcal{H}(\widehat{S})). \tag{2.20}$$

The squares $\widehat{S} \in \mathcal{S}_{m_2}$ with $\widehat{S} \subset \widehat{S}_n$ intersect only along their boundaries and their union is all of \widehat{S}_n . Therefore, the sum on the right-hand side of (2.20) is equal to the sum of

$$\mathfrak{c}(H, H') |\phi_{m_1, m_2}(H') - \phi_{m_1, m_2}(H)|^2$$

over all edges $\{H, H'\} \in \mathcal{EH}(\widehat{S}_n)$ except that edges are counted with multiplicity equal to the number of such squares \widehat{S} for which $\{H, H'\} \in \mathcal{EH}(\widehat{S})$. Since $H \sim H'$ implies that

$H \cap H' \neq \emptyset$, any edge $\{H, H'\}$ which is counted more than once must have $H, H' \in \mathcal{EH}(\partial\widehat{S})$ for some $\widehat{S} \in \mathcal{S}_{m_2}$. Both ϕ_{m_1} and ϕ_{m_2} coincide with ϕ_0 on the boundary of each such square \widehat{S} , so $|\phi_{m_1, m_2}(H') - \phi_{m_1, m_2}(H)|^2 = 0$ for each such edge $\{H, H'\}$. Therefore, (2.20) equals $\text{Energy}(\phi_{m_1, m_2}; \mathcal{H}(\widehat{S}_n))$. Combining this with (2.18) shows that (2.16) holds. The limiting value is uniformly bounded above by Lemma 2.12 and the triangle inequality for the Dirichlet norm. \square

We now argue that the specific Dirichlet energy $\mathbb{E}[\text{SE}_{m, m'}]$ has to be small when m and m' are large. Essentially, the reason for this is that the Dirichlet energy of ϕ_0 gets “used up” in the increments corresponding to indices smaller than m and m' .

LEMMA 2.14. *For each $\varepsilon > 0$, there exists $m_\varepsilon \in \mathbb{N}$ such that*

$$\mathbb{E}[\text{SE}_{m_\varepsilon, m}] \leq \varepsilon \quad \text{for all } m \geq m_\varepsilon. \tag{2.21}$$

Proof. For each $k \in \mathbb{N}$, the function $\phi_{m-1, m}$ vanishes on $\mathcal{H}(\bigcup_{\widehat{S} \in \mathcal{S}_m} \partial\widehat{S})$ (since both ϕ_m and ϕ_{m-1} agree with ϕ_0 there) and is discrete harmonic on $\mathcal{H} \setminus \mathcal{H}(\bigcup_{\widehat{S} \in \mathcal{S}_m} \partial\widehat{S})$. Therefore, the discrete analog of Green’s formula (summation by parts) implies that $\phi_{m-1, m}$ and $\phi_{m'-1, m'}$ are orthogonal with respect to the Dirichlet inner product for $m \neq m'$. Hence, for $m \in \mathbb{N}$ and $n \geq m$,

$$\text{Energy}(\phi_{0, m}; \mathcal{H}(\widehat{S}_n)) = \sum_{j=1}^m \text{Energy}(\phi_{j-1, j}; \mathcal{H}(\widehat{S}_n)). \tag{2.22}$$

Dividing both sides of (2.22) by $|\widehat{S}_n|^2$, sending $n \rightarrow \infty$, and applying Lemma 2.13 shows that

$$\mathbb{E}[\text{SE}_{0, m}] = \sum_{j=1}^m \mathbb{E}[\text{SE}_{j-1, j}]. \tag{2.23}$$

By the last part of Lemma 2.13, this last sum is uniformly bounded above, independently of m . Therefore, $\sum_{j=1}^\infty \mathbb{E}[\text{SE}_{j-1, j}] < \infty$ so for any $\varepsilon > 0$ we can find $m_\varepsilon \in \mathbb{N}$ such that

$$\mathbb{E}[\text{SE}_{m_\varepsilon, m}] \leq \sum_{j=m_\varepsilon+1}^\infty \mathbb{E}[\text{SE}_{j-1, j}] \leq \varepsilon \quad \text{for all } m \geq m_\varepsilon. \tag{2.24}$$

\square

Proof of Proposition 2.11. We will show that $\phi_m - \phi_m(H_0)$ is Cauchy in probability with respect to the local uniform topology on \mathcal{H} using Lemma 2.14. To this end, fix $l \in \mathbb{N}$; we will work on \widehat{S}_l . Given $\varepsilon > 0$, let $m_\varepsilon \in \mathbb{N}$ be as in Lemma 2.14 and define $\text{SE}_{m_\varepsilon, m}^z$ as

in (2.17). Then, if $m \geq l \vee m_\varepsilon$,

$$\begin{aligned} \int_{\widehat{S}_l} \text{SE}_{m_\varepsilon, m}^z dz &= \frac{1}{2} \sum_{H \in \mathcal{H}(\widehat{S}_l)} \frac{\text{Area}(H \cap \widehat{S}_l)}{\text{Area}(H \cap \widehat{S}_m)} \sum_{\substack{H' \in \mathcal{H}(\widehat{S}_m) \\ H' \sim H}} \mathbf{c}(H, H') |\phi_{m_\varepsilon, m}(H') - \phi_{m_\varepsilon, m}(H)|^2 \\ &\geq \frac{1}{2} \sum_{H \in \mathcal{H}(\widehat{S}_l) \setminus \mathcal{H}(\partial \widehat{S}_l)} \sum_{\substack{H' \in \mathcal{H}(\widehat{S}_m) \\ H' \sim H}} \mathbf{c}(H, H') |\phi_{m_\varepsilon, m}(H') - \phi_{m_\varepsilon, m}(H)|^2 \\ &\geq \text{Energy}(\phi_{m_\varepsilon, m}; \mathcal{H}(\widehat{S}_l) \setminus \mathcal{H}(\partial \widehat{S}_l)). \end{aligned} \tag{2.25}$$

Dividing by $|\widehat{S}_l|^2$ in (2.25), taking expectations, and applying Lemma 2.3 and our choice of m_ε to bound the expectation of the integral over \widehat{S}_l , we get that

$$\sup_{m \geq m_\varepsilon} \mathbb{E} \left[\frac{1}{|\widehat{S}_l|^2} \text{Energy}(\phi_{m_\varepsilon, m}; \mathcal{H}(\widehat{S}_l) \setminus \mathcal{H}(\partial \widehat{S}_l)) \right] \leq \varepsilon. \tag{2.26}$$

The variation of $\phi_{m_\varepsilon, m}$ over any single edge $\{H, H'\}$ of $\mathcal{H}(\widehat{S}_l) \setminus \mathcal{H}(\partial \widehat{S}_l)$ is bounded above by $\mathbf{c}(H, H')^{-1/2}$ times the square root of the energy of $\phi_{m_\varepsilon, m}$ over $\mathcal{H}(\widehat{S}_l) \setminus \mathcal{H}(\partial \widehat{S}_l)$. Summing over all such edges and applying (2.26) shows that, for each fixed $l \in \mathbb{N}$,

$$\sup_{m \geq m_\varepsilon} \mathbb{E} \left[\frac{\max_{H \in \mathring{\mathcal{H}}(\widehat{S}_l)} |\phi_{m_\varepsilon, m}(H) - \phi_{m_\varepsilon, m}(H_0)|}{|\widehat{S}_l| \sum_{\{H, H'\} \in \mathcal{E} \mathring{\mathcal{H}}(\widehat{S}_l)} \mathbf{c}(H, H')^{-1/2}} \right] \leq \varepsilon^{1/2}, \tag{2.27}$$

where $\mathring{\mathcal{H}}(\widehat{S}_l)$ denotes the connected component of $\mathcal{H}(\widehat{S}_l) \setminus \mathcal{H}(\partial \widehat{S}_l)$ which contains H_0 . Since $\varepsilon > 0$ is arbitrary and l is fixed, the sequence $\{\phi_m - \phi_m(H_0)\}_{m \in \mathbb{N}}$ is Cauchy in probability with respect to the topology of uniform convergence on $\mathring{\mathcal{H}}(\widehat{S}_l)$. By the local finiteness of \mathcal{H} , a.s. each finite subgraph of \mathcal{H} is contained in $\mathring{\mathcal{H}}(\widehat{S}_l)$ for large enough $l \in \mathbb{N}$. Therefore, there is a function $\phi_\infty: \mathcal{H} \rightarrow \mathbb{C}$ such that (2.12) holds.

The function ϕ_∞ is a.s. discrete harmonic since each ϕ_m is discrete harmonic on $\mathcal{H}(\widehat{S}_m) \setminus \mathcal{H}(\partial \widehat{S}_m)$ and the pointwise limit of discrete harmonic functions is discrete harmonic. The scaling/translation property of ϕ_∞ follows from the analogous property for the ϕ_m 's. □

2.4. Sublinearity of the corrector

In this subsection we prove (1.17) of Theorem 1.6 for the discrete harmonic function ϕ_∞ of Proposition 2.11 (this corresponds to Theorem 1.6 in the special case when \mathcal{H} is the planar dual of an embedded lattice). We continue to use the notation of §2.3. We will first argue that the Dirichlet energy per unit area of $\phi_\infty - \phi_m$ is small when m is

large (Lemma 2.16), which is a consequence of Lemmas 2.7 and 2.14. We will then show (essentially using Hölder’s inequality) that this implies that the total variation of $\phi_\infty - \phi_m$ over most horizontal and vertical line segments of the square S_k is of smaller order than $|S_k|$ when m is large (but fixed) and k is very large (Lemma 2.17). It is easily seen from the definition of ϕ_m that for any fixed $m \in \mathbb{N}$, the maximum value of $|\phi_m - \phi_0|$ over S_k is $o_k(|S_k|)$ as $k \rightarrow \infty$ (Lemma 2.18). We will then conclude by combining our bounds for $\phi_\infty - \phi_m$ and $\phi_m - \phi_0$ and using the maximum principle for discrete harmonic functions to transfer from a bound for “most” horizontal and vertical line segments to a bound which holds for all points simultaneously (Lemma 2.19).

Extending the definition (2.15), for $m \in \mathbb{N}$ we set $\phi_{m,\infty} := \phi_\infty - \phi_m$ and define

$$SE_{m,\infty} := \sum_{\substack{H \in \mathcal{H} \\ H \sim H_0}} \frac{\mathbf{c}(H_0, H) |\phi_{m,\infty}(H) - \phi_{m,\infty}(H_0)|^2}{2\text{Area}(H_0)}. \tag{2.28}$$

LEMMA 2.15. *We have $\mathbb{E}[S_{m,\infty}] < \infty$ for each $m \in \mathbb{N}_0$ and*

$$\lim_{m \rightarrow \infty} \mathbb{E}[S_{m,\infty}] = 0. \tag{2.29}$$

Proof. It is clear from Proposition 2.11 that, for each fixed $m \in \mathbb{N}$, $SE_{m,m'} \rightarrow SE_{m,\infty}$ in probability as $m' \rightarrow \infty$. By Fatou’s lemma,

$$\mathbb{E}[SE_{m,\infty}] \leq \liminf_{m' \rightarrow \infty} \mathbb{E}[SE_{m,m'}],$$

which is finite by Lemma 2.13. By Lemma 2.14,

$$\lim_{m \rightarrow \infty} \liminf_{m' \rightarrow \infty} \mathbb{E}[SE_{m,m'}] = 0. \quad \square$$

From Lemma 2.15 we can deduce the following.

LEMMA 2.16. *Almost surely,*

$$\lim_{m \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{1}{|\widehat{S}_n|^2} \text{Energy}(\phi_{m,\infty}; \mathcal{H}(\widehat{S}_n)) = 0. \tag{2.30}$$

Proof. For $z \in \mathbb{C}$, let $SE_{m,\infty}^z$ be defined as in (2.28) with H_z in place of H_0 . Equivalently, $SE_{m,\infty}^z$ is defined in the same manner as $SE_{m,\infty}$ but with $(\mathcal{H} - z, \mathcal{D} - z)$ in place of $(\mathcal{H}, \mathcal{D})$. By Lemma 2.7, for each fixed $m \in \mathbb{N}$, a.s.

$$\lim_{n \rightarrow \infty} \frac{1}{|\widehat{S}_n|^2} \int_{\widehat{S}_n} SE_{m,\infty}^z dz = \mathbb{E}[SE_{m,\infty}]. \tag{2.31}$$

As in (2.25), for $m, n \in \mathbb{N}$,

$$\begin{aligned} \int_{\widehat{S}_n} \text{SE}_{m,\infty}^z dz &= \frac{1}{2} \sum_{H \in \mathcal{H}(\widehat{S}_n)} \frac{\text{Area}(H \cap \widehat{S}_n)}{\text{Area}(H)} \sum_{\substack{H' \in \mathcal{H} \\ H' \sim H}} \mathfrak{c}(H, H') |\phi_{m,\infty}(H') - \phi_{m,\infty}(H)|^2 \\ &\geq \text{Energy}(\phi_{m,\infty}; \mathcal{H}(\widehat{S}_n) \setminus \mathcal{H}(\partial \widehat{S}_n)). \end{aligned} \quad (2.32)$$

Combining this with (2.31) shows that a.s.

$$\limsup_{n \rightarrow \infty} \frac{1}{|\widehat{S}_n|^2} \text{Energy}(\phi_{m,\infty}; \mathcal{H}(\widehat{S}_n) \setminus \mathcal{H}(\partial \widehat{S}_n)) \leq \mathbb{E}[\text{SE}_{m,\infty}]. \quad (2.33)$$

Lemma 2.9 implies that a.s. for large enough $n \in \mathbb{N}$, each cell in $\mathcal{H}(\widehat{S}_n)$ has diameter at most $\frac{1}{2}|\widehat{S}_n|$, and hence is contained in the union of the four dyadic parents of \widehat{S}_n . Furthermore, (2.33) combined with Lemma 2.8 implies that a.s. for large enough $n \in \mathbb{N}$, the relation (2.33) holds with \widehat{S}_n replaced by any of the four dyadic parents of \widehat{S}_n . Consequently, a.s.

$$\limsup_{n \rightarrow \infty} \frac{1}{|\widehat{S}_n|^2} \text{Energy}(\phi_{m,\infty}; \mathcal{H}(\widehat{S}_n)) \leq 4\mathbb{E}[\text{SE}_{m,\infty}] \quad (2.34)$$

which, together with Lemma 2.15, implies (2.30). \square

Lemma 2.16 gives us an upper bound for the discrete Dirichlet energy of $\phi_\infty - \phi_m$ when m is large. The following lemma allows us to transfer this to a pointwise bound.

LEMMA 2.17. *There is a deterministic constant $C > 0$ such that the following is true a.s. For $k \in \mathbb{N}$ and $r \in [0, 1]$, let $P_k(r)$ be the horizontal line segment connecting the left and right boundaries of the square S_k whose distance from the lower boundary of S_k is $r|S_k|$. Then, for each large enough $k \in \mathbb{N}$ and each function $f: \mathcal{H}(S_k) \rightarrow \mathbb{C}$,*

$$\int_0^1 \sum_{\{H, H'\} \in \mathcal{EH}(P_k(r))} |f(H) - f(H')| dr \leq C \text{Energy}(f; \mathcal{H}(S_k))^{1/2}. \quad (2.35)$$

The same is true if we replace horizontal line segments by vertical line segments.

Proof. By Lemma 2.10, there is a deterministic constant $C_0 > 0$ such that a.s. for sufficiently large $k \in \mathbb{N}$,

$$\frac{1}{|S_k|^2} \sum_{H \in \mathcal{H}(S_k)} \text{diam}(H)^2 \pi^*(H) \leq C_0, \quad (2.36)$$

where here we recall the definition of π^* from (1.4). Assume now that (2.36) holds. For $k \in \mathbb{N}$ and an edge $\{H, H'\} \in \mathcal{EH}(S_k)$, let $L_k(H, H')$ be the Lebesgue measure of the set of

$r \in [0, 1]$ for which H and H' both belong to $\mathcal{H}(P_k(r))$. Then, by interchanging integration and summation, we get

$$\int_0^1 \sum_{\{H, H'\} \in \mathcal{EH}(P_k(r))} |f(H) - f(H')| dr = \sum_{\{H, H'\} \in \mathcal{EH}(S_k)} L_k(H, H') |f(H) - f(H')|. \tag{2.37}$$

The quantity $L_k(H, H')$ is bounded above by the Lebesgue measure of the set of r for which $H \in \mathcal{H}(P_k(r))$, so

$$L_k(H, H') \leq \frac{2}{|S_k|} \text{diam}(H). \tag{2.38}$$

By (2.37), (2.38), and the Cauchy–Schwarz inequality, we see that if (2.36) holds then

$$\begin{aligned} & \int_0^1 \sum_{\{H, H'\} \in \mathcal{EH}(P_k(r))} |f(H) - f(H')| dr \\ & \leq \frac{2}{|S_k|} \sum_{\{H, H'\} \in \mathcal{H}(S_k)} \text{diam}(H) |f(H) - f(H')| \\ & \leq 2 \left(\frac{1}{|S_k|^2} \sum_{\{H, H'\} \in \mathcal{H}(S_k)} \text{diam}(H)^2 \mathfrak{c}(H, H')^{-1} \right)^{1/2} \\ & \quad \times \left(\sum_{\{H, H'\} \in \mathcal{H}(S_k)} \mathfrak{c}(H, H') |f(H) - f(H')|^2 \right)^{1/2} \\ & \leq 2 \left(\frac{1}{|S_k|^2} \sum_{H \in \mathcal{H}(S_k)} \text{diam}(H)^2 \pi^*(H) \right)^{1/2} \text{Energy}(f; \mathcal{H}(S_k))^{1/2} \\ & \leq 2C_0^{1/2} \text{Energy}(f; \mathcal{H}(S_k))^{1/2}. \end{aligned} \tag{2.39}$$

This concludes the proof with $C = 2C_0^{1/2}$. Obviously, the case of vertical line segments follows from the same argument. □

The preceding lemmas allow us to approximate ϕ_∞ by ϕ_m for a large value of $m \in \mathbb{N}$. To complement this, we have the following bound for $\phi_m - \phi_0$ for a fixed value of m .

LEMMA 2.18. *For each fixed $m \in \mathbb{N}$, a.s.*

$$\limsup_{k \rightarrow \infty} \frac{1}{|S_k|} \max_{H \in \mathcal{H}(S_k)} |\phi_m(H) - \phi_0(H)| = 0. \tag{2.40}$$

Proof. By Lemma 2.9, a.s.

$$\limsup_{k \rightarrow \infty} \frac{1}{|S_k|} \max_{H \in \mathcal{H}(S_k)} \text{diam}(H) = 0. \tag{2.41}$$

By (2.1), the interior of each $\widehat{S} \in \mathcal{S}_m$ contains at most m cells of \mathcal{H} . Consequently, the center point of \widehat{S} can be joined to $\partial\widehat{S}$ by a Euclidean path which intersects at most $m+1$ cells of $\mathcal{H}(\widehat{S})$. Since \widehat{S} is a square, it follows that the side length of \widehat{S} is at most $m+1$ times the maximal diameter of the cells which intersect \widehat{S} , so we can multiply the estimate (2.41) by $m+1$ to get that a.s.

$$\limsup_{k \rightarrow \infty} \frac{1}{|S_k|} \max_{\substack{\widehat{S} \in \mathcal{S}_m \\ \widehat{S} \subset S_k}} \text{diam}(\widehat{S}) = 0. \tag{2.42}$$

By definition, ϕ_0 and ϕ_m agree on $\mathcal{H}(\partial\widehat{S})$ for each $\widehat{S} \in \mathcal{S}_m$. Moreover, ϕ_m is discrete harmonic on $\mathcal{H}(\widehat{S}) \setminus \mathcal{H}(\partial\widehat{S})$ for each such square \widehat{S} . By combining (2.41) and (2.42) with the maximum principle for $\phi_m|_{\mathcal{H}(\widehat{S})}$ and the fact that $|\phi_0(H) - z| \leq \text{diam}(H)$ for each $z \in H \in \mathcal{H}$, we obtain (2.40). \square

LEMMA 2.19. *Almost surely,*

$$\limsup_{k \rightarrow \infty} \frac{1}{|S_k|} \max_{H \in \mathcal{H}(S_k)} |\phi_\infty(H) - \phi_0(H)| = 0. \tag{2.43}$$

Proof. For $k \in \mathbb{N}$, define the horizontal line segments $P_k(r) \subset S_k$ for $r \in [0, 1]$ as in Lemma 2.17. By Lemmas 2.16 and 2.17, a.s.

$$\lim_{m \rightarrow \infty} \limsup_{k \rightarrow \infty} \frac{1}{|S_k|} \int_0^1 \sum_{\{H, H'\} \in \mathcal{E}\mathcal{H}(P_k(r))} |\phi_{m,\infty}(H) - \phi_{m,\infty}(H')| dr = 0. \tag{2.44}$$

Given $\varepsilon > 0$, choose a (random) $m_\varepsilon \in \mathbb{N}$ sufficiently large that for $m \geq m_\varepsilon$, the limsup in (2.44) is at most $\frac{1}{100}\varepsilon^3$. By Markov’s inequality, a.s. for $m \geq m_\varepsilon$ it holds for large enough $k \in \mathbb{N}$, the Lebesgue measure of the set of $r \in [0, 1]$ for which the corresponding sum in (2.44) is at most $\frac{1}{100}|S_k|\varepsilon^2$ is at least $1 - \varepsilon$. Therefore, for large enough $k \in \mathbb{N}$ we can find a finite collection \mathcal{P}_k of at most $2/\varepsilon$ horizontal line segments joining the left and right boundaries of S_k such that each point of S_k lies within Euclidean distance $\varepsilon|S_k|$ of one of these line segments and

$$\frac{1}{|S_k|} \sum_{P \in \mathcal{P}_k} \sum_{\{H, H'\} \in \mathcal{E}\mathcal{H}(P)} |\phi_{m_\varepsilon,\infty}(H) - \phi_{m_\varepsilon,\infty}(H')| \leq \frac{\varepsilon}{2}. \tag{2.45}$$

Similarly, after possibly increasing m_ε it is a.s. the case that for large enough $k \in \mathbb{N}$, we can find an $\varepsilon|S_k|$ -dense collection \mathcal{P}'_k of at most $2/\varepsilon$ vertical line segments joining the top and bottom boundaries of S_k such that (2.45) holds with \mathcal{P}'_k in place of \mathcal{P}_k .

By Lemma 2.18, it is a.s. the case that for large enough $k \in \mathbb{N}$, we have

$$|S_k|^{-1} \max_{H \in \mathcal{H}(S_k)} |\phi_{m_\varepsilon}(S) - \phi_0(S)| \leq \frac{1}{4}\varepsilon.$$

Combining this with (2.45) and its analog for \mathcal{P}'_k , we see that a.s. for large enough $k \in \mathbb{N}$,

$$\frac{1}{|S_k|} \sum_{P \in \mathcal{P}_k} \sum_{\{H, H'\} \in \mathcal{EH}(P)} |\phi_{0,\infty}(H) - \phi_{0,\infty}(H')| \leq \varepsilon \tag{2.46}$$

and the same holds with \mathcal{P}'_k in place of \mathcal{P}_k . Each $P \in \mathcal{P}_k \cup \mathcal{P}'_k$ is a horizontal or vertical line segment, so assumption (3) implies that $\mathcal{H}(P)$ (viewed as a graph) is connected for each $P \in \mathcal{P}_k \cup \mathcal{P}'_k$, i.e., any two cells in $\mathcal{H}(P)$ can be joined by a path of edges in $\mathcal{EH}(P)$. Hence

$$\sum_{\{H, H'\} \in \mathcal{EH}(P)} |\phi_{0,\infty}(H) - \phi_{0,\infty}(H')|$$

is an upper bound for the maximum of $|\phi_{0,\infty}(H_1) - \phi_{0,\infty}(H_2)|$ over all $H_1, H_2 \in \mathcal{H}(P)$. By this and (2.46), we get that a.s. for each large enough $k \in \mathbb{N}$,

$$\frac{1}{|S_k|} \max \left\{ |\phi_{0,\infty}(H_1) - \phi_{0,\infty}(H_2)| : H_1, H_2 \in \mathcal{H} \left(\bigcup_{P \in \mathcal{P}_k \cup \mathcal{P}'_k} P \right) \right\} \leq 2\varepsilon. \tag{2.47}$$

We now want to extend from (2.47) to a bound which holds for all $H_1, H_2 \in \mathcal{H}(S_k)$. Recall that $\phi_{0,\infty} = \phi_\infty - \phi_0$. Since the collections \mathcal{P}_k and \mathcal{P}'_k are $\varepsilon|S_k|$ -dense, a.s. for large enough $k \in \mathbb{N}$ each $H \in \mathcal{H}(S_k)$ which is not contained in the Euclidean $\varepsilon|S_k|$ -neighborhood of ∂S_k intersects a rectangle of diameter at most $2\varepsilon|S_k|$ whose boundary is formed by segments of four elements of $\mathcal{P}_k \cup \mathcal{P}'_k$. Since ϕ_∞ is discrete harmonic, it attains its maximum and minimum values on this rectangle on the boundary of the rectangle. Furthermore, by Lemma 2.9 and the definition of ϕ_0 , the maximum and minimum values of ϕ_0 on each such rectangle differ from the maximum and minimum values of ϕ_0 on the boundary of the rectangle by at most a quantity which is of order $o_k(|S_k|)$ a.s. as $k \rightarrow \infty$, uniformly over all of the rectangles. Combining these facts with (2.47) shows that a.s. for each large enough $k \in \mathbb{N}$,

$$|\phi_{0,\infty}(H_1) - \phi_{0,\infty}(H_2)| \leq 4\varepsilon|S_k|$$

for each $H_1, H_2 \in \mathcal{H}(S_k)$ which are not contained in the Euclidean $\varepsilon|S_k|$ -neighborhood of ∂S_k .

We will now deal with the cells $H \in \mathcal{H}(S_k)$ which are close to ∂S_k . Let S'_k be a dyadic ancestor of S_k of side length $4|S_k|$ which contains S_k in its interior. By the preceding paragraph and Lemma 2.8, it is a.s. the case for each large enough $k \in \mathbb{N}$,

$$|\phi_{0,\infty}(H_1) - \phi_{0,\infty}(H_2)| \leq 16\varepsilon|S_k|$$

for each $H_1, H_2 \in \mathcal{H}(S'_k)$ which are not contained in the Euclidean $\varepsilon|S'_k|$ -neighborhood of $\partial S'_k$. Since $S_k \subset S'_k$ and S_k lies at Euclidean distance at least $\varepsilon|S'_k|$ from $\partial S'_k$, we get that a.s. for each large enough $k \in \mathbb{N}$, $|\phi_{0,\infty}(H_1) - \phi_{0,\infty}(H_2)| \leq 16\varepsilon|S_k|$ for each $H \in \mathcal{H}(S_k)$.

We have

$$|\phi_{0,\infty}(H_0)| = |\phi_0(H_0)| = \frac{1}{\text{Area}(H_0)} \left| \int_{H_0} z \, dz \right| \leq \text{diam}(H_0).$$

From this and the preceding paragraph (applied with one of H_1 or H_2 equal to H_0), we get that a.s. for each large enough $k \in \mathbb{N}$, $|\phi_\infty(H) - \phi_0(H)| \leq 17\varepsilon|S_k|$ for each $H_1, H_2 \in \mathcal{H}(S_k)$ which are not contained in the Euclidean $\varepsilon|S_k|$ -neighborhood of ∂S_k . Since $\varepsilon > 0$ is arbitrary, this implies (2.43). \square

Proof of (1.17) of Theorem 1.16. Lemmas 2.8 and 2.19 together imply that for each $\varepsilon > 0$, it a.s. holds for large enough $k \in \mathbb{N}$ that $|\phi_\infty(H) - \phi_0(H)| \leq \varepsilon|S_k|$ for each cell H which intersects one of the four dyadic parents of S_k . Since $0 \in S_k$, the union of these four dyadic parents contains $B_{|S_k|}(0)$. Since $|S_k| \rightarrow \infty$ as $k \rightarrow \infty$, this implies (1.6). \square

Although it is not needed in the proof of (1.17), we record here the following lemma which will be used in §3.2.

LEMMA 2.20. *The function ϕ_∞ has finite specific Dirichlet energy, i.e.,*

$$\mathbb{E} \left[\sum_{\substack{H \in \mathcal{H} \\ H \sim H_0}} \frac{\mathfrak{c}(H_0, H) |\phi_\infty(H)|^2}{\text{Area}(H_0)} \right] < \infty.$$

Proof. By Lemma 2.15 we have $\mathbb{E}[\text{SE}_{0,\infty}] < \infty$ so we only need to show that

$$\mathbb{E}[\text{SE}_0] < \infty \quad \text{where} \quad \text{SE}_0 := \sum_{\substack{H \in \mathcal{H} \\ H \sim H_0}} \frac{\mathfrak{c}(H_0, H) |\phi_0(H) - \phi_0(H_0)|^2}{\text{Area}(H_0)}. \tag{2.48}$$

To prove (2.48), we use Lemma 2.7 to get that a.s.

$$\mathbb{E}[\text{SE}_0] = \lim_{k \rightarrow \infty} \frac{1}{|S_k|^2} \int_{S_k} \sum_{\substack{H \in \mathcal{H} \\ H \sim H_z}} \frac{\mathfrak{c}(H_z, H) |\phi_0(H) - \phi_0(H_z)|^2}{\text{Area}(H_z)} \, dz. \tag{2.49}$$

By breaking up the integral into a sum over the regions $H \cap S_k$ and applying Lemma 2.9, we find that the integral on the right-hand side of (2.49) is a.s. bounded above by twice the Dirichlet energy of ϕ_0 over the union of the four dyadic parents of S_k for large enough $k \in \mathbb{N}$. The bound (2.48) therefore follows from (2.49) combined with Lemmas 2.8 and 2.12. \square

2.5. Random walk on \mathcal{H} is recurrent

In this subsection we will establish the following, which is part of Theorem 1.16 and is a generalization of Theorem 1.7.

PROPOSITION 2.21. *In the setting of Theorem 1.16, the simple random walk on \mathcal{H} is a.s. recurrent.*

The proof of Proposition 2.21 uses only the results of §2.2. By a standard criterion, it suffices to produce a sequence of finite vertex sets $\{V_r\}_{r \in \mathbb{N}} \subset \mathcal{H}$ containing H_0 , whose union is all of \mathcal{H} , with the property that the effective resistance from zero to $\mathcal{H} \setminus V_r$ tends to ∞ as $r \rightarrow \infty$. By Dirichlet’s principle (see, e.g., [LP, Exercise 2.13]), to prove this it suffices to find functions $\mathcal{H} \rightarrow [0, 1]$ which are identically equal to one on $\mathcal{H} \setminus V_r$ and vanish at H_0 whose discrete Dirichlet energy tends to zero as $r \rightarrow \infty$.

We will construct the sets V_r and the corresponding functions using the dyadic system \mathcal{D} . For $k \in \mathbb{Z}$ and $r \in \mathbb{N}_0$, let $B_k^{(r)}$ be the Euclidean ball of side length $2^r |S_k|$ with the same center at S_k . For $r \geq 2$, let $\mathbf{g}_k^{(r)}: B_k^{(r)} \setminus B_k^{(1)} \rightarrow [0, 1]$ be the continuous function which equals zero on $\partial B_k^{(1)}$, equals one on $\partial B_k^{(r)}$, and is discrete harmonic on the rest of $B_k^{(r)} \setminus B_k^{(1)}$. That is,

$$\mathbf{g}_k^{(r)}(z) = \frac{\log |z - v_k| - \log(2|S_k|)}{\log(2^{r-1})} \quad \text{for all } z \in B_k^{(r)} \setminus B_k^{(1)}, \tag{2.50}$$

where v_k is the center of S_k . Also let $\mathbf{f}_k^{(r)}: \mathcal{H} \rightarrow [0, 1]$ be the function which vanishes on $\mathcal{H}(B_k^{(1)})$, equals one on $\mathcal{H} \setminus \mathcal{H}(B_k^{(1)})$, and for other cells

$$H \in \mathcal{H}(B_k^{(r)} \setminus B_k^{(1)}) \setminus \mathcal{H}(\partial(B_k^{(r)} \setminus B_k^{(1)}))$$

satisfies $\mathbf{f}_k^{(r)}(H) = \mathbf{g}_k^{(r)}(z)$ for some (arbitrary) choice of $z \in H \cap (B_k^{(r)} \setminus B_k^{(1)})$. By the discussion just above (applied with $V_r = \mathcal{H}(B_{k_r}^{(r)})$ for a large random k_r), to prove Proposition 2.21, it suffices to establish the following.

LEMMA 2.22. *There is a deterministic constant $C > 0$ such that for each $r \in \mathbb{N}$, it a.s. holds for large enough $k \in \mathbb{N}$ that*

$$\text{Energy}(\mathbf{f}_k^{(r)}; \mathcal{H}) \leq \frac{C}{r}. \tag{2.51}$$

Proof. The proof is an extension of that of Lemma 2.12 with a few extra complications. Each of the balls $B_k^{(r)}$ above is contained in the union of four dyadic ancestors of S_k of side length $2^{r+1}|S_k|$. It therefore follows from Lemma 2.8 combined with Lemmas 2.9 and 2.10, respectively, that there exists a deterministic constant $C_0 > 0$ (independent of r) such that a.s. for large enough $k \in \mathbb{N}$,

$$\text{diam}(H) \leq \frac{1}{100} |S_k| \quad \text{for all } H \in \mathcal{H}(B_k^{(r)}) \tag{2.52}$$

and

$$\frac{1}{2^{2j}|S_k|^{2j}} \sum_{H \in \mathcal{H}(B_k^{(j)})} \text{diam}(H)^2 \pi(H) \leq C_0 \quad \text{for all } j \in [1, r]_{\mathbb{Z}}. \tag{2.53}$$

Henceforth assume that k is chosen sufficiently large that (2.52) and (2.53) hold.

We first observe that, by direct calculation and since $\log(2^{r-1}) \asymp r$, the function of (2.50) satisfies

$$|\nabla \mathfrak{g}_k^{(r)}(z)| \preceq \frac{1}{r|z|} \tag{2.54}$$

with a universal implicit constant. Now consider $j \in [1, r]_{\mathbb{Z}}$ and an edge

$$\{H, H'\} \in \mathcal{EH}(B_k^{(j)} \setminus B_k^{(j-1)}).$$

By the mean value theorem, for appropriate points $z \in H$ and $z' \in H'$,

$$|\mathfrak{f}_k^{(r)}(H) - \mathfrak{f}_k^{(r)}(H')| \leq \int_0^1 |\nabla \mathfrak{g}_k^{(r)}(P_{z,z'}(t)) \cdot P_{z,z'}(t)| dt \tag{2.55}$$

where $P_{z,z'}(t)$ denotes a Euclidean path from z to z' in $B_k^{(r)} \setminus B_k^{(1)}$ of minimal Euclidean length.

Since $P_{z,z'}(t)$ has minimal Euclidean length, it follows that its total length is

$$\int_0^1 |P_{z,z'}(t)| dt \leq 2(\text{diam}(H) + \text{diam}(H')).$$

Furthermore, by (2.52) we have $|P_{z,z'}(t)| \asymp 2^j |S_k|$ for $t \in [0, 1]$, so, by (2.54),

$$|\nabla \mathfrak{g}_k^{(r)}(P_{z,z'}(t))| \preceq r^{-1} 2^{-j} |S_k|^{-1}.$$

Plugging these estimates into (2.55) gives

$$|\mathfrak{f}_k^{(r)}(H) - \mathfrak{f}_k^{(r)}(H')| \preceq \frac{\text{diam}(H) + \text{diam}(H')}{r^2 2^{2j} |S_k|} \quad \text{for all } \{H, H'\} \in \mathcal{EH}(B_k^{(j)} \setminus B_k^{(j-1)}). \tag{2.56}$$

Squaring both sides of (2.56), multiplying by $\mathfrak{c}(H, H')$, and summing over all such edges $\{H, H'\}$, then applying (2.53) gives

$$\text{Energy}(\mathfrak{f}_k^{(r)}; \mathcal{H}(B_k^{(j)} \setminus B_k^{(j-1)})) \leq \frac{1}{r^2 2^{2j} |S_k|^2} \sum_{H \in \mathcal{H}(B_k^{(j)})} \text{diam}(H)^2 \pi(H) \preceq \frac{1}{r^2}$$

with a deterministic, r -independent implicit constant. Summing this over all $j \in [1, r]_{\mathbb{Z}}$ now yields (2.51) since $\mathfrak{f}_k^{(r)}$ is constant on each of the two connected components of $\mathcal{H} \setminus \mathcal{H}(B_k^{(r)} \setminus B_k^{(1)})$. □

3. Convergence of random walk to Brownian motion

Continue to assume that \mathcal{H} is a cell configuration satisfying the hypotheses of Theorem 1.16 and \mathcal{D} is an independent uniform dyadic system. In this section we will show that simple random walk on \mathcal{H} converges to Brownian motion, thereby completing the proof of Theorem 1.16. To do this, we will show that the image of simple random walk under ϕ_∞ converges to Brownian motion (Proposition 3.5), then apply Theorem 1.6. The proof of this convergence will use the martingale central limit theorem (see Theorem 3.6 for a precise statement of the version we use). In order to prove the necessary quadratic variation convergence conditions, we will need analogs of some of the ergodic theory results in §2.2 when we average along the path of the walk instead of over a large square. Such results are established in §3.1. We prove Theorem 1.16 in §3.2. §3.3 and §3.4, respectively, are devoted to the proof that the rate of convergence in Theorem 1.16 is uniform on compact subsets of \mathbb{C} (Theorem 3.10) and the proofs of the other variants of our main results.

3.1. Ergodic averages for simple random walk

Let Y be a 2-sided discrete time simple random walk on the graph \mathcal{H} , started at $Y_0 = H_0$. In this subsection we will prove an exact resampling property for Y analogous to Lemma 2.3 (Lemma 3.1) and a corresponding convergence statement for ergodic averages analogous to Lemma 2.7 (Lemma 3.2).

For this purpose we first need to choose a particular continuous-time parametrization X of Y with the property that X spends $\text{Area}(H)/\pi(H)$ units of time at the cell H whenever this cell is hit (recall that $\pi(H)$ is as in (1.4)). In order to avoid making zero a special time we first sample a uniform random variable θ_X from $[0, \text{Area}(H_0)/\pi(H_0)]$ and declare that $n_t = 0$ and $X_t = H_0$ for $t \in [-\theta_X, \text{Area}(H_0)/\pi(H_0) - \theta_X]$. For $t \geq \text{Area}(H_0)/\pi(H_0) - \theta_X$, we set

$$X_t = Y_{n_t} \quad \text{where } n_t := \min \left\{ n \in \mathbb{N} : \sum_{j=0}^n \frac{\text{Area}(Y_j)}{\pi(Y_j)} \geq t + \theta_X \right\} \quad (3.1)$$

and for $t \leq -\theta_X$, we define n_t and X_t analogously but with Y run backward, rather than forward, and the sum going from $-n$ to -1 instead of 0 to n .

For $j \in \mathbb{N}$, let τ_j be the j th smallest time after zero at which X jumps to a new vertex and for $j \in \mathbb{N}_0$, let τ_{-j} be the $(j+1)$ st largest time before zero at which X jumps to a new vertex (so $\tau_0 = -\theta$). Note that $X_t = Y_{n_t} = X_{\tau_{n_t}}$.

It will be convenient to view X_t as a process taking values in \mathbb{R}^2 rather than in \mathcal{H} . Conditional on (\mathcal{H}, X) , sample for each $j \in \mathbb{Z} \setminus \{0\}$ a point \widehat{X}_{τ_j} uniformly from Lebesgue measure on the cell X_{τ_j} (independently for different cells and different values of j). Also

set $\widehat{X}_0=0$. Define

$$\widehat{X}_t := \widehat{X}_{\tau_j} \quad \text{for all } t \in (\tau_{j-1}, \tau_j], \quad \text{so that } H_{\widehat{X}_t} = X_t \tag{3.2}$$

where in the second formula we recall that H_z for $z \in \mathbb{C}$ denotes the cell containing z .

We will now state and prove a translation invariance property for the walk \widehat{X} analogous to Lemma 2.3. To this end, we let $\widetilde{\mathcal{D}}$ be a 1-dimensional uniform dyadic system (Remark 2.4), independent from $(\mathcal{H}, \mathcal{D}, \widehat{X})$, and we let $\{I_k\}_{k \in \mathbb{Z}}$ be the sequence of intervals in $\widetilde{\mathcal{D}}$ which contain the origin. For $a > 0$, we let \hat{I}_a be the largest interval in $\widetilde{\mathcal{D}}$ containing zero with the property that

$$\sum_{j \in \mathbb{Z}} |[\tau_{j-1}, \tau_j] \cap \hat{I}_a| \leq a. \tag{3.3}$$

The interval \hat{I}_a is the 1-dimensional analog of the square \widehat{S}_m^0 from (2.1).

LEMMA 3.1. *The law of $(\mathcal{H}, \mathcal{D}, \widehat{X}, \widetilde{\mathcal{D}})$ satisfies the following translation invariance property. Let $a \geq 0$ and conditional on everything else, let t be sampled uniformly from \hat{I}_a . Then the 4-tuple*

$$(\mathcal{H} - \widehat{X}_t, \mathcal{D} - \widehat{X}_t, \widehat{X}_{\cdot + t} - \widehat{X}_t, \widetilde{\mathcal{D}} - t) \tag{3.4}$$

agrees in law with $(\mathcal{H}, \mathcal{D}, \widehat{X}, \widetilde{\mathcal{D}})$ modulo a global re-scaling, i.e., there is a random $C > 0$ such that

$$(C(\mathcal{H} - \widehat{X}_t), C(\mathcal{D} - \widehat{X}_t), C(\widehat{X}_{\cdot / C^2 + t} - \widehat{X}_t), C^2(\widetilde{\mathcal{D}} - t)) \stackrel{d}{=} (\mathcal{H}, \mathcal{D}, \widehat{X}, \widetilde{\mathcal{D}}). \tag{3.5}$$

Proof. We will prove the lemma by considering random walk on $\mathcal{H}(\widehat{S}_m)$ with reflected boundary conditions and sending $m \rightarrow \infty$. Let X^m , \widehat{X}^m , and τ_j^m be defined in the same manner as \widehat{X} and τ_j above except that we start with a simple random walk Y^m on $\mathcal{H}(\widehat{S}_m)$ (with reflected boundary conditions) started from H_0 instead of a simple random walk on \mathcal{H} started from H_0 ; we replace $\text{Area}(Y_j)/\pi(Y_j)$ by $\text{Area}(Y_j^m \cap \widehat{S}_m)/\pi^{\mathcal{H}(\widehat{S}_m)}(Y_j^m)$ in (3.1), where $\pi^{\mathcal{H}(\widehat{S}_m)}(H)$ is the sum of the conductances of the edges adjacent to H in $\mathcal{H}(\widehat{S}_m)$; and we sample the points $\widehat{X}_{\tau_j^m}^m$ from $X_{\tau_j^m}^m \cap \widehat{S}_m$ instead of from $X_{\tau_j^m}$. It is easily seen that, conditional on $(\mathcal{H}, \mathcal{D})$, the stationary distribution for X^m is the one which assigns probability proportional to $\text{Area}(H \cap \widehat{S}_m)$ to each cell $H \in \mathcal{H}(\widehat{S}_m)$. Equivalently, the stationary distribution of \widehat{X}^m is uniform on \widehat{S}_m .

As $t \rightarrow \infty$ or $t \rightarrow -\infty$, the conditional distribution of \widehat{X}_t^m given $(\mathcal{H}, \mathcal{D})$ converges in the total variation sense to the stationary distribution. Consequently, if we let t_T be sampled uniformly from Lebesgue measure on $[-T, T]$ (independently from everything

else), then as $T \rightarrow \infty$ the conditional distribution of $\widehat{X}_{t_T}^m$ given $(\mathcal{H}, \mathcal{D})$ converges in the total variation sense to the uniform measure on \widehat{S}_m . Hence, as $T \rightarrow \infty$ the joint law of

$$(\mathcal{H}, \mathcal{D}, \widehat{X}_{t_T}^m, \widehat{X}_{+t_T}^m,)$$

converges in the total variation sense to the law of the 4-tuple consisting of \mathcal{H}, \mathcal{D} , a point z sampled uniformly from Lebesgue measure on \widehat{S}_m , and the walk $\widehat{X}^{m,z}$ which is defined in the same manner as \widehat{X}^m but with $\widehat{X}_0^m = z$.

By Lemma 2.3, there is a random $C > 0$ such that with z as above,

$$(C(\mathcal{H} - z), C(\mathcal{D} - z), C(\widehat{X}_{\cdot/C^2}^{m,z} - z)) \stackrel{d}{=} (\mathcal{H}, \mathcal{D}, \widehat{X}^m). \tag{3.6}$$

Note that to get the scaling of the time parametrization of $\widehat{X}^{m,z}$, we use that replacing \mathcal{H} by $C\mathcal{H}$ for $C > 0$ scales areas of cells by C^2 , which in turn scales the time parametrization of $\widehat{X}^{m,z}$ by $1/C^2$.

From (3.6) and the above total variation convergence, we infer that there are random variables $C_T > 0$ such that

$$(C_T(\mathcal{H} - \widehat{X}_{t_T}^m), C_T(\mathcal{D} - \widehat{X}_{t_T}^m), C_T(\widehat{X}_{\cdot/C_T^2+t_T}^m - \widehat{X}_{t_T}^m)) \rightarrow (\mathcal{H}, \mathcal{D}, \widehat{X}^m)$$

in the total variation sense. By the 1-dimensional analog of Lemma 2.1, for each $T > 0$ the dyadic system $C_T^2(\widetilde{\mathcal{D}} - t_T)$ is a uniform dyadic system independent from $(\mathcal{H}, \mathcal{D}, \widehat{X}^m, t_T)$. Therefore, we have the total variation convergence

$$(C_T(\mathcal{H} - \widehat{X}_{t_T}^m), C_T(\mathcal{D} - \widehat{X}_{t_T}^m), C_T(\widehat{X}_{\cdot/C_T^2+t_T}^m - \widehat{X}_{t_T}^m), C_T^2(\widetilde{\mathcal{D}} - t_T)) \rightarrow (\mathcal{H}, \mathcal{D}, \widehat{X}^m, \widetilde{\mathcal{D}}). \tag{3.7}$$

The next several paragraphs of the proofs are similar to the proof that condition (2) implies condition (5) in Appendix B. Let $\widehat{I}_a^{t_T}$ be the largest interval of $\widetilde{\mathcal{D}}$ containing t_T for which $\sum_{j \in \mathbb{Z}} |[\tau_{j-1}, \tau_j] \cap \widehat{I}_a^{t_T}| \leq a$ (i.e., the analog of \widehat{I}_a with t_T in place of zero). We claim that

$$\lim_{T \rightarrow \infty} \mathbb{P}[\widehat{I}_a^{t_T} \subset [-T, T]] = 1. \tag{3.8}$$

Indeed, from (3.7) the law of the random interval $C_T^2(\widehat{I}_a^{t_T} - t_T)$ converges in the total variation sense to the law of \widehat{I}_a , so $C_T^2|\widehat{I}_a^{t_T}|$ is a tight random variable. On the other hand, $C_T(\widehat{S}_m - \widehat{X}_{t_T}^m) \rightarrow \widehat{S}_m$ in the total variation sense, and $\widehat{X}_{t_T}^m \subset \widehat{S}_m$. From this, we get that C_T and C_T^{-1} are tight random variables. Therefore, $|\widehat{I}_a^{t_T}|$ is a tight random variable. Since t_T is sampled uniformly from $[-T, T]$, the probability that $|t_T| \geq T - r$ goes to zero as $T \rightarrow \infty$ for any fixed $r > 0$. From this and the tightness of $|\widehat{I}_a^{t_T}|$, we obtain (3.8).

On the event $\{\widehat{I}_a^{t_T} \subset [-T, T]\}$, the conditional law of t_T given $(\mathcal{H}, \mathcal{D}, \widetilde{\mathcal{D}}, \widehat{X}^m, \widehat{I}_a^{t_T})$ is uniform on $\widehat{I}_a^{t_T}$. From (3.8), we get that as $T \rightarrow \infty$, the total variation distance between

the conditional law of t_T given $(\mathcal{H}, \mathcal{D}, \tilde{\mathcal{D}}, \hat{X}^m, \hat{I}_a^{t_T})$ and the uniform measure on $\hat{I}_a^{t_T}$ tends to zero as $T \rightarrow \infty$. By this and (3.7), if we condition on $(\mathcal{H}, \mathcal{D}, \tilde{\mathcal{D}}, \hat{X}^m, t_T)$ and sample t'_T uniformly from $\hat{I}_a^{t_T}$, then the total variation convergence (3.7) holds with t'_T in place of t_T and possibly different scaling factors $C'_T > 0$ in place of C_T .

On the other hand, the interval $C_T^2(\hat{I}_a^{t_T} - t_T)$ is determined by the left 4-tuple in (3.7) in the same manner that \hat{I}_a is determined by the right 4-tuple in (3.7). Consequently, we have the total variation convergence

$$\begin{aligned} (C_T(\mathcal{H} - \hat{X}_{t_T}^m), C_T(\mathcal{D} - \hat{X}_{t_T}^m), C_T(\hat{X}_{\cdot/C_T^2+t_T}^m - \hat{X}_{t_T}^m), C_T^2(\tilde{\mathcal{D}} - t_T), C_T^2(t'_T - t_T)) \\ \rightarrow (\mathcal{H}, \mathcal{D}, \hat{X}^m, \tilde{\mathcal{D}}, t) \end{aligned} \tag{3.9}$$

where t is sampled uniformly from \hat{I}_a . In particular, we have the total variation convergence

$$\begin{aligned} (C_T(\mathcal{H} - \hat{X}_{t'_T}^m), C_T(\mathcal{D} - \hat{X}_{t'_T}^m), C_T(\hat{X}_{\cdot/C_T^2+t'_T}^m - \hat{X}_{t'_T}^m), C_T^2(\tilde{\mathcal{D}} - t'_T)) \\ \rightarrow (\mathcal{H} - \hat{X}_t^m, \mathcal{D} - \hat{X}_t^m, \hat{X}_{\cdot+t}^m - \hat{X}_t^m, \tilde{\mathcal{D}} - t). \end{aligned} \tag{3.10}$$

From this and the analog of (3.7) with t'_T in place of t_T , we obtain (3.5) with \hat{X}^m instead of \hat{X} . We now conclude by sending $m \rightarrow \infty$. \square

Next, we will prove an analog of Lemma 2.7 for averages along segments of the walk \hat{X} .

LEMMA 3.2. *Let $F = F(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}})$ be a real-valued measurable function of*

$$(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}})$$

which is invariant under global re-scaling in the sense that

$$F(C\mathcal{H}, C\mathcal{D}, C\hat{X}_{\cdot/C^2}, C^2\tilde{\mathcal{D}}) = F(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}}) \tag{3.11}$$

for each $C > 0$. For $t \in \mathbb{R}$, let F_t denote F applied to the shifted 4-tuple of (3.4). If either $\mathbb{E}[|F|] < \infty$ or $F \geq 0$ a.s., then a.s.

$$\lim_{a \rightarrow \infty} \frac{1}{|\hat{I}_a|} \int_{\hat{I}_a} F_t dt = \mathbb{E}[F]. \tag{3.12}$$

Proof. We start out as in the proof of Lemma 2.7. Throughout the proof, we assume that $\mathbb{E}[|F|] < \infty$ (the case $F \geq 0$ and $\mathbb{E}[F] = \infty$ is treated exactly as in Lemma 2.7). By analogy with Definition 2.5, for $a \geq 0$ let $\tilde{\mathcal{F}}_a$ be the σ -algebra generated by the measurable functions of $(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}})$ whose values are unchanged if we replace $(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}})$ by the

4-tuple in (3.4) for any $t \in \hat{I}_a$ and then apply a global re-scaling to the function as in (3.11). By Lemma 3.1,

$$\frac{1}{|\hat{I}_a|} \int_{\hat{I}_a} F_t dt = \mathbb{E}[F | \tilde{\mathcal{F}}_a].$$

By the backward martingale convergence theorem, the limit in (3.12) exists a.s. and in L^1 . Call the limiting random variable $A = A(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}})$. We need to show that A is equal to a deterministic constant a.s.

The random variable A is $\bigcap_{a>0} \tilde{\mathcal{F}}_a$ -measurable, so by the definition of $\tilde{\mathcal{F}}_a$ and since $\bigcup_{a>0} \hat{I}_a = \mathbb{R}$, a.s.

$$A(C(\mathcal{H} - \tilde{X}_t), C(\mathcal{D} - \tilde{X}_t), C(\hat{X}_{\cdot/C^2+t} - \hat{X}_t), C^2(\tilde{\mathcal{D}} - t)) = A(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}}) \tag{3.13}$$

for all $C > 0$ and all $t \in \mathbb{R}$.

Recall from Proposition 2.21 that a.s. the random walk on \mathcal{H} is recurrent. If we condition on $(\mathcal{H}, \mathcal{D}, \tilde{\mathcal{D}})$, then the excursions of the walk X away from X_0 are i.i.d. random variables. Under this conditional law, A belongs to the tail σ -algebra for these random variables and the independent uniform random points of the cells of \mathcal{H} which we used to define \hat{X}_t . By the Kolmogorov zero-one law, A is a.s. determined by $(\mathcal{H}, \mathcal{D}, \tilde{\mathcal{D}})$. By a slight abuse of notation, we now write $A = A(\mathcal{H}, \mathcal{D}, \tilde{\mathcal{D}})$.

We will now argue that A is a.s. determined by $\tilde{\mathcal{D}}$. To this end, let

$$\tilde{A}(\mathcal{H}, \mathcal{D}) := \int A(\mathcal{H}, \mathcal{D}, \tilde{\mathcal{D}}) d\tilde{m}(\tilde{\mathcal{D}})$$

where \tilde{m} is the law of $\tilde{\mathcal{D}}$. If $C > 0$ and $t \in \mathbb{R}$ are random and independent from $\tilde{\mathcal{D}}$, but allowed to depend on $(\mathcal{H}, \mathcal{D})$, then since $\tilde{\mathcal{D}}$ is independent from $(\mathcal{H}, \mathcal{D})$ and by the 1-dimensional analog of Lemma 2.1, $C^2(\tilde{\mathcal{D}} - t)$ is a uniform dyadic system independent from $(\mathcal{H}, \mathcal{D})$. By this and (3.13),

$$\begin{aligned} \tilde{A}(C(\mathcal{H} - \hat{X}_t), C(\mathcal{D} - \hat{X}_t)) &= \int A(C(\mathcal{H} - \hat{X}_t), C(\mathcal{D} - \hat{X}_t), \tilde{\mathcal{D}}) d\tilde{m}(\tilde{\mathcal{D}}) \\ &= \int A(C(\mathcal{H} - \hat{X}_t), C(\mathcal{D} - \hat{X}_t), C^2(\tilde{\mathcal{D}} - t)) d\tilde{m}(\tilde{\mathcal{D}}) \\ &= \tilde{A}(\mathcal{H}, \mathcal{D}). \end{aligned} \tag{3.14}$$

Now let $R > 0$ be random and independent from $\tilde{\mathcal{D}}$, but allowed to depend on $(\mathcal{H}, \mathcal{D})$. Since the random walk on \mathcal{H} is recurrent and \mathcal{H} is connected, it is a.s. the case that for large enough $t > 0$ (depending on $(\mathcal{H}, \mathcal{D}, R)$), the Lebesgue measure on $B_R(0)$ is absolutely continuous with respect to the conditional law of \hat{X}_t given $(\mathcal{H}, \mathcal{D})$. From this and (3.14),

we get that if $C > 0$ is random and independent from $\tilde{\mathcal{D}}$, but allowed to depend on $(\mathcal{H}, \mathcal{D})$, then for Lebesgue-a.e. $z \in \mathbb{C}$,

$$\tilde{A}(C(\mathcal{H} - z), C(\mathcal{D} - z)) = \tilde{A}(\mathcal{H}, \mathcal{D}). \tag{3.15}$$

If w is a point in the origin-containing cell H_0 , then the definition of \hat{X} implies that we can couple the process \hat{X}_t with the analogous process \hat{X}'_t defined with $(\mathcal{H} - w, \mathcal{D} - w)$ in place of $(\mathcal{H}, \mathcal{D})$ in such a way that the following is true. The associated discrete time walks Y and Y' satisfy $Y'_j = Y_j - w$ for all $j \in \mathbb{Z}$. Moreover, we have $\hat{X}_{t+s} = \hat{X}'_{t+s'} - w$ for all $t \geq 0$, where s and s' are the first times that \hat{X} and \hat{X}' , respectively jump away from the origin-containing cell; and an analogous relation holds for negative times. Since A is the limit of

$$\frac{1}{|\hat{I}_a|} \int_{\hat{I}_a} F_t dt$$

as $t \rightarrow \infty$, we infer that the value of A , and hence also the value of \tilde{A} , is unaffected if we replace $(\mathcal{H}, \mathcal{D})$ by $(\mathcal{H} - w, \mathcal{D} - w)$. From this and the fact that the boundaries of the cells of \mathcal{H} have zero Lebesgue measure, we see that the fact that (3.15) holds a.s. for Lebesgue-a.e. z implies that it in fact holds a.s. for every choice of z .

By (3.15) and Lemma 2.6, \tilde{A} is a.s. equal to a deterministic constant, so A is a.s. determined by $\tilde{\mathcal{D}}$. By (3.13), the value of A is unaffected if we replace $\tilde{\mathcal{D}}$ by $C(\tilde{\mathcal{D}} - t)$ for any $C > 0$ and $t \in \mathbb{R}$. The same argument as at the end of Lemma 2.6 now shows that A is equal to a deterministic constant a.s. □

To end this subsection, we record analogs of some of the lemmas from §2.2 in the setting of Lemma 3.2. We start with a lemma which allows us to prove statements for intervals of $\tilde{\mathcal{D}}$ which do not contain the origin, analogous to Lemma 2.8.

LEMMA 3.3. *Let $E(S) = E(S, \mathcal{H})$ be an event depending on an interval $I \subset \mathbb{R}$ and the 4-tuple $(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}})$ of Lemma 3.2. Suppose that a.s. $E(I_k)$ occurs for each large enough $k \in \mathbb{N}$. Then, for each $l \in \mathbb{N}$, it is a.s. the case that for large enough $k \in \mathbb{N}$, $E(I)$ occurs for each dyadic ancestor I of I_k of length at most $2^l |I_k|$ and each dyadic descendant I of I_k of length at least $2^{-l} |I_k|$.*

Proof. This follows from exactly the same argument use to prove Lemma 2.8. □

Similarly to Lemma 2.9, we also find that the maximum size of the intervals $[\tau_{j-1}, \tau_j]$ (on which X is constant) which intersect I_k is a.s. $o_k(|I_k|)$.

LEMMA 3.4. *Almost surely, for each $\varepsilon \in (0, 1)$ it holds for large enough $k \in \mathbb{N}$ that*

$$\tau_j - \tau_{j-1} \leq \varepsilon |I_k| \quad \text{for all } j \in \mathbb{N} \quad \text{with } [\tau_{j-1}, \tau_j] \cap I_k \neq \emptyset.$$

Proof. We set

$$\tilde{\mathcal{K}}_\varepsilon = \min\{k \in \mathbb{N}_0 : \tau_1 - \tau_0 \leq \varepsilon |I_k|\}$$

then argue exactly as in the proof of Lemma 2.9, with Lemma 3.2 used in place of Lemma 2.7. □

3.2. Proof of Theorem 1.16

Continue to assume we are in the setting of §3.1, so that $\{Y_j\}_{j \in \mathbb{Z}}$ is a 2-sided simply random walk on \mathcal{H} started from $Y_0=0$ and $X_t=Y_{n_t}$ is obtained by re-parameterizing Y so that it spends $\text{Area}(H)/\pi(H)$ units of time at each cell H that it hits. In this subsection we will prove the following proposition, then use it in conjunction with (1.17) (which was established in §2.4) to conclude the proof of Theorem 1.16.

PROPOSITION 3.5. *There is a deterministic covariance matrix Σ such that a.s. as $\varepsilon \rightarrow 0$, the conditional law given $(\mathcal{H}, \mathcal{D})$ of $t \mapsto \varepsilon \phi_\infty(X_{t/\varepsilon^2})$ converges weakly to the law of Brownian motion with covariance matrix Σ started from zero, with respect to the local uniform topology.*

Since ϕ_∞ is discrete harmonic and by the optional stopping theorem, the process $t \mapsto \phi_\infty(X_t)$ is a local martingale. We can therefore prove Proposition 3.5 using the martingale functional central limit theorem. In particular, we will use the following theorem, which is proven in [R] (see [H, Theorem 5.1 (a)] for a statement in English).

THEOREM 3.6. ([R]) *Let $\{(M_n, \mathcal{F}_n)\}_{n \in \mathbb{N}}$ be continuous-time real-valued local martingales, not necessarily continuous. For $t \geq 0$, let $\Delta M_n(t) := M_n(t) - M_n(t^-)$ be the jump of M_n at time t and for $\delta > 0$, let*

$$\sigma^\delta[M_n](t) := \sum_{s \leq t} \Delta M_n(s)^2 \mathbf{1}_{(|\Delta M_n(s)| \geq \delta)}.$$

Also let $\tilde{\sigma}^\delta[M_n](t)$ be the compensator in the Doob–Meyer decomposition of $\sigma^\delta[M_n]$, so that $\sigma^\delta[M_n] - \tilde{\sigma}^\delta[M_n]$ is a local martingale. Suppose that we have the following convergence in probability for each $t \geq 0$:

$$\lim_{n \rightarrow \infty} \langle M_n \rangle_t = t \quad \text{and} \quad \lim_{n \rightarrow \infty} \tilde{\sigma}^\delta[M_n](t) = 0 \quad \text{for all } \delta > 0,$$

where $\langle M_n \rangle$ denotes the quadratic variation. Then M_n converges in law to standard linear Brownian motion in the local uniform topology.

We will apply Theorem 3.6 to the 1-dimensional projections $\mathbf{v} \cdot (\varepsilon \phi_\infty(X_{\cdot/\varepsilon^2}))$ for unit vectors $\mathbf{v} \in \mathbb{R}^2$, then use the Cramér–Wold theorem. We first check the condition concerning the convergence of quadratic variations. For a unit vector $\mathbf{v} \in \mathbb{R}^2$, let

$$V_t^\mathbf{v} := \langle \mathbf{v} \cdot \phi_\infty(X), \mathbf{v} \cdot \phi_\infty(X) \rangle_t = \sum_{j=1}^{n_t} (\mathbf{v} \cdot (\phi_\infty(Y_j) - \phi_\infty(Y_{j-1})))^2. \tag{3.16}$$

Also set

$$c_\mathbf{v} := \mathbb{E} \left[\frac{\pi(H_0)}{\text{Area}(H_0)} (\mathbf{v} \cdot \phi_\infty(Y_1))^2 \right]. \tag{3.17}$$

LEMMA 3.7. *For each unit vector \mathbf{v} , one has $0 < c_\mathbf{v} < \infty$.*

Proof. One obtains $c_\mathbf{v} < \infty$ by first taking a conditional expectation given $(\mathcal{H}, \mathcal{D})$ in (3.17), then applying Lemma 2.20. If $c_\mathbf{v}$ were equal to zero for some \mathbf{v} , then one would have $\mathbf{v} \cdot \phi_\infty(Y_1) = 0$ a.s. This together with Lemma 3.1 would then show that a.s. $\mathbf{v} \cdot (\phi_\infty(Y_j) - \phi_\infty(Y_{j-1})) = 0$ for every j . Since \mathcal{H} is connected and Y is recurrent, this means that a.s. the range of ϕ_∞ is contained in a straight line orthogonal to \mathbf{v} , which is impossible by Proposition 2.11. \square

LEMMA 3.8. *For each fixed unit vector $\mathbf{v} \in \mathbb{R}^2$, a.s.*

$$\lim_{T \rightarrow \infty} \frac{1}{T} V_T^\mathbf{v} = c_\mathbf{v}. \tag{3.18}$$

Proof. For $t \in \mathbb{R}$, let

$$F_t := \frac{\pi(X_t)}{\text{Area}(X_t)} (\mathbf{v} \cdot (\phi_\infty(Y_{n_t+1}) - \phi_\infty(Y_{n_t})))^2. \tag{3.19}$$

Then F_0 is scale invariant in the sense of Lemma 3.2, so we can apply that lemma to find that a.s.

$$\lim_{k \rightarrow \infty} \frac{1}{|I_k|} \int_{I_k} F_t dt = c_\mathbf{v}, \tag{3.20}$$

where here we recall that the I_k ’s are the intervals containing zero in the 1-dimensional dyadic system $\tilde{\mathcal{D}}$. Furthermore, by Lemma 3.4 it is a.s. the case that the maximal length of the intervals $[\tau_{j-1}, \tau_j]$ on which X is constant which intersect I_k is $o(|I_k|)$.

To obtain (3.18), we need to transfer from the random interval I_k to an interval of the form $[0, T]$ for fixed $T > 0$. We will do this by approximating $[0, T]$ by a union of small intervals I for which (3.20) holds with I in place of I_k . To this end, fix $l \in \mathbb{N}$ (which we will eventually send to ∞). By (3.20) and two applications of Lemma 3.3, it is a.s. the case that, for large enough $k \in \mathbb{N}$,

$$\left| \frac{1}{|I|} \int_I F_t dt - c_\mathbf{v} \right| \leq 2^{-2l} \quad \text{and} \quad \tau_{n_t} - \tau_{n_t-1} \leq 2^{-2l} |I_k| \quad \text{for all } t \in I, \tag{3.21}$$

for each interval I which is a dyadic descendant of either of the two dyadic parents of I_k and which has side length $2^{-l}|I_k|$. Note that the union of these intervals contains

$$[-|I_k|, |I_k|].$$

Give $T > 0$, let $k_T \in \mathbb{N}$ be chosen so that $T \in [\frac{1}{4}|I_{k_T}|, \frac{1}{2}|I_{k_T}|]$. Let $\underline{\mathcal{I}}_T$ (resp. $\bar{\mathcal{I}}_T$) be the union of those intervals I of the form described just after (3.21) which are contained in $[0, T]$ (resp. which either intersect $[0, T]$ or share an endpoint with such an interval which intersects $[0, T]$). Note that the intervals I contained in $\bar{\mathcal{I}}_T$ intersect only along their endpoints, $\underline{\mathcal{I}}_T \subset [0, T] \subset \bar{\mathcal{I}}_T$, and $\bar{\mathcal{I}}_T \setminus \underline{\mathcal{I}}_T$ is a union of at most four such intervals I .

The second inequality in (3.21) implies that a.s. for each large enough $T > 0$, the intervals $[\tau_0, \tau_1]$ and $[\tau_{n_T}, \tau_{n_T+1}]$ which contain the endpoints of $[0, T]$ are contained in $\bar{\mathcal{I}}_T$. Recalling that $\tau_j - \tau_{j-1} = \text{Area}(Y_j)/\pi(Y_j)$ and the definition 3.16 of $V_T^{\mathbf{v}}$, we can break integrals up into sums over intervals of the form $[\tau_{j-1}, \tau_j]$ to find that

$$\int_{\underline{\mathcal{I}}_T} F_t dt \leq V_T^{\mathbf{v}} \leq \int_{\bar{\mathcal{I}}_T} F_t dt. \quad (3.22)$$

By summing the first estimate in (3.21) over those intervals I which are contained in each of $\underline{\mathcal{I}}_T$ and $\bar{\mathcal{I}}_T$ (there are at most 2^{l+1} such intervals), we find that a.s., for large enough $T > 0$,

$$\frac{1}{|\underline{\mathcal{I}}_T|} \int_{\underline{\mathcal{I}}_T} F_t dt \geq c_{\mathbf{v}} - O_l(2^{-l}) \quad \text{and} \quad \frac{1}{|\bar{\mathcal{I}}_T|} \int_{\bar{\mathcal{I}}_T} F_t dt \leq c_{\mathbf{v}} + O_l(2^{-l}), \quad (3.23)$$

with the rate of the $O_l(2^{-l})$ depending only on $c_{0,0}$. Since the lengths of $\underline{\mathcal{I}}_T$, $[0, T]$, and $\bar{\mathcal{I}}_T$ each differ by at most $2^{-l+2}|I_k| \leq 2^{-l+4}T$, we can plug (3.23) into (3.22) and send $l \rightarrow \infty$ to obtain (3.18). \square

We next need to check the condition on $\tilde{\sigma}^\delta[M_n]$ from Theorem 3.6. This will be accomplished using the following lemma.

LEMMA 3.9. *Fix a unit vector $\mathbf{v} \in \mathbb{R}^2$. For $j \in \mathbb{Z}$ and $r > 0$, let*

$$\Delta_j^r := (\mathbf{v} \cdot (\phi_\infty(Y_j) - \phi_\infty(Y_{j-1})))^2 \mathbf{1}_{(|\mathbf{v} \cdot (\phi_\infty(Y_j) - \phi_\infty(Y_{j-1}))| \geq r)} \quad (3.24)$$

and

$$\begin{aligned} \tilde{\Delta}_j^r &:= \mathbb{E}[\Delta_j^r \mid \mathcal{H}, \mathcal{D}, Y|_{(-\infty, j-1]}] \\ &= \frac{1}{\pi(Y_{j-1})} \sum_{\substack{H \in \mathcal{H} \\ H \sim Y_{j-1}}} \mathbf{c}(Y_{j-1}, H) (\mathbf{v} \cdot (\phi_\infty(H) - \phi_\infty(Y_{j-1})))^2 \mathbf{1}_{(|\mathbf{v} \cdot (\phi_\infty(H) - \phi_\infty(Y_{j-1}))| \geq r)}. \end{aligned} \quad (3.25)$$

Almost surely, for each $\delta > 0$,

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{j=1}^{n_T} \Delta_j^{\delta T^{1/2}} = \lim_{T \rightarrow \infty} \sum_{j=1}^{n_T} \tilde{\Delta}_j^{\delta T^{1/2}} = 0. \quad (3.26)$$

Proof. We will prove that the second limit in (3.26) is a.s. equal to zero, then comment on the adaptations needed to treat the first limit. For $a \geq 0$, let $\hat{I}_a \in \tilde{\mathcal{D}}$ be as in (3.3) and let

$$\tilde{F}^a := \frac{\pi(H_0)}{\text{Area}(H_0)} \tilde{\Delta}_1^{\delta|\hat{I}_a|^{1/2}} = \frac{1}{\text{Area}(H_0)} \sum_{\substack{H \in \mathcal{H} \\ H \sim H_0}} \mathfrak{c}(H_0, H) (\mathbf{v} \cdot \phi_\infty(H))^2 \mathbb{1}_{(|\mathbf{v} \cdot \phi_\infty(H)| \geq \delta|\hat{I}_a|^{1/2})}, \tag{3.27}$$

where here we set $|\hat{I}_0|=0$, so that the indicator is identically equal to 1 in this case. By Lemma 2.20, we have $\mathbb{E}[\tilde{F}^0] < \infty$. Since a.s. $|\hat{I}_a| \rightarrow \infty$ as $a \rightarrow \infty$, for each $\alpha > 0$ there is a deterministic $a_* = a_*(\delta, \alpha) > 0$ such that $\mathbb{E}[\tilde{F}^{a_*}] \leq \alpha$. The random variable \tilde{F}^{a_*} is a scale-invariant functional of $(\mathcal{H}, \mathcal{D}, \hat{X}, \tilde{\mathcal{D}})$ in the sense of Lemma 3.2 (note that spatially scaling by C scales the time parametrization of X by $1/C^2$, and hence the length of \hat{I}_a by C^2). Therefore, Lemma 3.2 applied to this random variable implies that a.s.

$$\limsup_{a \rightarrow \infty} \frac{1}{|\hat{I}_a|} \int_{\hat{I}_a} \sum_{\substack{H \in \mathcal{H} \\ H \sim X_t}} \frac{\mathfrak{c}(X_t, H)}{\text{Area}(X_t)} (\mathbf{v} \cdot (\phi_\infty(H) - \phi_\infty(X_t)))^2 \times \mathbb{1}_{(|\mathbf{v} \cdot (\phi_\infty(H) - \phi_\infty(X_t))| \geq \delta|\hat{I}_a|^{1/2})} dt \leq \alpha.$$

Breaking up the integral into a sum over intervals of the form $[\tau_{j-1}, \tau_j]$ which are contained in \hat{I}_a , and noting that each such interval has length $\text{Area}(Y_{j-1})/\pi(Y_{j-1})$, we get that a.s.

$$\limsup_{a \rightarrow \infty} \frac{1}{|\hat{I}_a|} \sum_{j: [\tau_{j-1}, \tau_j] \subset \hat{I}_a} \tilde{\Delta}_j^{\delta|\hat{I}_a|^{1/2}} \leq \alpha.$$

By Lemma 3.3, the same is true with \hat{I}_a replaced by the union of the dyadic parents of \hat{I}_a , which in turn contains $[0, \hat{I}_a]$. Combining this with Lemma 3.4 (to deal with the intervals containing zero and T) and sending $\alpha \rightarrow 0$ shows that the second limit in (3.26) is zero.

To treat the first limit in (3.26), we apply the same argument as above with the random variable \tilde{F}^a from (3.27) replaced by

$$F^a := \frac{\pi(H_0)}{\text{Area}(H_0)} \Delta_1^{\delta|\hat{I}_a|^{1/2}}.$$

We note that $\mathbb{E}[F^0]$ can be seen to be finite by first taking a conditional expectation given $(\mathcal{H}, \mathcal{D})$, then applying Lemma 2.20. □

Proof of Proposition 3.5. Fix a unit vector $\mathbf{v} \in \mathbb{R}^2$. We will check the hypotheses of Theorem 3.6 under the conditional law given $(\mathcal{H}, \mathcal{D})$ for the local martingales

$$M_\varepsilon^{\mathbf{v}}(t) := \varepsilon(\mathbf{v} \cdot \phi_\infty(X_{t/\varepsilon^2})).$$

Note that $M_\varepsilon^{\mathbf{v}}$ is a martingale under the conditional law given $(\mathcal{H}, \mathcal{D})$ since ϕ_∞ is discrete harmonic.

Lemma 3.8 shows that for each $t \geq 0$, a.s.

$$\lim_{\varepsilon \rightarrow 0} \langle M_\varepsilon^{\mathbf{v}}, M_\varepsilon^{\mathbf{v}} \rangle_t = \lim_{\varepsilon \rightarrow 0} \varepsilon^2 V_{t/\varepsilon^2}^{\mathbf{v}} = c_{\mathbf{v}} t.$$

Furthermore, in the notation of Theorem 3.6 and Lemma 3.9, we have

$$\sigma^\delta[M_\varepsilon^{\mathbf{v}}](t) = \varepsilon^2 \sum_{j=1}^{n_{t/\varepsilon^2}} \Delta_j^{\delta/\varepsilon},$$

so, by the discrete Doob decomposition and the optional stopping theorem, the compensator is given by

$$\tilde{\sigma}^\delta[M_\varepsilon^{\mathbf{v}}](t) = \varepsilon^2 \sum_{j=1}^{n_{t/\varepsilon^2}} \tilde{\Delta}_j^{\delta/\varepsilon}.$$

Therefore, Lemma 3.9 shows that a.s. $\lim_{\varepsilon \rightarrow 0} \tilde{\sigma}^\delta[M_\varepsilon^{\mathbf{v}}](t) = 0$ for each $t \geq 0$.

Hence, we can apply Theorem 3.6 to find that for each fixed unit vector $\mathbf{v} \in \mathbb{R}^2$, a.s. the conditional distribution given $(\mathcal{H}, \mathcal{D})$ of $M_\varepsilon^{\mathbf{v}}$ converges in law as $\varepsilon \rightarrow 0$ to standard linear Brownian motion with quadratic variation $c_{\mathbf{v}} dt$. Almost surely, this holds for Lebesgue-a.e. choice of \mathbf{v} simultaneously. By the Cramér–Wold theorem and (3.17), we obtain the proposition statement for an appropriate choice of covariance matrix Σ . This covariance matrix is positive definite since each $c_{\mathbf{v}}$ is positive (Lemma 3.7). \square

Proof of Theorem 1.16. Combine Proposition 3.5 with Proposition 2.21 and (1.17) (which has already been proven). \square

3.3. Convergence is uniform in the starting point

In this subsection we prove a straightforward extension of Theorem 1.16 which shows that in fact random walk on $\varepsilon\mathcal{H}$ converges to Brownian motion modulo time parametrization *uniformly* over all choices of starting point in a compact subset of \mathbb{C} . This result is needed in [GMS3].

THEOREM 3.10. *For $z \in \mathbb{C}$, let Y^z denote the simple random walk on \mathcal{H} started from H_z . For $j \in \mathbb{N}_0$, let \hat{Y}_j^z be an arbitrarily chosen point of the cell Y_j^z and extend \hat{Y}^z from \mathbb{N}_0 to $[0, \infty)$ by piecewise linear interpolation. Let Σ be the covariance matrix from Theorem 1.16. For each fixed compact set $A \subset \mathbb{C}$, it is a.s. the case that as $\varepsilon \rightarrow 0$, the maximum over all $z \in A$ of the Prokhorov distance between the conditional law of $\varepsilon \hat{Y}^{z/\varepsilon}$ given \mathcal{H} and the law of Brownian motion started from z with covariance matrix Σ , with respect to the topology on curves modulo time parametrization (recall (1.8)), tends to zero.*

An important consequence of Theorem 3.10 is the convergence of discrete harmonic functions to their continuous counterparts. In the statement of the following corollary and in the rest of this subsection, when we refer to “Euclidean harmonic” functions we mean functions which are harmonic with respect to the Brownian motion with covariance matrix Σ , rather than the standard 2-dimensional Brownian motion.

COROLLARY 3.11. *Let $D \subset \mathbb{C}$ be an open domain bounded by a Jordan curve, let*

$$f: \bar{D} \longrightarrow \mathbb{R}$$

be Euclidean harmonic on D and continuous on \bar{D} , and for $\varepsilon > 0$ let

$$f^\varepsilon: \mathcal{H}(\varepsilon^{-1}\bar{D}) \longrightarrow \mathbb{R}$$

be discrete harmonic on $\mathcal{H}(\varepsilon^{-1}\bar{D}) \setminus \mathcal{H}(\varepsilon^{-1}\partial D)$ with boundary data chosen so that, for

$$H \in \mathcal{H}(\varepsilon^{-1}\partial D),$$

we have $f^\varepsilon(H) = f(z)$ for some (arbitrary) $z \in H \cap \bar{D}$. View f^ε as a function on $\varepsilon^{-1}\bar{D}$ which is constant on each cell. Then a.s. $\varepsilon f^\varepsilon(\varepsilon^{-1}\cdot) \rightarrow f$ uniformly as $\varepsilon \rightarrow 0$.

Theorem 3.10 will be extracted from Theorem 1.16 by using the re-sampling property of Lemma 2.3 to get uniform convergence of random walk started from a large (but finite) number of points to Brownian motion, then using continuity considerations to transfer to uniform convergence at all points simultaneously. The following lemma is needed to compare the random walks started from two nearby points.

LEMMA 3.12. *Let G be a connected graph and let $A \subset \mathcal{V}(G)$ be a set such that the simple random walk started from any vertex of G a.s. hits A in finite time. For $x \in \mathcal{V}(G)$, let X^x be the simple random walk started from x and let τ^x be the first time X^x hits A . For $x, y \in \mathcal{V}(G) \setminus A$,*

$$d_{TV}(X_{\tau^x}^x, X_{\tau^y}^y) \leq 1 - \mathbb{P}[X^x \text{ disconnects } y \text{ from } A \text{ before time } \tau^x], \tag{3.28}$$

where d_{TV} denotes the total variation distance.

Proof. The lemma is a consequence of Wilson’s algorithm. Let T be the uniform spanning tree of G with all of the vertices of A wired to a single point. For $x \in \mathcal{V}(G)$, let L^x be the unique path in T from x to A . For a path P in G , write $\text{LE}(P)$ for its chronological loop erasure. By Wilson’s algorithm [W], we can generate the union $L^x \cup L^y$ in the following manner.

- (1) Run X^y until time τ^y and generate the loop erasure $\text{LE}(X^y|_{[0, \tau^y]})$.
- (2) Conditional on $X^y|_{[0, \tau^y]}$, run X^x until the first time $\tilde{\tau}^x$ that it hits either $\text{LE}(X^y|_{[0, \tau^y]})$ or A .
- (3) Set $L^x \cup L^y = \text{LE}(X^y|_{[0, \tau^y]}) \cup \text{LE}(X^x|_{[0, \tilde{\tau}^x]})$.

Note that $L^y = \text{LE}(X^y|_{[0, \tau^y]})$ in the above procedure. Applying the above procedure with the roles of x and y interchanged shows that $L^x \stackrel{d}{=} \text{LE}(X^x|_{[0, \tau^x]})$. When we construct $L^x \cup L^y$ as above, the points where L^x and L^y hit A coincide provided X^x hits $\text{LE}(X^y|_{[0, \tau^y]})$ before A , which happens in particular if X^x disconnects y from A before time τ^x . We thus obtain a coupling of $\text{LE}(X^x|_{[0, \tau^x]})$ and $\text{LE}(X^y|_{[0, \tau^y]})$ such that the probability that these two loop erasures hit A at the same point is at least

$$\mathbb{P}[X^x \text{ disconnects } y \text{ from } A \text{ before time } \tau^x].$$

We now obtain (3.28) by observing that $X_{\tau^x}^x$ is the same as the point where $\text{LE}(X^x|_{[0, \tau^x]})$ first hits A , and similarly with y in place of x . □

Recall the dyadic system \mathcal{D} from §2.1 and the sequence $\{S_k\}_{k \in \mathbb{Z}}$ of squares containing the origin in this system. We also use the following notation.

Notation 3.13. For a square S , write S^P for the union of the four dyadic parents of S (equivalently, the square with the same center as S with side length $3|S|$).

LEMMA 3.14. *Almost surely, the supremum over all $z \in S_k$ of the Prokhorov distance between the conditional law given $(\mathcal{H}, \mathcal{D})$ of the point where the linearly interpolated walk \widehat{Y}^z exits S_k^P and the Euclidean harmonic measure on ∂S_k^P as viewed from z , with respect to $|S_k|^{-1}$ times the Euclidean metric on ∂S_k^P , tends to zero as $k \rightarrow \infty$.*

Proof. Fix $N \in \mathbb{N}$ and, for $m > 0$, let z_1^m, \dots, z_N^m be sampled uniformly and independently from Lebesgue measure on the square \widehat{S}_m of (2.1). By Theorem 1.16 and Lemma 2.3, if we take the walks $\widehat{Y}^{z_j^m}$ for $j \in [1, N]_{\mathbb{Z}}$ to be conditionally independent given $(\mathcal{H}, \mathcal{D})$, then a.s. the joint conditional law given $(\mathcal{H}, \mathcal{D})$ of

$$t \mapsto |\widehat{S}_m|^{-1} (\widehat{Y}^{z_j^m} - z_j^m)$$

for $j=1, \dots, N$ converges to the joint law of N independent Brownian motions with covariance matrix Σ started from zero with respect to the local topology on curves modulo time parametrization.

Since a Brownian motion started from zero has probability tending to one as $\delta \rightarrow 0$ to disconnect $B_\delta(0)$ from ∞ before leaving $B_1(0)$, if we are given $\delta > 0$, we can find $N = N(\delta) \in \mathbb{N}$ such that a.s., for each large enough $m > 0$ the following is true.

(1) For each $j \in [1, N]_{\mathbb{Z}}$, the Prokhorov distance between the conditional law given $(\mathcal{H}, \mathcal{D})$ of the point where $\widehat{Y}^{z_j^m}$ exits \widehat{S}_m^P and the Euclidean harmonic measure on $\partial \widehat{S}_m^P$ as viewed from z_j^m , with respect to $|\widehat{S}_m|^{-1}$ times the Euclidean distance, is at most δ .

(2) With conditional probability at least $1 - \delta$ given $(\mathcal{H}, \mathcal{D})$, we have

$$\widehat{S}_m \subset \bigcup_{j=1}^N B_{\delta|\widehat{S}_m|}(z_j^m).$$

(3) For each $j \in [1, N]_{\mathbb{Z}}$, it holds with conditional probability at least $1 - \delta$ given $(\mathcal{H}, \mathcal{D})$ that $Y^{z_j^m}$ disconnects $\mathcal{H}(B_{\delta|\widehat{S}_m|}(z_j^m))$ from ∞ before hitting $\mathcal{H}(\partial\widehat{S}_m^{\mathbb{P}})$.

By Lemma 3.12, the condition (3) implies that the total variation distance between the discrete harmonic measure on $\mathcal{H}(\partial\widehat{S}_m^{\mathbb{P}})$ as viewed from $H_{z_j^m}$ and the harmonic measure on $\mathcal{H}(\partial\widehat{S}_m^{\mathbb{P}})$ as viewed from any $z \in B_{\delta|\widehat{S}_m|}(z_j^m)$ is at most δ . Combined with condition (1) and the continuity of Euclidean harmonic measure with respect to the base point, this implies that for any $j \in [1, N]_{\mathbb{Z}}$ and any $z \in B_{\delta|\widehat{S}_m|}(z_j^m)$, the Prokhorov distance between the conditional law given $(\mathcal{H}, \mathcal{D})$ of the point where \widehat{Y}^z exits $\widehat{S}_m^{\mathbb{P}}$ and the Euclidean harmonic measure on $\partial\widehat{S}_m^{\mathbb{P}}$ as viewed from z is at most $o_{\delta}(1)$, with the $o_{\delta}(1)$ deterministic and independent of k . Finally, combining this with condition (2) and sending $\delta \rightarrow 0$ yields the statement of the lemma. \square

Proof of Theorem 3.10. Fix $l \in \mathbb{N}$ (we will eventually let $l \rightarrow \infty$). By Lemma 3.14 and two applications of Lemma 2.8, the following holds a.s. Let $\mathcal{D}_{k,l}$ be the set of squares which are dyadic descendants of at least one of the four dyadic parents of S_k and which have side length at least $2^{-l}|S_k|$. Then, for $S \in \mathcal{D}_{k,l}$, the maximum over all $z \in S$ of the Prokhorov distance between the conditional law given $(\mathcal{H}, \mathcal{D})$ of the point where the linearly interpolated walk \widehat{Y}^z exits $S^{\mathbb{P}}$ and the Euclidean harmonic measure on $\partial S^{\mathbb{P}}$ as viewed from z with respect to $|S_k|^{-1}$ times the Euclidean distance tends to zero as $k \rightarrow \infty$.

For $z \in S_k^{\mathbb{P}}$, let $\tau_0^z := z$ and inductively let τ_n^z for $n \in \mathbb{N}$ be the first time that \widehat{Y}^z exits $S^{\mathbb{P}}$, for S the square of \mathcal{D} with side length $2^{-l}|S_k|$ which contains $\widehat{Y}_{\tau_{n-1}^z}^z$; or $\tau_n^z = \tau_{n-1}^z$ if \widehat{Y}^z is not in $S_k^{\mathbb{P}}$. Similarly, define σ_n^z for $n \in \mathbb{N}_0$ with Brownian motion \mathcal{B}^z started z in place of \widehat{Y}^z . The preceding paragraph and the strong Markov property shows that a.s. for each fixed $N \in \mathbb{N}$, the maximum over all $z \in S_k^{\mathbb{P}}$ and all $n \in [0, N]_{\mathbb{Z}}$ of the Prokhorov distance between the conditional laws given $(\mathcal{H}, \mathcal{D})$ of $\widehat{Y}_{\tau_n^z}^z$ and $\mathcal{B}_{\sigma_n^z}^z$ (with respect to $|S_k|^{-1}$ times the Euclidean metric) tends to zero as $k \rightarrow \infty$. We have $|\widehat{Y}_t^z - \widehat{Y}_{\tau_{n-1}^z}^z| \leq 2^{-l+2}|S_k|$ for $t \in [\tau_{n-1}^z, \tau_n^z]$, and similarly for \mathcal{B}^z . By Brownian scaling, the law of the largest n for which $\sigma_{n-1}^z \neq \sigma_n^z$ can be bounded uniformly independently of z and k .

It follows that a.s. it holds uniformly over all $z \in S_k^{\mathbb{P}}$ that the Prokhorov distance between the conditional law given $(\mathcal{H}, \mathcal{D})$ of \widehat{Y}^z , stopped upon exiting $S_k^{\mathbb{P}}$ and \mathcal{B}^z , stopped upon exiting $S_k^{\mathbb{P}}$, each scaled by $|S_k|^{-1}$, is $o_l(1) + o_k(1)$, with the rate of the $o_l(1)$ deterministic and independent of k .

We now conclude by sending $l \rightarrow \infty$ and noting that $B_r(0) \subset S_k^{\mathbb{P}}$ for $r \leq |S_k|$. \square

3.4. Proof of Theorems 1.5, 1.11, and 1.13

We will now prove the other variants of our main theorem.

Proof of Theorems 1.5–1.7. The combination of these theorems is a special case of Theorem 1.16, since the adjacency graph of faces of an embedded lattice is a cell configuration satisfying the connectivity along lines hypothesis (3). \square

Proof of Theorem 1.11. Let \mathcal{M} be an embedded lattice satisfying the hypotheses of Theorem 1.11. We will construct a cell configuration $\mathring{\mathcal{H}}$ which satisfies the hypotheses of Theorem 1.16 and whose cells correspond to the vertices of \mathcal{M} .

We first divide each face $H \in \mathcal{FM}$ into cells corresponding to the vertices in ∂H . One way of doing this is as follows. For each such face H , let w_H be sampled uniformly from Lebesgue measure on H , and let x_1, \dots, x_n be the vertices lying on ∂H , in counterclockwise cyclic order. Let $\psi: H \rightarrow \mathbb{D}$ be a conformal map sending w_H to zero. For $j \in [1, n]_{\mathbb{Z}}$, let P_j be a radial line segment in \mathbb{D} from zero to the midpoint of the counterclockwise arc of $\partial \mathbb{D}$ from $\psi(x_j)$ to $\psi(x_{j-1})$ (where here we set $x_0 = x_n$). We then let

$$\mathring{H}_{x_j, H} \subset H$$

be the closure of the image under ψ^{-1} of the slice of \mathbb{D} lying between P_{j-1} and P_j (where here we set $P_0 = P_n$).

For $x \in \mathcal{VM}$, we define the cell

$$\mathring{H}_x := \bigcup_{H \in \mathcal{FM}: x \in \partial H} \mathring{H}_{x, H}.$$

We then set

$$\mathring{\mathcal{H}} := \{ \mathring{H}_x : x \in \mathcal{VM} \}$$

and declare that $\mathring{H}_x \sim \mathring{H}_y$ if and only if x and y are joined by an edge in \mathcal{M} , which (by the above definition of the $\mathring{H}_{x, H}$'s) is the case if and only if \mathring{H}_x and \mathring{H}_y intersect along a non-trivial boundary arc. We also set

$$\mathfrak{c}(\mathring{H}_x, \mathring{H}_y) := \mathfrak{c}(x, y) \quad \text{and} \quad \pi(\mathring{H}_x) := \pi(x).$$

We now check the hypotheses of Theorem 1.16 for the cell configuration $\mathring{\mathcal{H}}$. The translation invariance modulo scaling of $\mathring{\mathcal{H}}$ follows from the translation invariance modulo scaling hypothesis for \mathcal{M} and the fact that we defined the cells in each face in a manner depending only on the face (to avoid continuity issues, we use one of the definitions which does not involve limits, e.g., the one given in Lemma 2.3). Ergodicity modulo scaling for $\mathring{\mathcal{H}}$ is then immediate from ergodicity modulo scaling for \mathcal{M} . To check hypothesis (2) for $\mathring{\mathcal{H}}$, we observe that the translation invariance modulo scaling of \mathcal{M} shows that, if we

condition on \mathcal{M} and sample z uniformly from the face H_0 , then $\mathcal{M} - z \stackrel{d}{=} \mathcal{M}$. Consequently, we can average over H_0 to find that if \mathring{H}_0 is the cell of $\mathring{\mathcal{H}}$ containing 0, then

$$\begin{aligned} \mathbb{E} \left[\frac{\text{diam}(\mathring{H}_0)^2 \pi(\mathring{H}_0)}{\text{Area}(\mathring{H}_0)} \right] &= \mathbb{E} \left[\sum_{x \in \mathcal{V} \mathcal{M} \cap \partial H_0} \frac{\text{diam}(\mathring{H}_x)^2 \pi(x) \text{Area}(\mathring{H}_x \cap h_0)}{\text{Area}(\mathring{H}_x) \text{Area}(H_0)} \right] \\ &\leq \mathbb{E} \left[\sum_{x \in \mathcal{V} \mathcal{M} \cap \partial H_0} \frac{\text{Outrad}(x)^2 \pi(x)}{\text{Area}(H_0)} \right], \end{aligned}$$

which is finite by our finite expectation hypothesis for \mathcal{M} . Similarly with π^* in place of π . Hypothesis (3) for \mathcal{H} follows since one can check from the above construction that if $H_x \cap H_y \neq \emptyset$ but $H_x \not\sim H_y$, then $H_x \cap H_y$ is a finite collection of points of the form w_H defined above. The theorem statement now follows from Theorem 1.16. \square

The hardest part of the proof of Theorem 1.13 is the following analog of Lemma 2.9.

LEMMA 3.15. *In the setting of Theorem 1.13, a.s.*

$$\lim_{r \rightarrow \infty} \frac{1}{r} \max_{\substack{e \in \mathcal{EM} \\ e \cap B_r(0) \neq \emptyset}} \text{diam}(e) = \lim_{r \rightarrow \infty} \frac{1}{r} \max_{\substack{e^* \in \mathcal{EM}^* \\ e^* \cap B_r(0) \neq \emptyset}} \text{diam}(e^*) = 0. \tag{3.29}$$

We note that Lemma 3.15 implies that also the maximum diameter of the primal or dual faces which intersect $B_r(0)$ tends to zero since any two points on the boundary of such a face are contained in a connected union of two primal edges and two dual edges.

Lemma 3.15 will be proven in Appendix C. The proof is much harder than that of Lemma 2.9 since we need to rule out the possibility of large edges with small conductances. In the setting of Theorem 1.16, we have a bound for $\mathbb{E}[\text{diam}(H_0)/\text{Area}(H_0)]$, since $\pi(H_0) + \pi^*(H_0) \leq 1$, but we have no such bound in the setting of Theorem 1.13. Once Lemma 3.15 is established, Theorem 1.13 follows from essentially the same argument used to prove Theorem 1.16. We now give a brief account of the (entirely cosmetic) adaptations necessary.

Proof of Theorem 1.13. All of the proofs in §2.1 and §2.2 carry over verbatim except that we replace the quantities appearing in (1.5) with the quantities appearing in (1.12) and we apply Lemma 3.15 instead of Lemma 2.9. The arguments of §2 through §2.5 are exactly the same except that in the proof of Lemma 2.17 we sum over edges of \mathcal{M} whose corresponding dual edges cross the line segment $P_k(r)$ instead of over cells of \mathcal{H} which intersect this line segment. The set of all such edges is necessarily connected since if a dual edge crosses $P_k(r)$, then the faces on either side of the edge must intersect $P_k(r)$. The argument of §3 is identical, with the area of the face of \mathcal{M}^* containing a vertex used in place of the area of the cell in the parametrization of the walk X . \square

4. Applications to planar maps and Liouville quantum gravity

Our results are motivated by applications to the convergence of embedded random planar maps toward Liouville quantum gravity, which are explored in more depth in the papers [GMS2] and [GMS3]. Here we provide a brief overview of such applications.

There are various ways of embedding a planar map M into \mathbb{C} (i.e., drawing the vertices and edges) which can be thought of as canonical, in some sense. Particularly relevant for us is the *Tutte embedding* (also known as *harmonic* or *barycentric* embedding) where the position of each vertex is the average of the positions of its neighbors, equivalently the embedding function is discrete harmonic. In the setting of Theorem 1.6, the function ϕ_∞ is a version of the Tutte embedding of the dual map of \mathcal{M} . Other examples of embeddings include circle packing, Riemann uniformization, and various types of square tilings.

It is a long-standing conjecture that for a large class of random planar maps—including uniform maps (or triangulations, quadrangulations, etc.) and maps sampled with probability proportional to the partition functions of various statistical mechanics models—the image of the map under any one of the above embeddings should converge to γ -Liouville quantum gravity (γ -LQG) with parameter $\gamma \in (0, 2]$ depending on the universality class of the random planar map. More precisely, the counting measure on vertices of the map, appropriately re-scaled, should converge to a version of the γ -LQG measure, a random fractal measure on the plane constructed e.g. in [DS] (see also [Ka], [RV]). The idea that LQG should describe (in some sense) the scaling limit of random planar maps dates back at least to Polyakov in the 1980’s [P1], [P2]. Various precise conjectures along these lines are stated, e.g., in [DS], [S], [DMS], [HRV], [C]. We will not review the definition of LQG here. See the above references for more detail.

For each $\gamma \in (0, 2)$, there is a certain special LQG surface called the *zero-quantum cone* whose corresponding γ -LQG measure μ satisfies an exact continuum analog of translation invariance modulo scaling [DMS, Proposition 4.13 (ii)] (see also [GMS3, Proposition 3.1]). Roughly speaking, a zero-quantum cone is an infinite-volume γ -LQG surface centered at a “Lebesgue typical” point. One can use the zero-quantum cone to construct embedded lattices or cell configurations which approximate LQG and are translation invariant modulo scaling in the sense of Definition 1.2. For example, we could sample a dyadic system \mathcal{D} independent from μ and look at the embedded lattice whose faces are the maximal dyadic squares of \mathcal{D} with μ -mass smaller than one. Other LQG surfaces are not exactly translation invariant modulo scaling, but locally look like the zero-quantum cone in the sense of local absolute continuity.

It is easy to construct embeddings of random planar maps which converge to LQG if the embedding is not required to be defined in an intrinsic way with respect to the map. As a trivial example, if we wanted to embed a map with n vertices we could sample n

i.i.d. points from the γ -LQG measure, send one vertex to each of these points, then draw in the edges arbitrarily. More useful examples arise from the theory developed in [DMS], as we discuss just below.

The results of the present paper tell us that if we can find *some* a-priori embedding of a random planar map or its dual into \mathbb{C} under which (a) the embedded map converges to γ -LQG and (b) the hypotheses of one of our main theorems are satisfied, then the Tutte embedding will be close to the a-priori embedding at large scales, and hence the map will also converge to γ -LQG under the Tutte embedding. Moreover, one also gets convergence of random walk on the map to Brownian motion under both the a-priori embedding and the Tutte embedding.

The above approach is used in [GMS3] to show that the Tutte embedding of the *mated-CRT maps*, a family of random planar maps constructed by gluing together a pair of correlated continuum random trees (CRT's), converges to γ -LQG, with $\gamma \in (0, 2)$ determined by the correlation parameter. This constitutes the first rigorous proof that embedded random planar maps converge to LQG.

The needed a-priori embedding comes from the main result of [DMS], which gives a correspondence between pairs of correlated CRT's and LQG surfaces decorated by space-filling Schramm–Loewner evolution (SLE_κ) curves for $\kappa = 16/\gamma^2$. In particular, if μ is an appropriate version of the γ -LQG measure and η is a whole-plane space-filling SLE_κ sampled independently from μ and then parameterized so that it traverses one unit of μ -mass in one unit of time, then [DMS, Theorem 1.9] implies that the mated-CRT map is isomorphic to the adjacency graph of unit μ -mass “cells” $\eta([x-1, x])$ for $x \in \mathbb{Z}$. Theorem 1.16 is used to show that the random walk on this adjacency graph of cells converges to Brownian motion modulo time parametrization, and hence that the a-priori embedding is close to the Tutte embedding. (It is expected — but not yet proven — that the random walk on this adjacency graph of cells converges uniformly to *Liouville Brownian motion*, the “quantum” time parametrization of Brownian motion constructed in [GRV], [Be].) We also note that the paper [GMS1] proves more quantitative versions of many of the estimates in this paper (e.g., for the Dirichlet energy and modulus of continuity of discrete harmonic functions) in the special case of the cell configuration associated with a mated-CRT map.

Another application of the above approach appears in [GMS2], which considers a random planar map whose vertices correspond to the cells of a Poisson–Voronoi tessellation of the Brownian map (or some other Brownian surface, like the Brownian disk), with two vertices considered to be adjacent if the corresponding cells intersect. As in [GMS3], it is proven that the Tutte embeddings of these maps converge to $\sqrt{\frac{8}{3}}$ -Liouville quantum gravity in the appropriate sense when the intensity of the Poisson point process tends

to ∞ . In this case, the a-priori embedding comes from the results of [MS1]– [MS3], which show that there is a way of embedding of a Brownian surface into the plane under which it is equivalent to $\sqrt{\frac{8}{3}}$ -Liouville quantum gravity (but do not give a direct description of this embedding in terms of the Brownian surface). We also show that “Brownian motion on a Brownian surface” (an object that was previously constructed in a somewhat indirect way) can be obtained as the scaling limit of simple random walk on these cells. In [AFS], these results are extended to the case when the Brownian map is replaced by a γ -LQG surface for general $\gamma \in (0, 2)$, viewed as a metric measure space.

As we discuss in more detail in [GMS2] and [GMS3], there are many additional potential applications of our results in the theory of random planar maps and LQG along these same lines.

Appendix A. Specific Dirichlet energy

Let \mathcal{M} be an embedded lattice (Definition 1.1) which is translation invariant modulo scaling (Definition 1.2). Recall that the *specific Dirichlet energy* associated with \mathcal{M} is the “expected amount of discrete Dirichlet energy per unit area” for the function $\mathcal{V}(\mathcal{M}) \rightarrow \mathbb{C}$ which takes each vertex of \mathcal{M} to its position in \mathbb{C} .

As mentioned in the introduction, there are various ways to make this definition precise, depending on how one chooses the origin-containing region and how each edge’s energy is “localized”. In this appendix, we formalize this family of definitions (which includes the definition presented as (1.12) within the statement of Theorem 1.13). We then use mass transport to briefly explain why the different definitions are all equivalent.

To make the notion of specific Dirichlet energy precise, the Dirichlet energy

$$c(e)|x - y|^2$$

associated with each edge $e = \{x, y\} \in \mathcal{E}(\mathcal{M})$ needs to be spread out in space in some way. That is, we assume we have a way to associate to each edge a measure $\nu_{\mathcal{M}}^e$ on the plane of total mass $c(e)|x - y|^2$, where the map from (\mathcal{M}, e) to $\nu_{\mathcal{M}}^e$ commutes with dilation and translation in the obvious way:

$$\nu_{C(\mathcal{M}-z)}^{C(e-z)}(\cdot) = C^2 \nu_{\mathcal{M}}^e(C^{-1}\cdot).$$

We set

$$\nu = \nu_{\mathcal{M}} = \sum_{e \in \mathcal{E}(\mathcal{M})} \nu_{\mathcal{M}}^e.$$

One concrete way to do this is to let $\nu_{\mathcal{M}}^e$ spread half its mass uniformly over each of the two faces incident to e . Then, the density of ν on a face H is given by

$$\frac{1}{2} \sum c(e) \text{diam}(e)^2 / \text{Area}(H),$$

where the sum is over all edges on the boundary of H . Another example is to let $\nu_{\mathcal{M}}^e$ be a sum of two point masses (at the two endpoints of e), or a single point mass at some specified internal point on e .

Definition A.1. If the law of \mathcal{M} is translation invariant modulo scaling, then the *specific Dirichlet energy* of \mathcal{M} is defined as $\mathbb{E}[\nu(B)/\text{Area}(B)]$, where ν is as described above and B is the origin-containing block in any block partition $\mathcal{B}(\mathcal{M})$ of the sort defined in condition (3) of Definition 1.2 such that a.s. the ν measure of the boundary of any block is zero.⁽⁴⁾ Equivalent definitions of the specific Dirichlet energy (obtained as special cases) are as follows.

- (1) The expected density of ν at zero (for any choice of ν that is a.s. absolutely continuous with respect to Lebesgue measure).
- (2) With H_0 the origin-containing face,

$$\mathbb{E} \left[\sum_{e=\{x,y\} \in \mathcal{E}\mathcal{M} \cap \partial H_0} c(e) \frac{\text{diam}(e)^2}{\text{Area}(H_0)} \right].$$

Now, we will sketch an argument for the fact that Definition A.1 does not depend on the specific way the measures ν and the blocks B are defined. We first observe that the definition is not changed if we replace each $\nu_{\mathcal{M}}^e$ with a measure that spreads its mass *uniformly* over each block B to which it assigns measure, but otherwise assigns the same measure to each block B . Such a change would not impact the value of $\nu(B)$ for the origin-containing block B referenced in Definition A.1. We next observe that if we had a different measure $\tilde{\nu}$ with this property (for a different block partition), we could define a mass transport function G_e such that (for a given edge e) any amount of mass associated with $\nu_{\mathcal{M}}^e$ is spread evenly according to the measure $\tilde{\nu}_{\mathcal{M}}^e$. In other words,

$$\int_A \int_B G_e(a, b) da db = \nu_{\mathcal{M}}^e(A) \tilde{\nu}_{\mathcal{M}}^e(B).$$

Summing these G_e over all edges e , one obtains a mass transport function F on the whole plane. The difference in the density of $\tilde{\nu}$ and ν at zero is given by the net amount of flow in the transport function and since (by the assumption in Definition 1.2) this has expectation zero, the two definitions of specific free energy must agree.

⁽⁴⁾ This definition is independent of the choice of block partition \mathcal{B} and the energy location rule ν , as can easily be seen, e.g., from the mass transport principle (condition (4) of Definition 1.2).

Remark A.2. In Definition A.1, the Dirichlet energy associated with an edge

$$e = \{x, y\}$$

is taken to be

$$c(e)|x - y|^2.$$

However, in some of the proofs involving specific Dirichlet energy (and related quantities) we need bounds for $c(e) \text{diam}(e)^2$ instead, which is larger than $c(e)|x - y|^2$, unless the edges are straight lines. Hence, the finiteness conditions in those theorem statements involve bounds for diameters rather than just distance between vertices.

Appendix B. Equivalence of translation invariance modulo scaling definitions

In this appendix we prove Lemma 2.3, which asserts the equivalence of the conditions of Definition 1.2 plus the extra condition of Lemma 2.3 in the setting of a cell configuration (analogous equivalence results for embedded lattices or pairs of primal/dual embedded lattices follow from identical arguments). We will prove that each of the conditions of Definition 1.2 is equivalent to the dyadic system resampling condition (condition (5) from Lemma 2.3).

Proof of (1) \Rightarrow (2). Fix $m > 0$. Let \mathcal{H}_n, z_n , and C_n be as in condition (1) with cell configurations in place of lattices, so that each \mathcal{H}_n is a cell configuration with the property that the union of its cells is a disk or a square (instead of all of \mathbb{C}) and $C_n(\mathcal{H}_n - z_n) \rightarrow \mathcal{H}$ in law. For $n \in \mathbb{N}$, let \mathcal{D}_n be a uniform dyadic system independent from $(\mathcal{H}_n, z_n, C_n)$ and let $\widehat{S}_{m,n}^{z_n}$ be defined as in (2.1) with \mathcal{H}_n in place of \mathcal{H} and $z = z_n$, or $\widehat{S}_{m,n}^{z_n} = \mathbb{C}$ if no square satisfying the condition of (2.1) exists (which can happen since \mathcal{H}_n is finite). By Lemma 2.1, $C_n(\mathcal{D}_n - z_n)$ is a uniform dyadic system independent from $C_n(\mathcal{H}_n - z_n)$. The square $C_n(\widehat{S}_{m,n}^{z_n} - z_n)$ is determined by $(C_n(\mathcal{H}_n - z_n), C_n(\mathcal{D}_n - z_n))$ in the same continuous manner that \widehat{S}_m is determined by $(\mathcal{H}, \mathcal{D})$. Thus,

$$(C_n(\mathcal{H}_n - z_n), C_n(\mathcal{D}_n - z_n), C_n(\widehat{S}_{m,n}^{z_n} - z_n)) \rightarrow (\mathcal{H}, \mathcal{D}, \widehat{S}_m) \tag{B.1}$$

in law, with respect to the topology on cell configurations on the first coordinate and any choice of topology on the other two coordinates. It follows from (B.1) that, with probability tending to one as $n \rightarrow \infty$, the square $\widehat{S}_{m,n}^{z_n}$ is contained in the union of the cells of \mathcal{H}_n . Since z_n is a uniform sample from this union, it follows that the total variation distance between the conditional law of z_n given $(\mathcal{H}_n, \mathcal{D}_n, \widehat{S}_{m,n}^{z_n})$ and the uniform measure on $\widehat{S}_{m,n}^{z_n}$ a.s. converges to zero as $n \rightarrow \infty$. Combining this with (B.1) shows that condition (5) holds. \square

Proof of (5) \Rightarrow (1). For $m > 0$, define a cell configuration on \widehat{S}_m by

$$\mathcal{H}_m := \{H \cap \widehat{S}_m : H \in \mathcal{H}, H \subset \widehat{S}_m\} \cup \{L_m\} \quad \text{for } L_m := \overline{\widehat{S}_m \setminus \bigcup_{\substack{H \in \mathcal{H} \\ H \subset \widehat{S}_m}} H}$$

Then, \mathcal{H}_m is a collection of cells whose union is the square \widehat{S}_m . We endow $\mathcal{H}_m \cap \mathcal{H}$ with the adjacency relation and the conductances it inherits from \mathcal{H} and we declare that $H \sim L_m$ for each $H \in \mathcal{H}_m \cap \mathcal{H}$ with $H \cap L_m \neq \emptyset$.

For $m \in \mathbb{N}$, let z_m be sampled uniformly from Lebesgue measure on \widehat{S}_m . By condition (5), for each $m > 0$ there exists $C_m > 0$ such that $C_m(\mathcal{H} - z_m) \stackrel{d}{=} \mathcal{H}$. We claim that $C_m(\mathcal{H}_m - z_m) \rightarrow \mathcal{H}$ in law. Indeed, if $l > 0$ and $\widehat{S}_l^{z_m} \subset \widehat{S}_m \setminus L_m$, then $\mathcal{H}(\widehat{S}_l^{z_m}) = \mathcal{H}_m(\widehat{S}_l^{z_m})$, i.e., the analogs of $\mathcal{H}(\widehat{S}_l)$ for the scaled/translated cell configurations $C_m(\mathcal{H} - z_m)$ and $C_m(\mathcal{H}_m - z_m)$ are identical. Since the union of \widehat{S}_l over all $l > 0$ is all of \mathbb{C} and

$$C_m(\mathcal{H} - z_m) \stackrel{d}{=} \mathcal{H},$$

it suffices to show that, for each fixed $l > 0$, one has

$$\lim_{m \rightarrow \infty} \mathbb{P}[\widehat{S}_l^{z_m} \subset \widehat{S}_m \setminus L_m] = 1.$$

By the definition of L_m , this will follow from

$$\lim_{m \rightarrow \infty} \mathbb{P} \left[\bigcup_{H \in \mathcal{H}(\widehat{S}_l^{z_m})} H \subset \widehat{S}_m \right] = 1. \quad (\text{B.2})$$

To prove (B.2), we observe that condition (5) implies that

$$\frac{1}{|\widehat{S}_m|} \text{diam} \left(\bigcup_{H \in \mathcal{H}(\widehat{S}_l^{z_m})} H \right) \stackrel{d}{=} \frac{1}{|\widehat{S}_m|} \text{diam} \left(\bigcup_{H \in \mathcal{H}(\widehat{S}_l)} H \right),$$

which tends to zero in probability as $m \rightarrow \infty$ (l fixed). Since z_m is sampled uniformly from \widehat{S}_m (so is unlikely to be close to $\partial \widehat{S}_m$), this implies (B.2). \square

Proof of (2) \Rightarrow (5). The proof is similar to the analogous implication started from condition (1). Fix $m > 0$. Let z_n be sampled uniformly from U_n , as in condition (2), and let $C_n > 0$ be chosen so that $C_n(\mathcal{H} - z_n) \rightarrow \mathcal{H}$ in law. By Lemma 2.1, $C_n(\mathcal{D} - z_n)$ is a uniform dyadic system independent from $C_n(\mathcal{H} - z_n)$. Using the notation (2.1), the square $C_n(\widehat{S}_m^{z_n} - z_n)$ is determined by $(C_n(\mathcal{H} - z_n), C_n(\mathcal{D} - z_n))$ in the same continuous manner that \widehat{S}_m is determined by $(\mathcal{H}, \mathcal{D})$. Thus,

$$(C_n(\mathcal{H} - z_n), C_n(\mathcal{D} - z_n), C_n(\widehat{S}_m^{z_n} - z_n)) \rightarrow (\mathcal{H}, \mathcal{D}, \widehat{S}_m) \quad (\text{B.3})$$

in law, with respect to the topology on cell configurations (1.15) on the first coordinate, any choice of topology on the second coordinate, and the Hausdorff distance on the last coordinate.

We will now argue that

$$\lim_{n \rightarrow \infty} \mathbb{P}[\widehat{S}_m^{z_n} \subset U_n] = 1. \tag{B.4}$$

Note that this is not obvious, since we do not assume that the maximum diameter of the cells of \mathcal{H} which intersect U_n is $o(\text{diam}(U_n))$, so it is a priori possible that $\widehat{S}_m^{z_n}$ could be comparable in size to U_n (although we will show that this is not the case). Since \mathcal{H} is locally finite and the number of cells which intersect U_n tends to ∞ as $n \rightarrow \infty$, the convergence $C_n(\mathcal{H} - z_n) \rightarrow \mathcal{H}$ implies that $C_n \text{diam}(U_n) \rightarrow \infty$ in probability (for otherwise there would be a subsequence of n 's along which $C_n(\mathcal{H} - z_n)$ has uniformly positive probability to contain an arbitrarily large number of cells which intersect some fixed compact set). Since \widehat{S}_m has finite side length, it follows from (B.3) that the laws of the random variables $C_n|\widehat{S}_m^{z_n}|$ are tight. Combining the preceding two sentences shows that $|\widehat{S}_m^{z_n}|/\text{diam}(U_n) \rightarrow 0$ in probability, and thus, since z_n is sampled uniformly from U_n (so is unlikely to be close to ∂U_n), we get (B.4).

By (B.4) and since z_n is sampled uniformly from U_n , a.s. the total variation distance between the conditional law of z_n given $(\mathcal{H}, \mathcal{D}, \widehat{S}_m^{z_n})$ and the uniform measure on $\widehat{S}_m^{z_n}$ tends to zero as $n \rightarrow \infty$. Combining this with (B.3) shows that condition (5) holds. \square

Proof of (5) \Rightarrow (2). Sample a uniform dyadic system \mathcal{D} independent from \mathcal{H} and take $U_n = \widehat{S}_n$ (as in (2.1)) for $n \in \mathbb{N}$. \square

Proof of (3) \Rightarrow (5). Let \mathcal{D} be a dyadic system independent from \mathcal{H} and apply condition (3) to the block decomposition $\{\widehat{S}_m^z : z \in \mathbb{C}\}$, with \widehat{S}_m^z as in (2.1). \square

Proof of (5) \Rightarrow (3). Given a block decomposition \mathcal{B} as in condition (3), let B_z for $z \in \mathbb{C}$ be the (a.s. unique) block of \mathcal{B} containing the z . Also, let \mathcal{D} be a uniform dyadic system independent from $(\mathcal{H}, \mathcal{B})$. By condition (5) (and the scaling condition on block decompositions), if $m > 0$ and we sample z_m uniformly from \widehat{S}_m , then

$$(\mathcal{H} - z_m, \mathcal{B} - z_m, \mathcal{D} - z_m)$$

agrees in law with $(\mathcal{H}, \mathcal{B}, \mathcal{D})$ modulo spatial scaling. Since the block B_0 is a compact set and the union of the \widehat{S}_m 's is all of \mathbb{C} , this shows that

$$\mathbb{P}[B_{z_m} \subset \widehat{S}_m] = \mathbb{P}[B_0 \subset \widehat{S}_m] \rightarrow 1 \quad \text{as } m \rightarrow \infty. \tag{B.5}$$

If we condition on $(\mathcal{H}, \mathcal{D}, \mathcal{B})$ and sample w_m uniformly from Lebesgue measure on

$$B_{z_m} \cap \widehat{S}_m,$$

then the conditional law of w_m given $(\mathcal{H}, \mathcal{B}, \mathcal{D})$ is uniform on \widehat{S}_m . Thus, by condition (5) $\mathcal{H} - w_m$ agrees in law with \mathcal{H} modulo scaling. On the other hand, (B.5) implies that the total variation distance between the law of (\mathcal{H}, w_m) and the law of (\mathcal{H}, w) tends to zero as $m \rightarrow \infty$. Therefore, $\mathcal{H} - w$ agrees in law with \mathcal{H} modulo scaling. \square

Proof of (4) \Rightarrow (5). To prove condition (5), it suffices to show that for each non-negative measurable function $G(\mathcal{H}, w)$ of a cell configuration and a point in \mathbb{C} which is invariant under translation and scaling in the sense that

$$G(C(\mathcal{H} - z), C(w - z)) = G(\mathcal{H}, w)$$

for each $C > 0$ and $z \in \mathbb{C}$, we have

$$\mathbb{E} \left[\frac{1}{|\widehat{S}_m|^2} \int_{\widehat{S}_m} G(\mathcal{H}, z) dz \right] = \mathbb{E}[G(\mathcal{H}, 0)] \quad \text{for all } m > 0. \tag{B.6}$$

To see why this is sufficient, suppose that (B.6) holds and we are given an arbitrary function \widehat{G} from the space of cell configurations to $[0, \infty)$. Define

$$G(\mathcal{H}, w) := \widehat{G} \left(\frac{\mathcal{H} - w}{\text{diam}(H_w)} \right),$$

where H_w is a cell of \mathcal{H} which contains w (chosen by some deterministic convention if there is more than one). If z_m is sampled uniformly from \widehat{S}_m , then (B.6) applied with this choice of G shows that

$$\mathbb{E} \left[\widehat{G} \left(\frac{\mathcal{H}}{\text{diam}(H_0)} \right) \right] = \mathbb{E} \left[\widehat{G} \left(\frac{\mathcal{H} - z_m}{\text{diam}(H_{z_m})} \right) \right].$$

Thus, $\mathcal{H} \stackrel{d}{=} \mathcal{H} - z_m$ modulo scaling.

To prove (B.6), we fix $m \geq 0$ and define a function on cell configurations with two marked points by

$$F(\mathcal{H}, w_0, w_1) := \mathbb{E} [|\widehat{S}_m^{w_0}|^{-2} G(\mathcal{H}, w_1) \mathbf{1}_{(w_1 \in \widehat{S}_m^{w_0})} | \mathcal{H}].$$

Note that

$$\{w_1 \in \widehat{S}_m^{w_0}\} = \{\widehat{S}_m^{w_1} = \widehat{S}_m^{w_0}\}.$$

Then, F is covariant with respect to translation and scaling in the sense of (1.2), so we can apply condition (4) to F to find that (B.6) holds. \square

Proof of (5) ⇒ (4). For $m \in \mathbb{N}$, let z_m be sampled uniformly from Lebesgue measure on \widehat{S}_m , as in condition (5). This condition tells us that there is a random $C_m > 0$ such that

$$(C_m(\mathcal{H} - z_m), C_m(\mathcal{D} - z_m)) \stackrel{d}{=} (\mathcal{H}, \mathcal{D}).$$

If we condition on $(C_m(\mathcal{H} - z_m), C_m(\mathcal{D} - z_m))$, then the conditional law of the point $-C_m z_m$ (which corresponds to the origin) is uniform on $C_m(\widehat{S}_m - z_m)$. Thus,

$$(C_m(\mathcal{H} - z_m), C_m(\widehat{S}_m - z_m), -C_m z_m) \stackrel{d}{=} (\mathcal{H}, \widehat{S}_m, z_m).$$

Hence, if $F(\mathcal{H}, w_1, w_2)$ is any non-negative measurable function of a cell configuration and a pair of points, then

$$\begin{aligned} \mathbb{E} \left[\int_{\widehat{S}_m} F(\mathcal{H}, 0, z) dz \right] &= \mathbb{E} [|\widehat{S}_m|^2 F(\mathcal{H}, 0, z_m)] \\ &= \mathbb{E} [C_m^2 |\widehat{S}_m|^2 F(C_m(\mathcal{H} - z_m), 0, -C_m z_m)] \\ &= \mathbb{E} \left[\int_{C_m(\widehat{S}_m - z_m)} F(C_m(\mathcal{H} - z_m), 0, z) dz \right] \\ &= \mathbb{E} \left[C_m^2 \int_{\widehat{S}_m} F(C_m(\mathcal{H} - z_m), 0, C_m(u - z_m)) du \right], \end{aligned} \tag{B.7}$$

where in the last line we have made the change of variables $z = C_m(u - z_m)$, $dz = C_m^2 du$. If F satisfies (1.2), then

$$C_m^2 F(C_m(\mathcal{H} - z_m), 0, C_m(u - z_m)) = F(\mathcal{H}, z_m, u)$$

so (B.7) gives

$$\begin{aligned} \mathbb{E} \left[\int_{\widehat{S}_m} F(\mathcal{H}, 0, z) dz \right] &= \mathbb{E} \left[\int_{\widehat{S}_m} F(\mathcal{H}, z_m, u) du \right] \\ &= \mathbb{E} \left[\frac{1}{|\widehat{S}_m|^2} \int_{\widehat{S}_m \times \widehat{S}_m} F(\mathcal{H}, v, u) du dv \right], \end{aligned} \tag{B.8}$$

here using that z_m is uniform on \widehat{S}_m conditional on $(\mathcal{H}, \mathcal{D})$. The same argument shows that the right-hand side of (B.8) is also equal to $\mathbb{E} [\int_{\widehat{S}_m} F(\mathcal{H}, z, 0) dz]$, so we get

$$\mathbb{E} \left[\int_{\widehat{S}_m} F(\mathcal{H}, 0, z) dz \right] = \mathbb{E} \left[\int_{\widehat{S}_m} F(\mathcal{H}, z, 0) dz \right] \quad \text{for all } m > 0.$$

Sending $m \rightarrow \infty$ now shows that (1.3) holds. □

Appendix C. No macroscopic edges in Theorem 1.13

In this appendix we prove Lemma 3.15, which says that the hypotheses of Theorem 1.13 imply that the maximal diameter of the edges of the embedded lattice \mathcal{M} and of its dual \mathcal{M}^* which intersect $B_r(0)$ is a.s. of order $o_r(r)$. The proof of this assertion is much more difficult than the proof of the analogous assertion in the settings of the other theorems (see Lemma 2.9), since we only have bounds for edge diameters weighted by conductances, so one needs to rule out the possibility of macroscopic edges with small conductances.

We emphasize that nothing in this appendix is needed for the proofs of our other main results (Theorems 1.5, 1.11, and 1.16) or in the companion paper [GMS3].

Throughout this appendix we identify the real axis with \mathbb{R} . We also fix a uniform dyadic system \mathcal{D} independent from $(\mathcal{M}, \mathcal{M}^*)$ (§2.1) and let $\{S_k\}_{k \in \mathbb{Z}}$ be its sequence of origin-containing squares.

The basic idea of the proof is to study the edges of \mathcal{M}^* whose corresponding primal edges intersect \mathcal{M} . As we will explain just below, simple geometric considerations imply that this set of dual edges contains a bi-infinite simple path $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$. The main step of the proof is to show that this path does not deviate too far from \mathbb{R} , in the sense that, for each $\varepsilon > 0$, it holds for large enough $k \in \mathbb{N}$ that $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ admits a sub-path which joins the left and right boundaries of S_k and stays in the $\varepsilon|S_k|$ -neighborhood of $S_k \cap \mathbb{R}$ (Lemma C.7). This will be accomplished in two main steps: we first show that a.s. for large enough k , each point of $S_k \cap \mathbb{R}$ is contained in a small (side length of order $o_k(|S_k|)$) square of \mathcal{D} which intersects one of the $\mathbf{e}^*(j)$'s and satisfies an additional technical property (Lemma C.5). We will then argue that the excursions of the path $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ between the times when it hits these squares cannot be macroscopic (Lemma C.6). Both of these statements are proven using variants of the argument used to prove Lemma 2.9.

Once we know that the maximum deviation of the path $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ from \mathbb{R} is macroscopic, we will use translation invariance modulo scaling to construct a “grid” of paths in \mathcal{M}^* which approximate horizontal and vertical lines in S_k such that each of the complementary connected components of the union of these paths has diameter $o_k(|S_k|)$. Since no dual edge can cross any of the paths, this will conclude the proof of Lemma 3.15.

Since we will mostly be working with \mathbb{R} rather than \mathbb{C} , we will first set up an analog of Lemma 2.7 for averages along \mathbb{R} .

Definition C.1. Recall the squares $\widehat{S}_m = \widehat{S}_m^0$ for $m > 0$ from (2.1) (defined with faces of \mathcal{M} in place of cells of \mathcal{H}). For $m \in \mathbb{N}$, let $\mathcal{F}_m^{\mathbb{R}}$ be the σ -algebra generated by the measurable functions of $F = F(\mathcal{H}, \mathcal{D})$ which satisfy

$$F(C(\mathcal{H} - z), C(\mathcal{D} - z)) = F(\mathcal{H}, \mathcal{D}) \quad \text{for all } z \in \widehat{S}_m \cap \mathbb{R} \text{ and all } C > 0. \tag{C.1}$$




Figure 7. *Left:* A primal edge/dual edge pair (e, e^*) and the corresponding rectangle R_e , defined as in (C.4). *Right:* Illustration of the construction of the path $\{\epsilon^*(j)\}_{j \in \mathbb{Z}}$. The faces H of \mathcal{M} which intersect the real axis are shown in black and the corresponding first intersection points $\tilde{\tau}(H)$ are shown in blue. The edges $\tilde{\epsilon}^*(i)$ for $i \in \mathbb{Z}$ are shown in red. These edges form a tree. We take $\{\epsilon^*(j)\}_{j \in \mathbb{Z}}$ to be the lowest bi-infinite path in this tree.

Also set

$$\mathcal{F}_\infty^{\mathbb{R}} := \bigcap_{m > 0} \mathcal{F}_m^{\mathbb{R}}. \tag{C.2}$$

Unlike in the setting of §2.2, we do *not* know that the tail σ -algebra $\mathcal{F}_\infty^{\mathbb{R}}$ is trivial since functions in $\mathcal{F}_\infty^{\mathbb{R}}$ are only invariant under translation along the line \mathbb{R} .

LEMMA C.2. *Let*

$$F = F(\mathcal{M}, \mathcal{M}^*, \mathcal{D})$$

be a measurable function on the space of embedded lattice/dual lattice/dyadic system pairs which is scale invariant (i.e., $F(C\mathcal{M}, C\mathcal{M}^, C\mathcal{D}) = F(\mathcal{M}, \mathcal{M}^*, \mathcal{D})$ for each $C > 0$). If either $\mathbb{E}[|F|] < \infty$ or $F \geq 0$ a.s., then a.s.*

$$\lim_{k \rightarrow \infty} \frac{1}{|S_k|} \int_{S_k \cap \mathbb{R}} F(\mathcal{M} - x, \mathcal{M}^* - x, \mathcal{D} - x) dx = \mathbb{E}[F | \mathcal{F}_\infty^{\mathbb{R}}]. \tag{C.3}$$

Proof. The analog of Lemma 2.3 in our setting implies that if x is sampled uniformly from $\widehat{S}_m \cap \mathbb{R}$, then $(\mathcal{M} - x, \mathcal{M}^* - x)$ agrees in law with $(\mathcal{M}, \mathcal{M}^*)$ modulo scaling. The lemma therefore follows from the backward martingale convergence theorem via exactly the same proof as Lemma 2.7. □

We will now prove a re-formulation of our finite expectation hypothesis (1.12) which will be more convenient for our purposes. For $e \in \mathcal{EM}$, define the rectangle

$$R_e := \left[\min_{z \in e^*} \operatorname{Re} z, \max_{z \in e^*} \operatorname{Re} z \right] \times \left[\min_{z \in e} \operatorname{Im} z, \max_{z \in e} \operatorname{Im} z \right]. \tag{C.4}$$

That is, the width (resp. height) of R_e is the same as the width (resp. height) of e^* (resp. e). See Figure 7, left. Note that e and e^* are not necessarily contained in R_e , but that R_e intersects \mathbb{R} if and only if e intersects \mathbb{R} . For $z \in \mathbb{C}$, let

$$\mathcal{E}_z := \{e \in \mathcal{EM} : R_e \ni z\}. \tag{C.5}$$

LEMMA C.3. *In the notation introduced just above, we have*

$$\mathbb{E} \left[\sum_{e \in \mathcal{E}_0} \frac{\text{diam}(e) \text{diam}(e^*)}{\text{Area}(R_e)} \right] < \infty.$$

In particular, $\mathbb{E}[\#\mathcal{E}_0] < \infty$.

Proof. We will use the mass transport principle (condition (4) of Definition 1.2 with $(\mathcal{M}, \mathcal{M}^*)$ in place of \mathcal{M}) and the Cauchy–Schwarz inequality to deduce the lemma from the finite expectation hypothesis (1.12). For $w_0, w_1 \in \mathbb{C}$, define

$$F(\mathcal{M}, \mathcal{M}^*, w_0, w_1) := \sum_{e \in \mathcal{E} \mathcal{M} \cap \partial H_{w_0}} \frac{\text{diam}(e^*)^2 \mathbf{c}^*(e^*)}{\text{Area}(H_{w_0}) \text{Area}(R_e)} \mathbb{1}_{(w_1 \in R_e)}.$$

Then F is covariant with respect to scaling and translation in the sense of (1.2). We compute

$$\int_{\mathbb{C}} F(\mathcal{M}, \mathcal{M}^*, 0, z) dz = \sum_{e \in \mathcal{E} \mathcal{M} \cap \partial H_0} \frac{\text{diam}(e^*)^2 \mathbf{c}^*(e^*)}{\text{Area}(H_0)}$$

and

$$\int_{\mathbb{C}} F(\mathcal{M}, \mathcal{M}^*, z, 0) dz = 2 \sum_{e \in \mathcal{E}_0} \frac{\text{diam}(e^*)^2 \mathbf{c}^*(e^*)}{\text{Area}(R_e)},$$

where here we use that each edge $e \in \mathcal{E}_0$ is contained in exactly two faces $H \in \mathcal{F} \mathcal{M}$. By the mass transport principle and the hypothesis (1.12), it therefore follows that

$$2\mathbb{E} \left[\sum_{e \in \mathcal{E}_0} \frac{\text{diam}(e^*)^2 \mathbf{c}^*(e^*)}{\text{Area}(R_e)} \right] = \mathbb{E} \left[\sum_{e \in \mathcal{E} \mathcal{M} \cap \partial H_0} \frac{\text{diam}(e^*)^2 \mathbf{c}^*(e^*)}{\text{Area}(H_0)} \right] < \infty. \tag{C.6}$$

Similarly,

$$2\mathbb{E} \left[\sum_{e \in \mathcal{E}_0} \frac{\text{diam}(e)^2 \mathbf{c}(e)}{\text{Area}(R_e)} \right] = \mathbb{E} \left[\sum_{e \in \mathcal{E} \mathcal{M}^* \cap \partial H_0^*} \frac{\text{diam}(e)^2 \mathbf{c}(e)}{\text{Area}(H_0^*)} \right] < \infty. \tag{C.7}$$

Since $\mathbf{c}(e) = \mathbf{c}^*(e^*)^{-1}$, we can now apply the Cauchy–Schwarz inequality to get

$$\mathbb{E} \left[\sum_{e \in \mathcal{E}_0} \frac{\text{diam}(e) \text{diam}(e^*)}{\text{Area}(R_e)} \right] \leq \mathbb{E} \left[\sum_{e \in \mathcal{E}_0} \frac{\text{diam}(e^*)^2 \mathbf{c}^*(e^*)}{\text{Area}(R_e)} \right]^{1/2} \mathbb{E} \left[\sum_{e \in \mathcal{E}_0} \frac{\text{diam}(e)^2 \mathbf{c}(e)}{\text{Area}(R_e)} \right]^{1/2} < \infty.$$

This implies that $\mathbb{E}[\#\mathcal{E}_0] < \infty$, since $\text{Area}(R_e) \leq \text{diam}(e) \text{diam}(e^*)$. □




Figure 8. *Left:* The path $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ (red) and the squares $\tilde{S}_x \in \mathcal{D}$ for $x \in \mathbb{R}$ which it intersects (black). The set \mathcal{A} consists of excursions of the path away from the union of the squares \tilde{S}_x . *Right:* The squares \tilde{S}_x for $x \in \mathbb{R}$ are chosen in such a way that $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ cannot cross between the vertical lines containing the left and right boundaries of \tilde{S}_x without hitting \tilde{S}_x . This condition is needed to rule out pathological behavior of the sort shown in the figure.

Let us now define the path we will consider. See Figure 7, right, for an illustration. For each face H of \mathcal{M} with $H \cap \mathbb{R} \neq \emptyset$, let $\tilde{\mathfrak{r}}(H)$ be the smallest $x \in \mathbb{R}$ with $x \in H$. We can enumerate

$$\{H \in \mathcal{FM} : H \cap \mathbb{R} \neq \emptyset\} = \{H(i)\}_{i \in \mathbb{Z}}$$

in such a way that $\tilde{\mathfrak{r}}(H(i-1)) \leq \tilde{\mathfrak{r}}(H(i))$ for each $i \in \mathbb{Z}$ and $H(0) = H_0$ is the origin-containing face. Let $\tilde{\mathfrak{e}}(i)$ be the (a.s. unique, by translation invariance modulo scaling and since \mathcal{VM} is countable) edge containing $\tilde{\mathfrak{r}}(i)$ and let $\tilde{\mathfrak{e}}^*(i)$ be the corresponding dual edge.

Since $\tilde{\mathfrak{e}}(i)$ lies on the boundary of $H(i)$ and some face $H(i')$ with $i' \leq i-1$, the dual edge $\tilde{\mathfrak{e}}^*(i)$ joins the dual vertex corresponding to $H(i)$ and the dual vertex corresponding to $H(i')$. If we orient $\tilde{\mathfrak{e}}^*(i)$ toward the dual vertex corresponding to $H(i)$, then each of the edges $\tilde{\mathfrak{e}}^*(i)$ for $i \in \mathbb{Z}$ is oriented from a dual vertex with a lower index to a dual vertex with a higher index. Consequently, the collection of edges $\{\tilde{\mathfrak{e}}^*(i) : i \in \mathbb{Z}\}$ has no cycles, so is a tree. By planarity, this tree admits a unique bi-infinite simple oriented path which lies below every other such bi-infinite simple oriented path. (This path can be found via a depth-first search where we always start from the rightmost unexplored edge emanating from the current vertex.)

Let $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ be this bi-infinite path. For $j \in \mathbb{Z}$ let $\mathfrak{e}(j)$ be the primal edge which crosses $\mathbf{e}^*(j)$ and let $\mathfrak{r}(j)$ be the leftmost $x \in \mathbb{R}$ with $x \in \mathfrak{e}(j)$. We assume that the enumeration is chosen so that $0 \in [\mathfrak{r}(0), \mathfrak{r}(1)]$.

We emphasize that the above definition of the path $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ is determined by $(\mathcal{M}, \mathcal{M}^*)$ in a translation- and scale-invariant manner, modulo index shifts, i.e., replacing $(\mathcal{M}, \mathcal{M}^*)$ by $(C(\mathcal{M}-x), C(\mathcal{M}^*-x))$ for $C > 0$ and $x \in \mathbb{R}$ gives us the path

$$\{C(\mathfrak{e}(j)-x)\}_{j \in \mathbb{Z}},$$

modulo an index shift.

We now introduce certain special squares of \mathcal{D} which we will use to “pin down” the path $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ to \mathbb{R} .

LEMMA C.4. *Almost surely, for each $x \in \mathbb{R}$ there exists a square S of \mathcal{D} containing x with the following property: Every path of edges $e^* \in \mathcal{EM}^*$ whose corresponding dual edges e intersect \mathbb{R} and which crosses the vertical line through the left boundary of S and the vertical line through the right boundary of S must contain a dual edge which intersects S .*

Proof. Lemma C.3 shows that a.s. \mathcal{E}_0 is a finite set, and this, together with translation invariance modulo scaling (e.g., using the condition given in Lemma 2.3) shows that a.s. $\#\mathcal{E}_x < \infty$ for each $x \in \mathbb{R}$. Therefore, a.s. for each $x \in \mathbb{R}$ there exists a square S of \mathcal{D} containing x such that

$$\bigcup_{e \in \mathcal{E}_x} R_e \subset S.$$

We claim that the condition in the statement of the lemma holds for this choice of S . Indeed, the width of each rectangle R_e is the same as the width of e^* and the height of R_e is the same as the height of e , so if $e \in \mathcal{E}(\mathbb{R})$ such that e^* intersects the vertical line through x , then $x \in R_e$ and hence $R_e \subset S$. Therefore every path of edges e^* for $e \in \mathcal{E}(\mathbb{R})$ which intersects this vertical line must pass through S . \square

For $x \in \mathbb{R}$, let \mathring{S}_x be the smallest square containing x which satisfies the condition of Lemma C.4 (chosen by some arbitrary deterministic convention in the case when the square is not unique). Note that \mathring{S}_x necessarily intersects one of the edges $\mathfrak{e}^*(j)$ for $j \in \mathbb{Z}$. In fact, every sub-path of $\{\mathfrak{e}^*(j)\}_{j \in \mathbb{Z}}$ which crosses between the vertical lines containing the left and right boundaries of \mathring{S}_x must include an edge which intersects \mathring{S}_x . We also note that $\mathring{S}_x = \mathring{S}_y$ whenever $y \in \mathring{S}_x \cap \mathbb{R}$ and y is not one of the two endpoints of $\mathring{S}_x \cap \mathbb{R}$. We write

$$\mathring{S} := \{\mathring{S}_x : x \in \mathbb{R}\}. \tag{C.8}$$

We will now show that none of the squares \mathring{S}_x are macroscopic.

LEMMA C.5. *Almost surely,*

$$\lim_{k \rightarrow \infty} |S_k|^{-1} \max_{x \in S_k \cap \mathbb{R}} |\mathring{S}_x| = 0. \tag{C.9}$$

Proof. The proof is similar to that of Lemma 2.9 but since we are working with squares (so our sets cannot be “long and skinny”) no moment bound is required. The square \mathring{S}_0 a.s. has finite side length, so a.s. we can find a random $\mathcal{F}_\infty^{\mathbb{R}}$ -measurable $m_\varepsilon > 0$ such that

$$\mathbb{P}[|\mathring{S}_0| > \varepsilon |\widehat{S}_{m_\varepsilon}| | \mathcal{F}_\infty^{\mathbb{R}}] \leq \varepsilon^2.$$

Applying Lemma C.2 to the function $F = \mathbb{1}_{(|\mathring{S}_0| > \varepsilon |\widehat{S}_{m_\varepsilon}|)}$ shows that a.s.

$$\limsup_{m \rightarrow \infty} \frac{1}{|\widehat{S}_m|} \int_{\widehat{S}_m \cap \mathbb{R}} \mathbb{1}_{(|\mathring{S}_x| > \varepsilon |\widehat{S}_m|)} dx \leq \varepsilon^2, \tag{C.10}$$

where here we use that if $m \geq m_\varepsilon$, then $\widehat{S}_{m_\varepsilon}^x \subset \widehat{S}_m$ for each $x \in \widehat{S}_m$ and so

$$\{|\dot{S}_x| > \varepsilon | \widehat{S}_m|\} \Rightarrow \{|\dot{S}_x| > \varepsilon | \widehat{S}_{m_\varepsilon}^x|\}.$$

For Lebesgue-a.e. $x, y \in \widehat{S}_m \cap \mathbb{R}$, we have $y \in \dot{S}_x \Rightarrow \dot{S}_x = \dot{S}_y$. Hence if there is an $x \in \widehat{S}_m \cap \mathbb{R}$ with $\dot{S}_x \cap \mathbb{R} \subset \widehat{S}_m$ and $|\dot{S}_x| > \varepsilon | \widehat{S}_m|$, then the integral in (C.10) is at least ε . From this and the fact that a.s. each S_k is one of the \widehat{S}_m 's, we get that a.s. for large enough $k \in \mathbb{N}$,

$$\max\{|\dot{S}_x| : x \in S_k \cap \mathbb{R}, \dot{S}_x \cap \mathbb{R} \subset S_k\} \leq \varepsilon |S_k|. \tag{C.11}$$

Using (C.11) and Lemma 2.8, one obtains (C.9) via the same argument used at the end of the proof of Lemma 2.9 to deal with the squares which intersect the endpoints of the interval $S_k \cap \mathbb{R}$. \square

Lemma C.5 ensures that our path $\{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$ gets close to each point of $S_k \cap \mathbb{R}$, but we still need to rule out the possibility of large excursions between the times when it hits the \dot{S}_x 's.

Recalling the set of squares \mathring{S} from (C.8), let $\{J_n\}_{n \in \mathbb{Z}}$ be the ordered bi-infinite sequence of times $j \in \mathbb{Z}$ for which $\mathbf{e}^*(j) \in \mathring{S}$ for some $\mathring{S} \in \mathring{\mathcal{S}}$, enumerated so that $0 \in [J_0, J_1]$. Let

$$\mathcal{A} := \left\{ \bigcup_{j=J_{n-1}+1}^{J_n} \mathbf{e}^*(j) : n \in \mathbb{Z} \right\} \tag{C.12}$$

be the set of excursions of $j \mapsto \mathbf{e}^*(j)$ away from the union of the squares in \mathring{S} . For $e \in \mathcal{EM}$ with $e^* \in \{\mathbf{e}^*(j)\}_{j \in \mathbb{Z}}$, let $A(e) \in \mathcal{A}$ be the unique such excursion with $e^* \in A(e)$.

LEMMA C.6. *Almost surely,*

$$\lim_{k \rightarrow \infty} |S_k|^{-1} \max\{\text{diam}(A) : A \in \mathcal{A} \text{ and } A \cap S_k \neq \emptyset\} = 0. \tag{C.13}$$

Proof. The proof is similar to that of Lemma 2.9. By Lemma C.3, a.s.

$$\mathbb{E} \left[\sum_{e \in \mathcal{E}_0} \frac{\text{diam}(e) \text{diam}(e^*)}{\text{Area}(R_e)} \mid \mathcal{F}_\infty^{\mathbb{R}} \right] < \infty \tag{C.14}$$

and a.s. $\max_{e \in \mathcal{E}_0} \text{diam}(A(e)) < \infty$. So, for each $\varepsilon > 0$, we can a.s. find a random $m_\varepsilon > 0$ such that $\mathbb{E}[D_{m_\varepsilon}(0) \mid \mathcal{F}_\infty^{\mathbb{R}}] \leq \varepsilon^2$, where

$$D_m(x) = D_m(x, \varepsilon) := \sum_{e \in \mathcal{E}_x} \frac{\text{diam}(e) \text{diam}(e^*)}{\text{Area}(R_e)} \mathbb{1}_{(\text{diam}(A(e)) \geq \varepsilon | \widehat{S}_m|)}.$$

Applying Lemma C.2 with $F = D_{m_\varepsilon}(0)$ shows that a.s.

$$\limsup_{m \rightarrow \infty} \frac{1}{| \widehat{S}_m|} \int_{\widehat{S}_m \cap \mathbb{R}} D_m(x) dx \leq \varepsilon^2, \tag{C.15}$$

where here we use that if $m \geq m_\varepsilon$, then $\widehat{S}_{m_\varepsilon}^x \subset \widehat{S}_m$ for each $x \in \widehat{S}_m$ and so for $e \in \mathcal{E}_x$,

$$\{\text{diam}(A(e)) > \varepsilon|\widehat{S}_m|\} \Rightarrow \{\text{diam}(A(e)) > \varepsilon|\widehat{S}_{m_\varepsilon}^x|\}.$$

By breaking up the integral (C.15) as a sum of integrals over $R_e \cap \mathbb{R}$ for $e \in \mathcal{EM}$ with $R_e \cap \widehat{S}_m \cap \mathbb{R} \neq \emptyset$, and recalling that each S_k is one of the \widehat{S}_m 's, we get that a.s.

$$\limsup_{k \rightarrow \infty} \frac{1}{|S_k|} \sum_{\substack{e \in \mathcal{EM} \\ R_e \cap S_k \cap \mathbb{R} \neq \emptyset}} \frac{\text{diam}(e) \text{Len}(R_e \cap S_k \cap \mathbb{R}) \text{diam}(e^*)}{\text{Area}(R_e)} \mathbf{1}_{(\text{diam}(A(e)) > \varepsilon|S_k|)} \leq \varepsilon^2, \tag{C.16}$$

where here Len denotes 1-dimensional Lebesgue measure. If $R_e \cap \mathbb{R} \subset S_k$, then

$$\text{Len}(R_e \cap S_k \cap \mathbb{R})$$

is the width of R_e . Since the height of R_e is at most $\text{diam}(e)$, in this case the corresponding summand in (C.16) is at least $\text{diam}(e^*) \mathbf{1}_{(\text{diam}(A(e)) > \varepsilon|S_k|)}$. If $A \in \mathcal{A}$ with $A \subset S_k$, then $R_e \cap \mathbb{R} \subset S_k$ for each $e \in \mathcal{EM}$ with $A(e) = A$ (this is because the width of R_e is the same as the width of e^*). If also $\text{diam}(A) > \varepsilon|S_k|$, then the preceding discussion shows that the sum (C.16) is at least

$$\sum_{\substack{e \in \mathcal{EM} \\ e \subset A}} \text{diam}(e^*) \geq \text{diam}(A) \geq \varepsilon|S_k|.$$

Consequently, (C.16) implies that a.s., for large enough $k \in \mathbb{N}$,

$$\text{diam}(A) \leq \varepsilon|S_k| \quad \text{for all } A \in \mathcal{A} \text{ with } A \subset S_k. \tag{C.17}$$

We now conclude the proof by applying Lemma 2.8 as at the end of the proof of Lemma 2.9. □

LEMMA C.7. *Almost surely, for each fixed $\varepsilon \in (0, 1)$ and each large enough $k \in \mathbb{N}$, there is a path of edges $e^* \in \mathcal{EM}^*$ with $\text{diam}(e^*) \leq \varepsilon|S_k|$ from the left boundary of S_k to the right boundary of S_k which is contained in the $\varepsilon|S_k|$ -neighborhood of $S_k \cap \mathbb{R}$.*

For the proof of Lemma C.7, we will need a slight extension of Lemma C.6. For an excursion set $A = \bigcup_{j=J_{n-1}+1}^{J_n} \mathbf{e}^*(j) \in \mathcal{A}$, we set $A' := \bigcup_{j=J_{n-1}}^{J_n-1} \mathbf{e}^*(j)$, so that the first edge of A' , rather than the last edge, intersects a square in \mathcal{S} . The same argument as in the proof of Lemma C.6 shows that a.s.

$$\lim_{k \rightarrow \infty} \max\{\text{diam}(A') : A \in \mathcal{A}, A' \cap S_k \neq \emptyset\} = 0.$$

Proof of Lemma C.7. For $k \in \mathbb{N}$, let \mathcal{J}_k^R be the smallest $j \in \mathbb{Z}$ for which $\mathbf{e}^*(j)$ intersects the vertical line through the right boundary of S_k , and \mathcal{J}_k^L be the last time j before \mathcal{J}_k^R for which $\mathbf{e}^*(j)$ intersects the vertical line through the left boundary of S_k . Then $\mathbf{e}^*|_{[\mathcal{J}_k^L, \mathcal{J}_k^R]_{\mathbb{Z}}}$ is a path of edges of \mathcal{EM}^* between these two horizontal lines and each edge of this path except for $\mathbf{e}^*(\mathcal{J}_k^L)$ and $\mathbf{e}^*(\mathcal{J}_k^R)$ lies strictly between these two horizontal lines.

By Lemma C.5, a.s. for each large enough $k \in \mathbb{N}$, each of the squares \mathring{S}_x for $x \in S_k \cap \mathbb{R}$ is properly contained in S_k . By the definition of \mathring{S}_x (i.e., the property described in Lemma C.4) and the preceding paragraph, the path $\mathbf{e}^*|_{[\mathcal{J}_k^L, \mathcal{J}_k^R]_{\mathbb{Z}}}$ must include an edge which intersects each such square \mathring{S}_x , and no edge of $\mathbf{e}^*|_{[\mathcal{J}_k^L+1, \mathcal{J}_k^R-1]_{\mathbb{Z}}}$ can intersect a square of the form \mathring{S}_y for $y \in \mathbb{R} \setminus S_k$. From this and the definition (C.12) of \mathcal{A} , it follows that a.s. for large enough $k \in \mathbb{N}$,

$$\bigcup_{j=\mathcal{J}_k^L}^{\mathcal{J}_k^R} \mathbf{e}^*(j) \subset \mathbb{A}_k, \tag{C.18}$$

where \mathbb{A}_k is the union of all of the of sets of the form A or A' for $A \in \mathcal{A}$ which intersect \mathring{S}_x for some $x \in S_k \cap \mathbb{R}$. Note that we need to allow for sets of the form A' , since the element of \mathcal{A} containing $\mathbf{e}^*(\mathcal{J}_k^R)$ need not intersect \mathring{S}_x for any $x \in S_k \cap \mathbb{R}$.

By Lemma C.6 and the extension of Lemma C.6 mentioned just above, a.s. the maximum diameter of the sets A or A' in the union defining \mathbb{A}_k is of order $o_k(|S_k|)$ as $k \rightarrow \infty$. In particular, by combining this with (C.18) we get that a.s. the maximal diameter of the edges of $\mathbf{e}^*|_{[\mathcal{J}_k^L, \mathcal{J}_k^R]_{\mathbb{Z}}}$ is $o_k(|S_k|)$. Since each set in the union defining \mathbb{A}_k intersects \mathring{S}_x for some $x \in S_k \cap \mathbb{R}$ by definition, Lemma C.5 implies that a.s.

$$|S_k|^{-1} \max_{z \in \mathbb{A}_k} \text{dist}(z, S_k \cap \mathbb{R}) \rightarrow 0$$

as $k \rightarrow \infty$. It therefore follows from (C.18) that a.s. for large enough $k \in \mathbb{N}$, each edge of $\mathbf{e}^*|_{[\mathcal{J}_k^L, \mathcal{J}_k^R]_{\mathbb{Z}}}$ is contained in the $\varepsilon|S_k|$ -neighborhood of $S_k \cap \mathbb{R}$, as required. \square

Proof of Lemma 3.15. By Lemmas 2.3 and C.7, if z_k is sampled uniformly from S_k , then a.s. the conditional probability given $(\mathcal{M}, \mathcal{M}^*, \mathcal{D})$ that there is a path of edges $e^* \in \mathcal{EM}^*$ with $\text{diam}(e^*) \leq \varepsilon|S_k|$ from the left boundary of S_k to the right boundary of S_k which is contained in the $(\varepsilon/100)|S_k|$ -neighborhood of $S_k \cap (\mathbb{R} + z_k)$ tends to one as $k \rightarrow \infty$. A symmetric argument shows that the same is true with “top” and “bottom” in place of “left” and “right” and $\mathbb{R} + z_k$ replaced by the vertical line through z_k .

Applying this to a large, ε -dependent number N of uniformly random points sampled from S_k , we find that it is a.s. the case that for large enough $k \in \mathbb{N}$, we can find a finite collection E_k^* of edges of \mathcal{EM}^* such that each $e^* \in E_k^*$ has diameter at most ε and each

connected component of $S_k \setminus \bigcup_{e^* \in E_k^*} e^*$ has diameter at most ε (see Lemma 2.19 for a similar argument). Since no two dual edges can cross, a.s. for large enough $k \in \mathbb{N}$ the maximal diameter of the edges of \mathcal{EM}^* which are contained in S_k and lie at Euclidean distance at least ε from ∂S_k is at most ε . Applying Lemma 2.8 to deal with the edges near the boundary (via exactly the same argument as at the end of the proof of Lemma 2.9) shows that a.s. the maximal diameter of the edges of \mathcal{EM}^* which intersect S_k is of order $o_k(|S_k|)$ $k \rightarrow \infty$. A symmetric argument shows that the same is true of the edges which of \mathcal{EM} which intersect S_k . Using Lemma 2.8 again to transfer from the squares S_k to the disks $B_r(0)$ concludes the proof. \square

References

- [AB] AIZENMAN, M. & BURCHARD, A., Hölder regularity and dimension bounds for random curves. *Duke Math. J.*, 99 (1999), 419–453.
- [AL] ALDOUS, D. & LYONS, R., Processes on unimodular random networks. *Electron. J. Probab.*, 12 (2007), no. 54, 1454–1508.
- [ABDH] ANDRES, S., BARLOW, M. T., DEUSCHEL, J.-D. & HAMBLY, B. M., Invariance principle for the random conductance model. *Probab. Theory Related Fields*, 156 (2013), 535–580.
- [AFS] ANG, M., FALCONET, H. & SUN, X., Volume of metric balls in Liouville quantum gravity. *Electron. J. Probab.*, 25 (2020), Paper No. 160, 50 pp.
- [Ba] BARLOW, M. T., Diffusions on fractals, in *Lectures on Probability Theory and Statistics* (Saint-Flour, 1995), Lecture Notes in Math., 1690, pp. 1–121. Springer, Berlin–Heidelberg, 1998.
- [BD] BARLOW, M. T. & DEUSCHEL, J.-D., Invariance principle for the random conductance model with unbounded conductances. *Ann. Probab.*, 38 (2010), 234–276.
- [BF] BEN AROUS, G. & FRIBERGH, A., Biased random walks on random graphs, in *Probability and Statistical Physics in St. Petersburg*, Proc. Sympos. Pure Math., 91, pp. 99–153. Amer. Math. Soc., Providence, RI, 2016.
- [Be] BERESTYCKI, N., Diffusion in planar Liouville quantum gravity. *Ann. Inst. Henri Poincaré Probab. Stat.*, 51 (2015), 947–964.
- [BB] BERGER, N. & BISKUP, M., Quenched invariance principle for simple random walk on percolation clusters. *Probab. Theory Related Fields*, 137 (2007), 83–120.
- [Bi] BISKUP, M., Recent progress on the random conductance model. *Probab. Surv.*, 8 (2011), 294–373.
- [BDG] BISKUP, M., DING, J. & GOSWAMI, S., Return probability and recurrence for the random walk driven by two-dimensional Gaussian free field. *Comm. Math. Phys.*, 373 (2020), 45–106.
- [BP] BISKUP, M. & PRESCOTT, T. M., Functional CLT for random walk among bounded random conductances. *Electron. J. Probab.*, 12 (2007), no. 49, 1323–1348.
- [C] CURIE, N., A glimpse of the conformal structure of random planar maps. *Comm. Math. Phys.*, 333 (2015), 1417–1463.
- [DFGW] DE MASI, A., FERRARI, P. A., GOLDSTEIN, S. & WICK, W. D., An invariance principle for reversible Markov processes. Applications to random motions in random environments. *J. Statist. Phys.*, 55 (1989), 787–855.

- [DMS] DUPLANTIER, B., MILLER, J. & SHEFFIELD, S., Liouville quantum gravity as a mating of trees. *Astérisque*, 427 (2021), viii+257.
- [DS] DUPLANTIER, B. & SHEFFIELD, S., Liouville quantum gravity and KPZ. *Invent. Math.*, 185 (2011), 333–393.
- [D] DURRETT, R., *Probability: Theory and Examples*. Cambridge Series in Statistical and Probabilistic Mathematics, 31. Cambridge Univ. Press, Cambridge, 2010.
- [GRV] GARBAN, C., RHODES, R. & VARGAS, V., Liouville Brownian motion. *Ann. Probab.*, 44 (2016), 3076–3110.
- [G] GRIMMETT, G., *Percolation*. Grundlehren der mathematischen Wissenschaften, 321. Springer, Berlin–Heidelberg, 1999.
- [GGN] GUREL-GUREVICH, O. & NACHMIAS, A., Recurrence of planar graph limits. *Ann. of Math.*, 177 (2013), 761–781.
- [GMS1] GWYNNE, E., MILLER, J. & SHEFFIELD, S., Harmonic functions on mated-CRT maps. *Electron. J. Probab.*, 24 (2019), Paper No. 58, 55.
- [GMS2] — The Tutte embedding of the Poisson-Voronoi tessellation of the Brownian disk converges to $\sqrt{8/3}$ -Liouville quantum gravity. *Comm. Math. Phys.*, 374 (2020), 735–784.
- [GMS3] — The Tutte embedding of the mated-CRT map converges to Liouville quantum gravity. *Ann. Probab.*, 49 (2021), 1677–1717.
- [H] HELLAND, I. S., Central limit theorems for martingales with discrete or continuous time. *Scand. J. Statist.*, 9 (1982), 79–94.
- [HRV] HUANG, Y., RHODES, R. & VARGAS, V., Liouville quantum gravity on the unit disk. *Ann. Inst. Henri Poincaré Probab. Stat.*, 54 (2018), 1694–1730.
- [Ka] KAHANE, J.-P., Sur le chaos multiplicatif. *Ann. Sci. Math. Québec*, 9 (1985), 105–150. Comments and corrections in *Ann. Sci. Math. Québec*, 10 (1986), 117–118.
- [KV] KIPNIS, C. & VARADHAN, S. R. S., Central limit theorem for additive functionals of reversible Markov processes and applications to simple exclusions. *Comm. Math. Phys.*, 104 (1986), 1–19.
- [KLO] KOMOROWSKI, T., LANDIM, C. & OLLA, S., *Fluctuations in Markov Processes*. Grundlehren der mathematischen Wissenschaften, 345. Springer, Berlin–Heidelberg, 2012.
- [Ku] KUMAGAI, T., *Random Walks on Disordered Media and their Scaling Limits*. Lecture Notes in Mathematics, 2101. Springer, Cham, 2014.
- [L] LEE, J. R., Conformal growth rates and spectral geometry on distributional limits of graphs. *Ann. Probab.*, 49 (2021), 2671–2731.
- [LP] LYONS, R. & PERES, Y., *Probability on Trees and Networks*. Cambridge Series in Statistical and Probabilistic Mathematics, 42. Cambridge Univ. Press, New York, 2016.
- [MS1] MILLER, J. & SHEFFIELD, S., Liouville quantum gravity and the Brownian map I: the QLE(8/3, 0) metric. *Invent. Math.*, 219 (2020), 75–152.
- [MS2] — Liouville quantum gravity and the Brownian map II: Geodesics and continuity of the embedding. *Ann. Probab.*, 49 (2021), 2732–2829.
- [MS3] — Liouville quantum gravity and the Brownian map III: The conformal structure is determined. *Probab. Theory Related Fields*, 179 (2021), 1183–1211.
- [M] MURUGAN, M., Quasisymmetric uniformization and heat kernel estimates. *Trans. Amer. Math. Soc.*, 372 (2019), 4177–4209.
- [P1] POLYAKOV, A. M., Quantum geometry of bosonic strings. *Phys. Lett. B*, 103 (1981), 207–210.
- [P2] — Quantum geometry of fermionic strings. *Phys. Lett. B*, 103 (1981), 211–213.

- [PRS] PROCACCIA, E. B., ROSENTHAL, R. & SAPOZHNIKOV, A., Quenched invariance principle for simple random walk on clusters in correlated percolation models. *Probab. Theory Related Fields*, 166 (2016), 619–657.
- [R] REBOLLEDO, R., Remarques sur la convergence en loi des martingales vers des martingales continues. *C. R. Acad. Sci. Paris Sér. A-B*, 285 (1977), A517–A520.
- [RV] RHODES, R. & VARGAS, V., Gaussian multiplicative chaos and applications: a review. *Probab. Surv.*, 11 (2014), 315–392.
- [S] SHEFFIELD, S., Conformal weldings of random surfaces: SLE and the quantum gravity zipper. *Ann. Probab.*, 44 (2016), 3474–3545.
- [SS] SIDORAVICIUS, V. & SZNITMAN, A. S., Quenched invariance principles for walks on clusters of percolation or among random conductances. *Probab. Theory Related Fields*, 129 (2004), 219–244.
- [W] WILSON, D. B., Generating random spanning trees more quickly than the cover time, in *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing* (Philadelphia, PA, 1996), pp. 296–303. ACM, New York, 1996.

EWAIN GWYNNE

Department of Mathematics
University of Chicago
Eckhart Hall
5734 S University Ave
Chicago, IL 60637
U.S.A.
ewain@uchicago.edu

JASON MILLER

Statistical Laboratory
University of Cambridge
Wilberforce Road
Cambridge CB3 0WB
U.K.
jpmiller@statslab.cam.ac.uk

SCOTT SHEFFIELD

Department of Mathematics
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02478
U.S.A.
sheffield@math.mit.edu

Received September 14, 2018