



A mating-of-trees approach for graph distances in random planar maps

Ewain Gwynne¹ · Nina Holden² · Xin Sun³

Received: 5 July 2018 / Revised: 9 March 2020
© The Author(s) 2020

Abstract

We introduce a general technique for proving estimates for certain random planar maps which belong to the γ -Liouville quantum gravity (LQG) universality class for $\gamma \in (0, 2)$. The family of random planar maps we consider are those which can be encoded by a two-dimensional random walk with i.i.d. increments via a mating-of-trees bijection, and includes the uniform infinite planar triangulation (UIPT; $\gamma = \sqrt{8/3}$); and planar maps weighted by the number of different spanning trees ($\gamma = \sqrt{2}$), bipolar orientations ($\gamma = \sqrt{4/3}$), or Schnyder woods ($\gamma = 1$) that can be put on the map. Using our technique, we prove estimates for graph distances in the above family of random planar maps. In particular, we obtain non-trivial upper and lower bounds for the cardinality of a graph distance ball consistent with the Watabiki (Prog Theor Phys Suppl 114:1–17, 1993) prediction for the Hausdorff dimension of γ -LQG and we establish the existence of an exponent for certain distances in the map. The basic idea of our approach is to compare a given random planar map M to a *mated-CRT map*—a random planar map constructed from a correlated two-dimensional Brownian motion—using a strong coupling (Zaitsev in ESAIM Probab Stat 2:41–108, 1998) of the encoding walk for M and the Brownian motion used to construct the mated-CRT map. This allows us to deduce estimates for graph distances in M from the estimates for graph distances in the mated-CRT map which we proved (using continuum theory) in a previous work. In the special case when $\gamma = \sqrt{8/3}$, we instead deduce estimates for the $\sqrt{8/3}$ -mated-CRT map from known results for the UIPT. The arguments of this paper do not directly use SLE/LQG, and can be read without any knowledge of these objects.

✉ Ewain Gwynne
eg558@cam.ac.uk

¹ University of Cambridge, Cambridge, UK

² ETH Zürich, Zürich, Switzerland

³ Columbia University, New York City, USA

Mathematics Subject Classification 60D05 Geometric probability · 60G50 Sums of independent random variables; random walks

Contents

1	Introduction	· · · · ·
1.1	Overview	· · · · ·
1.2	Background and context	· · · · ·
1.3	Basic notation	· · · · ·
1.4	Main results	· · · · ·
1.5	A stronger coupling theorem	· · · · ·
1.6	Comparison of graph metric balls	· · · · ·
2	Comparing distances via strong coupling	· · · · ·
3	Combinatorial arguments	· · · · ·
3.1	Spanning-tree decorated planar maps	· · · · ·
3.2	Site percolation on the UIPT	· · · · ·
3.3	Bipolar-oriented planar maps and Schnyder woods	· · · · ·
4	Open questions	· · · · ·
	References	· · · · ·

1 Introduction

1.1 Overview

A *planar map* is a graph embedded in the plane, viewed modulo orientation-preserving homeomorphisms.¹ Random planar maps are a natural model of discrete random surfaces and are of fundamental importance in mathematical physics, probability, and combinatorics. Many interesting random planar maps converge in various topologies to γ -Liouville quantum gravity (LQG) surfaces—a family of continuum random fractal surfaces—with the parameter $\gamma \in (0, 2]$ depending on the random planar map model. Such random planar maps are said to belong to the γ -LQG universality class.

The $\sqrt{8/3}$ -LQG universality class contains uniform random planar maps, including uniform triangulations, quadrangulations, and uniform maps with unconstrained face degree. The γ -LQG universality class for $\gamma \neq \sqrt{8/3}$ contains random planar maps sampled with probability proportional to the partition function of some statistical mechanics model on the map, e.g., the uniform spanning tree ($\gamma = \sqrt{2}$), bipolar orientations ($\gamma = \sqrt{4/3}$), the Ising model ($\gamma = \sqrt{3}$), or the Schnyder wood ($\gamma = 1$). It is expected (and in some cases known) that the universality class is affected only by changing the statistical mechanics model, not by changing the microscopic features of the planar map such as constraints on face or vertex degree.

The goal of this paper is to set up a general framework for proving estimates for a certain family of random planar maps in the γ -LQG universality class for $\gamma \in (0, 2)$ and apply this framework to obtain estimates for graph distances on these maps (see Sect. 1.4 for precise statements). The family of planar maps we consider are those

¹ It is more common in the planar map literature to define a planar map to be a graph embedded in the 2-sphere, rather than in the plane. We choose to embed in the plane since we will only be working with infinite planar maps and finite planar maps with boundary, so it is convenient to have a notion of “infinity”.

which can be encoded by a two-dimensional random walk with i.i.d. increments via a so-called *mating-of-trees bijection*. It includes the uniform infinite planar triangulation (UIPT) as well as infinite-volume limits of planar maps sampled with probability proportional to the number of different spanning trees, bipolar orientations, or Schnyder woods that can be put on the map.

Using a strong coupling of the encoding walk with Brownian motion [47, 64], we will obtain a coupling of each map in our family with the γ -mated-CRT map—a specific random planar map in the γ -LQG universality class constructed from a correlated two-dimensional Brownian motion (see Sect. 1.2.1 for a precise definition). The mated-CRT map, in turn, is directly connected to SLE and LQG, so can be studied using continuum theory (we will not use this theory directly in the present paper, however).

To estimate graph distances in our family of random planar maps, we will show that in the above coupling the mated-CRT map and the other map are roughly isometric up to a polylogarithmic factor with high probability (Theorem 1.5). We will then deduce estimates for distances in the other map from the estimates for distances in the mated-CRT map which were proven in [35]. See Fig. 1 for a schematic illustration of our approach.

In particular, for each of the maps in our family, we obtain upper and lower bounds for the cardinality of a graph-metric ball consistent with the prediction of Watabiki [63] for the dimension of γ -LQG (Theorem 1.6); and establish the existence of an exponent χ which describes various distances in the planar map and which, at least for $\gamma \leq \sqrt{2}$, we expect to be the scaling exponent for the graph distance needed to get a non-trivial scaling limit in the Gromov-Hausdorff topology (Theorem 1.7). To our knowledge, our results constitute the first non-trivial bounds for distances for planar maps weighted by spanning trees, bipolar orientations, or Schnyder woods. Since distances in uniform maps are already well understood, in the case when $\gamma = \sqrt{8/3}$ we instead obtain sharper bounds for distances in the $\sqrt{8/3}$ -mated-CRT map than those in [35] (Theorem 1.8).

The results of this paper can also be used to prove other results for random planar maps. For example, our coupling theorems are used in [33, 37] to solve several open problems concerning the simple random walk on the UIPT and on the other random planar maps considered in this paper (including computing the exponents for the return probability and graph distance displacement of the walk). Our results are also used in [24] to relate the metric ball volume exponent for random planar maps to various exponents related to LQG—arising from the Liouville heat kernel, Liouville graph distance, and Liouville first passage percolation—and to obtain new bounds for these exponents. More generally, since there is a wide range of tools available for studying mated-CRT maps (due to their connection to SLE/LQG) we expect that the techniques developed in this paper could potentially be useful whenever one needs to prove estimates for random planar maps which can be encoded by mating-of-trees bijections.

In the rest of this section, we provide some background on distances in random planar maps, define mated-CRT maps, introduce some notation, and state our main results. Section 2 contains the core of our argument, where we apply the strong coupling theorem of [64] to compare graph distances for mated-CRT maps and other random




Fig. 1 Schematic illustration of how the results of this paper fit together with other works (blue indicates what is done in the present work). As explained in Sect. 1.2.1, the mated-CRT map can be constructed from a two-dimensional correlated Brownian motion, or equivalently from a γ -LQG surface decorated by a space-filling SLE_κ for $\kappa = 16/\gamma^2$ due to the results of [25]. In [35], we proved estimates for distances in mated-CRT maps using a combination of Brownian motion and SLE/LQG techniques. In the present paper, we transfer these estimates to other random planar map models which can be bijectively encoded by a two-dimensional random walk with i.i.d. increments. This is done by coupling the encoding walk with the Brownian motion used to define the mated-CRT map using [64]; and arguing that in such a coupling, the other random planar map and the mated-CRT map are roughly isometric up to a polylogarithmic factor (color figure online)

planar maps. In Sect. 3 we review the mating-of-trees bijections for several particular random planar maps, then use the results of Sect. 2 to deduce our main results (since the bijections for the various maps are defined in slightly different ways, we need to treat the maps individually). In Sect. 4, we discuss some open problems.

1.2 Background and context

1.2.1 Mating-of-trees bijections and mated-CRT maps

Many random planar map models in the γ -LQG universality class for $\gamma \in (0, 2)$ can be encoded by a two-dimensional random walk via a bijection of *mating-of-trees* (a.k.a. *peanosphere*) type. The reason for the name is that these bijections can be interpreted as gluing together two discrete trees (corresponding to the two coordinates of a two-dimensional random walk) to construct a planar map decorated by a space-filling (“peano”) curve.

The first mating-of-trees bijection is that of Mullin [58] (explained more explicitly in [15]), which encodes a random spanning tree-decorated planar map by means of a nearest-neighbor random walk in \mathbb{Z}^2 . Sheffield’s hamburger-cheeseburger bijection [62] (which is a re-formulation of a bijection in [15]) is a generalization of Mullin’s bijection which encodes a planar map decorated by the critical Fortuin Kasteleyn (FK) model [30] with parameter $q \in (0, 4)$. There are also mating-of-trees bijections for bipolar-oriented planar maps [46], uniform random triangulations decorated by site percolation [14, 16], random planar maps decorated by a Schnyder wood [51], and random planar maps decorated by an activity-weighted spanning tree [36].

For $\gamma \in (0, 2)$, the γ -*mated-CRT map* is a random planar map in the γ -LQG universality class, constructed by means of a semi-continuous variant of a mating-of-trees bijection with a correlated two-dimensional Brownian motion in place of a




Fig. 2 Left: To construct the mated-CRT map \mathcal{G} geometrically, one can draw the graph of L (red) and the graph of $C - R$ (blue) for some large constant $C > 0$ chosen so that the parts of the graphs over some time interval of interest do not intersect. One then divides the region between the graphs into vertical strips (boundaries shown in orange) and identifies each strip with the horizontal coordinate $x \in \mathbb{Z}$ of its rightmost point. Vertices $x_1, x_2 \in \mathbb{Z}$ are connected by an edge if and only if the corresponding strips either share a boundary segment (i.e., $|x_1 - x_2| = 1$) or are connected by a horizontal line segment which lies under the graph of L or above the graph of $C - R$. One such segment is shown in green in the figure for each pair of vertices for which this latter condition holds. Right: One can draw the graph \mathcal{G} in the plane by connecting two vertices $x_1, x_2 \in \mathbb{Z}$ by an edge by an arc above (resp. below) the blue line if the corresponding strips are connected by a horizontal segment above (resp. below) the graph of L (resp. $C - R$); and connecting each pair of consecutive vertices of \mathbb{Z} by an edge. This gives \mathcal{G} a planar map structure under which it is a triangulation. Note that the two pictures do not correspond to the same mated-CRT map realization (color figure online)

random walk. Suppose $\gamma \in (0, 2)$ and $Z = (L, R)$ is a two-sided, two-dimensional Brownian motion with variances and covariances

$$\text{Var}(L_t) = \text{Var}(R_t) = |t| \quad \text{and} \quad \text{Cov}(L_t, R_t) = -\cos(\pi\gamma^2/4)|t|, \quad \forall t \in \mathbb{R}. \quad (1.1)$$

Note that the correlation of L and R ranges over $(-1, 1)$ as γ ranges over $(0, 2)$.

We define the mated-CRT map \mathcal{G} to be the graph with vertex set \mathbb{Z} , with two vertices $x_1, x_2 \in \mathbb{Z}$ with $x_1 < x_2$ connected by an edge if and only if either

$$\begin{aligned} \left(\inf_{t \in [x_1-1, x_1]} L_t \right) \vee \left(\inf_{t \in [x_2-1, x_2]} L_t \right) &\leq \inf_{t \in [x_1, x_2-1]} L_t \quad \text{or} \\ \left(\inf_{t \in [x_1-1, x_1]} R_t \right) \vee \left(\inf_{t \in [x_2-1, x_2]} R_t \right) &\leq \inf_{t \in [x_1, x_2-1]} R_t. \end{aligned} \quad (1.2)$$

Note that (1.2) holds in particular if $|x_1 - x_2| = 1$. The vertices x_1 and x_2 are connected by two edges if both conditions in (1.2) hold but $|x_1 - x_2| > 1$. See Fig. 2 for an illustration of this definition and an explanation of why \mathcal{G} is a planar map (in fact, a triangulation).²

The mated-CRT map is a discretized mating of the continuum random trees (CRT's) [4–6] associated with L and R . As we will see in Sect. 3, this adjacency condition (1.2)

² The papers [35, 39] consider the mated-CRT maps \mathcal{G}^ϵ with vertex set $\epsilon\mathbb{Z}$, with each vertex x corresponding to an interval of the form $[x - \epsilon, x]$. In this notation $\mathcal{G} = \mathcal{G}^1$. By Brownian scaling the law of \mathcal{G}^ϵ does not depend on ϵ .

is an exact continuum analogue of the condition for two vertices to be adjacent in terms of the encoding walk in various discrete mating-of-trees bijections.

It follows from [25, Theorem 1.9] (and is explained in more detail in [35]) that the mated-CRT map is directly connected to SLE-decorated Liouville quantum gravity. We will not directly use this relationship here, but we briefly mention it for the sake of context. Suppose h is the distribution corresponding to a γ -quantum cone (a type of γ -LQG surface) and η is an independent whole-plane space-filling SLE $_{\kappa}$ from ∞ to ∞ for $\kappa = 16/\gamma^2$, as constructed in [57, Sections 1.2.3 and 4.3]. If we parametrize $\eta : \mathbb{R} \rightarrow \mathbb{C}$ in such a way that the γ -LQG mass of $\eta([s, t])$ equals $t - s$ whenever $s < t$, then two vertices $x_1, x_2 \in \mathbb{Z}$ of \mathcal{G} are joined by an edge in \mathcal{G} , i.e., (1.2) holds, if and only if the ‘‘cells’’ $\eta([x_1 - 1, x_1])$ and $\eta([x_2 - 1, x_2])$ intersect along a non-trivial connected boundary arc (this adjacency graph of cells is sometimes called the *structure graph* associated with (h, η)). This gives an embedding of the mated-CRT map into \mathbb{C} which can be analyzed using SLE and LQG estimates.

Using a combination of the above relationship between mated-CRT maps and SLE/LQG and Brownian motion techniques, the paper [35] proves a number of estimates for graph distances in the mated-CRT map \mathcal{G} , valid for all $\gamma \in (0, 2)$. This paper will transfer these estimates to other random planar maps. We also mention the recent paper [39], which proves that the mated-CRT map under the so-called *Tutte embedding* converges to γ -LQG by showing that it is close to the a priori embedding (coming from SLE-decorated LQG) discussed above. We will not need the results of this latter paper here.

1.2.2 Graph distances in random planar maps

To put our results in context, we provide here some background about graph distances in random planar maps.

Uniformly random planar maps (i.e., those in the $\sqrt{8/3}$ -LQG universality class) can be analyzed by means of the Schaeffer bijection [61] and variants thereof (e.g., [13]), which, unlike mating-of-trees bijections, describe graph distances in a simple way. In particular, for a number of different types of uniformly random planar maps (including k -angulations for $k = 3$ or $k \geq 4$ even) it is known that if we sample a map of size n and re-scale graph distances by $n^{-1/4}$ then the re-scaled metric converges in law with respect to the Gromov-Hausdorff topology to the *Brownian map* [49, 52], a continuum metric measure space constructed via a continuum analogue of the Schaeffer bijection. It is shown in [54–56] that the Brownian map is equivalent, as a metric measure space, to a certain type of $\sqrt{8/3}$ -LQG surface. Hence the metric space structure of uniformly random planar maps is well understood.

For random planar maps in the γ -LQG universality class for $\gamma \neq \sqrt{8/3}$, however, the graph distance is not at all well understood. In this case there is no known bijective encoding analogous to the Schaeffer bijection which describes distances in a simple way. However, a candidate for the Gromov-Hausdorff limit, the so-called γ -LQG metric has recently been constructed in [38].

In fact, even the basic exponents for distances in random planar maps outside the $\sqrt{8/3}$ -LQG universality class are unknown. Such exponents include the exponent for the diameter of a random planar map of size n or the exponent for the cardinality of a

metric ball of radius n in an infinite-volume random planar map (which is expected to be its reciprocal). This latter exponent should be the same as the Hausdorff dimension d_γ of the γ -LQG metric. It was predicted by Watabiki in [63] that

$$d_\gamma = 1 + \frac{\gamma^2}{4} + \frac{1}{4}\sqrt{(4 + \gamma^2)^2 + 16\gamma^2}. \quad (1.3)$$

The prediction (1.3) appears to match closely with numerical simulations [1] and there are several rigorous upper and lower bounds for quantities associated with γ -LQG which are consistent with this prediction [3, 9, 24, 28, 35, 41, 53].

However, (1.3) is known to be false for small values of γ . Indeed, Ding and Goswami [23] proved an estimate for various discrete approximations of the LQG metric which shows that $2/d_\gamma \leq 1 - c \frac{\gamma^{4/3}}{\log \gamma^{-1}}$ for small enough γ , whereas (1.3) gives $2/d_\gamma = 1 - O_\gamma(\gamma^2)$. It follows from the results of [24, 28] that one has an analogous estimate for the γ -mated CRT map when γ is small, which then extends to other random planar maps (in particular, the biased bipolar-oriented maps in 4 below) using the results of the present paper. This then contradicts (1.3) for random planar maps. It is not clear whether (1.3) is true in the regime when γ is not close to 0. See [23, 24, 35] for further discussion of Watabiki's prediction and the problem of computing d_γ .

1.3 Basic notation

We write \mathbb{N} for the set of positive integers and $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$.

For $a, b \in \mathbb{R}$ with $a < b$, we define the discrete intervals $[a, b]_{\mathbb{Z}} := [a, b] \cap \mathbb{Z}$ and $(a, b)_{\mathbb{Z}} := (a, b) \cap \mathbb{Z}$.

If a and b are two quantities depending on a variable x , we write $a = O_x(b)$ (resp. $a = o_x(b)$) if a/b remains bounded (resp. tends to 0) as $x \rightarrow 0$ or as $x \rightarrow \infty$ (the regime we are considering will be clear from the context). We write $a = o_x^\infty(b)$ if $a = o_x(b^s)$ for every $s \in \mathbb{R}$.

1.3.1 Notation for graphs

For a planar graph G , we write $\mathcal{V}(G)$, $\mathcal{E}(G)$, and $\mathcal{F}(G)$, respectively, for the set of vertices, edges, and faces of G , respectively. We write $v_1 \sim v_2$ in G if $v_1, v_2 \in \mathcal{V}(G)$ is connected by an edge in G . For $v \in \mathcal{V}(G)$, we write $\deg(v; G)$ for the degree of v (i.e., the number of edges with v as an endpoint).

A *path* in G is a function $P : [0, n]_{\mathbb{Z}} \rightarrow \mathcal{V}(G)$ for some $n \in \mathbb{N}$ such that $P(i) \sim P(i + 1)$ for $i = 0, \dots, n - 1$. The *length* $|P|$ of P is the integer n .

For sets A, B consisting of vertices and/or edges of G , we write $\text{dist}(A, B; G)$ for the graph distance from A to B in G , i.e., the length of the shortest path from a vertex in A or the endpoint of an edge in A to a vertex or an endpoint of an edge in B . If $A = \{x\}$ and/or $B = \{y\}$ is a singleton, we drop the set brackets.

For a set A consisting of vertices and/or edges of G and $r > 0$, we define the *graph metric ball* $B_r(A; G)$ to be the subgraph of G consisting of the set of vertices of G

which lie at graph distance at most r from A and the set of edges of G which join two such vertices (the set brackets are omitted if $A = \{x\}$ is a singleton).

We write $\text{diam}(A; G) := \max_{x, y \in \mathcal{V}(A)} \text{dist}(x, y; G)$ for the graph-distance diameter of A with respect to the graph metric on G . We abbreviate $\text{diam}(G) := \text{diam}(G; G)$.

1.4 Main results

We will consider the following infinite random planar maps M with a distinguished oriented root edge e_0 , each decorated by statistical mechanics models, which are the local limits of corresponding finite random planar maps in the Benjamini–Schramm topology [20].

1. $\gamma = \sqrt{2}$ and (M, e_0, T) is the uniform infinite spanning-tree decorated map, which can be constructed via the infinite-volume version of Mullin’s bijection [15, 21, 58, 62].
2. $\gamma = \sqrt{8/3}$ and (M, e_0, θ) is the uniform infinite planar triangulation of type II [10] (i.e., it allows multiple edges but not self-loops) decorated by a critical ($p = 1/2$ [8]) site percolation configuration.
3. $\gamma = \sqrt{4/3}$ and (M, e_0, \mathcal{O}) is the uniform infinite bipolar-oriented map, constructed via the infinite-volume version of the bijection in [46].³
4. $\gamma < \sqrt{2}$ and (M, e_0, \mathcal{O}) is an infinite bipolar-oriented map with one of the other face degree distributions considered in [46, Section 2.3] for which the face degree distribution has an exponential tail and the correlation between the coordinates of the encoding walk is $-\cos(\pi\gamma^2/4)$ (e.g., an infinite bipolar-oriented k -angulation for $k \geq 3$ —in which case $\gamma = \sqrt{4/3}$ —or one of the bipolar-oriented maps with biased face degree distributions considered in [46, Remark 1]).
5. $\gamma = 1$ and (M, e_0, T) is the uniform infinite Schnyder wood-decorated triangulation, as constructed in [51].

The reason why we consider these five cases is that in each of the above cases, the map (M, e_0) and its associated statistical mechanics model can be encoded by a two-sided two-dimensional random walk $\mathcal{Z} : \mathbb{Z} \rightarrow \mathbb{Z}^2$ with i.i.d. increments via a mating-of-trees bijection (these bijections are reviewed in Sect. 3), so we can apply the two-dimensional version of the KMT strong coupling theorem [47, 64] to couple \mathcal{Z} with the correlated Brownian motion Z of (1.1).

We emphasize that all of our results can be extended to any other random planar map model which can be encoded by a mating-of-trees bijection wherein the encoding walk can be strongly approximated by Brownian motion—the models listed above are just the most natural ones for which we presently know this to be the case. Section 2 considers a larger (but perhaps less natural) class of random planar map models for which slightly stronger versions of our results hold. See Sect. 4 for open problems and additional discussion concerning expanding the set of models for which the results of this paper apply.

³ We give a careful proof that this and other infinite-volume limits of bipolar-oriented planar maps exist in Sect. 3.3.2. See Proposition 3.16.

As in Sect. 1.2.1, in each of the above four cases we let $Z = (L, R)$ be a correlated two-dimensional Brownian motion correlation $-\cos(\pi\gamma^2/4)$, as in (1.1), and we let \mathcal{G} be the mated-CRT map constructed from Z as in (1.2).

Just below, we will state a coupling result (Theorem 1.5) which allows us to compare M and \mathcal{G} . The strong coupling theorem [64] only allows us to compare random walk and Brownian motion on a finite time interval, so to state our result we need to describe finite (approximate) subgraphs of M and \mathcal{G} which correspond to finite time intervals. We start by defining such graphs for \mathcal{G} .

Definition 1.1 For $n \in \mathbb{N}$ we write \mathcal{G}_n for the subgraph of \mathcal{G} whose vertex set is $[-n, n]_{\mathbb{Z}}$ and whose edge set consists of all of the edges of \mathcal{G} between two such vertices (\mathcal{G}_n is called $\mathcal{G}^1|_{[-n, n]}$ in [35]). We also write $\partial\mathcal{G}_n$ for the subgraph of \mathcal{G}_n consisting of all vertices of \mathcal{G}_n which are connected by an edge to a vertex of $\mathcal{G} \setminus \mathcal{G}_n$; and all edges of \mathcal{G}_n which join two such vertices.

We now discuss the analogue of Definition 1.1 for the map M . In each of the above situations, the corresponding bijection gives for each $n \in \mathbb{N}$ a planar map M_n associated with the random walk increment $\mathcal{Z}|_{[-n, n]_{\mathbb{Z}}}$. The maps M_n are defined slightly different in each of the four cases and will be defined precisely in Sect. 3.

The map M_n is not necessarily a subgraph of M (see, however, Remark 1.2), but there is a distinguished boundary⁴ ∂M_n and an “almost” inclusion function

$$\iota_n : M_n \rightarrow M \quad \text{which is injective on } M_n \setminus \partial M_n, \quad (1.4)$$

i.e., a function from $\mathcal{V}(M_n) \cup \mathcal{E}(M_n)$ to $\mathcal{V}(M) \cup \mathcal{E}(M)$. The injectivity of ι_n on $M_n \setminus \partial M_n$ means that we can canonically identify $M_n \setminus \partial M_n$ with a subgraph of M . We emphasize that ι_n is not injective on all of M_n in case 2: see Remark 1.2. The map M_n possesses a canonical root vertex which is mapped to \mathbf{v} by ι_n , and which (by a slight abuse of notation) we identify with \mathbf{v} .

To set up a correspondence between the vertex sets $\mathcal{V}(M_n)$ and $\mathcal{V}(\mathcal{G}_n) = [-n, n]_{\mathbb{Z}}$, we will define in Sect. 3 for each $n \in \mathbb{N}$ functions

$$\begin{aligned} \phi_n : \mathcal{V}(M_n) &\rightarrow [-n, n]_{\mathbb{Z}} \quad \text{and} \quad \psi_n : [-n, n]_{\mathbb{Z}} \rightarrow \mathcal{V}(M_n), \quad \text{with} \\ \phi_n(\mathbf{v}) &= 0 \quad \text{and} \quad \psi_n(0) = \mathbf{v}, \end{aligned} \quad (1.5)$$

where \mathbf{v} denotes the initial endpoint of e_0 . Roughly speaking, the vertex $\psi_n(i)$ corresponds to the i th step of the walk $\mathcal{Z}|_{[-n, n]_{\mathbb{Z}}}$ in the bijective construction of (M, e_0, T) from \mathcal{Z} and ϕ_n is “close” to being the inverse of ψ_n . However, the construction of M_n from $\mathcal{Z}|_{[-n, n]_{\mathbb{Z}}}$ does not set up an exact bijection between $[-n, n]_{\mathbb{Z}}$ and the vertex set of M_n , so the functions ϕ_n and ψ_n are neither injective nor surjective. See Fig. 3 for an illustration of the above definitions.

Remark 1.2 In cases 1 and 3 through 5 above, each of the planar maps M_n can be canonically identified with a subgraph of M and one can take ι_n to be the identity

⁴ Recall that a planar map with boundary is a planar map G with a distinguished face (the *external face*), in which case the boundary ∂G is the set of vertices and edges on the boundary of this face.




Fig. 3 Illustration of the map M_n , the “almost inclusion” map $i_n : M_n \rightarrow M$, and the maps $\phi_n : V(M_n) \rightarrow V(\mathcal{G}_n)$ and $\psi_n : V(\mathcal{G}_n) \rightarrow V(M_n)$. Note that two edges of ∂M_n get identified when we apply i_n . Theorem 1.5 asserts that ϕ_n and ψ_n are rough isometries up to polylogarithmic factors

map. Moreover, there are functions $\phi : V(M) \rightarrow \mathbb{Z}$ and $\psi : \mathbb{Z} \rightarrow V(M)$ for which $\phi_n = \phi|_{V(M_n)}$ and $\psi_n = \psi|_{[-n, n]_{\mathbb{Z}}}$. However, in case 2 M_n cannot be identified with a subgraph of M since it may be the case that one or more pairs of vertices or edges of M_n get identified together to a single edge of M at a later step of the bijection; see Fig. 11.

Remark 1.3 For each of the random planar maps considered here, for $i \in \mathbb{Z}$ the translated encoding walk $j \mapsto \mathcal{Z}_{j+i} - \mathcal{Z}_i$ has the same law as \mathcal{Z} , so encodes rooted planar map with the same law as (M, e_0) , decorated by a statistical mechanics model. It is easy to see from the definitions of the bijections, as reviewed in Sect. 3, that the random planar map encoded by the translated walk is the same as M but with a different choice of root edge which we denote by $\lambda(i)$ (the statistical mechanics model may not be the same). The function $\lambda : \mathbb{Z} \rightarrow \mathcal{E}(M)$ is a bijection in cases 2 through 5 and is two-to-one in case 1. It is tempting to try to define M_n to be the subgraph of M whose edge set is $\lambda([-n, n]_{\mathbb{Z}})$ and whose vertex set consists of the endpoints of edges in $\lambda([-n, n]_{\mathbb{Z}})$. This definition does not work in general, however, since the resulting graph is not necessarily connected in cases 3 and 4 and the resulting graph may have pairs of vertices which get identified at a later step in case 2, as in Remark 1.2. However, with our definitions of M_n one can check that in each of cases 1 through 4 above,

$$\lambda^{-1}(\iota_n(\mathcal{E}(M_n) \setminus \mathcal{E}(\partial M_n))) \subset [-n, n]_{\mathbb{Z}} \quad \text{and} \quad \lambda([-n, n-1]_{\mathbb{Z}}) \subset \iota_n(\mathcal{E}(M_n)). \quad (1.6)$$

In the case 5 of the Schnyder wood-decorated map, (1.6) does not hold with our definitions since we treat such maps as a special case of bipolar-oriented maps for the sake of convenience (see Sect. 3.3.4). However, we expect that one can use the bijection from [51] to give a different definition of M_n for which (1.6) holds.

The main result which allows us to compare the graph metrics on M and \mathcal{G} is Theorem 1.5 just below, which says that these two maps are roughly isometric (up to a polylogarithmic factor), in the following sense.

Definition 1.4 Let (X_1, d_1) and (X_2, d_2) be metric spaces. For $a, b, c > 0$, a map $\phi : X_1 \rightarrow X_2$ is called a *rough isometry* with parameters (a, b, c) if

$$a^{-1}d_1(x, y) - b \leq d_2(\phi(x), \phi(y)) \leq ad_1(x, y) + b, \quad \forall x, y \in X_1 \quad (1.7)$$

and for each $z \in X_2$, there is an $x \in X_1$ such that

$$d_2(z, \phi(x)) \leq c. \quad (1.8)$$

Theorem 1.5 Suppose we are in one of the five settings listed above. For each $A > 0$, there is a constant $C = C(A) > 0$ such that for each $n \in \mathbb{N}$, there is a coupling of Z and (M, e_0, T) such that with probability $1 - O_n(n^{-A})$, the map ϕ_n of (1.5) is a rough isometry from M_n to \mathcal{G}_n (each equipped with their respective graph distances) with parameters $a = C(\log n)^4$, $b = 2$, and $c = C(\log n)^4$; and the map ψ_n is a rough isometry from \mathcal{G}_n to M_n with these same values of a, b , and c .

In cases 1 and 3 through 5 above, Theorem 1.5 will be used to deduce estimates for the graph metric on M from the estimates for the graph metric on the mated-CRT map \mathcal{G} from [35] (in case 2, these estimates for M are weaker than the ones obtained from a Schaeffer-type bijection or from peeling). We now state the results we can obtain for the graph metric on M . Our first main result concerns the cardinality of graph metric balls in M , and is a planar map analogue of [35, Theorem 1.10].

Theorem 1.6 In each of the settings listed above, the following is true. Let

$$d_- := \frac{2\gamma^2}{4 + \gamma^2 - \sqrt{16 + \gamma^4}} \quad \text{and} \quad d_+ := 2 + \frac{\gamma^2}{2} + \sqrt{2}\gamma. \quad (1.9)$$

Then for each $u > 0$, there exists $c = c(u) > 0$ such that with probability $1 - O_n(n^{-c})$, we have the following bounds for the number of vertices in graph metric balls of M centered at the root edge:

$$n^{d_- - u} \leq \#\mathcal{V}(B_n(e_0; M)) \leq n^{d_+ + u}. \quad (1.10)$$

As noted in [35], it is expected that typically $\#\mathcal{V}(B_n(e_0; M)) = n^{d_\gamma + o_n(1)}$, where d_γ is the Hausdorff dimension of the γ -LQG metric (in fact, it can be proven that this is the case by combining [24, Theorem 1.6] and [43, Corollary 1.7], both of which were established after this paper). Hence Theorem 1.6 gives upper and lower bounds for d_γ . See Fig. 4 for a graph of these bounds and a table giving their values for various special planar map models.

The papers [24, 41] build on the results of the present paper to prove new upper and lower bounds for $\#\mathcal{V}(B_n(e_0; M))$ which are sharper than those of Theorem 1.6 for most values of γ . For example, the bound in the case of spanning-tree weighted maps is improved to $3.550408 \leq d_{\sqrt{2}} \leq 3.63299$ [41, Corollary 2.5]. The proof of these new bounds relies crucially on the results of the present paper both to transfer from the mated-CRT map to other random planar maps and to obtain that the volume




Fig. 4 Left: graph of the upper and lower bounds for the cardinality of a graph metric ball from Theorem 1.6 (blue and red) and the Watabiki prediction (1.3) for the Hausdorff dimension of γ -LQG (green). This graph was produced using Mathematica, and also appears in [35]. Right: Table of the exact and approximate values of the known upper/lower bounds and the Watabiki prediction for several special planar map models (color figure online)

growth exponent for the $\sqrt{8/3}$ -mated-CRT map is 4 (Theorem 1.8), which is used in combination with the monotonicity (in γ) of various exponents to obtain the new bounds.

In [35, Theorem 1.12], it is shown that for $\gamma \in (0, 2)$, there exists an exponent $\chi = \chi(\gamma) > 0$ which can be defined as the limit

$$\chi = \lim_{n \rightarrow \infty} \frac{\log \mathbb{E} [\text{diam} (\mathcal{G}_n)]}{\log n} \quad (1.11)$$

and which satisfies

$$\frac{1}{2 + \gamma^2/2 + \sqrt{2}\gamma} \vee \left(1 - \frac{2}{\gamma^2}\right) \leq \chi \leq \frac{1}{2}. \quad (1.12)$$

Our next result shows that this same exponent χ also describes distances in the other random planar maps considered above.

Theorem 1.7 *Suppose we are in one of the settings listed above and let χ be as in (1.11) (for our given choice of γ). For $u > 0$ we have (with the notation $o_n^\infty(n)$ as in Sect. 1.3),*

$$\mathbb{P} [\text{diam} (M_n) \leq n^{\chi+u}] \geq 1 - o_n^\infty(n) \quad (1.13)$$

and

$$\mathbb{P} [\text{diam} (M_n) \geq n^{\chi-u}] \geq n^{-o_n(1)}. \quad (1.14)$$

In particular, for each $p > 0$,

$$\mathbb{E} [\text{diam} (M_n)^p] = n^{\chi p + o_n(1)}. \quad (1.15)$$

We note that (1.14) does not tell us that $\text{diam} (M_n) \geq n^{\chi-o_n(1)}$ with probability tending to 1 as $n \rightarrow \infty$, although we expect this to be the case. In the special case

of spanning-tree weighted planar maps ($\gamma = \sqrt{2}$), a much stronger version of this statement is proven in [42, Theorem 3.1].

In case 2, certain exponents for graph distances in the UIPT are already known (see, e.g., [7]), so we can deduce estimates for graph distances in the $\sqrt{8/3}$ -mated-CRT map from our coupling result Theorem 1.5. For example, we get the correct exponent for the cardinality of a metric ball.

Theorem 1.8 *Suppose $\gamma = \sqrt{8/3}$, so the correlation of L and R is $1/2$. For each $u > 0$, the graph metric ball in the mated-CRT map satisfies*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[n^{4-u} \leq \#\mathcal{V}(B_n(0; \mathcal{G})) \leq n^{4+u} \right] = 1. \quad (1.16)$$

1.5 A stronger coupling theorem

Our proof of Theorem 1.5 yields a stronger statement than just the existence of rough isometries between the maps \mathcal{G}_n and M_n , which we state here.

Theorem 1.9 *Suppose we are in one of the five settings listed at the beginning of Sect. 1.4. For each $A > 0$, there is a constant $C = C(A) > 0$ such that for each $n \in \mathbb{N}$, there is a coupling of Z and (M, e_0, T) such that with probability $1 - O_n(n^{-A})$, the following is true (with ϕ_n and ψ_n as in (1.5)).*

1. *For each $v_1, v_2 \in \mathcal{V}(M_n)$ with $v_1 \sim v_2$ in M_n , there is a path $P_{v_1, v_2}^{\mathcal{G}}$ from $\phi_n(v_1)$ to $\phi_n(v_2)$ in \mathcal{G}_n with $|P_{v_1, v_2}^{\mathcal{G}}| \leq C(\log n)^4$; and each $i \in [-n, n]_{\mathbb{Z}}$ is hit by a total of at most $C(\log n)^7$ of the paths $P_{v_1, v_2}^{\mathcal{G}}$ for $v_1, v_2 \in \mathcal{V}(M_{[-n, n]})$ with $v_1 \sim v_2$.*
2. *For each $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 \sim i_2$ in \mathcal{G}_n , there is a path P_{i_1, i_2}^M from $\psi_n(i_1)$ to $\psi_n(i_2)$ in M_n with $|P_{i_1, i_2}^M| \leq C(\log n)^4$; and each $v \in \mathcal{V}(M_n)$ is hit by a total of at most $C(\log n)^7$ of the paths P_{i_1, i_2}^M for $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 \sim i_2$.*
3. *We have $\text{dist}(\psi_n(\phi_n(v)), v; M_n) \leq C(\log n)^4$ for each $v \in \mathcal{V}(M_n)$ and $\text{dist}(\phi_n(\psi_n(i)), i; \mathcal{G}_n) \leq C(\log n)^4$ for each $i \in [-n, n]_{\mathbb{Z}}$.*

Theorem 1.9 is strictly stronger than Theorem 1.5. Indeed, Theorem 1.9 trivially implies Theorem 1.5 (we record this fact as Lemma 1.10 below for the sake of reference). However, the upper bounds on the total number of the paths $P_{v_1, v_2}^{\mathcal{G}}$ or P_{i_1, i_2}^M which hit a given vertex in Theorem 1.9 are not implied by Theorem 1.5. These bounds will be important in [37].

Lemma 1.10 *Let $A > 0$ and $n \in \mathbb{N}$. Any coupling of Z with (M, e_0, T) which satisfies the conditions of Theorem 1.9 for this choice of A and n also satisfies the conditions of Theorem 1.5 for this choice of A and n , with the same value of C .*

Proof The second condition from Definition 1.4 with $c = C(\log n)^4$ for either ϕ_n or ψ_n is immediate from condition 3 of Theorem 1.9, so we just need to check the first condition (concerning the amount by which each of ϕ_n and ψ_n distort distances).

We first argue that the conditions of Theorem 1.9 imply that

$$\text{dist}(\phi_n(v_1), \phi_n(v_2); \mathcal{G}_n) \leq C(\log n)^4 \text{dist}(v_1, v_2; M_n), \quad \forall v_1, v_2 \in \mathcal{V}(M_n). \quad (1.17)$$

Indeed, suppose we are given $v_1, v_2 \in \mathcal{V}(M_n)$ and let P be a geodesic in $\mathcal{V}(M_n)$ from v_1 to v_2 . Condition 1 from Theorem 1.9 implies that

$$\text{dist}(\phi_n(P(k-1)), \phi_n(P(k)); \mathcal{G}_n) \leq C(\log n)^4,$$

for each $k \in [1, |P|]_{\mathbb{Z}}$. Summing over all such k and applying the triangle inequality yields (1.17). Similarly, condition 2 from Theorem 1.9 implies that

$$\text{dist}(\psi_n(i_1), \psi_n(i_2); M_n) \leq C(\log n)^4 \text{dist}(i_1, i_2; \mathcal{G}_n), \quad \forall i_1, i_2 \in [-n, n]_{\mathbb{Z}}. \quad (1.18)$$

Combining (1.18) (applied with $i_1 = \phi_n(v_1)$ and $i_2 = \phi_n(v_2)$) and condition 3 from Theorem 1.9 and using the triangle inequality gives that for $v_1, v_2 \in \mathcal{V}(M_n)$,

$$\text{dist}(\phi_n(v_1), \phi_n(v_2); \mathcal{G}_n) \geq C^{-1}(\log n)^{-4} \text{dist}(v_1, v_2; M_n) - 2. \quad (1.19)$$

Similarly, combining (1.17) and condition 3 from Theorem 1.9 gives

$$\text{dist}(\psi_n(i_1), \psi_n(i_2); M_n) \geq C^{-1}(\log n)^{-4} \text{dist}(i_1, i_2; \mathcal{G}_n) - 2. \quad (1.20)$$

Combining (1.17), (1.18), (1.19), and (1.20) gives the first condition from Definition 1.4 with $a = C(\log n)^4$ and $b = 2$. \square

1.6 Comparison of graph metric balls

In order to use Theorem 1.5 (or Theorem 1.9) to compare graph metric balls in M and \mathcal{G} , we need to make sure that such balls are contained in M_n and \mathcal{G}_n , respectively, with high probability. This is the purpose of the present subsection. In particular, we will establish the following lemma, which is sufficient for our purposes.

Lemma 1.11 *Suppose we are in one of the five settings listed at the beginning of Sect. 1.4. For each $A > 0$, there exists $K = K(A) > 0$ such that with probability at least $1 - O_n(n^{-A})$, we have $B_n(0; \mathcal{G}) \subset \mathcal{G}_{n^K}$ and the map ι_{n^K} of (1.4) restricts to a graph isomorphism from $B_n(v; M_{n^K})$ to $B_n(v; M)$.*

To prove Lemma 1.11 we will need the following lemma.

Lemma 1.12 *Suppose $m < n$ are positive integers such that*

$$\max \left\{ \min_{i \in [m+1, n]_{\mathbb{Z}}} \mathcal{L}_i, \min_{i \in [-n, -m-1]_{\mathbb{Z}}} \mathcal{L}_i \right\} < \min_{i \in [-m-1, m+1]_{\mathbb{Z}}} \mathcal{L}_i \quad (1.21)$$

and the same holds with \mathcal{R} in place of \mathcal{L} . Then the map ϕ_n of (1.5) satisfies

$$\phi_n(\partial M_n) \cap [-m, m]_{\mathbb{Z}} = \emptyset. \quad (1.22)$$

In each of the five cases we consider, Lemma 1.12 is an easy consequence of the definitions in Sect. 3. We will check the lemma separately for each case in the appropriate subsection of Sect. 3.

Proof of Lemma 1.11 In light of (1.4), we only need to find $K = K(A)$ as in the statement of the lemma such that

$$\mathbb{P} [B_n(0; \mathcal{G}) \subset \mathcal{G}_{n^K} \text{ and } B_n(v; M_{n^K}) \subset M_{n^K} \setminus \partial M_{n^K}] \geq 1 - O_n(n^{-A}). \quad (1.23)$$

By [35, Corollary 3.2], there exists $K > 2$ such that with probability at least $1 - O_n(n^{-A})$,

$$B_{n^2}(0; \mathcal{G}) \subset \mathcal{G}_{n^K}. \quad (1.24)$$

We will now deduce from (1.24), Lemma 1.12, and Theorem 1.9 that after possibly increasing K , we also have $\mathbb{P} [B_n(v; M_{n^{K^2}}) \subset M_{n^{K^2}} \setminus \partial M_{n^{K^2}}] \geq 1 - O_n(n^{-A})$. Since (1.24) implies the analogous statement with K^2 in place of K , this will imply (1.23) with K^2 in place of K .

Recall that each of the two coordinates of \mathcal{Z} is a one-dimensional random walk started at zero with i.i.d. increments having an exponential tail at ∞ . Just below, we will explain using basic random walk estimates that if K is chosen to be sufficiently large, in a manner depending only on A , then with probability at least $1 - O_n(n^{-A})$,

$$\max \left\{ \min_{i \in [n^K+1, n^{K^2}]_{\mathbb{Z}}} \mathcal{L}_i, \min_{i \in [-n^{K^2}-1, -n^K]_{\mathbb{Z}}} \mathcal{L}_i \right\} < \min_{i \in [-n^K-1, n^K+1]_{\mathbb{Z}}} \mathcal{L}_i \quad (1.25)$$

and the same holds with \mathcal{R} in place of \mathcal{L} . One way to justify this is as follows. Let B be a standard linear Brownian motion. Using the reflection principle to estimate the running minima of B , one gets that if K is chosen to be sufficiently large, in a manner depending only on A , then with probability at least $1 - O_n(n^{-A})$,

$$\max \left\{ \inf_{t \in [n^K+1, n^{K^2}]_{\mathbb{Z}}} B_t, \inf_{t \in [-n^{K^2}-1, -n^K]_{\mathbb{Z}}} B_t \right\} < \inf_{t \in [-n^K-1, n^K+1]_{\mathbb{Z}}} B_t - n. \quad (1.26)$$

We then deduce (1.25) from (1.26) and the KMT coupling theorem [47] (see [50, Theorem 7.1.1] for the precise version which we use here).

By Lemma 1.12, if (1.25) holds, then

$$\phi_{n^{K^2}}(\partial M_{n^{K^2}}) \cap [-n^K, n^K]_{\mathbb{Z}} = \emptyset. \quad (1.27)$$

By Theorem 1.5, there is a $C = C(A) > 1$ such that with probability at least $1 - O_n(n^{-A})$,

$$\text{dist}\left(v, \mathbb{v}; M_{n^{K^2}}\right) \geq C^{-1}(\log n)^{-4} \text{dist}\left(\phi_{n^{K^2}}(v), 0; \mathcal{G}_{n^{K^2}}\right) - 2, \quad \forall v \in \mathcal{V}(\mathcal{M}_{n^{K^2}}). \quad (1.28)$$

Henceforth assume that (1.24), (1.27), and (1.28) all hold, which happens with probability at least $1 - O_n(n^{-A})$. If $v \in \partial M_{n^{K^2}}$, then by (1.27), $\phi_{n^{K^2}}(v) \notin [-n^K, n^K]_{\mathbb{Z}} = \mathcal{V}(\mathcal{G}_{n^K})$. By (1.24), this means that

$$\text{dist}\left(\phi_{n^{K^2}}(v), 0; \mathcal{G}_{n^{K^2}}\right) \geq \text{dist}\left(\phi_{n^{K^2}}(v), 0; \mathcal{G}\right) \geq n^2. \quad (1.29)$$

By (1.28), it therefore follows that for large enough n ,

$$\text{dist}\left(v, \mathbb{v}; M_{n^{K^2}}\right) \geq C^{-1}(\log n)^{-4}n^2 - 2 \geq n + 1. \quad (1.30)$$

That is, $v \notin B_n(\mathbb{v}; M_{n^{K^2}})$. This gives the lemma statement with K^2 in place of K . \square

Remark 1.13 All of the results in this and the previous two subsections remain true with M replaced by its dual map M_* . One can even use the same coupling in Theorems 1.5 and 1.9 for both M and M_* simultaneously. This is because one can formulate the bijections used in this paper in terms of M_* rather than M , and then use similar arguments to the ones in Sect. 3 to transfer from Theorem 2.1 below to estimates for M_* instead of estimates for M . Similar considerations hold if, instead of M_* , we consider, e.g., the so-called *radial quadrangulation* $\mathcal{Q} = \mathcal{Q}(M)$ whose vertex set is $\mathcal{V}(M) \cup \mathcal{V}(M_*)$ with two vertices joined by an edge if and only if they correspond to a face of M and a vertex on the boundary of this face; or the dual map \mathcal{Q}_* of \mathcal{Q} .

2 Comparing distances via strong coupling

In this section, we prove a variant of Theorem 1.9 which will be the main technical input in the proofs of our main theorems. This theorem compares the mated-CRT map \mathcal{G} to a random planar map \mathcal{H} constructed from a general two-sided random walk in \mathbb{R}^2 with i.i.d. increments. Roughly speaking, this planar map is constructed via a simpler bijection than the ones corresponding to the four cases in Sect. 1.4, where the functions ϕ_n and ψ_n of (1.5) can be taken to be restrictions of globally defined inverse bijections $\phi : \mathcal{V}(\mathcal{H}) \rightarrow \mathbb{Z}$ and $\psi : \mathbb{Z} \rightarrow \mathcal{V}(\mathcal{H})$, so the vertex set of \mathcal{H} can be identified with $\mathbb{Z} = \mathcal{V}(\mathcal{G})$. As we will explain in Sect. 3, for several particular choices of \mathcal{Z} the graph \mathcal{H} is a close approximation of one of the random planar maps considered in Sect. 1.4.

Let $\mathcal{Z} = (\mathcal{L}, \mathcal{R}) : \mathbb{Z} \rightarrow \mathbb{R}^2$ be a two-sided two-dimensional random walk normalized so that $\mathcal{Z}_0 = 0$. We assume that the increments $\mathcal{Z}_j - \mathcal{Z}_{j-1}$ of \mathcal{Z} are i.i.d. with mean-zero; that there is a constant $c > 0$ such that the increment distribution satisfies

$$\mathbb{E} \left[\exp \left((\mathcal{Z}_j - \mathcal{Z}_{j-1}) \cdot \xi \right) \right] < \infty, \quad \forall \xi \in \mathbb{R}^2 \text{ with } |\xi| \leq c; \quad (2.1)$$

and that the walk is truly two-dimensional in the sense that the correlation $\rho := \text{Corr}(\mathcal{L}_j - \mathcal{L}_{j-1}, \mathcal{R}_j - \mathcal{R}_{j-1})$ belongs to $(-1, 1)$. We note that the hypothesis (2.1) is precisely the condition needed to apply the strong coupling result [64, Theorem 1.3].

We define an infinite random planar map $\mathcal{H} = \mathcal{H}(\mathcal{Z})$ via the following discrete analogue of the formula (1.2) defining the mated-CRT map. The vertex set of \mathcal{H} is \mathbb{Z} , and for $i_1, i_2 \in \mathbb{Z}$ with $i_1 < i_2$ we declare that i_1 and i_2 are connected by an edge in \mathcal{H} if and only if either

$$\begin{aligned} \mathcal{L}_{i_1-1} \vee \mathcal{L}_{i_2} &< \min_{j \in [i_1, i_2-1] \cap \mathbb{Z}} \mathcal{L}_j \quad \text{or} \quad \mathcal{R}_{i_1-1} \vee \mathcal{R}_{i_2} \\ &< \min_{j \in [i_1, i_2-1] \cap \mathbb{Z}} \mathcal{R}_j \quad \text{or} \quad i_2 - i_1 = 1. \end{aligned} \quad (2.2)$$

(The arguments in this subsection are robust with respect to small modifications of (2.2); see Remark 2.5).

As in Definition 1.1, for $n \in \mathbb{N}$ we define \mathcal{H}_n to be the subgraph of \mathcal{H} whose vertex set is $[-n, n] \cap \mathbb{Z}$, with two vertices connected by an edge in \mathcal{H}_n if and only if they are connected by an edge in \mathcal{H} .

Let $\gamma \in (0, 2)$ be chosen so that $\rho = -\cos(\pi\gamma^2/4)$, so that the Brownian motion Z of (1.1) for this choice of γ has correlation ρ .

Theorem 2.1 *For each $n \in \mathbb{N}$ and each $A > 0$, there is a $C > 0$ depending on A and the particular law of \mathcal{Z} and a coupling of \mathcal{Z} with the Brownian motion Z (and thereby the mated-CRT map \mathcal{G}) such that for each $n \in \mathbb{N}$, the following is true with probability at least $1 - O_n(n^{-A})$. For each $i_1, i_2 \in [-n, n] \cap \mathbb{Z}$ with $i_1 \sim i_2$ in \mathcal{H} , there is a path $P_{i_1, i_2}^{\mathcal{G}}$ from i_1 to i_2 in \mathcal{G}_n with $|P_{i_1, i_2}^{\mathcal{G}}| \leq C(\log n)^3$; and each $i \in [-n, n] \cap \mathbb{Z}$ is hit by a total of at most $C(\log n)^6$ of the paths $P_{i_1, i_2}^{\mathcal{G}}$. Moreover, the same is true with \mathcal{H} and \mathcal{G} interchanged. In particular, for each $i_1, i_2 \in [-n, n] \cap \mathbb{Z}$,*

$$C^{-1}(\log n)^{-3} \text{dist}(i_1, i_2; \mathcal{G}_n) \leq \text{dist}(i_1, i_2; \mathcal{H}_n) \leq C(\log n)^3 \text{dist}(i_1, i_2; \mathcal{G}_n). \quad (2.3)$$

The proof of Theorem 2.1 proceeds by coupling Z and \mathcal{Z} using the two-dimensional variant of the KMT coupling theorem [64, Theorem 1.3] then comparing the adjacency conditions 1.2 and (2.2). See Fig. 5 for an illustration.

The formula (2.2) is unaffected if we rescale each of the coordinates \mathcal{L} and \mathcal{R} by a (possibly different) constant, so we can assume without loss of generality that these coordinates are normalized so that $\text{Var}(\mathcal{L}_j - \mathcal{L}_{j-1}) = \text{Var}(\mathcal{R}_j - \mathcal{R}_{j-1}) = 1$.

By the multi-dimensional strong coupling theorem [64, Theorem 1.3] (the higher dimensional analogue of [47]), there are constants $b_0, b_1 > 0$, depending only on the law of the increments of \mathcal{Z} , and a coupling of \mathcal{Z} with Z such that

$$\mathbb{E} \left[\exp \left(b_0 \max_{j \in [-n, n] \cap \mathbb{Z}} |\mathcal{Z}_j - Z_j| \right) \right] \leq O_n(n^{b_1}). \quad (2.4)$$

By the Chebyshev inequality, (2.4) implies that there is a constant $C_0 > 0$, depending on A and the law of the increments of \mathcal{Z} , such that except on an event of probability $O_n(n^{-A})$ we have $\max_{j \in [-n, n] \cap \mathbb{Z}} |\mathcal{Z}_j - Z_j| \leq C_0 \log n$. Henceforth fix such a coupling




Fig. 5 Illustration of the proof of Theorem 2.1. If $i_1, i_2 \in \mathbb{Z}$ are adjacent in \mathcal{G} , then by (1.2) either we can draw a horizontal line segment (dark green) under the graph of L connecting (t_1, L_{t_1}) and (t_2, L_{t_2}) for some $t_1 \in [i_1 - 1, i_1]$ and $t_2 \in [i_2 - 1, i_2]$; or the same holds with R in place of L . By the strong coupling result [64, Theorem 1.3] we can couple Z and \mathcal{Z} such that they differ by at most $O_n(\log n)$ on $[-n, n]_{\mathbb{Z}}$ with high probability. The purple segment is obtained by translating the dark green segment upward by $O_n(\log n)$. The blue and red dots correspond to times $j \in [i_1, i_2]_{\mathbb{Z}}$ for which (j, L_j) lies below the purple line. For each such j , (j, L_j) lies within distance $O_n(\log n)$ of the dark green segment (in particular, in the light green region), so Lemma 2.3 implies that there are at most $O_n((\log n)^3)$ blue dots. By (2.2), successive blue or red dots are connected by edges of \mathcal{H} , which gives a path in \mathcal{H} from i_1 to i_2 of length at most $O_n((\log n)^3)$. Similar considerations hold with the roles of \mathcal{H} and \mathcal{G} interchanged (color figure online)

and constants $C_1, C_2 \geq C_0 \vee 1$ to be chosen later in a manner depending only on γ , A , and the law of the increments of \mathcal{Z} .

We will show that the conditions in the statement of Theorem 2.1 are satisfied on an event E^n depending on Z and \mathcal{Z} (c.f. Lemma 2.4). In particular, we let $E^n = E^n(C_0, C_1, C_2)$ be the event that the following is true.

1. $\max_{j \in [-n, n]_{\mathbb{Z}}} |\mathcal{Z}_j - Z_j| \leq C_0 \log n$.
2. For each $i \in [-n, n]_{\mathbb{Z}}$, we have $\sup_{s, t \in [i-1, i]} |Z_t - Z_s| \leq \log n$.
3. For each pair of integers (i_1, i_2) satisfying $-n \leq i_1 < i_2 \leq n$ and

$$\inf_{t \in [i_1, i_2]} (L_t - L_{i_1}) \geq -6C_1 \log n \quad \text{and} \quad |L_{i_2} - L_{i_1}| \leq 6C_1 \log n, \quad (2.5)$$

we have

$$\# \left\{ j \in [i_1, i_2]_{\mathbb{Z}} : \inf_{t \in [j-1, j]} (L_t - L_{i_1}) \leq 7C_1 \log n \right\} \leq C_1^3 (\log n)^3. \quad (2.6)$$

and the same holds with R in place of L .

4. For each $j \in [-n, n]_{\mathbb{Z}}$, the number of pairs $(i_1, i_2) \in [-n, n]_{\mathbb{Z}}^2$ with $i_1 < i_2$ for which (2.5) holds and j belongs to the set in (2.6) is at most $C_2(\log n)^6$; and the same holds with R in place of L .

The condition (2.5) says that $[i_1, i_2]$ is in some sense “close” to being an excursion interval for L , in the sense that the minimum of $L_t - L_{i_1}$ over this interval is not much less than zero and the difference $|L_{i_2} - L_{i_1}|$ is small. The condition (2.6) requires the L does not get close to zero too many times during any such interval. See also Fig. 5.

Before checking that the conditions in the theorem statement are satisfied on E^n , we show that E^n occurs with high probability.

Lemma 2.2 *If the constants $C_1, C_2 \geq C_0 \vee 1$ are chosen sufficiently large, in a manner depending only on γ , A , and the law of \mathcal{Z} , then*

$$\mathbb{P}[E^n] \geq 1 - O_n(n^{-A}). \quad (2.7)$$

For the proof of Lemma 2.2, we will need the following elementary lemma about Brownian motion.

Lemma 2.3 *Let B be a standard linear Brownian motion and for $r \geq 1$, let $T_r := \inf\{t \geq 0 : B_t \leq -r\}$. For each $c > 0$, there are constants $a_0, a_1 > 0$ depending only on c such that for $r \geq 1$ and $s > 0$,*

$$\mathbb{P}\left[\#\left\{j \in (0, T_r]_{\mathbb{Z}} : \inf_{t \in [j-1, j]} B_t \leq cr\right\} > sr^2\right] \leq a_0 e^{-a_1 s}. \quad (2.8)$$

Proof Let $\tau_0 = 0$ and for $k \in \mathbb{N}$ inductively let τ_k be the smallest $t \geq \tau_{k-1} + r^2$ for which $B_t \leq cr$. Also let K be the smallest $k \in \mathbb{N}$ for which $\tau_k \geq T_r$, i.e., $K = \inf\{k \in \mathbb{N} : \inf_{t \in [\tau_{k-1}, \tau_k]} B_t \leq -r\}$. Then for $r \geq 1$,

$$\#\left\{j \in (0, T_r]_{\mathbb{Z}} : \inf_{t \in [j-1, j]} B_t \leq r\right\} \leq K(r^2 + 2) \leq 3Kr^2.$$

By the strong Markov property and Brownian scaling, there is a $p = p(c) \in (0, 1)$ such that for each $k \in \mathbb{N}$,

$$\mathbb{P}\left[\inf_{t \in [\tau_{k-1}, \tau_k]} B_t \leq -r \mid B|_{[0, \tau_{k-1}]}\right] \geq p.$$

Hence for $s > 0$, $\mathbb{P}[K \geq s] \leq (1 - p)^{\lfloor s \rfloor}$, whence (2.8) holds. \square

Proof of Lemma 2.2 Condition 1 holds except on an event of probability $O_n(n^{-A})$ by our choice of coupling and condition 2 holds except on an event of sub-polynomial probability in n by the Gaussian tail bound, the reflection principle, and a union bound over all $i \in [-n, n]_{\mathbb{Z}}$.

By Lemma 2.3 applied with $r = 6C_1 \log n$, $c = 7/6$, and $s = \frac{1}{36}C_1 \log n$, there is a constant $a_1 = a_1(\gamma) > 0$ such that for each $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 \leq i_2$ and each choice of $C_1 > 0$, the probability of that (2.5) holds but (2.6) fails is at most $O_n(n^{-a_1 C_1})$. If we choose C_1 sufficiently large, then by a union bound over all such pairs (i_1, i_2) , we find that condition 3 holds except on an event of probability $O_n(n^{-A})$.

Now we turn our attention to condition 4, which will also be obtained using Lemma 2.3. If $(i_1, i_2) \in [-n, n]_{\mathbb{Z}}^2$ for which (2.5) holds and $j \in [i_1, i_2]_{\mathbb{Z}}$, then for some universal constant $K > 0$,

$$\inf_{t \in [i_1, j]} (L_t - L_{i_1}) \geq -KC_1 \log n \quad \text{and} \quad \inf_{t \in [j, i_2]} (L_t - L_{i_2}) \geq -KC_1 \log n. \quad (2.9)$$

If also j belongs to the set in (2.6) for this choice of (i_1, i_2) , then also $\inf_{t \in [j-1, j]} (L_t - L_{i_1}) \leq 7C_1 \log n$. If these properties are satisfied and furthermore condition 2 in the definition of E^n holds, then L_j differs from $\inf_{t \in [j-1, j]} L_t$ by at most $\log n$ so $|L_j - L_{i_1}| \vee |L_j - L_{i_2}|$ is bounded above by a universal constant times $C_1 \log n$. Hence, for a possibly larger universal constant K ,

$$\inf_{t \in [i_1, j]} (L_t - L_j) \geq -KC_1 \log n, \quad \inf_{t \in [j, i_2]} (L_t - L_j) \geq -KC_1 \log n, \\ \text{and } (L_{i_1} - L_j) \vee (L_{i_2} - L_j) \leq KC_1 \log n. \quad (2.10)$$

By Lemma 2.3, applied to each of the Brownian motions $t \mapsto L_{t+j} - L_j$ and $t \mapsto L_{-t+j} - L_j$ and with s equal to a large enough constant times $\log n$ and $r = KC_1 \log n$, there is a constant $C_2 = C_2(A) > 0$ such for each fixed $j \in \mathbb{Z}$, the probability that the number of pairs $(i_1, i_2) \in \mathbb{Z}^2$ for which (2.10) holds is larger than $C_2(\log n)^6$ is at most $O_n(n^{-A-1})$. By a union bound over all $j \in [-n, n]_{\mathbb{Z}}$, we see that the probability that condition 2 holds but condition 4 fails is at most $O_n(n^{-A})$. Combining the four preceding paragraphs shows that (2.7) holds. \square

Lemma 2.4 *Let $C_1, C_2 \geq C_0 \vee 1$ be the constants from Lemma 2.2. There is a constant $C \geq 1$, depending only on C_1, C_2 , and C_0 such that if the event E^n defined just above Lemma 2.2 occurs, then the first part of the conclusion of Theorem 2.1 is satisfied. That is, for each $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 \sim i_2$ in \mathcal{H} , there is a path $P_{i_1, i_2}^{\mathcal{G}}$ from i_1 to i_2 in \mathcal{G}_n with $|P_{i_1, i_2}^{\mathcal{G}}| \leq C(\log n)^3$; and each $i \in [-n, n]_{\mathbb{Z}}$ is contained in at most $C(\log n)^6$ of the paths $P_{i_1, i_2}^{\mathcal{G}}$. Moreover, the same is true with \mathcal{H} and \mathcal{G} interchanged.*

Proof Throughout the proof we assume that E^n occurs.

First consider a pair of vertices $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 < i_2$ and $i_1 \sim i_2$ in \mathcal{H} . We will construct a path $P_{i_1, i_2}^{\mathcal{G}}$ from i_1 to i_2 in \mathcal{G}_n with length at most $C_1^3(\log n)^3$.

By (2.2) either $|i_1 - i_2| = 1$, $\mathcal{L}_{i_1-1} \vee \mathcal{L}_{i_2} < \min_{j \in [i_1, i_2-1]_{\mathbb{Z}}} \mathcal{L}_j$, or the same holds with \mathcal{R} in place of \mathcal{L} . If $|i_1 - i_2| = 1$ we take $P_{i_1, i_2}^{\mathcal{G}}$ to be the length-1 path in \mathcal{G} from i_1 to i_2 . We will construct $P_{i_1, i_2}^{\mathcal{G}}$ in the case when $\mathcal{L}_{i_1-1} \vee \mathcal{L}_{i_2} < \min_{j \in [i_1, i_2-1]_{\mathbb{Z}}} \mathcal{L}_j$; the construction when this holds with \mathcal{R} in place of \mathcal{L} is similar. By conditions 1 and 2 in the definition of E^n and since $C_0 \leq C_1$,

$$\inf_{t \in [i_1, i_2]} (L_t - L_{i_1}) \geq -3C_1 \log n \quad \text{and} \quad |L_{i_2} - L_{i_1}| \leq 3C_1 \log n. \quad (2.11)$$

By condition 3 in the definition of E^n , the set

$$\left\{ j \in [i_1, i_2]_{\mathbb{Z}} : \inf_{t \in [j-1, j]} (L_t - L_{i_1}) \leq 3C_1 \log n \right\} \quad (2.12)$$

has cardinality at most $C_1^3(\log n)^3$. By (1.2), any two consecutive elements of the set (2.12) are connected by an edge in \mathcal{G} . Hence this set is connected in \mathcal{G} . Since the set (2.12) contains i_1 and i_2 , the vertices i_1 and i_2 can be connected by a path $P_{i_1, i_2}^{\mathcal{G}}$ contained in this set which has length at most $C_1^3(\log n)^3$.

Using condition 4 in the definition of E^n , we see that each $j \in [-n, n]_{\mathbb{Z}}$ is contained in at most $C_2(\log n)^6$ of the sets (2.12) for $i_1 < i_2$ such that $i_1 \sim i_2$ in \mathcal{H} , so each such j is hit by at most $C_2(\log n)^6$ of the paths $P_{i_1, i_2}^{\mathcal{G}}$. This gives the statement of the lemma for adjacent vertices in \mathcal{H} .

We next prove the reverse relationship. Assume $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 < i_2$ and $i_1 \sim i_2$ in \mathcal{G} . We will construct a path $P_{i_1, i_2}^{\mathcal{H}}$ in \mathcal{H}_n from i_1 to i_2 with length at most $C_1^3(\log n)^3$. The argument is similar to the one for adjacent vertices of \mathcal{H} given above.

By (1.2), the condition that $i_1 \sim i_2$ in \mathcal{G} implies that either

$$\left(\inf_{t \in [i_1-1, i_1]} L_t \right) \vee \left(\inf_{t \in [i_2-1, i_2]} L_t \right) \leq \inf_{t \in [i_1, i_2-1]} L_t \quad (2.13)$$

or the same holds with R in place of L . Assume without loss of generality that we are in the former setting. By (2.13) and condition 2 in the definition of E^n ,

$$\inf_{t \in [i_1, i_2]} (L_t - L_{i_1}) \geq -2 \log n \quad \text{and} \quad |L_{i_1} - L_{i_2}| \leq 2 \log n. \quad (2.14)$$

By this and condition 1 in the definition of E^n ,

$$\inf_{t \in [i_1, i_2]} (\mathcal{L}_t - \mathcal{L}_{i_1}) \geq -4C_1 \log n \quad \text{and} \quad |\mathcal{L}_{i_2} - \mathcal{L}_{i_1}| \leq 4C_1 \log n. \quad (2.15)$$

Moreover, by condition 2 in the definition of E^n we have $|L_j - L_{j-1}| \leq \log n$ for each $j \in [-n, n]_{\mathbb{Z}}$ so by condition 1 in the definition of E^n ,

$$\begin{aligned} & \{j \in [i_1, i_2]_{\mathbb{Z}} : \mathcal{L}_j \wedge \mathcal{L}_{j-1} - \mathcal{L}_{i_1} \leq 4C_1 \log n\} \\ & \subset \{j \in [i_1, i_2]_{\mathbb{Z}} : L_j - L_{i_1} \leq 7C_1 \log n\}. \end{aligned} \quad (2.16)$$

By (2.14) and condition 3 in the definition of E^n , the set on the right side of (2.16) has cardinality at most $C_1^3(\log n)^3$. By (2.15) the set on the left side in (2.16) contains i_1 and i_2 . If j_1 and j_2 are consecutive elements of the left set in (2.16) then either $j_2 = j_1 + 1$ or

$$\begin{aligned} \mathcal{L}_{j_1-1} & \leq 4C_1 \log n + \mathcal{L}_{i_1}, \quad \mathcal{L}_{j_2} \leq 4C_1 \log n + \mathcal{L}_{i_1}, \quad \text{and} \\ & \min_{j \in [j_1, j_2-1]_{\mathbb{Z}}} \mathcal{L}_j > 4C_1 \log n + \mathcal{L}_{i_1}, \end{aligned}$$

so by (2.2) j_1 and j_2 are connected by an edge in \mathcal{H} . Hence i_1 and i_2 can be connected by a path $P_{i_1, i_2}^{\mathcal{H}}$ in \mathcal{H} with length at most $C_1^3(\log n)^3$ which is contained in the set on the left side of (2.16).

Using condition 4 in the definition of E^n , we see that each $j \in [-n, n]_{\mathbb{Z}}$ is contained in at most $C_2(\log n)^6$ of the sets on the right side of (2.16) for $i_1 < i_2$ such that (2.14) holds, so each such j is contained in at most $C_2(\log n)^6$ of the paths $P_{i_1, i_2}^{\mathcal{H}}$.

Consequently, the statement of the lemma holds with $C = C_2 \vee C_1^3$. \square

Proof of Theorem 2.1 The first statement of the lemma (concerning the existence of paths satisfying the desired properties) is immediate from Lemmas 2.2 and 2.4. To deduce (2.3) from this statement, suppose $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ and let \widehat{P} be a \mathcal{H}_n -geodesic from i_1 to i_2 . By concatenating paths of length at most $C(\log n)^3$ in \mathcal{G}_n between the vertices of \mathcal{G}_n corresponding to the vertices traversed by \widehat{P} , we obtain a path of length at most $C(\log n)^3 |\widehat{P}|$ in \mathcal{G}_n from i_1 to i_2 , which gives the lower bound in (2.3). We similarly obtain the upper bound in (2.3). \square

Remark 2.5 The arguments in this subsection still work almost verbatim if we slightly modify the adjacency condition (2.2), e.g., by inserting ± 1 in various places. The reason for using (2.2) is that this particular adjacency condition is closely connected to the mating-of-trees bijections for spanning tree-decorated maps and for site percolation on the UIPT (see Sects. 3.1 and 3.2). In the case of bipolar-oriented and Schnyder wood-decorated maps (Sect. 3.3), we will apply Theorem 2.1 in the case when \mathcal{H} is defined with (2.2) replaced by (3.15) below.

3 Combinatorial arguments

In this section we review the bijective encodings of each of the random planar maps M listed in Sect. 1.4 by means of a certain two-dimensional random walk. We then compare M to the random planar map \mathcal{H} constructed in Sect. 2 from this same two-dimensional random walk; and deduce the strong coupling result Theorem 1.9 from this and Theorem 2.1. Our other main results will be easy consequences of Theorem 1.9 and the results of [35].

Since the comparison between M and \mathcal{H} relies on the fine geometric properties of the bijection, each of the cases needs to be treated separately. The reader may wish to read only one of the subsections of this section to get a general idea of the sort of arguments involved.

We start in Sect. 3.1 by treating the simplest case—that of the infinite spanning tree-decorated map (case 1), which is encoded by a simple random walk on \mathbb{Z}^2 via the Mullin bijection. We review this encoding, show how the map \mathcal{H} from Sect. 2 arises from the bijection (Proposition 3.3) then prove our main results in this case. In Sect. 3.2, we treat the case of site percolation on the UIPT (case 2) by relating the bijection from [14, 16] to the Mullin bijection.

In Sect. 3.3, we treat the case of the uniform infinite bipolar-oriented map (case 3) using the bijection of [46]. Along the way, we check carefully that this infinite map exists and is the local limit of the finite bipolar-oriented maps considered in [46] (unlike in the other cases, the existence of this local limit has not previously been established rigorously). We then treat the case of more general bipolar-oriented maps (case 4) and deduce the case of Schnyder-wood decorated maps (case 5) from the result for bipolar-oriented maps plus a bijection relating Schnyder wood with a special type of bipolar-oriented maps [29]. Section 3.3 can be read independently of Sects. 3.1 and 3.2.

Several places in this section we will use the following notion of *submap* of a planar map.




Fig. 6 Left: The subset of the triangulation $\mathcal{Q} \cup T \cup T_*$ consisting of the triangles in $t([a, b]_{\mathbb{Z}})$, with edges of \mathcal{Q} (resp. T , T_*) shown in black (resp. red, blue). The green curve indicates the order in which the triangles are hit by t . The edges of triangles of T which do not belong to $t([a, b]_{\mathbb{Z}})$, but which are adjacent to triangles in $t([a, b]_{\mathbb{Z}})$, are shown as dotted lines. If $a = -n$ and $b = n$, then the vertices of M_n are the red vertices and the edges of M_n are the edges of M which join these vertices (this includes the red edges in the figure plus some additional edges of M which cross blue edges in the figure). Right: The correspondence between M and \mathcal{Z} . If $i \in \mathbb{Z}$ and $t(i)$ has an edge in the red tree T , then $\mathcal{Z}_i - \mathcal{Z}_{i-1}$ is equal to $(1, 0)$ or $(-1, 0)$ according to whether the other triangle which shares this same edge of T is hit by t before or after time i . The other coordinate of \mathcal{Z} is defined symmetrically. Bottom: The steps of \mathcal{Z} corresponding to this segment of t if we assume that each of the exterior triangles (i.e., those with dotted edges) is hit by t before each of the non-dotted triangles (color figure online)

Definition 3.1 A planar map M' is a *submap* of a planar map M if, with $F_\infty \subset \mathcal{F}(M')$ denoting the set of faces of M' containing ∞ , we have

$$\mathcal{V}(M') \subset \mathcal{V}(M), \quad \mathcal{E}(M') \subset \mathcal{E}(M), \quad \mathcal{F}(M') \setminus F_\infty \subset \mathcal{F}(M).$$

The *boundary* of M' is the set of $v \in \mathcal{V}(M)$ and $e \in \mathcal{E}(M')$ which is incident to a face in F_∞ .

3.1 Spanning-tree decorated planar maps

3.1.1 Mullin bijection

The first mating-of-trees bijection to be discovered is the Mullin bijection [58], which encodes a spanning-tree decorated map by a nearest-neighbor walk in \mathbb{Z}^2 . This bijection is explained in more detail in [15], and is also equivalent to the $p = 0$ case of Sheffield's hamburger-cheeseburger bijection [62]. Here we will review the infinite-volume version of the Mullin bijection, which is also explained in [18, 21, 40] in the more general setting of the hamburger-cheeseburger bijection with arbitrary $p \in [0, 1]$. See Fig. 6 for an illustration.

Let (M, e_0, T) be the *uniform infinite spanning-tree decorated map*, which is the Benjamini–Schramm [20] local limit of uniformly random triples consisting of a planar map with an oriented root edge and a distinguished spanning tree. This infinite-volume limit is shown to exist in [21, 62].

Let M_* be the dual map of M and let T_* be the dual spanning tree of T , so that $\mathcal{E}(T_*)$ is the set of edges of M_* which do not cross edges of T . Also let $\mathcal{Q} = \mathcal{Q}(M)$ be the *radial quadrangulation*, whose vertex set is $\mathcal{V}(M) \cup \mathcal{V}(M_*)$, with two vertices of \mathcal{Q} connected by an edge if and only if they correspond to a face of M , i.e., a vertex of M_* , and a vertex of M incident to that face (the number of edges is equal to the multiplicity of the vertex as a prime end on the boundary of the face). We declare that the root edge of \mathcal{Q} is the edge $\textcircled{e}_0 \in \mathcal{E}(\mathcal{Q})$ whose primal endpoint coincides with the initial endpoint of e_0 and which is the first edge in clockwise order after e_0 with this property.

Each face of \mathcal{Q} is crossed diagonally by an edge of M and an edge of M_* , exactly one of which belongs to $T \cup T_*$. Hence the graph $\mathcal{Q} \cup T \cup T_*$ is a triangulation with the same vertex set as \mathcal{Q} . Let \mathcal{T} be the planar dual of this triangulation, so that \mathcal{T} is the adjacency graph of triangles of $\mathcal{Q} \cup T \cup T_*$, with two triangles considered adjacent if they share an edge. We declare that the root edge of \mathcal{T} is the edge of \mathcal{T} which crosses \textcircled{e}_0 , oriented so that the primal (resp. dual) endpoint of \textcircled{e}_0 is to its left (resp. right).

There is a unique path $\mathbf{t} : \mathbb{Z} \rightarrow \mathcal{V}(\mathcal{T})$ which hits each triangle (vertex) of \mathcal{T} exactly once; hits the initial and terminal points of the root edge of \mathcal{T} at times 0 and 1, respectively; and does not cross T or T_* in the sense that $\mathbf{t}(i)$ and $\mathbf{t}(i-1)$ share an edge belonging to \mathcal{Q} for each $i \in \mathbb{Z}$.

We define a walk $\mathcal{Z} = (\mathcal{L}, \mathcal{R}) : \mathbb{Z} \rightarrow \mathbb{R}^2$ with increments in $\{(0, 1), (1, 0), (-1, 0), (0, -1)\}$ as follows. Define $\mathcal{Z}_0 := 0$. Suppose $i \in \mathbb{Z}$ and the triangle $\mathbf{t}(i) \in \mathcal{V}(\mathcal{T})$ has an edge in T on its boundary (the other two boundary edges are in \mathcal{Q}). There is one other triangle of \mathcal{T} with this same edge of T on its boundary. If this triangle is hit by \mathbf{t} before (resp. after) time i , we set $\mathcal{Z}_i - \mathcal{Z}_{i-1} = (1, 0)$ (resp. $(-1, 0)$). Symmetrically, if $\mathbf{t}(i)$ has a boundary edge in T_* and the other triangle of \mathcal{T} sharing this boundary edge is hit before (resp. after) time i , we set $\mathcal{Z}_i - \mathcal{Z}_{i-1} = (0, 1)$ (resp. $(0, -1)$). Then \mathcal{Z} has the law of a standard nearest-neighbor simple random walk in \mathbb{Z}^2 .

We next state and prove a proposition which will allow us to transfer from the results of Sect. 2 to results for the infinite spanning-tree weighted map. For the statement, we will use the following definition.

Definition 3.2 Let G and G' be graphs and let $f : \mathcal{V}(G) \rightarrow \mathcal{V}(G')$. We say that f is a *graph isomorphism modulo multiplicity* if f is a bijection and two distinct vertices $v, w \in \mathcal{V}(G)$ are connected by at least one edge in G if and only if $f(v)$ and $f(w)$ are connected by at least one edge in G' (i.e., adjacency, but not necessarily the number of edges between two vertices, is preserved).

Since we are primarily interested in graph distances, there is no difference for our purposes between a true graph isomorphism and a graph isomorphism modulo multiplicity. However, some of the graphs we consider naturally have multiple edges whereas the graphs \mathcal{H} from Sect. 2 by definition do not, so we often get isomorphisms modulo multiplicity instead of true isomorphisms.

Proposition 3.3 *The function $i \mapsto t(i)$ is an isomorphism modulo multiplicity from the graph \mathcal{H} of Sect. 2 with the walk \mathcal{Z} as above and adjacency defined as in (2.2); to the graph of triangles \mathcal{T} .*

Proof For $i \in \mathbb{Z}$, the triangle $t(i)$ has two edges in \mathcal{Q} which are shared by $t(i-1)$ and $t(i+1)$, respectively, and one edge in either T or T_* . Let e_i be this third edge. For each edge e of T , there are precisely two values of $i \in \mathbb{Z}$ for which $e_i = e$, and the corresponding triangles $t(i)$ are adjacent in \mathcal{T} . The ordering of the edges of the tree T obtained from $i \mapsto e_i$ by ignoring the values of i for which $e_i \in T_*$ is the same as the contour (depth-first) ordering of T : indeed, this follows immediately from the definition of the path of triangles $i \mapsto t(i)$. The same holds with T_* in place of T . By combining this with the above definitions of \mathcal{L} and \mathcal{R} , we find that \mathcal{L} and \mathcal{R} coincide with the contour functions of the trees T and T_* [48, Section 1] except that they have extra constant steps (which do not affect the trees). From this, it follows that for $i_1 < i_2$, the condition that $e_{i_1} = e_{i_2}$ is equivalent to the adjacency condition (2.2) for \mathcal{H} . Since the three neighbors of $t(i)$ in \mathcal{T} are $t(i-1)$, $t(i+1)$, and the other triangle with e_i on its boundary, we obtain the statement of the lemma. \square

3.1.2 The inverse construction and the sewing procedure

In the above discussion, we only explained how to produce the walk \mathcal{Z} from the decorated map (M, e_0, T) . We can also go in the reverse direction and produce (M, e_0, T) from \mathcal{Z} by first constructing \mathcal{T} from \mathcal{Z} via (2.2) (see Proposition 3.3) then constructing (M, e_0, T) from \mathcal{T} . Note that this uses the fact that we can tell which edges of the dual map $\mathcal{Q} \cup T \cup T_*$ of \mathcal{T} belong to each of \mathcal{Q} , T , and T_* since edges of \mathcal{Q} correspond to edges of \mathcal{T} between consecutive vertices. We emphasize that the procedure for constructing (M, e_0, T) works for any bi-infinite walk $\mathcal{Z} = (\mathcal{L}, \mathcal{R})$ in \mathbb{Z}^2 with nearest neighbor steps such that

$$\liminf_{n \rightarrow -\infty} \mathcal{L}_n = \liminf_{n \rightarrow \infty} \mathcal{L}_n = \liminf_{n \rightarrow -\infty} \mathcal{R}_n = \liminf_{n \rightarrow \infty} \mathcal{R}_n = -\infty, \quad (3.1)$$

not just a.s. We refer to [21, Section 4.1] for more details on the inverse bijection.

The construction of (M, e_0, T) from \mathcal{Z} is uniquely characterized by an inductive relationship between certain submaps of the triangulation $\mathcal{Q} \cup T \cup T_*$, called the *sewing procedure*, which we now describe (see Fig. 7 for an illustration). For $i \in \mathbb{Z}$, let $\widehat{M}_{-\infty, i}$ be the submap of $\mathcal{Q} \cup T \cup T_*$ consisting of the vertices and edges on the boundaries of the triangles in $t((-\infty, i]_{\mathbb{Z}})$. For $i \in \mathbb{Z}$, let e_i^h (the “head”) be the edge shared by the triangles $t(i-1)$ and $t(i)$. Then $\widehat{M}_{-\infty, i}$ is an infinite triangulation with boundary, and the external face has infinitely many edges lying to the left and right of e_i^h . We can recover $(\widehat{M}_{-\infty, i+1}, e_{i+1}^h)$ from $(\widehat{M}_{-\infty, i}, e_i^h)$ by attaching a certain triangle (corresponding to $t(i+1)$) with one edge equal to e_i^h and one edge equal to e_{i+1}^h . If $\mathcal{Z}_{i+1} - \mathcal{Z}_i = (1, 0)$ (resp. $(0, 1)$), this triangle has two edges in the external face of $\widehat{M}_{-\infty, i}$ and e_{i+1}^h is the right (resp. left) one of these two edges if we face outward toward the external face. If $\mathcal{Z}_{i+1} - \mathcal{Z}_i = (-1, 0)$ (resp. $(0, -1)$), then one of the edges of the new triangle coincides with the edge of $\widehat{M}_{-\infty, i}$ immediately to the left (resp.




Fig. 7 Illustration of the inverse Mullin bijection as described in Sect. 3.1.2. The leftmost panel shows the outer boundary of the map $\widehat{M}_{-\infty, i}$, which is a subgraph of the triangulation $\mathcal{Q} \cup T \cup T_*$ (with the external face lying above the line). The next four panels show how to construct $\widehat{M}_{-\infty, i+1}$ from $\widehat{M}_{-\infty, i}$ depending on the walk increment $\mathcal{Z}_{i+1} - \mathcal{Z}_i$

right) of e_i^h , and the remaining edge lies in the external face of $\widehat{M}_{-\infty, i}$ and equals e_{i+1}^h . The edges of the tree T (resp. T_*) are precisely the edges of $\mathcal{Q} \cup T \cup T_*$ which are not crossed by \mathbf{t} and which lie to the left (resp. right) of the path \mathbf{t} . It follows from the inverse bijection described above [21, 62] that, given a walk \mathcal{Z} with i.i.d. increments there a.s. exists a unique tuple (M, e_0, T) which can be constructed via the sewing procedure just described.

3.1.3 Proofs of main theorems in the case of spanning-tree decorated maps

In this subsection we prove our main results in case 1, when $\gamma = \sqrt{2}$ and (M, e_0) is an infinite-volume spanning-tree weighted planar map. Throughout, we define the quadrangulation \mathcal{Q} , the graph of triangles \mathcal{T} , and the bijective path $\mathbf{t} : \mathbb{Z} \rightarrow \mathcal{V}(\mathcal{T})$ as in Sect. 3.1.1. We also let $Z = (L, R)$ be the Brownian motion as in (1.1) for $\gamma = \sqrt{2}$, so that L and R are independent; and we let \mathcal{G} be the associated mated-CRT map.

Our first task is to define the maps M_n and the functions ϕ_n and ψ_n from Sect. 1.4 in this setting. In keeping with Remark 1.2, we will first define functions $\phi : \mathcal{V}(M) \rightarrow \mathbb{Z}$ and $\psi : \mathbb{Z} \rightarrow \mathcal{V}(M)$. For $v \in \mathcal{V}(M)$, we choose $\phi(v) \in \mathbb{Z}$ to be the integer with the smallest absolute value for which the triangle $\mathbf{t}(\phi(v)) \in \mathcal{V}(\mathcal{T})$ has the vertex v on its boundary (in the case of a tie, we choose i to have a positive sign). For $i \in \mathbb{Z}$, we define $\psi(i) \in \mathcal{V}(M)$ to be one of the vertices of M on the boundary of the triangle $\mathbf{t}(i)$, chosen via some arbitrary deterministic convention in the case when there are two such vertices. We require $\psi(0)$ to be the initial endpoint of e_0 (i.e., the primal endpoint of e_0), in order to ensure that the last condition in (1.5) is satisfied.

For $n \in \mathbb{N}$, we define M_n the vertex set of M_n to be $\phi^{-1}([-n, n]_{\mathbb{Z}})$. Equivalently, $\mathcal{V}(M_n)$ consists of all of the vertices of M on the boundaries of the triangles in $\mathbf{t}([-n, n]_{\mathbb{Z}})$. The edge set of M_n is defined to be the set of edges e of M such that either e lies on the boundary of a triangle in $\mathbf{t}([-n, n]_{\mathbb{Z}})$; or e is contained in the union of two triangles of $\mathbf{t}([-n, n]_{\mathbb{Z}})$ (if the endpoints of e are in $\mathcal{V}(M_n)$, this latter condition is equivalent to requiring that the dual edge of e lies on the boundary of a triangle in $\mathbf{t}([-n, n]_{\mathbb{Z}})$). We define ∂M_n to be the subgraph consisting of those vertices and edges of M_n which lie on the boundary of the unbounded connected component of M_n (viewed as a subgraph of M) and we define $\iota_n : M_n \rightarrow M$ to be the inclusion map. Set

$$\phi_n := \phi|_{\mathcal{V}(M_n)} \quad \text{and} \quad \psi_n := \psi|_{[-n, n]_{\mathbb{Z}}}. \quad (3.2)$$

Note that ϕ_n maps $\mathcal{V}(M_n)$ into $[-n, n]_{\mathbb{Z}}$ by definition and ψ_n maps $[-n, n]_{\mathbb{Z}}$ into $\mathcal{V}(M_n)$ since the above definitions of ϕ and ψ show that $|\phi(\psi(i))| \leq |i|$ for each $i \in \mathbb{Z}$.

Let us now prove our main coupling theorems with this choice of ϕ_n and ψ_n . By Proposition 3.3, if we let \mathcal{Z} be the simple random walk on \mathbb{Z}^2 which encodes (M, e_0, T) via Mullin's bijection, then the map \mathcal{H} of Theorem 2.1 with this choice of \mathcal{Z} is isomorphic to T via $i \mapsto t(i)$. The basic idea of the proof is to deduce Theorem 1.9 from the analogous statements for T from Theorem 2.1 and straightforward geometric arguments. We first need the following degree bound, which is the reason why we have $(\log n)^4$ in Theorem 1.9 as compared to the $(\log n)^3$ in Theorem 2.1.

Lemma 3.4 *For $A > 0$, there exists $C = C(A) > 0$ such that for each $n \in \mathbb{N}$, it holds with probability at least $1 - O_n(n^{-A})$ that each vertex of M_n lies on the boundary of at most $C \log n$ triangles of T .*

Proof It follows from [21, Lemma 6] (see the proof of recurrence in [21, Section 4.2]) and the stationary increments property of the walk \mathcal{Z} that there are universal constants $c_0, c_1 > 0$ such that for $i \in \mathbb{Z}$ and $k \in \mathbb{N}$, the probability that the triangle $t(i)$ has a vertex of M with degree (in M) at least k on its boundary is at most $c_0 e^{-c_1 k}$. By a union bound over all $i \in [-n, n]_{\mathbb{Z}}$, we can find $C = C(A) > 0$ as in the statement of the lemma such that

$$\mathbb{P} \left[\max_{v \in \mathcal{V}(M_n)} \deg(v; M) \leq \frac{1}{2} C \log n \right] \geq 1 - O_n(n^{-A}). \quad (3.3)$$

If $v \in \mathcal{V}(M)$, then each quadrilateral of \mathcal{Q} with v on its boundary is bisected by a unique edge of M with v as an endpoint. Consequently, the number of triangles of T with v on their boundaries is at most $2 \deg(v; M)$. Combining this with (3.3) concludes the proof. \square

Proof of Theorems 1.5 and 1.9 in case 1 In light of Lemma 1.10, we only need to prove Theorem 1.9.

We first define a high-probability event on which we will show that the conditions in the theorem statement are satisfied. Let $C_0 > 0$ be the constant from Theorem 2.1 for our given choice of A and for \mathcal{Z} the simple random walk on \mathbb{Z}^2 . Recall from Proposition 3.3 that the corresponding graph \mathcal{H} from that proposition is isomorphic to the graph of triangles T via $i \mapsto t(i)$. Hence Theorem 2.1 shows that with probability at least $1 - O_n(n^{-A})$, the following is true (where here, in obvious notation, we define \mathcal{T}_n to be the graph $t(\mathcal{H}_n)$).

- (A) For each $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $t(i_1) \sim t(i_2)$ in T , there is a path $\tilde{P}_{i_1, i_2}^{\mathcal{G}}$ from i_1 to i_2 in \mathcal{G}_n with $|\tilde{P}_{i_1, i_2}^{\mathcal{G}}| \leq C_0(\log n)^3$; and each $j \in [-n, n]_{\mathbb{Z}}$ is contained in at most $C_0(\log n)^6$ of the paths $\tilde{P}_{i_1, i_2}^{\mathcal{G}}$.
- (B) For each $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 \sim i_2$ in \mathcal{G} , there is a path \tilde{P}_{i_1, i_2}^T from $t(i_1)$ to $t(i_2)$ in \mathcal{T}_n with $|\tilde{P}_{i_1, i_2}^T| \leq C_0(\log n)^3$; and each $j \in [-n, n]_{\mathbb{Z}}$ is contained in at most $C_0(\log n)^6$ of the paths \tilde{P}_{i_1, i_2}^T .

By Lemma 3.4, there exists $C_1 = C_1(A) > 0$ such that with probability at least $1 - O_n(n^{-A})$,

(C) Each vertex of M_n lies on the boundary of at most $C_1 \log n$ triangles of \mathcal{T} .

Henceforth assume that conditions (A), (B), and (C) above are all satisfied. We will check the conditions of the theorem statement in order.

Proof of condition 1 Suppose $v_1, v_2 \in \mathcal{V}(M_n)$ with $v_1 \sim v_2$. We will construct a path $P_{v_1, v_2}^{\mathcal{G}}$ in \mathcal{G}_n between $\phi(v_1)$ and $\phi(v_2)$.

By the definition of $\mathcal{E}(M_n)$, either the edge of M_n joining v_1 and v_2 lies on the boundary of a triangle of \mathcal{T}_n or this edge is contained in the union of two triangles of \mathcal{T}_n . In either case, we can find $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ such that v_1 is a vertex of $\mathbf{t}(i_1)$, v_2 is a vertex of $\mathbf{t}(i_2)$, and the triangles $\mathbf{t}(i_1)$ and $\mathbf{t}(i_2)$ are either identical or they share an edge (i.e., they are adjacent in \mathcal{T}).

By the definition of ϕ , the triangle $\mathbf{t}(\phi(v_1))$ also has v_1 as a vertex. The set of triangles in $\mathbf{t}([-n, n]_{\mathbb{Z}})$ with v_1 as a vertex is connected in \mathcal{T}_n . Consequently, there exists a path in \mathcal{T}_n from $\mathbf{t}(i_1)$ to $\mathbf{t}(\phi(v_1))$ with length at most the total number of triangles of \mathcal{T} with v_1 on their boundaries. Similar considerations hold with v_2 in place of v_1 .

By the preceding paragraph and condition (C) above, there is a $d \leq 2C_1 \log n$ and integers $j_0, \dots, j_d \in [-n, n]_{\mathbb{Z}}$ such that $j_0 = \phi(v_1)$, $j_d = \phi(v_2)$, and $\mathbf{t}(j_k) \sim \mathbf{t}(j_{k-1})$ in \mathcal{T} for each $k \in [1, d]_{\mathbb{Z}}$. For each such k , let $\tilde{P}_{j_{k-1}, j_k}^{\mathcal{G}}$ be the path from j_{k-1} to j_k afforded by condition (A) above and let $P_{v_1, v_2}^{\mathcal{G}}$ be the concatenation of all d of these paths. Then $P_{v_1, v_2}^{\mathcal{G}}$ is a path from $\phi(v_1)$ to $\phi(v_2)$ in \mathcal{G}_n with length at most $C_0 d (\log n)^3 \leq 2C_0 C_1 (\log n)^4$.

Each of the triangles $\mathbf{t}(j_k)$ for $k \in [0, d]_{\mathbb{Z}}$ has either v_1 or v_2 on its boundary. Therefore, for each $j \in [-n, n]_{\mathbb{Z}}$ the number of pairs $(v_1, v_2) \in [-n, n]_{\mathbb{Z}} \times [-n, n]_{\mathbb{Z}}$ for which $v_1 \sim v_2$ in M and j is one of the corresponding j_k 's is at most the sum of the degrees of the (at most two) vertices of M on the boundary of $\mathbf{t}(j)$, which is at most $2C_1 \log n$ by condition (C) above.

Consequently, each of the paths $\tilde{P}_{j, j'}^{\mathcal{G}}$ for $j, j' \in [-n, n]_{\mathbb{Z}}$ with $\mathbf{t}(j) \sim \mathbf{t}(j')$ in \mathcal{T}_n is a sub-path of at most $2C_0 \log n$ of the paths $P_{v_1, v_2}^{\mathcal{G}}$ for $v_1, v_2 \in \mathcal{V}(M)$. Since the total number of the $\tilde{P}_{j, j'}^{\mathcal{G}}$'s which hit each specified vertex of \mathcal{G}_n is at most $C_0 (\log n)^6$, the total number of the $P_{v_1, v_2}^{\mathcal{G}}$'s which hit each such vertex is at most $2C_0 C_1 (\log n)^7$. This gives condition 1 in the theorem statement with $C = 2C_0 C_1$. \square

Proof of condition 2 Suppose $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 \sim i_2$ in \mathcal{G} . Let $\tilde{P} = \tilde{P}_{i_1, i_2}^{\mathcal{T}}$ be the path from $\mathbf{t}(i_1)$ to $\mathbf{t}(i_2)$ in $\mathbf{t}([-n, n]_{\mathbb{Z}})$ as in condition (B) above. We will construct a path P_{i_1, i_2}^M in M_n between the vertices $\psi(i_1)$ and $\psi(i_2)$ of M .

Set $v_0 = \psi(i_1)$, let $v_{|\tilde{P}|} = \psi(i_2)$, and for $k \in [1, |\tilde{P}| - 1]_{\mathbb{Z}}$ let v_k be a vertex of the triangle $\tilde{P}(k)$ which belongs to $\mathcal{V}(M_n)$. For $k \in [1, |\tilde{P}|]_{\mathbb{Z}}$, the triangles $\tilde{P}(k-1)$ and $\tilde{P}(k)$ share an edge. This implies that the vertices v_{k-1} and v_k can be connected by a path of length at most 2 in M_n which hits only vertices of $\tilde{P}(k-1)$ or $\tilde{P}(k)$ (see Fig. 8). If we let P_{i_1, i_2}^M be the concatenation of these $|\tilde{P}|$ paths, then P_{i_1, i_2}^M is a path from v_0 to $v_{|\tilde{P}|}$ in M_n with length at most $2|\tilde{P}| \leq 2C_0 (\log n)^3$.




Fig. 8 Illustration of the seven different situations that can arise when two triangles of \mathcal{T} share a vertex. Vertices and edges of M (resp. M_*) are shown in red (resp. blue). Edges of \mathcal{Q} are shown in black. The dotted red edge is an edge of M which does not belong to the tree T . Note that in each case, there is a path of length at most 2 in M from any given red vertex of the first triangle to any given red vertex of the second triangle. Note that if the two triangles are in \mathcal{T}_n , then by definition the red vertices and edges in the figure are in M_n (color figure online)




Fig. 9 Illustration of the proof of Lemma 1.12 in case 1. The triangles in the path P in \mathcal{T} which separates $t([-m, m]_Z)$ from $\mathcal{T} \setminus t([-n, n]_Z)$ are shown in light green. The edges on the boundary of the union of the triangles in each of $t([-m, m]_Z)$, $t([-n-1, -m]_Z)$, and $t([m+1, n]_Z)$ are shown in red, blue, or black according to whether they below to T , T_* , or Q . Other edges are not shown (color figure online)

It remains to bound the number of paths P_{i_1, i_2}^M for $i_1, i_2 \in [-n, n]_Z$ with $i_1 \sim i_2$ in \mathcal{G} which hit a fixed vertex v of M_n . Indeed, if P_{i_1, i_2}^M hits v then the above construction shows that $\tilde{P}_{i_1, i_2}^{\mathcal{T}}$ must hit one of the triangles of \mathcal{T} which has v as a vertex. By condition (C), the number of such triangles is at most $2C_1 \log n$. By condition (B) above, the total number of the paths P_{i_1, i_2}^M which hit each such triangle is at most $C_0(\log n)^6$. Therefore, v is hit by at most $2C_0C_1(\log n)^7$ of the P_{i_1, i_2}^M 's, which gives condition 2 in the theorem statement with $C = 2C_0C_1$. \square

Proof of condition 3 Since the vertices v and $\psi(\phi(v))$ lie on the boundary of the same triangle of \mathcal{T} by definition, we obviously have $\text{dist}(\psi(\phi(v)), v; M_n) \leq 1$. On the other hand, the triangles $t(i)$ and $t(\phi(\psi(i)))$ share a common vertex of M , so by condition (C) the \mathcal{T}_n -graph distance between these triangles is at most $C_1 \log n$. By considering a path in \mathcal{T}_n between these two triangles and applying condition (A) to the successive pairs of triangles which it hits, we get $\text{dist}(\phi(\psi(i)), i; \mathcal{G}_n) \leq C_0C_1(\log n)^4$. \square

Proof of Lemma 1.12 in case 1 See Fig. 9 for an illustration of the proof. Recall that the adjacency graph of triangles \mathcal{T} is isomorphic to the graph \mathcal{H} of (2.2) via $i \mapsto t(i)$. Furthermore, the condition on the minimum of \mathcal{L} (resp. \mathcal{R}) in (2.2) holds if and only if $t(i_1)$ and $t(i_2)$ share an edge of T (resp. T_*). From this, we infer that if the random walk condition (1.21) and its analog for \mathcal{R} hold, then there exists $i_+, j_+ \in [m+1, n]_Z$ and

$i_-, j_- \in [-n, -m-1]_{\mathbb{Z}}$ such that $t(i_+), t(i_-)$ share an edge of T and $t(j_+), t(j_-)$ share an edge of T_* . Since consecutive triangles of \mathcal{T} share a side, each of $t([m+1, n]_{\mathbb{Z}})$ and $t([-n, -m-1]_{\mathbb{Z}})$ is connected in \mathcal{T} . Hence we can find a path of triangles P_+ in $t([m+1, n]_{\mathbb{Z}})$ from $t(i_+)$ to $t(j_+)$ and a path of triangles P_- in $t([-n, -m-1]_{\mathbb{Z}})$ from $t(j_-)$ to $t(i_-)$. The concatenation P of the paths P_+ and P_- is a cycle in $t([-n, n]_{\mathbb{Z}})$.

We claim that P necessarily disconnects $t([-m, m]_{\mathbb{Z}})$ from $\mathcal{T} \setminus t([-n, n]_{\mathbb{Z}})$. Indeed, consider the set of edges of triangles in \mathcal{T} which lie on the boundary of both a triangle of $t([m+1, n]_{\mathbb{Z}})$ and a triangle of $t([-n, -m-1]_{\mathbb{Z}})$. This set has at most two connected components, one of which is a subset of T and the other of which is a subset of T_* . By definition, P crosses an edge of each of these two connected components (in particular, they are both non-empty). It follows that the union of the triangles in $t([-n, n]_{\mathbb{Z}}) \setminus t([-m, m]_{\mathbb{Z}})$ has the topology of a Euclidean annulus whose inner and outer boundaries are disconnected by P . Our claim thus follows.

The vertex set $\mathcal{V}(M_m)$ consists of all vertices on the boundaries of triangles in $t([-m, m]_{\mathbb{Z}})$, so the preceding paragraph implies that no vertex of M_m can lie on the boundary of a triangle which is not in $t([-n, n]_{\mathbb{Z}})$ (otherwise, such a triangle would have to cross P). If $v \in \partial M_n$ then v lies on the boundary of a triangle which is not in $t([-n, n]_{\mathbb{Z}})$, so $\phi(v)$ does not lie on the boundary of a triangle in $t([-m, m]_{\mathbb{Z}})$. Since $\phi(v) \in \mathbb{Z}$ is chosen so that v lies on the boundary of $t(\phi(v))$, we therefore have $\phi(v) \notin [-m, m]_{\mathbb{Z}}$. Hence $\phi(\partial M_n) \cap [-m, m]_{\mathbb{Z}} = \emptyset$. \square

Proof of Theorems 1.6 and 1.7 in case 1 Theorem 1.6 in the case of the spanning-tree weighted map is an immediate consequence of Theorem 1.5, Lemma 1.11, and [35, Theorem 1.10]. Theorem 1.7 follows from Theorem 1.5 and [35, Theorem 1.15]. \square

3.2 Site percolation on the UIPT

3.2.1 Bijection for site percolation on the UIPT

In this subsection we review the infinite-volume version of the bijection between site-percolated triangulations of type II (no self-loops, but multiple edges allowed) and *Kreweras walks*—those with steps in $\{(1, 0), (0, 1), (-1, -1)\}$ —which was first described in [16]. This bijection is based on an earlier bijection by Bernardi [14] between such walks and trivalent maps decorated by a depth-first search tree. We emphasize that the only part of [16] needed here is the definition of the bijection in [16, Section 2].

Let (M, e_0) be an instance of the uniform infinite planar triangulation (UIPT) of type II, rooted at a directed edge e_0 . Let P be an instance of critical site percolation on the map, i.e., P associates each $v \in \mathcal{V}(M)$ with either red or blue, uniformly and independently at random. In [16, Section 2.5] (based on [14]) it was proved that (M, e_0, P) can a.s. be encoded by a bi-infinite walk with increments $a = (1, 0)$, $b = (0, 1)$, $c = (-1, -1)$.⁵ Conversely, a bi-infinite walk $\mathcal{Z} = (\mathcal{Z}_k)_{k \in \mathbb{Z}}$ with $\mathcal{Z}_0 = 0$ and steps a, b, c chosen uniformly and independently at random, a.s. encodes an

⁵ To be precise, the correspondence between walks and triples (M, e_0, P) in [16] is for the case where the edge e_0 is undirected. The direction of e_0 does not play an important role in this section, and we may for example assume that the direction of e_0 is determined by an additional binary random variable.




Fig. 10 Illustration of the sewing procedure for the measure-preserving map between percolation-decorated triangulations and Kreweras walks (i.e., walks with steps a, b, c). The black arrows indicate that we glue e_{k-1}^h to an adjacent edge, i.e., the two edges are identified. See the main text for the precise description. Vertices are colored gray if their color is not yet determined, while they are colored black if their color has been determined, so they are either blue or red, but the color cannot be determined from the figure (color figure online)

instance of (M, e_0, P) . Throughout this section we let $(\Delta \mathcal{Z}_k)_{k \in \mathbb{N}}$ denote the steps of \mathcal{Z} , i.e.,

$$\Delta \mathcal{Z}_k := \mathcal{Z}_k - \mathcal{Z}_{k-1} \in \{a, b, c\}.$$

As in the case of the Mullin bijection (see Sect. 3.1.2), the construction of (M, e_0, P) from the walk \mathcal{Z} is uniquely characterized by a sewing procedure, which we now describe. Each $k \in \mathbb{Z}$ corresponds to a map $M_{-\infty, k}$ with a percolation configuration $P_{-\infty, k}$ on the inner vertices and a marked edge e_k^h . Again the sewing procedure describes how to obtain $(M_{-\infty, k}, P_{-\infty, k}, e_k^h)$ from $(M_{-\infty, k-1}, P_{-\infty, k-1}, e_{k-1}^h)$ by observing the step $\Delta \mathcal{Z}_k$ of the walk. The map M rooted at e_0 is the local limit of $M_{-\infty, k}$ rooted at e_0 as k goes to ∞ in the Benjamin-Schramm sense. We will see that each k can be associated with a unique element $\tilde{\lambda}(k) \in \mathcal{V}(M) \cup \mathcal{F}(M)$. More precisely, we have $\tilde{\lambda}(k) \in \mathcal{V}(M)$ if $\Delta \mathcal{Z}_k = c$ and $\tilde{\lambda}(k) \in \mathcal{F}(M)$ if $\Delta \mathcal{Z}_k \in \{a, b\}$. The root edge e_0 of the map is e_{-1}^h . For any $k \in \mathbb{Z}$ the map $M_{-\infty, k}$ has an infinite left frontier and an infinite right frontier. Note that the sewing procedure is only an inductive procedure, but, as for the Mullin bijection, it can be shown that given a walk \mathcal{Z} with i.i.d. increments there is a.s. a unique collection of triples $(M_{-\infty, k}, P_{-\infty, k}, e_k^h)$ such that $(M_{-\infty, k}, P_{-\infty, k}, e_k^h)$ is obtained from $(M_{-\infty, k-1}, P_{-\infty, k-1}, e_{k-1}^h)$ via the sewing procedure. We refer to [16, Section 2.5] for a more complete description, and to Fig. 10 for an illustration.

Given $M_{-\infty, k-1}$ and e_{k-1}^h , we will now explain how to construct $M_{-\infty, k}$ via the sewing procedure. See Fig. 10. If $\Delta \mathcal{Z}_k = a$ (resp. $\Delta \mathcal{Z}_k = b$), we add a triangle to the edge e_{k-1}^h . We let the new head e_k^h be the right (resp. left) side of the triangle. We define $\tilde{\lambda}(k) \in \mathcal{F}(M)$ to be the triangle we added. Next consider the case $\Delta \mathcal{Z}_k = c$. For an edge e on the frontier of the map we say that a face f is *behind* e if it is the unique face of the map for which e is on its boundary. Let e_L (resp. e_R) be the edge on the frontier of $M_{-\infty, k-1}$ immediately to the left (resp. right) of e_{k-1}^h , and let f_L (resp. f_R) be the face behind e_L (resp. e_R). If $\tilde{\lambda}^{-1}(f_L) < \tilde{\lambda}^{-1}(f_R)$ (resp. $\tilde{\lambda}^{-1}(f_L) > \tilde{\lambda}^{-1}(f_R)$) define $e_k^h = e_R$ (resp. $e_k^h = e_L$). Then identify e_{k-1}^h with e_L (resp. e_R), i.e., glue the two edges together. The common vertex of e_L (resp. e_R) and e_{k-1}^h is colored red (resp. blue), and we define $\tilde{\lambda}(k) \in \mathcal{V}(M)$ to be this vertex. We see from this description that the function \mathcal{L} (resp. \mathcal{R}) represents the net change in the length of the left (resp. right) frontier of the map, relative to the length at time 0.




Fig. 11 The figure shows M_n embedded into M (left) and M_n (right) in light blue. Notice that vertices which are adjacent in M may correspond to vertices of M_n which are far apart. The green path is part of the depth-first search (DFS) tree of the dual map M^* as defined in [16, Section 4.1] (this DFS tree is not used in the present paper). The DFS tree gives an ordering of the faces of M that is consistent with the ordering given by $\tilde{\lambda}$ (color figure online)

We remark that $M_{-\infty, k}$ is not a submap of M , since for every step $\Delta \mathcal{Z}_k = c$ in the sewing procedure we identify two edges and two vertices, and vertices and edges which are identified in the sewing procedure at times $> k$ are *not* identified in $M_{-\infty, k}$. However, as remarked above, there are natural maps from the set of vertices (resp. inner faces, edges) of $M_{-\infty, k}$ to the set of vertices (resp. faces, edges) of $M_{-\infty, k+1}$ and M . For faces this map is injective, while this is not the case for vertices and edges. We will often identify vertices (resp. faces, edges) of the various maps when this identification is well-defined. We remark that for $k \in \mathbb{Z}$, the edge $\lambda(k)$ of Remark 1.3 is the edge of M corresponding⁶ to e_k^h .

Given $n \in \mathbb{N}$ we now define the map M_n introduced in Sect. 1.4. We will let M_n be a particular submap of $M_{-\infty, n}$ (see Fig. 11). Let the vertex set $\mathcal{V}(M_n) \subset \mathcal{V}(M_{-\infty, n})$ of M_n be the set of all vertices which are either an endpoint of some edge e_{k-1}^h or e_k^h or on the boundary of some face $\tilde{\lambda}(k) \in \mathcal{F}(M_{-\infty, n})$ for $k \in [-n, n]_{\mathbb{Z}}$. Let $\mathcal{E}_1 \subset \mathcal{E}(M_{-\infty, n})$ be the set of edges $e = \{v_1, v_2\}$ which are on the boundary of some face $\tilde{\lambda}(k) \in \mathcal{F}(M_{-\infty, n})$. Let $\mathcal{E}_2 = \{e_{k-1}^h : k \in [-n, n]_{\mathbb{Z}}, \Delta \mathcal{Z}_k = c\} \subset \mathcal{E}(M_{-\infty, n})$ be the set of edges which were glued to another edge in the sewing procedure. Define $\mathcal{E}(M_n) = \mathcal{E}_1 \cup \mathcal{E}_2$. We let $\iota_n : M_n \rightarrow M$ be the function which sends each vertex, face, and edge of M_n to the corresponding vertex, face, or edge of M . We will define the mappings ϕ_n and ψ_n of (1.5) in (3.4) just below.

3.2.2 Proofs of main theorems in the UIPT case

To prove the theorems from Sect. 1 in case 2, we will make small local changes in the map M and the walk \mathcal{Z} defined above and observe that the resulting map and walk are related by Mullin's bijection as described in Sect. 3.1. In particular, as we will see, if we replace each of the c steps of the walk \mathcal{Z} by a $-a$ step followed by a $-b$ step and each of the corresponding “gluing steps” in the sewing procedure by a step which adds a certain subgraph consisting of two triangles to the map, then the resulting

⁶ Note that the function $\lambda : \mathbb{Z} \rightarrow \mathcal{E}(M)$ is *not* the same as the function $\eta_e : \mathbb{Z} \rightarrow \mathcal{E}(M)$ considered in [16]. For all $k \in \mathbb{Z}$, we have $\lambda(k-1) = \eta_e(k)$.




Fig. 12 Illustration of $\widehat{\alpha} : \mathbb{Z} \rightarrow \mathbb{Z}$, which describes the time change of $\widehat{\mathcal{Z}}$ relative to \mathcal{Z} (color figure online)

modified walk and planar map are related via Mullin's bijection (Sect. 3.1.1). This will allow us to reduce the proof of Theorem 1.9 in the UIPT case to the Mullin bijection case, which we have already treated. It is also possible to treat case 2 directly, without reference to the Mullin bijection, but the argument given here is shorter and simpler.

We first define a modified walk $\widehat{\mathcal{Z}}$ with steps in $\{-a, -b, a, b\}$ by replacing each step c in \mathcal{Z} by one step $-a$ followed by one step $-b$. More precisely, $\widehat{\mathcal{Z}}$ and the maps α_n and $\widehat{\alpha}_n$ describing the associated time changes are defined as follows. First define $\alpha : \mathbb{Z} \rightarrow \mathbb{Z}$ and $\widehat{\alpha} : \mathbb{Z} \rightarrow \mathbb{Z}$ by

$$\widehat{\alpha}(k) = \begin{cases} k + \#\{j \in \{0, \dots, k-1\} : \Delta\mathcal{Z}_j = c\} & \text{for } k \geq 0 \\ k - \#\{j \in \{k, \dots, -1\} : \Delta\mathcal{Z}_j = c\} & \text{for } k < 0, \end{cases}$$

$$\alpha(k) = \sup\{j \in \mathbb{Z} : \widehat{\alpha}(j) \leq k\}.$$

See Fig. 12 for an illustration. Let $\widehat{\alpha}_n = \widehat{\alpha}|_{[-n, n]\mathbb{Z}}$, and let $n_- = -\widehat{\alpha}_n(-n)$. If $\Delta\mathcal{Z}_n \neq c$ set $n_+ = \widehat{\alpha}_n(n)$, and if $\Delta\mathcal{Z}_n = c$ set $n_+ = \widehat{\alpha}_n(n) + 1$. Define $\alpha_n : [-n_-, n_+]\mathbb{Z} \rightarrow [-n, n]\mathbb{Z}$ by $\alpha_n = \alpha|_{[-n_-, n_+]\mathbb{Z}}$, and define $\widehat{\mathcal{Z}}$ by

$$\widehat{\mathcal{Z}}_k = \begin{cases} \mathcal{Z}_{\alpha(k)} & \text{if } \mathcal{Z}_{\alpha(k)} \in \{a, b\}, \\ -a & \text{if } \mathcal{Z}_{\alpha(k)} = c \text{ and } \alpha(k) = \alpha(k+1), \\ -b & \text{if } \mathcal{Z}_{\alpha(k)} = c \text{ and } \alpha(k) \neq \alpha(k+1). \end{cases}$$

Since the walk $\widehat{\mathcal{Z}}$ has nearest-neighbor steps, it determines an infinite triangulation \widehat{M} (analogous to the triangulation $\mathcal{Q} \cup T \cup T_*$ from Sect. 3.1.1), a map $\widehat{t} : \mathbb{Z} \rightarrow \mathcal{F}(\widehat{M})$, and an infinite triangulations with boundary $\widehat{M}_{-\infty, n}$ for $n \in \mathbb{N}$ (which is a submap of \widehat{M}) via the Mullin bijection sewing procedure described in Sect. 3.1.2.

More precisely, for steps $\Delta\widehat{\mathcal{Z}}_k \in \{a, b\}$ the sewing procedure is as before and $\widehat{t}(k)$ is the triangle which we add. If $\Delta\mathcal{Z}_k = -a$ (resp. $\Delta\mathcal{Z}_k = -b$) we instead add a triangle $\widehat{t}(k)$ adjacent to e_{k-1}^h , we glue the right (resp. left) side of $\widehat{t}(k)$ to the edge immediately to its right (resp. left) on the frontier, and we define e_k^h to be the (only) edge of $\widehat{t}(k)$ which is still on the frontier of the map.

We define \widehat{M}_n be the (finite) submap of $\widehat{M}_{-\infty, n}$ consisting of the vertices and edges on the boundaries of the $2n+1$ triangles added at times in $[-n, n]\mathbb{Z}$, as in Sect. 3.1.3.

Given $k \in [-n, n]\mathbb{Z}$ such that $\Delta\mathcal{Z}_k = c$, let \widehat{D}_k be the submap of \widehat{M}_n of boundary length 4 consisting of the two faces $\widehat{t}_n(\widehat{\alpha}_n(k))$ and $\widehat{t}_n(\widehat{\alpha}_n(k)+1)$ (which correspond to consecutive steps of the $\widehat{\mathcal{Z}}$ of the form $-a$ and $-b$) and the vertices and edges on their boundaries. We say that we *contract* \widehat{D}_k if we modify the map \widehat{M}_n by the following procedure. (See Fig. 13 for an illustration.) Let e_1, e_2, e_3, e_4 denote the edges on the




Fig. 13 Contraction of \widehat{D}_k (light blue), as defined before the statement of Lemma 3.5, and illustration (in purple) of the map $\beta_n : \mathcal{V}(\widehat{M}_n) \rightarrow \mathcal{V}(M_n)$ defined after Lemma 3.5 (color figure online)

boundary of \widehat{D}_k in counterclockwise order, such that $e_1 = \widehat{e}_{\widehat{\alpha}_n(k)-1}^h$ and $e_3 = \widehat{e}_{\widehat{\alpha}_n(k)+1}^h$. A contraction means that we remove the two faces $\widehat{t}_n(\widehat{\alpha}_n(k))$ and $\widehat{t}_n(\widehat{\alpha}_n(k)+1)$ from the map, and then we identify edges pairwise along the boundary of \widehat{D}_k . The edge between the two faces is removed, while the other faces on the boundary of the two faces are kept. Let f_L (resp. f_R) be the face of $\widehat{M}_{-\infty, n}$ which has e_4 (resp. e_2) on its boundary, but is not equal to $\widehat{t}_n(\widehat{\alpha}_n(k))$ or $\widehat{t}_n(\widehat{\alpha}_n(k)+1)$. If $\widetilde{\lambda}^{-1}(f_L) < \lambda^{-1}(f_R)$ (resp. $\widetilde{\lambda}^{-1}(f_L) > \lambda^{-1}(f_R)$) then we identify e_1 and e_4 (resp. e_1 and e_2), and e_2 and e_3 (resp. e_3 and e_4); notice that the two edges we identify will always have one common vertex and one not common vertex, and when we identify the edges we choose the orientation of the edges in the natural way, such that the common (resp. not common) vertices are identified.

Lemma 3.5 *Let $\widehat{\mathcal{H}}_n$ denote the map with vertex set $[-n_-, n_+]_{\mathbb{Z}}$ and adjacency condition defined by (2.2) with the walk $\widehat{\mathcal{Z}}$ in place of \mathcal{Z} .*

- (i) *The map $k \mapsto \widehat{t}(k)$ is a graph isomorphism modulo multiplicity (recall Definition 3.2) from $\widehat{\mathcal{H}}_n$ to the dual of \widehat{M}_n .*
- (ii) *If we contract all the submaps \widehat{D}_k of \widehat{M}_n for $k \in [-n, n]_{\mathbb{Z}}$ such that $\Delta \mathcal{Z}_k = c$, then we get a map which is isomorphic to M_n .*

Proof (i) This is immediate by Proposition 3.3, and since $\widehat{\mathcal{Z}}$ and \widehat{M}_n (in the limit as $n \rightarrow \infty$) are related by Mullin's bijection.

(ii) Consider two walks \mathcal{Z}' and \mathcal{Z}'' with steps $\Delta \mathcal{Z}'$ and $\Delta \mathcal{Z}''$, respectively, in $\{a, b, c, -a, -b\}$. Assume \mathcal{Z}'' is obtained from \mathcal{Z}' by choosing a $k_0 \in \mathbb{Z}$ such that $\Delta \mathcal{Z}'_{k_0} = c$, and replacing this step c by two steps $-a$ and $-b$. In other words, we have $\Delta \mathcal{Z}_k'' = \Delta \mathcal{Z}_k'$ for $k < k_0$, $\Delta \mathcal{Z}_{k_0}'' = -a$, $\Delta \mathcal{Z}_{k_0+1}'' = -b$, and $\Delta \mathcal{Z}_k'' = \Delta \mathcal{Z}_{k-1}'$ for $k > k_0 + 1$. Let M' (resp. M'') denote the map we get when applying the sewing procedure described above to \mathcal{Z}' (resp. \mathcal{Z}''), and let D'' be the submap of M'' of boundary length 4 corresponding to the faces added in steps k_0 and $k_0 + 1$ of the sewing procedure. Since we can proceed by induction, in order to conclude it is sufficient to show that when we contract D'' , we get the map M' . This property is immediate by the sewing procedure applied to \mathcal{Z}' and \mathcal{Z}'' , respectively. \square

From Lemma 3.5(ii) we see that there exists a natural map $\beta_n : \mathcal{V}(\widehat{M}_n) \rightarrow \mathcal{V}(M_n)$ given by contracting all of the \widehat{D}_k 's. See Fig. 13. We define $\widehat{\beta}_n : \mathcal{V}(M_n) \rightarrow \mathcal{V}(\widehat{M}_n)$ in an arbitrary way such that $\beta_n \circ \widehat{\beta}_n$ is the identity map $\mathcal{V}(M_n) \rightarrow \mathcal{V}(M_n)$, and such that $\widehat{\beta}_n$ sends ∂M_n to $\partial \widehat{M}_n$. Note that it is possible to find a map $\widehat{\beta}_n$ satisfying the latter property by the construction in the proof of Lemma 3.5(ii). Let $\mathcal{F}(\widehat{M}_n)$ denote

the set of bounded faces of \widehat{M}_n , i.e., $\mathcal{F}(\widehat{M}_n) \setminus \mathring{\mathcal{F}}(\widehat{M}_n)$ consist of the single unbounded face of \widehat{M}_n . Let $\widehat{\phi}_n : \mathcal{V}(\widehat{M}_n) \rightarrow \mathring{\mathcal{F}}(\widehat{M}_n)$ map $v \in \mathcal{V}(\widehat{M}_n)$ to some arbitrarily chosen face of \widehat{M}_n which has v on its boundary. Let $\widehat{\psi}_n : \mathring{\mathcal{F}}(\widehat{M}_n) \rightarrow \mathcal{V}(\widehat{M}_n)$ map $f \in \mathring{\mathcal{F}}(\widehat{M}_n)$ to some arbitrarily chosen vertex on its boundary. We also require that $\phi_n, \widehat{\psi}_n, \widehat{\beta}_n$ are defined such that the maps ϕ_n and ψ_n defined just below satisfy the requirements on the root in (1.5) (note that by the definition of the contraction operation, it is possible to satisfy both this requirement and the second requirement on $\widehat{\beta}_n$ specified right above). By Lemma 3.5(i), we may identify $\mathring{\mathcal{F}}(\widehat{M}_n)$ with $[-n_-, n_+]_{\mathbb{Z}}$ in a natural way, so we may view $\widehat{\psi}_n$ and $\widehat{\phi}_n$ as maps $\widehat{\psi}_n : [-n_-, n_+]_{\mathbb{Z}} \rightarrow \mathcal{V}(\widehat{M}_n)$ and $\widehat{\phi}_n : \mathcal{V}(\widehat{M}_n) \rightarrow [-n_-, n_+]_{\mathbb{Z}}$. Define

$$\begin{aligned}\phi_n : \mathcal{V}(M_n) &\rightarrow [-n, n]_{\mathbb{Z}} \text{ by } \phi_n := \alpha_n \circ \widehat{\phi}_n \circ \widehat{\beta}_n \quad \text{and} \\ \psi_n : [-n, n]_{\mathbb{Z}} &\rightarrow \mathcal{V}(M_n) \text{ by } \psi_n := \beta_n \circ \widehat{\psi}_n \circ \widehat{\alpha}_n.\end{aligned}\quad (3.4)$$

Proof of Lemma 1.12 in case 2 By the definition of the maps ϕ_n it is sufficient to show the following to conclude the proof of the lemma:

- (i) $\widehat{\beta}_n$ sends ∂M_n to $\partial \widehat{M}_n$;
- (ii) α_n sends the complement of $[-m_-, m_+]_{\mathbb{Z}}$ to the complement of $[-m, m]_{\mathbb{Z}}$;
- (iii) if (1.21) holds then (1.21) also holds with $\widehat{\mathcal{L}}, \widehat{\mathcal{R}}, n_{\pm}, m_{\pm}$ instead of $\mathcal{L}, \mathcal{R}, n, m$;
- (iv) the map \widehat{M}_n satisfies Lemma 1.12, with (1.22) replaced by $\widehat{\phi}_n(\partial \widehat{M}_n) \cap [-m_-, m_+]_{\mathbb{Z}} = \emptyset$.

It follows from the definition of $\widehat{\beta}_n$ that (i) holds, and it follows from the definition of α and $\widehat{\mathcal{Z}}$ that (ii) and (iii), respectively, hold. Property (iv) follows from the argument for case 1 since the sewing procedure for \widehat{M}_n and the definition of $\widehat{\phi}_n$ are the same as for case 1. \square

Proof of Theorem 1.9 in case 2 Define the planar maps \mathcal{H} and $\{\mathcal{H}_n\}_{n \in \mathbb{N}}$ as in (2.2) and the discussion just after with \mathcal{Z} the encoding walk for M . Throughout the proof, we fix $A > 0$ and $n \in \mathbb{N}$ and we couple \mathcal{Z} with the correlated Brownian motion Z in such a way that the conclusion of Theorem 2.1 is satisfied for this choice of \mathcal{H} .

To prove condition 1 from the theorem statement, we will argue that for sufficiently large $C > 1$, depending only on A , the following is true for each $n \in \mathbb{N}$.

- (i) For any $v_1 \sim v_2$ in $\mathcal{V}(M_n)$ we can find a path $P_{v_1, v_2}^{\widehat{M}_n}$ from $\widehat{\beta}_n(v_1)$ to $\widehat{\beta}_n(v_2)$ in \widehat{M}_n of length at most 3, and each edge of \widehat{M}_n is on at most 10 of these paths. (Notice that when we bound the number of paths hitting each $i \in \mathcal{V}(\mathcal{G}_n)$ below, it is sufficient to bound the number of paths hitting each edge, rather than each vertex, in this step.)
- (ii) With probability at least $1 - O_n(n^{-A})$, for any $\widehat{v}_1 \sim \widehat{v}_2$ in $\mathcal{V}(\widehat{M}_n)$ we can find a path $P_{\widehat{v}_1, \widehat{v}_2}^{\widehat{H}_n}$ from $\phi_n(\widehat{v}_1)$ to $\phi_n(\widehat{v}_2)$ in \widehat{H}_n of length at most $C \log n$, and each $i \in [-n_-, n_+]_{\mathbb{Z}}$ is hit by at most $C \log n$ of the paths $P_{\widehat{v}_1, \widehat{v}_2}^{\widehat{H}_n}$.
- (iii) For any $i_1 \sim i_2$ in \widehat{H}_n we can find a path $P_{i_1, i_2}^{\mathcal{H}_n}$ from $\alpha_n(i_1)$ to $\alpha_n(i_2)$ in \mathcal{H}_n of length at most 2, and each vertex i of \mathcal{H}_n is hit by at most 10 of the paths $P_{i_1, i_2}^{\mathcal{H}_n}$.

(iv) With probability at least $1 - O_n(n^{-A})$, for any $i_1 \sim i_2$ in \mathcal{H}_n we can find a path $P_{i_1, i_2}^{\mathcal{G}_n}$ from i_1 to i_2 in \mathcal{G}_n of length at most $C(\log n)^3$, and each $i \in [-n, n]_{\mathbb{Z}}$ is hit by at most $C(\log n)^6$ of the paths $P_{i_1, i_2}^{\mathcal{G}_n}$.

Before we prove (i)–(iv), we will argue that they imply condition 1 upon replacing C by $100C^2$. Indeed, if the events in (i)–(iv) occur (which happens with probability at least $1 - O_n(n^{-A})$) and $v_1 \sim v_2$ in $\mathcal{V}(M_n)$ we can construct the path $P_{v_1, v_2}^{\mathcal{G}_n}$ from $\phi_n(v_1)$ to $\phi_n(v_2)$ in \mathcal{G} by considering the path $P_{v_1, v_2}^{\widehat{M}_n}$ from (i), replacing each step $(\widehat{v}_1, \widehat{v}_2)$ by the path $P_{\widehat{v}_1, \widehat{v}_2}^{\widehat{\mathcal{H}}_n}$ from (ii), replacing each step of the resulting path by the appropriate path from (iii), then replacing each step of this last path by the appropriate path from (iv). It is immediate from (i)–(iv) that the length of the path will be at most $3 \cdot C \log n \cdot 2 \cdot C(\log n)^3 = 6C^2(\log n)^4$, and that each i is hit by at most $10 \cdot C \log n \cdot 10 \cdot C(\log n)^6 = 100C^2(\log n)^7$ of the paths $P_{v_1, v_2}^{\widehat{M}_n}$.

The property (i) follows from Lemma 3.5(ii), while (iii) is immediate from the adjacency condition (2.2) for \mathcal{H}_n and $\widehat{\mathcal{H}}_n$. It follows from Theorem 2.1 that we can couple so that (iv) holds. Since the degree of the root of the UIPT has an exponential tail [10, Lemma 4.2], we know that for large enough C with probability at least $1 - O_n(n^{-A})$,

$$\deg(v; M_n) \leq C \log n, \quad \forall v \in \mathcal{V}(M_n). \quad (3.5)$$

The property (ii) follows from Lemma 3.5(i) and (3.5), since the definition of ϕ_n implies that the face of \widehat{M}_n associated with $\phi_n(\widehat{v}_1)$ (resp. $\phi_n(\widehat{v}_2)$) has \widehat{v}_1 (resp. \widehat{v}_2) on its boundary, so we can find a path of faces for \widehat{M}_n from $\phi_n(\widehat{v}_1)$ to $\phi_n(\widehat{v}_2)$ of length at most $C \log n$ (c.f. the proof of condition 1 of Theorem 1.9 from Sect. 3.1.3 for a similar argument).

The proof of condition 2 is very similar, and we therefore only give a brief sketch. First we prove variants of (i)–(iv) above for the inverse maps. For example, instead of (i) above, we show that for any $\widehat{v}_1 \sim \widehat{v}_2$ in $\mathcal{V}(\widehat{M}_n)$ we can find a path $P_{\widehat{v}_1, \widehat{v}_2}^{M_n}$ from $\beta_n(\widehat{v}_1)$ to $\beta_n(\widehat{v}_2)$ in M_n of length at most 3, and each edge of M_n is on at most 10 of these paths. The new variants of (i)–(iv) are proved as before, and they imply condition 2 similarly as in the proof of condition 1.

To prove condition 3 we will only explain how to bound $\text{dist}(\phi_n \circ \psi_n(i), i; \mathcal{G}_n)$, since the bound for $\text{dist}(\psi_n \circ \phi_n(v), v; M_n)$ is proven in a similar way. Assume the events in (i)–(iv) and (3.5) are satisfied, and choose an arbitrary $i \in [-n, n]_{\mathbb{Z}}$. Define

$$\widehat{i} = \widehat{\alpha}_n(i), \quad \widehat{v} = \widehat{\psi}_n(\widehat{i}), \quad v = \beta_n(\widehat{v}), \quad \widehat{v}' = \widehat{\beta}_n(v), \quad \widehat{i}' = \widehat{\phi}_n(\widehat{v}'), \quad i' = \alpha_n(\widehat{i}').$$

By the definition of contraction and since $\widehat{v}' = \widehat{\beta}_n(\beta_n(\widehat{v}))$, we have $\text{dist}(\widehat{v}, \widehat{v}'; \widehat{M}_n) \leq 2$. By the triangle inequality,

$$\text{dist}(\widehat{i}, \widehat{i}'; \widehat{\mathcal{H}}_n) \leq \text{dist}(\widehat{i}, \widehat{\phi}_n(\widehat{v}); \widehat{\mathcal{H}}_n) + \text{dist}(\widehat{\phi}_n(\widehat{v}), \widehat{i}'; \widehat{\mathcal{H}}_n).$$

The first term on the right side is bounded by $C \log n$, since it follows from the definition of $\widehat{\psi}_n$ and $\widehat{\phi}_n$ that the faces $\widehat{t}(\widehat{i})$ and $\widehat{t}(\widehat{\phi}_n(\widehat{\psi}_n(\widehat{i})))$ have a common vertex, so the

distance between the two faces in the dual of \widehat{M}_n (equivalently, $\widehat{\mathcal{H}}_n$) is at most $C \log n$ by (3.5). The second term on the right side is bounded by $2C \log n$, since (ii) and $\text{dist}(\widehat{v}, \widehat{v}'; \widehat{M}_n) \leq 2$ imply that $\text{dist}(\widehat{\phi}_n(\widehat{v}), \widehat{\phi}_n(\widehat{v}'); \widehat{\mathcal{H}}_n) \leq 2C \log n$. We conclude that $\text{dist}(\widehat{i}, \widehat{i}'; \widehat{\mathcal{H}}_n) \leq 3C \log n$. Using this bound, (iii), the definition of $\alpha_n, \widehat{\alpha}_n$, and $i' = \alpha_n(\widehat{i}')$, we get further

$$\text{dist}(i, i'; \mathcal{H}_n) \leq \text{dist}(i, \alpha_n(\widehat{i}); \mathcal{H}_n) + \text{dist}(\alpha_n(\widehat{i}), i'; \mathcal{H}_n) \leq 1 + 2 \cdot 3C \log n.$$

Combining this with (iv) concludes the proof: $\text{dist}(i, i'; \mathcal{G}_n) \leq (1+6C \log n)(\log n)^3 \leq 7C(\log n)^4$, where we assume that n is chosen sufficiently large in the last step. \square

Proof of Theorems 1.6, 1.7, and 1.8 in case 2 By [7, Theorem 1.2], the metric ball $B_n(e_0; M)$ of radius n centered at the root edge in the UIPT satisfies

$$\#\mathcal{V}(B_n(e_0; M)) = n^{4+o_n(1)} \quad \text{with probability tending to 1 as } n \rightarrow \infty, \quad (3.6)$$

which is stronger than the UIPT case of Theorem 1.6. (Observe that the weaker variant stated in Theorem 1.6 can also be proved by proceeding as in the UST case considered above.) Theorem 1.7 follows from Theorem 1.9 and [35, Theorem 1.15]. Theorem 1.8 is immediate from (3.6) together with the UIPT versions of Theorem 1.9 and Lemma 1.11. \square

3.3 Bipolar-oriented planar maps and Schnyder woods

In this section we study the case of planar maps decorated by a bipolar orientation or Schnyder wood. We first do the necessary combinatorial arguments for finite bipolar-oriented maps in Sect. 3.3.1. Then we move to the infinite-volume setting in Sect. 3.3.2, and explain why the *uniform infinite bipolar-oriented map*, the local limit of finite bipolar-oriented maps, exists (the analogous arguments in the case of the other maps considered in this paper have already been carried out elsewhere). We then prove Theorem 1.9 in the case of the uniform infinite bipolar-oriented map (case 3) in Sect. 3.3.3. Finally in Sect. 3.3.4 we treat the case of bipolar-oriented maps with other face degree distributions (case 4) and show that Schnyder wood-decorated maps (case 5) can be treated as a special case of bipolar-oriented maps.

3.3.1 The mating-of-trees bijection for bipolar-oriented maps

An *orientation* on a graph is an assignment of a direction to each edge. A vertex is called a sink (resp. source) if there are no outgoing (resp. incoming) edges incident to the vertex. Sources and sinks are also called *poles*. A *bipolar orientation* of G with specified source and sink is an acyclic orientation with no source or sink except at the specified poles. A *bipolar-oriented planar map* is a planar map M , with a marked face f_0 and a bipolar orientation \mathcal{O} on M whose source and sink are both on the boundary of f_0 . Bipolar-oriented planar maps have rich combinatorial structures and numerous applications in algorithms. For an overview of the graph theoretic perspective on




Fig. 14 In the first panel, we embed a bipolar-oriented map (M, f_0, \mathcal{O}) so that each edge is pointing to the north direction. We use blue and red vertices to represent their corresponding north and south poles. The second panel illustrates the (blue) NW tree and (red) SE tree. On the third panel, there is a path λ , drawn in green, as the interface between the NW and SE trees. The path λ induces a total ordering of $\mathcal{E}(M)$ which we still denote by λ . The fourth panel lists the two cases in the definition of λ and $\tilde{\lambda}$. In both cases $\lambda(t-1)$ is colored green and $\lambda(t)$ is colored purple. In case 1 (top figure) $\tilde{\lambda}(t)$ is the boldfaced vertex and in case 2 (bottom figure) $\tilde{\lambda}(t)$ is the face (color figure online)

bipolar orientations, we refer to [22] and references therein. See also [12, 31] and the references therein for enumerative and bijective results.

The plane \mathbb{R}^2 is naturally associated with the notion of east, west, south and north. It is known [2] that any bipolar-oriented planar map can be embedded into \mathbb{R}^2 in such a way that every edge is north oriented and the marked face is unbounded. See Fig. 14 for an example. Given a bipolar-oriented map (M, f_0, \mathcal{O}) , we always embed it this way so that \mathcal{O} can be read off from the embedding and the source and sink can be identified as the south pole and north pole, respectively, of the map. We therefore denote the source and sink by $S(M)$ and $N(M)$, respectively. For any edge $e \in \mathcal{E}(M)$ directed as in \mathcal{O} , we also write $S(e)$ and $N(e)$, respectively, for the initial and terminal edge of e .

Let $\mathcal{V}(M) := \mathcal{V}(M) \setminus \{N(M), S(M)\}$ be the set of non-polar vertices of M and $\mathring{\mathcal{F}}(M) := \mathcal{F}(M) \setminus \{f_0\}$ the set of *bounded* faces of M . We call a bipolar-oriented map consisting of a simple cycle (separating a single bounded face from the unbounded face) a *bipolar-oriented face*. The edges of a bipolar-oriented face can be divided into those which lie on the west and east boundary arcs connected the two poles, which we call *west edges* and *east edges*, respectively. We can think of any face in $\mathring{\mathcal{F}}(M)$ as a bipolar-oriented face. We call the clockwise (resp. counterclockwise) arc on f_0 from $S(M)$ to $N(M)$ the west (resp. east) boundary of M . An edge e is on both the west and east boundaries of M if and only if removing e separates M into two connected components.

We will now review the mating-of-trees bijection for finite bipolar-oriented maps from [46]. For any $v \in \mathcal{V}(M) \setminus \{N(M)\}$, there is a unique edge incident to v which is oriented away from v and is the furthest west among all such edges. We call this edge the *NW (northwest) edge* of v . Similarly, there is a unique edge incident to v which is oriented toward v and is the furthest east among all such edges. We call this edge the *SE (southeast) edge* of v . The *NW tree* \mathcal{T}_{NW} of (M, f_0, \mathcal{O}) is the directed planar tree that can be drawn as follows: each NW edge is entirely in \mathcal{T}_{NW} ; for each other edge,

a segment containing the head of the edge is in \mathcal{T}_{NW} . Similarly, the SE tree \mathcal{T}_{SE} can be drawn as follows: each SE edge is entirely in \mathcal{T}_{SE} ; for each other edge, a segment containing the tail of the edge is in \mathcal{T}_{SE} . Once \mathcal{T}_{NW} and \mathcal{T}_{SE} are drawn, we can draw an *interface* path λ winding between them from $\mathbf{S}(M)$ to $\mathbf{N}(M)$. The path λ introduces an ordering on $\mathcal{E}(M)$, i.e. a mapping (still denoted by λ) from $[0, \#\mathcal{E}(M) - 1]_{\mathbb{Z}}$ to $\mathcal{E}(M)$. It also defines a bijection $\tilde{\lambda} : [1, \#\mathcal{E}(M) - 1]_{\mathbb{Z}} \rightarrow \mathring{\mathcal{V}}(M) \cup \mathring{\mathcal{F}}(M)$. We now give a formal definition of λ , $\tilde{\lambda}$.

Definition 3.6 Set $\lambda(0)$ to be the NW edge of $\mathbf{S}(M)$. Inductively, for $1 \leq t \leq |\mathcal{E}(M)| - 1$, given $\lambda(t - 1)$:

1. If $\lambda(t - 1)$ is the SE edge of $\mathbf{N}(\lambda(t - 1))$, let $\lambda(t)$ be the NW edge of $\mathbf{N}(\lambda(t - 1))$ and set $\tilde{\lambda}(t) := \mathbf{S}(\lambda(t)) \in \mathcal{V}(M)$.
2. Otherwise, let $\tilde{\lambda}(t)$ be the unique face where $\lambda(t - 1)$ is a west edge and $\lambda(t)$ be the unique east edge of $\tilde{\lambda}(t)$ where $\mathbf{S}(\lambda(t)) = \mathbf{S}(\tilde{\lambda}(t))$.

See Fig. 14 for an illustration of \mathcal{T}_{NW} , \mathcal{T}_{SE} , λ , $\tilde{\lambda}$.

Suppose (M, f_0, \mathcal{O}) has $m + 1$ edges on the west boundary and $n + 1$ edges on the east boundary. We can associate (M, f_0, \mathcal{O}) with a walk $\mathcal{Z} = (\mathcal{L}, \mathcal{R})$ on \mathbb{Z}^2 , starting from $(m, 0)$ when $t = 0$ and ending at $(0, n)$ when $t = \#\mathcal{E}(M) - 1$. For $t \in [1, \#\mathcal{E}(M) - 1]_{\mathbb{Z}}$, if $\tilde{\lambda}(t) \in \mathring{\mathcal{V}}(M)$, we define the increment $\Delta \mathcal{Z}_t := \mathcal{Z}_t - \mathcal{Z}_{t-1}$ to be $(-1, 1)$. If $\tilde{\lambda}(t) \in \mathring{\mathcal{F}}(M)$, we set $\Delta \mathcal{Z}_t = (i, -j)$ where $i + 1$ and $j + 1$ are the number of the east and west edges of the face $\tilde{\lambda}(t)$ respectively. It is shown in [46, Theorem 2] that this procedure gives a bijection from bipolar-oriented maps to finite-length lattices walks on $\mathbb{Z}_{\geq 0}^2$ starting at the x -axis, ending at the y -axis, whose steps are in the set⁷

$$\{(-1, 1)\} \cup \{(i, -j) : i, j \in \mathbb{Z}_{\geq 0}\}. \quad (3.7)$$

Following [46, Section 2.2], we now describe the *sewing procedure* in this case, which is a way to dynamically build a bipolar-oriented map (M, f_0, \mathcal{O}) from the lattice path $(\mathcal{Z}_t)_{0 \leq t \leq \#\mathcal{E}(M)-1}$ of this type. For each $0 \leq k \leq \#\mathcal{E}(M) - 1$, we inductively associate $(\mathcal{Z}_t)_{0 \leq t \leq k}$ with a bipolar-oriented map $(M^k, f_0^k, \mathcal{O}^k)$ with a marked edge e^k on its east boundary. When $k = 0$, the map $(M^0, f_0^0, \mathcal{O}^0)$ is just a directed edge e^0 . For $0 < k \leq \#\mathcal{E}(M) - 1$, suppose $(M^{k-1}, f_0^{k-1}, \mathcal{O}^{k-1}, e^{k-1})$ is constructed.

1. If $\Delta \mathcal{Z}_k = (-1, 1)$ and $\mathbf{N}(e^{k-1}) \neq \mathbf{N}(M^{k-1})$, then $(M^k, f_0^k, \mathcal{O}^k) = (M^{k-1}, f_0^{k-1}, \mathcal{O}^{k-1})$ and e^k is the unique edge on the east boundary of M^{k-1} with $\mathbf{S}(e^k) = \mathbf{N}(e^{k-1})$.
2. If $\Delta \mathcal{Z}_k = (-1, 1)$ and $\mathbf{N}(e^{k-1}) = \mathbf{N}(M^{k-1})$, then $(M^k, f_0^k, \mathcal{O}^k, e^k)$ is obtained by attaching a north going edge e^k to $(M^{k-1}, f_0^{k-1}, \mathcal{O}^{k-1})$ at $\mathbf{N}(M^{k-1})$ so that $\mathbf{N}(M^k) = \mathbf{N}(e^k)$.
3. If $\Delta \mathcal{Z}_k = (i, -j)$ for some $i, j \geq 0$, to construct $(M^k, f_0^k, \mathcal{O}^k, e^k)$, we glue a bipolar-oriented face f^k with $i + 1$ east edges and $j + 1$ west edges to M^{k-1} , so that $\mathbf{N}(f^k) = \mathbf{N}(e^k)$, the west edges of f^k are east edges of f_0^k , and the east edges of f^k are *not* edges of M^{k-1} . The marked edge e^k is updated to be the unique east edge of f^k satisfying $\mathbf{S}(e^k) = \mathbf{S}(f^k)$.

⁷ The walk in [46] is written as $(\mathcal{R}, \mathcal{L})$, thus the set of steps differs from (3.7) by a reflection against $\{x = y\}$.




Fig. 15 We present ten snapshots of the sewing procedure for the bipolar-oriented map in Fig. 14 where $k \in \{1, 2, 3, 5, 6, 8, 9, 11, 12, 13\}$. The ten figures represent these $(M^k, f_0^k, \mathcal{O}^k)$'s in order with the corresponding incremental step below. $\Delta \mathcal{Z} = (-1, 1)$ corresponds to a Type 1 or Type 2 move, i.e. adding an edge. Here the Type 2 moves are displayed in gray. $\Delta \mathcal{Z} \neq (-1, 1)$ corresponds to a Type 3 move., i.e. adding a face. The marked edge e^k of $(M^k, f_0^k, \mathcal{O}^k)$ is always the current edge visited by the (green) path λ , i.e. $\lambda(k)$. The final map $(M^{13}, f_0^{13}, \mathcal{O}^{13})$ recovers Fig. 14 (color figure online)

From the construction, $\{(M^k, f_0^k, \mathcal{O}^k)\}_{0 \leq k \leq \#\mathcal{E}(M)-1}$ is an increasing family of bipolar-oriented maps with $S(M)$ as its common south pole. See Fig. 15 for an illustration. It is explained in [46] (and can be easily checked by induction) that the final map $(M^{\#\mathcal{E}(M)-1}, f_0^{\#\mathcal{E}(M)-1}, \mathcal{O}^{\#\mathcal{E}(M)-1})$ is the bipolar-oriented map (M, f_0, \mathcal{O}) producing \mathcal{Z} , thus each $(M^k, f_0^k, \mathcal{O}^k)$ can be considered as a submap of (M, f_0, \mathcal{O}) . Moreover, the number of edges from $S(e^k)$ to $S(M)$ (resp. $N(e^k)$ to $N(M)$) on the east boundary of $(M^k, f_0^k, \mathcal{O}^k)$ is equal to \mathcal{R}_k (resp. \mathcal{L}_k). From 1 and 2 in the Definition 3.6 we observe that

$$\begin{aligned} \lambda(k) &= e^k && \text{for all } k, \\ \tilde{\lambda}(k) &= f^k && \text{whenever } \tilde{\lambda}(k) \in \mathring{\mathcal{F}}(M), \\ \tilde{\lambda}(k) &= S(e^k) && \text{whenever } \tilde{\lambda}(k) \in \mathring{\mathcal{V}}(M). \end{aligned}$$

By the definition of \mathcal{Z} , λ and the sewing procedure, we get the following.

Lemma 3.7 *For two integers i, j , the edge $\lambda(i)$ is the NW edge of $N(\lambda(j))$ if and only if $i = \inf\{t > j : \mathcal{L}_t = \mathcal{L}_j - 1\}$. Similarly, $\lambda(i)$ is the SE edge of $S(\lambda(j))$ if and only if $i = \sup\{t < j : \mathcal{R}_t = \mathcal{R}_j - 1\}$.*

The following observation will allow us to apply Theorem 2.1 in the proof of Theorem 1.9 in the bipolar-oriented map case. In its statement we extend \mathcal{Z} to




Fig. 16 Illustration of the proof of Lemma 3.8. The left panel is an example of \mathcal{R} and τ_{\bullet}^{t-1} when $\Delta \mathcal{Z}_t \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\leq 0}$. We write τ_{\bullet}^{t-1} as τ_{\bullet} for simplification. In the right panel, the straight line represents the east boundary of M^{t-1} . The green edge is $\lambda(t-1)$ and the purple edge is $\lambda(t)$. The vertices to the south of $N(\lambda(t-1))$ are $\tilde{\lambda}(\tau_1), \tilde{\lambda}(\tau_2), \dots$ in order. At time t we add a face $\tilde{\lambda}(t)$ with west boundary length $\ell = 7$. We see that $S(\tilde{\lambda}(t)) = \tilde{\lambda}(\tau_7)$. Moreover, in the left panel for t, \mathcal{R} fixed, the time s satisfies the inequality (3.9) if and only if $s \in \{\tau_i\}_{1 \leq i \leq 7}$, which proves Assertion 1 (color figure online)

$[-1, \#\mathcal{E}(M)]_{\mathbb{Z}}$ and $\tilde{\lambda}$ to $[0, \#\mathcal{E}(M)]_{\mathbb{Z}}$ by setting

$$\mathcal{Z}_0 - \mathcal{Z}_{-1} = \mathcal{Z}_{\#\mathcal{E}(M)} - \mathcal{Z}_{\#\mathcal{E}(M)-1} = (-1, 1), \quad \tilde{\lambda}(0) = S(M), \quad \text{and} \quad \tilde{\lambda}(n) = N(M). \quad (3.8)$$

Lemma 3.8 For $s, t \in [0, \#\mathcal{E}(M)]_{\mathbb{Z}}$ with $s < t$, we have that

1. $\tilde{\lambda}(s)$ is the tail of a west edge on the boundary of $\tilde{\lambda}(t)$ if and only if

$$\mathcal{R}_{s-1} \vee (\mathcal{R}_t - 1) < \min_{s \leq j < t} \mathcal{R}_j \quad (3.9)$$

2. $\tilde{\lambda}(t)$ is the head of an east edge on the boundary of $\tilde{\lambda}(s)$ if and only if

$$\mathcal{L}_t \vee (\mathcal{L}_{s-1} - 1) < \min_{s-1 < j \leq t-1} \mathcal{L}_j \quad (3.10)$$

Proof See Fig. 16 for an illustration of the proof. For all $k \in [0, \#\mathcal{E}(M)]_{\mathbb{Z}}$, let $\tau_0^k = k$. For $i \geq 0$ and $\tau_i^k \geq 0$, we set $\sup \emptyset = -\infty$ by convention and inductively define

$$\tau_{i+1}^k = \sup\{t < \tau_i^k : \Delta \mathcal{Z}_t = (-1, 1) \text{ and } \mathcal{R}_t = \mathcal{R}_{\tau_i^k} - 1\}.$$

We first prove the following claim by an induction on k :

$S(e^k) = \tilde{\lambda}(\tau_1^k)$ and the vertices on the east boundary of M^k between $S(e^k)$ and $spole(M^k)$ are $\tilde{\lambda}(\tau_1^k), \tilde{\lambda}(\tau_2^k), \dots, \tilde{\lambda}(\tau_{\mathcal{R}_k+1}^k)$ in order. (3.11)

For $k = 0$, (3.11) holds by (3.8). For $1 \leq k \leq \#\mathcal{E}(M)$, assume (3.11) holds for $k - 1$. If $\Delta\mathcal{Z}_k = (-1, 1)$, then $\tau_i^k = \tau_{i-1}^{k-1}$ for $i \geq 1$ and by the definition of the sewing procedure we add an edge at time k (so we do not cover up any boundary vertices), which shows that (3.11) holds for k . If $\Delta\mathcal{Z}_k \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\leq 0}$, then at time k we add a face $\tilde{\lambda}(k)$ with west boundary length $\ell_k := \mathcal{R}_k - \mathcal{R}_{k-1}$. Therefore $\tau_i^k = \tau_{\ell_k+i-1}^{k-1}$ for $i \geq 1$ and $S(e^k) = \tilde{\lambda}(\tau_{\ell_k}^{k-1}) = \tilde{\lambda}(\tau_1^k)$. Thus (3.11) holds for k . This proves (3.11). Note that s, t satisfies (3.9) if and only if $s \in \{\tau_i^{t-1} : 1 \leq i \leq \ell_{t-1}\}$. Now Assertion 1 follows from (3.11).

By rotating \mathbb{R}^2 by 180 degrees, we reverse the direction of each edge on (M, f_0, \mathcal{O}) and swap the role of south and north as well as west and east. We can also reverse the path λ and the walk \mathcal{Z} . Then the role of \mathcal{R} is played by the time reversal of \mathcal{L} and vice versa. Here we note that after the time reversal, $(s - 1, s, t - 1, t)$ becomes $(t, t - 1, s, s - 1)$. This time reversal trick gives Assertion 2 from Assertion 1. \square

Remark 3.9 The time reversal trick in the proof of Lemma 3.8 will be used frequently.

3.3.2 The uniform infinite bipolar-oriented map

In this section, we construct an infinite planar map which will be the local limit of the uniform bipolar-oriented maps as the size tends to ∞ (see Proposition 3.16 and Definition 3.18.). Consider a bi-infinite simple random walk $\mathcal{Z} = (\mathcal{L}, \mathcal{R})$ with $\mathcal{Z}_0 = 0$ and the step distribution ν supported on (3.7) such that

$$\nu(i, -j) = 2^{-i-j-3} \text{ for } i, j \geq 0 \text{ or } i = j = -1. \quad (3.12)$$

As above, we write $\Delta\mathcal{Z}_t = \mathcal{Z}_t - \mathcal{Z}_{t-1}$ for all $t \in \mathbb{Z}$. For two integers $m < n$, we choose $K \in \mathbb{N}$ large enough so that $\mathcal{Z}([m, n]_{\mathbb{Z}}) + (K, K) \subset (0, \infty)^2$. Let $x = \mathcal{L}_m + \mathcal{R}_m + 2K$ and $y = \mathcal{L}_n + \mathcal{R}_n + 2K$ so that the straight line segment between $\mathcal{Z}_m + (K, K)$ (resp., $\mathcal{Z}_n + (K, K)$) and $(x, 0)$ (resp., $(0, y)$) has slope -1 . By concatenating the trajectory of $\mathcal{Z}|_{[m, n]_{\mathbb{Z}}} + (K, K)$ with these two line segments of slope -1 , each traversed by steps of the form $(-1, 1)$, we get a walk $\mathcal{Z}^{m, n, K}$ with steps in (3.7), starting from $(x, 0)$ and ending at $(0, y)$. Applying the sewing procedure in Sect. 3.3.1 to $\mathcal{Z}^{m, n, K}$ we obtain bipolar-oriented map $(M^{m, n, K}, f_0^{m, n, K}, \mathcal{O}^{m, n, K})$. We define $\lambda^{m, n, K}$ as λ in Definition 3.6 with $(M^{m, n, K}, f_0^{m, n, K}, \mathcal{O}^{m, n, K})$ in place of (M, f_0, \mathcal{O}) . See Fig. 17 for an illustration.

Since $\mathcal{Z}_t = \mathcal{Z}_{\mathcal{R}_m + K + t - m}^{m, n, K}$ for each $t \in [m, n]_{\mathbb{Z}}$, the map $M^{m, n, K}$ can be identified as a submap of $M^{m, n, K+1}$ with compatible edge orientations. In this identification, for each $t \in [m, n]_{\mathbb{Z}}$, $\lambda^{m, n, K}(\mathcal{R}_m + K + t - m)$ and $\lambda^{m, n, K+1}(\mathcal{R}_m + K + 1 + t - m)$ represent the same edge, which we denote by $\lambda^{m, n}(t)$. This edge corresponds to the step $\Delta\mathcal{Z}_{\mathcal{R}_m + K + t - m}^{m, n, K} = \Delta\mathcal{Z}_t$. If $M^{m, n, K}$ is removed from $M^{m, n, K+1}$, we are left with two disjoint edges, one oriented from $N(M^{m, n, K})$ to $N(M^{m, n, K+1})$, the other from $S(M^{m, n, K+1})$ to $S(M^{m, n, K})$. Taking the union $\bigcup_K M^{m, n, K}$, we obtain an infinite planar map $M^{m, n}$ with an edge orientation $\mathcal{O}^{m, n}$ where $\mathcal{O}^{m, n}|_{\mathcal{E}(M^{m, n, K})} = \mathcal{O}^{m, n, K}$ for large enough K . Moreover, $M^{m, n, K}$ is a submap of $M^{m, n}$ such that $\mathcal{E}(M^{m, n}) \setminus \mathcal{E}(M^{m, n, K})$ form two disjoint rays emanating from




Fig. 17 The walk $Z^{m,n,K}$ for $m = 1, n = 6, K = 4$ and its corresponding bipolar-oriented map $(M^{m,n,K}, f_0^{m,n,K}, O^{m,n,K})$, adjusted from the fifth figure in Fig. 15. The green edges are $\lambda^{m,n}(t)$ for $t \in [m, n]_{\mathbb{Z}}$. An edge of a certain color in the map corresponds to a step of the same color in the walk (color figure online)

$N(M^{m,n,K})$ and $S(M^{m,n,K})$, respectively. Concatenating the ray from $N(M^{m,n,K})$ to the NW tree of $M^{m,n,K}$, we obtain an infinite planar tree which we call the NW tree of $M^{m,n}$ that does not depend on K . Similarly, by concatenating the ray from $S(M^{m,n,K})$ to the SE tree of $M^{m,n,K}$, we obtain the SE tree of $M^{m,n}$.

By the time reversal trick in the proof of Lemma 3.8 (see also Remark 3.9), we have the following.

Lemma 3.10 For each $t \in \mathbb{Z}$, let $\bar{Z}_t = (\mathcal{R}_{-t}, \mathcal{L}_{-t})$. For $m < n \in \mathbb{Z}$, let $\bar{\mathcal{O}}^{m,n}$ be the orientation obtained by reversing each arrow in $\mathcal{O}^{m,n}$. Then $(M^{m,n}, f_0^{m,n}, \bar{\mathcal{O}}^{m,n})$ is isomorphic to $(\bar{M}^{-n,-m}, \bar{f}_0^{-n,-m}, \bar{\mathcal{O}}^{-n,-m})$, which is defined in the same way as $(M^{-n,-m}, f_0^{-n,-m}, \mathcal{O}^{-n,-m})$ with Z replaced by \bar{Z} .

Recall that $M^{m,n}$ can be embedded so that $\mathcal{O}^{m,n}$ is pointing upward for each edge. By Lemma 3.10, $(\bar{M}^{-n,-m}, \bar{\mathcal{O}}^{-n,-m})$ is obtained by rotating this picture by 180 degrees.

For $\ell \in [m, n]_{\mathbb{Z}}$, let $\{b_{\ell}^{m,n}(i)\}_{i \in \mathbb{Z}} \subset \mathcal{E}(M^{m,n})$ be such that $b_{\ell}^{m,n}(0) = \lambda^{m,n}(\ell)$, $b_{\ell}^{m,n}(i+1)$ is the NW edge of $N(b_{\ell}^{m,n}(i))$ for $i \geq 0$, and $b_{\ell}^{m,n}(i-1)$ is the SE edge of $S(b_{\ell}^{m,n}(i))$ for $i \leq 0$. Intuitively, we can think of $\{b_{\ell}^{m,n}(i)\}_{i > 0}$ as the branch of the NW tree of $M^{m,n}$ from $N(\lambda^{m,n}(\ell))$ to infinity, and $\{b_{\ell}^{m,n}(i)\}_{i < 0}$ as the branch of the SE tree of $M^{m,n}$ from $S(\lambda^{m,n}(\ell))$ to infinity.

Lemma 3.11 For $m' < m < n < n'$, $M^{m,n}$ is isomorphic to the submap of $M^{m',n'}$ bounded by the bi-infinite lines $b_m^{m',n'}$ and $b_n^{m',n'}$ in $M^{m,n}$. Under this identification, $\mathcal{O}^{m,n}|_{\mathcal{E}(M^{m,n})} = \mathcal{O}^{m,n}$. For each $t \in [m, n]_{\mathbb{Z}}$, we have $\lambda^{m,n}(t) = \lambda^{m',n'}(t)$ and $b_t^{m',n'} = b_t^{m,n}$.

Proof This can be checked straightforwardly from the definition of the sewing procedure in the finite volume setting. We leave the details to the reader. \square

Now for each $n \in \mathbb{N}$ we let $e^0 = \lambda^{-n,n}(0)$ be the root edge of $M^{-n,n}$. Taking the union $\bigcup_{n \in \mathbb{N}} M^{-n,n}$, we obtain an infinite planar map M with the root edge e^0 and




Fig. 18 We use the notations in the proof of Lemma 3.12. Left: An illustration of the times $\tau, \rho^\pm, T_\bullet, \hat{T}_\bullet$. The green step is $\Delta \mathcal{Z}_\tau$. Right: The corresponding picture for Nb_0 . The arrowed edge is e^0 . The green edge is $\lambda(\tau)$. The faces of Nb_0 correspond to the blue and red steps in the left panel. From Assertion 1 in Lemma 3.8, the time τ is first time when $S(\lambda(\tau)) = S(e^0)$. Moreover, ρ^+ (resp. ρ^-) is the unique time t such that $\tilde{\lambda}(t)$ is a face and $S(e^0)$ is on the west (resp. east) boundary of $\tilde{\lambda}(t)$ but is not a pole of $\tilde{\lambda}(t)$. Finally, $\{\tilde{\lambda}(T_i)\}_{\xi^- < i \leq \xi^+}$ (resp. $\{\tilde{\lambda}(\hat{T}_i)\}_{1 \leq i \leq \zeta}$) describes the set of faces where $S(e^0)$ is the south (resp. north) pole (color figure online)

an edge orientation \mathcal{O} such that $\mathcal{O}|_{\mathcal{E}(M^{-n,n})} = \mathcal{O}^{-n,n}$ for each $n \in \mathbb{N}$. For each $t \in \mathbb{Z}$, let $\lambda(t) = \lambda^{-n,n}(t)$ for $n > |t|$. This defines a function $\lambda : \mathbb{Z} \rightarrow \mathcal{E}(M)$. By definition λ is an injection. Lemma 3.14 below will show that λ is in fact a bijection.

So far there is no probability involved. Now we prove a few probabilistic lemmas ensuring that almost surely M does not have certain pathological behaviors. We will use the following notation. For integers $m < n$, we call the union of the east (resp., west) boundary of $M^{m,n,K}$ for all large enough K 's the west (resp. east) boundary of $M^{m,n}$.

Lemma 3.12 *With probability 1, there exists $n \in \mathbb{N}$ such that $S(e^0)$, i.e., the tail of e^0 , is not on the east or west boundary of $M^{-n,n}$.*

Proof Consider the following special times for \mathcal{Z} :

$$\begin{aligned} \tau &:= \max\{t \leq 0 : \Delta \mathcal{Z}_t = (-1, 1) \text{ and } \mathcal{R}_t = \mathcal{R}_0 = 0\}, \\ \rho^+ &:= \inf\{t \geq 0 : \mathcal{R}_t < 0\}, \quad \text{and} \quad \rho^- := \sup\{t < \tau : \Delta \mathcal{Z}_t \\ &\neq (-1, 1) \text{ and } \mathcal{L}_{t-1} \leq \mathcal{L}_{\tau-1}\}. \end{aligned} \quad (3.13)$$

These times are illustrated in Fig. 18. Note that all three times are finite almost surely.

Let K be so large that $M^{\rho^-, \rho^+, K}$ is well defined and let $(M', f'_0, \mathcal{O}') := (M^{\rho^-, \rho^+, K}, f_0^{\rho^-, \rho^+, K}, \mathcal{O}^{\rho^-, \rho^+, K})$. Now we define a function $\tilde{\lambda} : [\rho^-, \rho^+] \mathbb{Z} \rightarrow \mathring{\mathcal{V}}(M') \cup \mathring{\mathcal{F}}(M')$ with $\tilde{\lambda}$ as in Definition 3.6 but with a different interval of definition. For $t \in [\rho^-, \rho^+] \mathbb{Z}$, if $\lambda(t-1)$ is the SE edge of $N(\lambda(t-1))$, set $\tilde{\lambda}(t) := S(\lambda(t))$. Otherwise, let $\tilde{\lambda}(t)$ be the unique face where $\lambda(t-1)$ is a west edge. Here the notion of east/west/south/north is with respect to (M', f'_0, \mathcal{O}') . We define $\tilde{\lambda}$ this way so that Lemma 3.8 can be adapted to this situation straightforwardly.

Let Nb_0 be the subgraph of M' consisting of all faces of M' with $S(e^0)$ on their boundaries, and all vertices and edges of M' that are on the boundaries of such faces.

We claim that faces in Nb_0 all belong to $\mathring{\mathcal{F}}(M')$, which will imply Lemma 3.12 with $n = |\rho^-| \vee |\rho^+|$.

Let $\xi^+ = \#\{t \in (0, \rho^+) : \mathcal{R}_t = 0\}$ and $\xi^- = -\#\{t \in [\tau, 0) : \mathcal{R}_t = 0\}$. Let $\{T_i\}_{\xi^- \leq i \leq \xi^+}$ be the set of times in $[\tau, \rho^+]$ for which $\mathcal{R}_t = 0$, in increasing order. Note that $\tau = T_{\xi^-}$ and $T_0 = 0$. Let $\zeta := \#\{t \in (\rho^-, \tau - 1) : \mathcal{L}_t = \mathcal{L}_{\tau-1}\}$ and let $\{\widehat{T}_i\}_{1 \leq i \leq \zeta}$ represent this time set in decreasing order. Recall Lemma 3.8 and its proof. We make the following observations which imply the claim in the previous paragraph. We will justify them just below.

1. $\{\lambda(T_i)\}_{\xi^- \leq i \leq \xi^+}$ is equal to the set of outgoing edges in Nb_0 in clockwise order. In particular, $\lambda(\tau)$ is the NW edge of $S(\lambda(0))$.
2. The cardinality of the set of faces in M' with $S(e^0)$ as south pole equals $\xi^+ - \xi^-$. When the set is nonempty, it equals $\{\tilde{\lambda}(T_i)\}_{\xi^- < i \leq \xi^+}$.
3. $\tilde{\lambda}(\rho^+)$ is the unique face in M' with $S(e^0)$ on its west boundary but not as one of its poles.
4. The cardinality of the set of faces in M' with $S(e^0)$ as north pole equals ζ . When the set is nonempty, it equals $\{\tilde{\lambda}(\widehat{T}_i)\}_{1 \leq i \leq \zeta}$.
5. $\tilde{\lambda}(\rho^-)$ is the unique face in M' with $S(e^0)$ on its east boundary but not as one of its poles.

See Fig. 18 for an illustration of observations 1-5. Note that if $\Delta\mathcal{Z}_0 = (1, -1)$, then $\tau = 0$. Otherwise $\tau = \inf\{s : s \text{ and } 0 \text{ satisfy (3.9)}\}$. In both cases $S(\lambda(\tau)) = S(e^0)$. Since $\Delta\mathcal{Z}_\tau = (1, -1)$, condition 1 in the definition of λ implies that $\lambda(\tau)$ is the NW edge of $S(\lambda(0)) = S(e^0)$. Now $\{T_i\}_{\xi^- < i \leq \xi^+}$ is the set of times t satisfying $\tau = \inf\{s : s \text{ and } t \text{ satisfy (3.9)}\}$. Moreover, ρ^+ is the unique time $> \tau$ such that τ and ρ^+ satisfy (3.9) and $\mathcal{R}_{\rho^+} < \mathcal{R}_\tau$. This combined with Assertion 1 in Lemma 3.8 proves observations 1-3. Recall Remark 3.9 and Lemma 3.10. If we replace τ by $\tau - 1$ and \mathcal{R} by the time reversal of \mathcal{L} in the above argument, we obtain observations 4-5. \square

Lemma 3.13 *The following occurs with probability 1. Each vertex in M has finite degree. For each r , there exists $n = n(r) \in \mathbb{N}$ large enough such that $B_r(S(e^0); M)$ is a subgraph of $M^{-n,n}$ and has no vertices on the east or west boundary of $M^{-n,n}$.*

Proof By Lemma 3.12, the degree of $S(e^0)$ is finite. Note that $e^0 = S(\lambda(0))$. By the stationarity of \mathcal{Z} , Lemma 3.12 remains true if $S(e^0)$ is replaced by $S(\lambda(k))$ for each $k \in \mathbb{Z}$. Let v be an adjacent vertex of $S(e^0)$. Then by the proof of Lemma 3.12 there exists $k \in \mathbb{Z}$ such that $v = S(\lambda(k))$ or $v = N(\lambda(k))$. In the later case, we must have $v = S(\lambda(k'))$ where $k' = \inf\{t > k : \mathcal{L}_t = \mathcal{L}_k - 1\}$. Therefore we can always find $k \in \mathbb{Z}$ such that $v = S(\lambda(k))$. Thus there almost surely exists $n_v \in \mathbb{N}$ such that $v \in \mathcal{V}(M^{-n_v, n_v})$ and v is not on the east or west boundary of M^{-n_v, n_v} . Let $n(1)$ be the maximum over all such n_v 's. Then $n(1)$ satisfies the desired property for $B_1(S(e^0); M)$. Iterating this argument for each vertex in $B_1(S(e^0); M)$ gives the desired result for $B_2(S(e^0); M)$. The general statement follows from further iterations. \square

By Lemma 3.13, given each $v \in \mathcal{V}(M)$, by considering $M^{-n,n}$ with large enough n , we can talk about the SE and NW edges of v as in the finite volume case.

Lemma 3.14 For each edge e on M , there exists $n \in \mathbb{N}$ such that $e \in \mathcal{E}(M^{-n,n})$ and e has no endpoints on the east or west boundary of $M^{-n,n}$. In particular, there exists $k \in [n, n]_{\mathbb{Z}}$ such that $e = \lambda(k)$ and the two faces with e on their boundaries belong to $\mathring{\mathcal{F}}(M^{-n,n})$.

Proof The existence of $M^{-n,n}$ in the first statement immediately follows from Lemma 3.13. By the definition of east and west boundaries, the two faces in $\mathcal{F}(M^{-n,n})$ with e on their boundaries belong to $\mathring{\mathcal{F}}(M^{-n,n})$. Choose K large enough so that $M^{-n,n,K}$ is well defined. By Lemma 3.8 and the definition of $\mathcal{Z}^{-n,n,K}$, for each $t < \mathcal{R}_{-n} + K$ there is no face in $\mathring{\mathcal{F}}(M^{-n,n,K})$ with $\lambda^{-n,n,K}(t)$ being an edge. Therefore, $\lambda^{-n,n,K}(t)$ must be on the west boundary of $M^{-n,n,K}$. Similarly, $\lambda^{-n,n,K}(t)$ must be on the east boundary of $M^{-n,n,K}$ for each $t > \mathcal{R}_{-n} + K + 2n$. Therefore there exists $k \in [-n, n]_{\mathbb{Z}}$ such that $e = \lambda^{-n,n}(k) = \lambda(k)$. \square

By Lemma 3.14 λ is a bijection between \mathbb{Z} and $\mathcal{E}(M)$. Moreover, the degree of each face of M is bounded. Therefore, given a face of M we can define its poles and east/west boundaries as in the finite map case. Now for each $t \in \mathbb{Z}$, if $\lambda(t-1)$ is the SE edge of $N(\lambda(t-1))$, set $\tilde{\lambda}(t) := S(\lambda(t))$. Otherwise, let $\tilde{\lambda}(t)$ be the unique face where $\lambda(t-1)$ is a west edge. Then $\tilde{\lambda} : \mathbb{Z} \rightarrow \mathcal{V}(M) \cup \mathcal{F}(M)$ is a bijection. For each $t \in \mathbb{Z}$, both Items 1 and 2 in Definition 3.6 hold in this infinite volume setting. The following fact is an immediate consequence of this extension of definition.

Lemma 3.15 For $n \in \mathbb{N}$ and K large enough such that $M^{-n,n,K}$ is well defined, we have that $\tilde{\lambda}([-n, n]_{\mathbb{Z}}) \cap \mathcal{F}(M) = \mathring{\mathcal{F}}(M^{-n,n,K})$.

We adapt the notion of Benjamini–Schramm convergence [20] to the setting of oriented rooted graphs. Given planar maps G and G' with oriented root edges e and e' , respectively, the Benjamini–Schramm distance [20] from (G, e) to (G', e') is defined by

$$d_{\text{loc}} = \inf \{2^{-r} : r \in \mathbb{N}, B_r(o; G) \simeq B_r(o'; G')\},$$

where o and o' are the tail of e and e' respectively and $B_r(e; G) \simeq B_r(e'; G')$ means that the two metric balls are isomorphic as oriented rooted planar maps. Then d_{loc} defines a metric on the space of finite rooted planar maps. Let \mathfrak{M} be the Cauchy completion of the space of finite oriented rooted planar maps under d_{loc} . Elements in \mathfrak{M} which are not finite are infinite rooted oriented planar maps. The following proposition asserts that M is the Benjamini–Schramm limit of the uniform finite bipolar-oriented maps viewed from a typical edge.

Proposition 3.16 Let $(M(n), \mathcal{O}(n))$ be a uniformly chosen bipolar-oriented map with n edges, and let $e(n)$ be an edge uniformly chosen in $\mathcal{E}(M(n))$. Then as $n \rightarrow \infty$, $(M(n), e(n), \mathcal{O}(n))$ converges in law to (M, e^0, \mathcal{O}) defined above in the Benjamini–Schramm topology.

Proof Let (ξ_1, ξ_2) be independent of \mathcal{Z} so that $\mathbb{P}[(\xi_1, \xi_2) = (i, j)] = 2^{-i-j-2}$ for all $(i, j) \in \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}$. Let E_n be the event that $\{(\xi_1, 0) + \mathcal{Z}_t\}_{0 \leq t \leq n-1}$ is a walk starting at $(\xi_1, 0)$, ending at $(0, \xi_2)$ and staying in $\mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}$ during $[0, n-1]$. Then

conditioning on E_n and applying the sewing procedure to $\{(\xi_t, 0) + \mathcal{Z}_t\}_{0 \leq t \leq n-1}$, we obtain a sample of $(M(n), \mathcal{O}(n))$ (see [46, Remark 2]). Let λ_n be the edge-valued function associated with this sewing procedure and let U_n be a uniform sample from $[0, n-1]_{\mathbb{Z}}$. Then $(M(n), \lambda(U_n), \mathcal{O}(n))$ is a sample of $(M(n), e(n), \mathcal{O}(n))$.

Now let $\mathcal{Z}^n(\bullet)$ be sampled from the conditional law of $\mathcal{Z}(\bullet + U_n) - \mathcal{Z}(U_n)$ given E_n . Standard results for random walk in cones (see, e.g., [27, Theorem 4] and [26, Theorem 1]) yield that as long as $\mathbb{P}[E_n] > 0$, it holds that $\mathbb{P}[E_n]$ decays sub-exponentially in n . Therefore, Cramer's large deviation theorem implies that for all $N > 0$ the walks \mathcal{Z}^n and \mathcal{Z} can be coupled so that their restrictions to $[-N, N]$ agree with probability $1 - o_n(1)$ (see [62, Section 4.2] for a similar application of Cramer's theorem). Now Proposition 3.16 follows from Lemma 3.13. \square

For each $\ell \in \mathbb{Z}$, let $M_{\ell, \infty}$ be the submap of M obtained by taking the union $\bigcup_{n > \ell} M^{\ell, n}$, rooted at the edge $\lambda(\ell)$. Then $M_{\ell, \infty}$ has a bi-infinite boundary, which coincides with the west boundary of $M^{\ell, n}$ for each integer $n > \ell$. List these boundary edges as $\{b_{\ell}(i)\}_{i \in \mathbb{Z}}$ such that $b_{\ell}(0) = \lambda(\ell)$ and $N(b_{\ell}(i)) = S(b_{\ell}(i+1))$ for each $i \in \mathbb{Z}$. By Lemma 3.11, we have the following.

Lemma 3.17 *In the setting right above, almost surely, $b_{\ell}(i+1)$ is the NW edge of $N(b_{\ell}(i))$ for each $i \geq 0$, and $b_{\ell}(i-1)$ is the SE edge of $S(b_{\ell}(i))$ for each $i \leq 0$.*

For each $\ell \in \mathbb{Z}$, we can also consider $M_{-\infty, \ell} = \bigcup_{n < \ell} M^{n, \ell}$, where $\lambda(\ell)$ is assigned to be the root edge. Then $M_{-\infty, \ell}$ and $M_{\ell, \infty}$ share the same boundary.

Definition 3.18 We call (M, e^0, \mathcal{O}) the *uniform infinite bipolar-oriented map* (UIBOM). We call $\{M_{-\infty, \ell}\}_{\ell \in \mathbb{Z}}$ and $\{M_{\ell, \infty}\}_{\ell \in \mathbb{Z}}$ the forward and backward *sewing procedure* of (M, e^0, \mathcal{O}) , respectively.

Proposition 3.19 *For each $\ell \in \mathbb{Z}$, $\{\mathcal{Z}_k - \mathcal{Z}_{\ell}\}_{k \leq \ell}$ and $(M_{-\infty, \ell}, \lambda(\ell), \mathcal{O}|_{\mathcal{E}(M_{-\infty, \ell})})$ almost surely determine each other. The same statement holds for $M_{\ell, \infty}$ with $\{\mathcal{Z}_k - \mathcal{Z}_{\ell}\}_{k \leq \ell}$ replaced by $\{\mathcal{Z}_k - \mathcal{Z}_{\ell}\}_{k \geq \ell}$. In particular, \mathcal{Z} and (M, e^0, \mathcal{O}) almost surely determine each other.*

Proof For $n > \ell$, by definition $M^{\ell, n}$ is determined by $\{\mathcal{Z}_k - \mathcal{Z}_{\ell}\}_{k \in [\ell, n]_{\mathbb{Z}}}$. Therefore $\{\mathcal{Z}_k - \mathcal{Z}_{\ell}\}_{k \geq \ell}$ determines $(M_{\ell, \infty}, \lambda(\ell), \mathcal{O}|_{\mathcal{E}(M_{\ell, \infty})})$. Now given $(M_{\ell, \infty}, \lambda(\ell), \mathcal{O}|_{\mathcal{E}(M_{\ell, \infty})})$, let f be the face with $\lambda(\ell)$ on its west boundary. According to whether $N(\lambda(\ell)) = N(f)$, we can decide $\lambda(\ell+1)$. By Lemma 3.17, $(M_{\ell+1, \infty}, \lambda(\ell+1), \mathcal{O}|_{\mathcal{E}(M_{\ell+1, \infty})})$ is determined by $(M_{\ell, \infty}, \lambda(\ell), \mathcal{O}|_{\mathcal{E}(M_{\ell, \infty})})$. By induction, $(M_{k, \infty}, \lambda(k), \mathcal{O}|_{\mathcal{E}(M_{k, \infty})})$ is determined by $(M_{\ell, \infty}, \lambda(\ell), \mathcal{O}|_{\mathcal{E}(M_{\ell, \infty})})$ for all $k > \ell$. This allows us to recover $\{\Delta \mathcal{Z}_t\}_{t > \ell}$ hence $\{\mathcal{Z}_k - \mathcal{Z}_{\ell}\}_{k \in [\ell, n]_{\mathbb{Z}}}$. Using the time reversal trick in Lemma 3.10, we obtain the same result for $M_{-\infty, \ell}$ and $\{\mathcal{Z}_k - \mathcal{Z}_{\ell}\}_{k \leq \ell}$. \square

Given $v \in \mathcal{V}(M)$, we define Nb_v to be the subgraph of M consisting of all faces of M with v on their boundaries and all vertices and edges of M on the boundaries of such faces. In the proof of Lemma 3.12, if M' is identified as a subgraph of M , then $Nb_{S(e^0)}$ equals Nb_0 defined there. For $f \in \mathcal{F}(M)$, let Nb_f be subgraph of M given by the union of Nb_v 's where v is a boundary vertex of f , i.e., Nb_f is the set of all

faces of M which share a vertex with f and all vertices and edges which lie on the boundaries of such faces. Let

$$|\text{Nb}_x| := \#\mathcal{V}(\text{Nb}_x) + \#\mathcal{F}(\text{Nb}_x) + \#\mathcal{E}(\text{Nb}_x), \quad \forall x \in \mathcal{V}(M) \cup \mathcal{F}(M).$$

We conclude this section by a degree-bound type estimates for the UIBOM.

Lemma 3.20 *For all fixed $i \in \mathbb{Z}$, set $\text{Nb}_i := \text{Nb}_{S(\lambda(i))}$ and $\widetilde{\text{Nb}}_i := \text{Nb}_{\widetilde{\lambda}(i)}$. Then the laws of both $|\text{Nb}_i|$ and $|\widetilde{\text{Nb}}_i|$ do not depend on i and have exponential tails. In particular, the degrees of the two endpoints of the root edge of the UIBOM have exponential tails.*

For the proof of Lemma 3.20, we need the following elementary estimate.

Lemma 3.21 *Suppose $\{X_i\}_{i \geq 1}$ is a sequence of i.i.d. nonnegative random variables such that $\mathbb{E}[e^{bX_1}] < \infty$ for some $b > 0$. Let ξ be a nonnegative random variable (not necessarily independent from the X_i 's) with an exponential tail. Then there exists $c > 0$ only depending on the law of X_1 and ξ so that*

$$\mathbb{P}\left[\sum_{i=1}^{\xi} X_i \geq n\right] \leq e^{-cn}.$$

Proof By a union bound, for $\delta > 0$,

$$\mathbb{P}\left[\sum_{i=1}^{\xi} X_i \geq n\right] \leq \mathbb{P}[\xi \geq \delta n] + \mathbb{P}\left[\sum_{i=1}^{\lfloor \delta n \rfloor} X_i \geq n\right].$$

By choosing δ small enough according to the law of X_1 , the lemma follows from Cramer's Large Deviation Theorem. \square

Proof of Lemma 3.20 The stationarity of $\{\text{Nb}_i\}_{i \in \mathbb{Z}}$ and $\{\widetilde{\text{Nb}}_i\}_{i \in \mathbb{Z}}$ follows from the stationarity of \mathcal{Z} . So it suffices to focus on $i = 0$ in the rest of the proof.

To deal with Nb_0 , we retain the notation used in the proof of Lemma 3.12. By the Markov property of \mathcal{Z} , all of ξ^- , ξ^+ and ζ have geometric distributions. Moreover, conditioning on $\xi^- \neq 0$, we have that $\{\Delta \mathcal{Z}_{T_i}\}_{\xi^- < i \leq 0}$ are i.i.d and have exponential tails. The same also holds for $\{\Delta \mathcal{Z}_i\}_{1 \leq i \leq \xi^+}$ and $\{\Delta \mathcal{Z}_{\widehat{T}_i}\}_{1 \leq i \leq \zeta}$. (However, the laws of $\Delta \mathcal{Z}_{T_{-1}}$ and $\Delta \mathcal{Z}_{T_1}$ are not the same.) Now observations 1–5 in the proof of Lemma 3.12 combined with Lemma 3.21 and (3.12) imply that $|\text{Nb}_0|$ has an exponential tail. We obtain the same result for $\{\text{Nb}_{S(\lambda(i))}\}_{i \in \mathbb{Z}}$ by reversing the time direction and apply Lemma 3.10.

It remains to treat $\widetilde{\text{Nb}}_0$. Set $t_0 = 0$. For a positive integer k , we inductively define $t_k = \inf\{t > t_{k-1} : \mathcal{L}_t = \mathcal{L}_{t_{k-1}} - 1\}$. Then $\{t_k\}_{k \geq 0}$ are stopping times for the filtration $\{\mathcal{F}_\ell = \sigma(\{\mathcal{Z}_i\}_{i \leq \ell})\}_{\ell \in \mathbb{Z}}$. Moreover, $\lambda(t_k)$ is the NW edge of $\lambda(t_{k-1})$ for $k \geq 1$. Let $\text{Nb}_k^+ = (\text{Nb}_{S(\lambda(t_{k-1}))} \cup \text{Nb}_{S(\lambda(t_k))}) \cap M^{t_{k-1}, t_k}$. Then $\{\text{Nb}_k^+\}_{k \geq 1}$ is a sequence of independent identically distributed random variables. Recall the $M_{0, \infty}$ in Definition 3.18. On the

event $v \in \mathcal{V}(M)$, $\widetilde{\text{Nb}}_0 = \text{Nb}_0$ is handled in the previous paragraph. On the event $\widetilde{\lambda}(0) \in \mathcal{F}(M)$, let ℓ be the number of edges on the east boundary of $\widetilde{\lambda}(0)$, which has an exponential tail. Then $\{\text{Nb}_k^+\}_{1 \leq k \leq \ell}$ and $\text{Nb}_{S(\lambda(t_\ell))} \cap M_{0,\infty}$ cover $\widetilde{\text{Nb}}_0 \cap M_{0,\infty}$. Let $|\text{Nb}_k^+| = \#\mathcal{V}(\text{Nb}_k^+) + \#\mathcal{F}(\text{Nb}_k^+) + \#\mathcal{E}(\text{Nb}_k^+)$. Since $|\text{Nb}_k^+| \leq |\text{Nb}_{S(e^0)}| + |\text{Nb}_{N(e^0)}|$, we see that $|\text{Nb}_k^+|$ has an exponential tail. On the other hand, $\text{Nb}_{S(\lambda(t_\ell))} \cap M_{0,\infty}$ has the same law as $\text{Nb}_0 \cap M_{0,\infty}$, whose size has an exponential tail. By Lemma 3.21 and reversing the time as in Lemma 3.10, we can bound $|\widetilde{\text{Nb}}_0|$ by the sum of two random variables with exponential tails, thus proving Lemma 3.20. \square

Remark 3.22 A bipolar orientation on a map can be associated with a bipolar orientation on its dual map via rotating the orientation of each edge by 90 degrees. Therefore the UIBOM and its dual map agree in law. Thus bounds on $|\widetilde{\text{Nb}}_0|$ immediately follow from bounds on $|\text{Nb}_0|$. However, our proof of Lemma 3.20 is more robust and works for the more general setting in Sect. 3.3.4.

3.3.3 Proof of Theorem 1.9 for the UIBOM

From the discussion in Sect. 3.3.2, Assertions 1 and 2 in Lemma 3.8 hold for the UIBOM for all $s, t \in \mathbb{Z}$ with $s < t$. We still refer to this fact as Lemma 3.8 for the UIBOM. Let $\mathcal{Q} := \mathcal{Q}(M)$ be the radial map associated with M , i.e., the quadrangulation obtained from M by connecting the center of each face of M with its adjacent vertices and removing all edges of M . We identify $\mathcal{V}(M) \cup \mathcal{F}(M)$ with $\mathcal{V}(\mathcal{Q})$ so that $\widetilde{\lambda}$ is a bijection between \mathbb{Z} and $\mathcal{V}(\mathcal{Q})$. By Lemma 3.8 for the UIBOM, we have the following.

Proposition 3.23 *The graph \mathcal{Q} is isomorphic via $i \leftrightarrow \widetilde{\lambda}(i)$ to the graph $\widehat{\mathcal{H}}$ defined as follows. The vertex set of $\widehat{\mathcal{H}}$ is \mathbb{Z} . For $s, t \in \mathcal{V}(\widehat{\mathcal{H}})$ with $s < t$, the vertices s and t are connected by an edge in $\widehat{\mathcal{H}}$ if and only if either*

$$(\mathcal{L}_{s-1} - 1) \vee \mathcal{L}_t < \min_{s-1 < j \leq t-1} \mathcal{L}_j \quad \text{or} \quad \mathcal{R}_{s-1} \vee (\mathcal{R}_t - 1) < \min_{s \leq j < t} \mathcal{R}_j. \quad (3.14)$$

Remark 3.24 Although it is not needed later, we note that \mathcal{Q} is a quadrangulation with boundary which can be oriented as follows. Each edge $e \in \mathcal{E}(\mathcal{Q})$ must be of the form $\{v, f\}$ with $v \in \mathcal{V}(M)$ and $f \in \mathcal{F}(M)$. Orient e from f to v if $v = \text{N}(f)$ or $\text{S}(f)$ and from v to f otherwise. This combinatorial object is called a 2-orientation on quadrangulation (see e.g. [29, Section 5.2]) because every vertex has out-degree 2. This construction gives a general bijection between bipolar orientations and 2-orientations. When $M = \mathbb{Z}^2$, the 2-orientation on \mathcal{Q} is called a 6-vertex configuration. This special case is considered in [45], based on which a connection between SLE₁₂ and the 6-vertex model is found.

To use Theorem 2.1 and Remark 2.5 to prove Theorem 1.9, we define a planar map $\mathcal{H} = \mathcal{H}(\mathcal{Z})$ via the following discrete analogue of (1.2). The vertex set of \mathcal{H} is $\mathcal{V}(\mathcal{H}) = \mathbb{Z}$, and for $s, t \in \mathcal{V}(\mathcal{H})$ with $s < t$, we declare that s and t are connected by an edge in \mathcal{H} if and only if either

$$\begin{aligned}
(\mathcal{L}_{s-1} - 1) \vee \mathcal{L}_t &< \min_{s-1 < j \leq t-1} \mathcal{L}_j \quad \text{or} \quad \mathcal{R}_{s-1} \vee (\mathcal{R}_t - 1) \\
&< \min_{s \leq j < t} \mathcal{R}_j \quad \text{or} \quad t - s = 1.
\end{aligned} \tag{3.15}$$

By Proposition 3.23, the radial map \mathcal{Q} can be identified with a submap of \mathcal{H} via $i \leftrightarrow \tilde{\lambda}(i)$. Moreover, \mathcal{H} can be recovered from $\widehat{\mathcal{H}}$ in Proposition 3.23 by adding edges $\{t, t+1\}$ for all $t \in \mathbb{Z}$. Recall Definition 3.6 of λ and $\tilde{\lambda}$ and its infinite volume extension in Sect. 3.3.2. We have that \mathcal{H} is isomorphic via $\tilde{\lambda}$ to the graph $\overline{\mathcal{Q}}$ such that $\mathcal{V}(\overline{\mathcal{Q}}) = \mathcal{V}(\mathcal{Q})$ and $v, w \in \mathcal{V}(\overline{\mathcal{Q}})$ are connected in $\overline{\mathcal{Q}}$ if and only if one of the following three possibilities occurs:

1. v and w are connected in \mathcal{Q} ;
2. v, w are two vertices of M and are connected by an edge $e \in \mathcal{E}(M)$ so that e is the NW edge of $S(e)$ and the SE edge of $N(e)$ (see Item 1 in Definition 3.6);
3. v, w are two faces of M which share an edge e so that $S(e)$ is the south pole of one face and $N(e)$ is the north pole of the other face (see Item 2 in Definition 3.6).

For $n \in \mathbb{N}$, let \mathcal{H}_n be the subgraph of \mathcal{H} whose vertex set is $[-n, n]_{\mathbb{Z}}$, with two vertices connected by an edge in \mathcal{H}_n if and only if they are connected by an edge in \mathcal{H} . Let $\overline{\mathcal{Q}}_n$ be the image of \mathcal{H}_n under the isomorphism $\tilde{\lambda}$. Let M_n be the smallest connected subgraph of M containing each vertex in $\{\tilde{\lambda}(t)\}_{t \in [-n, n]_{\mathbb{Z}}}$ and all vertices and edges on the boundary of each face in $\{\tilde{\lambda}(t)\}_{t \in [-n, n]_{\mathbb{Z}}}$.

One can give a more direct description of M_n . Recall Definition 3.18. For $n \in \mathbb{N}$, let $\{b_n(i)\}_{i \in \mathbb{Z}}$ and $\{b_{-n}(i)\}_{i \in \mathbb{Z}}$ be as in Lemma 3.17 with $\ell = n$ and $\ell = -n$, respectively. Define

$$\begin{aligned}
\tau_n &= \inf\{t \geq n : \lambda(t) = b_{-n}(i) \text{ for some } i > 0\} & \text{and } v_{+,n} = S(\lambda(\tau_n)); \\
\tau_{-n} &= \sup\{t \leq -n : S(\lambda(t)) = S(b_n(i)) \text{ for some } i \leq 0\} & \text{and } v_{-,n} = S(\lambda(\tau_{-n})).
\end{aligned}$$

Both $v_{+,n}$ and $v_{-,n}$ are on the boundary of both $M_{-n,\infty}$ and $M_{-\infty,n}$. The segments between $v_{+,n}$ and $v_{-,n}$ on these two boundaries enclose a bipolar oriented map $(\widehat{M}_n, \widehat{f}_n, \widehat{\mathcal{O}}_n)$. In other words, \widehat{M}_n is the unique finite connected component of $M^{-n,n}$ after removing $\{v_{+,n}, v_{-,n}\}$. Moreover, $\mathcal{O}|_{\mathcal{E}(\widehat{M}_n)} = \widehat{\mathcal{O}}_n$, $v_{+,n} = N(\widehat{M}_n)$ and $v_{-,n} = S(\widehat{M}_n)$.

Lemma 3.25 M_n is isomorphic to \widehat{M}_n when \widehat{M}_n is viewed as a graph.

Proof We first show that \widehat{M}_n is a subgraph of M_n . Since both \widehat{M}_n and M_n are connected subgraphs of M , it suffices to show that $\mathcal{E}(\widehat{M}_n) \subset \mathcal{E}(M_n)$. Choose K large enough so that $M^{n,n,K}$ is well defined. Given $e \in \mathcal{E}(\widehat{M}_n)$, if e is on the boundary of a face $f \in \mathring{\mathcal{F}}(M^{n,n,K})$, by Lemma 3.15 there exists $i \in [-n, n]_{\mathbb{Z}}$ such that $\lambda(i) = f$. Therefore $e \in \mathcal{E}(M_n)$. Now we suppose e is not on the boundary of any face in $\mathring{\mathcal{F}}(M^{n,n,K})$. Then e must be on the boundary of both $M_{-n,\infty}$ and $M_{-\infty,n}$. Let $k \in \mathbb{Z}$ be such that $\lambda(k) = e$. If $k \geq n$, by the minimality of τ_n , we must have $\tau_n \leq k$. This contradicts that $v_{+,n} = S(\lambda(\tau_n))$ is the north pole of \widehat{M}_n . Therefore $k < n$. Similarly, if $k < -n$, we must have $\tau_{-n} \geq k$ which contradicts that $v_{-,n} = S(\lambda(\tau_{-n}))$ is the south pole of \widehat{M}_n . Therefore $k \geq -n$. Let i be such that $b_{-n}(i) = e$. Then we must

have $i \geq 0$. Therefore e is the NW edge of $S(e)$ and hence $S(e) = \tilde{\lambda}(k) \in \mathcal{V}(M_n)$. On the other hand, $\lambda(k+1)$ must be the NW edge of $N(e)$. Namely, $\lambda(k+1) = b_{-n}(i+1)$. Otherwise $\tilde{\lambda}(k+1)$ would be a face with e on its boundary. Therefore $N(e) \in \mathcal{V}(M_n)$, which concludes the proof that $\mathcal{E}(\widehat{M}_n) \subset \mathcal{E}(M_n)$.

It remains to show that M_n is a subgraph of \widehat{M}_n . By the definition of M_n above, it suffices to show that $\{\tilde{\lambda}(t)\}_{t \in [-n, n]_{\mathbb{Z}}} \subset \mathcal{V}(\widehat{M}_n) \cup \mathring{\mathcal{F}}(\widehat{M}_n)$. By Lemma 3.15, $\{\tilde{\lambda}(t)\}_{t \in [-n, n]_{\mathbb{Z}}} \cap \mathcal{F}(M) = \mathring{\mathcal{F}}(M^{-n, n}) = \mathring{\mathcal{F}}(\widehat{M}_n)$. Now suppose $j \in \mathbb{Z}$ is such that $v = \tilde{\lambda}(j) \in \mathcal{V}(M^{-n, n, K}) \setminus \mathcal{V}(\widehat{M}_n)$. Then v must be in one of the two infinite connected of $M^{-n, n}$ after removing $\{v_{+, n}, v_{-, n}\}$. Therefore $j > \tau_n \geq n$ or $j < \tau_{-n} \leq -n$. This gives $\{\tilde{\lambda}(t)\}_{t \in [-n, n]_{\mathbb{Z}}} \cap \mathcal{V}(M) \subset \mathcal{V}(\widehat{M}_n)$ and concludes the proof. \square

From now on we identify the graph M_n with the submap \widehat{M}_n of M in Lemma 3.25. Let $\iota_n : M_n \rightarrow M$ be the inclusion map. Define $\psi : \mathbb{Z} \rightarrow \mathcal{V}(M)$ by $\psi(i) = S(\lambda(i))$ and set $\psi_n := \psi|_{[-n, n]_{\mathbb{Z}}}$. Define

$$\phi_n(v) := \tilde{\lambda}^{-1}(v) \mathbf{1}_{v \neq S(e^0)} \quad \text{for } v \in \mathcal{V}(M_n) \text{ with } |\tilde{\lambda}^{-1}(v)| \leq n.$$

To define $\phi_n(v)$ for the case when $|\tilde{\lambda}^{-1}(v)| > n$ we proceed as follows. Suppose $v \in \mathcal{V}(M_n)$ and $\mathcal{F}(Nb_v) \cap \mathcal{F}(M_n) = \emptyset$. Then v cannot be on any inner face of M_n . Here recall that $\mathcal{F}(Nb_v)$ is a subset of $\mathcal{F}(M)$ and note that the outer face of M_n is not a face in M . So $\mathcal{F}(Nb_v) \cap \mathcal{F}(M_n)$ is just $\mathcal{F}(Nb_v) \cap \mathring{\mathcal{F}}(M_n)$. By the first paragraph in the proof of Lemma 3.25, $v \in \{\tilde{\lambda}(t)\}_{t \in [-n, n]_{\mathbb{Z}}}$. Therefore, for $v \in \mathcal{V}(M_n)$ with $|\tilde{\lambda}^{-1}(v)| > n$, we must have $\mathcal{F}(Nb_v) \cap \mathcal{F}(M_n) \neq \emptyset$. In this case, we let $\phi_n(v)$ be the integer with the smallest absolute value for which $\tilde{\lambda}(\phi_n(v)) \in \mathcal{F}(Nb_v) \cap \mathcal{F}(M_n)$ (in the case of a tie, we choose $\phi_n(v)$ to be positive). Now we have a function $\phi_n : \mathcal{V}(M_n) \rightarrow [-n, n]_{\mathbb{Z}}$. Since $\tilde{\lambda}(0) \in \mathcal{F}(Nb_0) \cap \mathcal{F}(M_n)$ when $\tilde{\lambda}(0) \in \mathcal{F}(M)$, we see that $\phi_n(S(e^0)) = 0$. Therefore ψ_n, ϕ_n, M_n satisfy (1.5).

Proof of Lemma 1.12 in case 3 We first prove that under the condition (1.21) of Lemma 1.12, there does not exist $v \in \mathcal{V}(\partial M_n)$ such that $k = \phi_n(v) \in [-m, m]_{\mathbb{Z}}$ and $|\tilde{\lambda}^{-1}(v)| \leq n$. Suppose by way of contradiction that there exists such a v . If v is on the west boundary of M_n , by the definition of ϕ_n , we have $v = \tilde{\lambda}(k)$ and $\lambda(k)$ is the NW edge of v . If v is on the west boundary of M_n then there exists $i > 0$ such that $\lambda(k) = b_{-n}(i)$. (Recall b_{-n} in Lemma 3.17.) Therefore, by Lemma 3.7 and Lemma 3.17, $k = \inf\{t > -n : \mathcal{L}_t = \mathcal{L}_{-n} - i\}$. Since $k \in [-m, m]_{\mathbb{Z}}$, this contradicts the assumption (1.21) for \mathcal{L} . Similarly, if v is on the east boundary of M_n , by the time reversal trick in Remark 3.9 and Lemma 3.10, there exists $i < 0$ such that $\lambda(k-1) = b_n(i)$ and $k-1 = \sup\{t < n : \mathcal{R}_t = \mathcal{R}_n + i\}$. This contradicts the assumption (1.21) for \mathcal{R} .

It remains to show that there exists no $v \in \mathcal{V}(\partial M_n)$ such that $k = \phi_n(v) \in [-m, m]_{\mathbb{Z}}$ and $|\tilde{\lambda}^{-1}(v)| > n$. Suppose by way of contradiction that there exists such a v . Then $\tilde{\lambda}(k)$ is a face in $\mathring{\mathcal{F}}(M_m)$ with v on its boundary. Suppose $\tilde{\lambda}^{-1}(v) > n$. Then v must be on the east boundary of M_n . Thus the north pole of $\lambda(k)$, which is $N(\lambda(k-1))$, must be on the east boundary of M_n . Let $\ell = \inf\{t > k-1 : \mathcal{L}_t = \mathcal{L}_{k-1} - 1\}$. By Lemma 3.7, $\lambda(\ell)$ is the NW edge of $N(\lambda(k-1))$. Since $\tilde{\lambda}^{-1}(v) > n$, we must have $\ell > n$. Therefore $\mathcal{L}_{k-1} = \min_{t \in [k-1, n]_{\mathbb{Z}}} \mathcal{L}_t$, which contradicts (1.21) for \mathcal{L} . Suppose

$\tilde{\lambda}^{-1}(v) < -n$. Then similarly, v and $S(\lambda(k))$ must be on the east boundary of M_n . This gives $\sup\{t < k : \mathcal{R}_t = \mathcal{R}_k - 1\} < -n$, which contradicts (1.21) for \mathcal{R} . \square

Recall the notations above Lemma 3.20. The following fact follows from Lemma 3.12 and the definition of $\overline{\mathcal{Q}}_n$.

Lemma 3.26 *Suppose $v \in \mathcal{V}(M_n)$ and $f_1, f_2 \in \mathcal{F}(\text{Nb}_v) \cap \mathcal{F}(M_n)$. Then f_1, f_2 are connected in $\overline{\mathcal{Q}}_n$ only using vertices and faces in Nb_v . (Recall the identification $\mathcal{V}(M) \cup \mathcal{F}(M) = \mathcal{V}(\overline{\mathcal{Q}})$.)*

Proof We leave the reader to check the following fact. Write $\mathcal{F}(\text{Nb}_v)$ as $\{\tilde{\lambda}(i_k) : 1 \leq k \leq \#\mathcal{F}(\text{Nb}_v)\}$ where $i_1 < \dots < i_{\#\mathcal{F}(\text{Nb}_v)}$. Then for any $1 \leq k \leq \#\mathcal{F}(\text{Nb}_v) - 1$, if $\tilde{\lambda}(i_k)$ and $\tilde{\lambda}(i_{k+1})$ are both in $\mathcal{F}(M_n)$ but do not satisfy condition 3 in the definition of $\overline{\mathcal{Q}}$, then they are connected in \mathcal{Q}_n only using vertices and faces in Nb_v . These facts together with the definition of $\overline{\mathcal{Q}}_n$ imply the statement of the lemma. \square

From here on the proof of Theorem 1.9 in this case is very similar to the case of the infinite spanning tree-decorated map (Sect. 3.1.3) but there are minor technical differences so we give the complete argument.

Let $C_1 > 0$ be the constant from Theorem 2.1 for A in Theorem 1.9 and for \mathcal{Z} above. As in the spanning tree case, by the isomorphism (modulo multiplicity) $\tilde{\lambda}$ between \mathcal{H} and $\overline{\mathcal{Q}}$ plus Remark 2.5, we see that with probability at least $1 - O_n(n^{-A})$, the following is true

- (A) For each $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $\tilde{\lambda}(i_1) \sim \tilde{\lambda}(i_2)$ in $\overline{\mathcal{Q}}$, there is a path $\tilde{P}_{i_1, i_2}^{\mathcal{G}}$ from i_1 to i_2 in \mathcal{G}_n with $|\tilde{P}_{i_1, i_2}^{\mathcal{G}}| \leq C_1(\log n)^3$; and each $j \in [-n, n]_{\mathbb{Z}}$ is contained in at most $C_1(\log n)^6$ of the paths $\tilde{P}_{i_1, i_2}^{\mathcal{G}}$.
- (B) For each $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 \sim i_2$ in \mathcal{G} , then there is a path $\tilde{P}_{i_1, i_2}^{\overline{\mathcal{Q}}}$ from $\tilde{\lambda}(i_1)$ to $\tilde{\lambda}(i_2)$ in $\overline{\mathcal{Q}}_n$ with $|\tilde{P}_{i_1, i_2}^{\overline{\mathcal{Q}}}| \leq C_1(\log n)^3$; and each $j \in [-n, n]_{\mathbb{Z}}$ is contained in at most $C_1(\log n)^6$ of the paths $\tilde{P}_{i_1, i_2}^{\overline{\mathcal{Q}}}$.

By basic properties of Brownian motion, $\deg(i; \mathcal{G})$ has an exponential tail. See [39, Lemma 2.2] for a proof. Combined with Lemma 3.20, a union bound implies that there exists $C_2 = C_2(A)$ such that with probability at least $1 - O_n(n^{-A})$,

- (C) $\max_{x \in \mathcal{V}(M_n) \cup \mathcal{F}(M_n)} |\text{Nb}_x| \leq C_2 \log n$ and $\max_{i \in [-n, n]_{\mathbb{Z}}} \deg(i; \mathcal{G}) \leq C_2 \log n$.

From here on we work on the event that (A)–(C) hold.

Proof of condition 1 For $v_1, v_2 \in \mathcal{V}(M_n)$ such that $v_1 \sim v_2$ in M_n , let e be the edge connecting them. If both of the two faces of M with e on their boundaries are not in M_n , then one can check that v_1, v_2 satisfy condition 2 in the definition of $\overline{\mathcal{Q}}$ and $v_1, v_2 \in \mathcal{V}(\mathcal{Q}_n)$. Indeed, if $\{v_1, v_2\}$ is not on a face, it must be of the form $\lambda(t)$ for some $t \in [-n, n]_{\mathbb{Z}}$ by the definition of M_n . This means that it is the NW edge of its tail. If it is not the SE edge of its head, then the face on its right will be visited at $t + 1$. By the definition of ϕ_n , we have that $\phi_n(v_i) = \tilde{\lambda}^{-1}(v_i)$ for $i = 1, 2$. In this case we let $P_{v_1, v_2}^{\mathcal{G}} = P_{\tilde{\lambda}^{-1}(v_1), \tilde{\lambda}^{-1}(v_2)}^{\mathcal{G}}$.

If there exists $f \in \mathcal{F}(M_n)$ with e on its boundary, then $f \in \mathcal{F}(\text{Nb}_{v_1}) \cap \mathcal{F}(\text{Nb}_{v_2})$. Now Lemma 3.26 yields that $\tilde{\lambda}(\phi_n(v_i))$ and f are connected by a path P_i in $\overline{\mathcal{Q}}_n$ using only vertices and faces in Nb_{v_i} for $i = 1, 2$. Let P_{v_1, v_2}^G be the concatenation of paths $\tilde{P}_{\tilde{\lambda}^{-1}(v), \tilde{\lambda}^{-1}(w)}^G$ where $\{v, w\}$ is an edge on P_1 or P_2 . Now

$$|P_{v_1, v_2}^G| \stackrel{(A)}{\leq} 2 \max_{f \in \mathcal{F}(M_n)} |\text{Nb}_f| \cdot C_1 (\log n)^3 \stackrel{(C)}{\leq} 2C_1 C_2 (\log n)^4. \quad (3.16)$$

Given $f \sim v$ in \mathcal{Q}_n , if $\{f, v\}$ is an edge on P_1 or P_2 defined above for some $v_1, v_2 \in \mathcal{V}(M)$, we must have $v_1, v_2 \in \mathcal{V}(\text{Nb}_f)$. Therefore for each $j \in [-n, n]_{\mathbb{Z}}$, the number of paths of the form P_{v_1, v_2}^G hitting j is at most

$$\max_{f \in \mathcal{F}(M_n)} |\text{Nb}_f| \cdot C_2 (\log n)^6 \stackrel{(C)}{\leq} C_1 C_2 (\log n)^7. \quad (3.17)$$

□

Proof of condition 2 Given a path P on $\overline{\mathcal{Q}}$ of length m , we write it as $\{\tilde{\lambda}(k_i)\}_{1 \leq i \leq m}$. In addition, if $\tilde{\lambda}(k_1) \in \mathcal{F}(M)$ (resp. $\tilde{\lambda}(k_m) \in \mathcal{F}(M)$), we add $\psi(k_1)$ (resp. $\psi(k_m)$) at the beginning (resp. end) of P . Denote this possibly augmented path by \overline{P} . Now we define a path $\Pi(P)$ on M_n in two steps:

1. for any two consecutive faces f_1, f_2 of M , insert a vertex $v \in \mathcal{V}(M_n)$ which is an endpoint of the edges shared by f_1, f_2 . Denote this path by \overline{P}' .
2. For any three consecutive vertices v_1, f, v_2 on \overline{P}' so that $v_1, v_2 \in \mathcal{V}(M)$ and $f \in \mathcal{F}(M)$, replace the segment $[v_1, f, v_2]$ with an arbitrary arc (say, clockwise) on f between v_1 and v_2 .

For each $i_1, i_2 \in [-n, n]_{\mathbb{Z}}$ with $i_1 \sim i_2$ in \mathcal{G} , set $P_{i_1, i_2}^M := \Pi(\tilde{P}_{i_1, i_2}^{\overline{\mathcal{Q}}})$. Then P_{i_1, i_2}^M connects $\psi_n(i_1)$ and $\psi_n(i_2)$ and

$$|P_{i_1, i_2}^M| \stackrel{(B)}{\leq} \max_{f \in \mathcal{F}(M_n)} (\#\text{Nb}_f) \cdot 2C_1 (\log n)^3 \stackrel{(C)}{\leq} 2C_1 C_2 (\log n)^7. \quad (3.18)$$

For each $v \in \mathcal{V}(M_n)$, in order for v to be hit by some P_{i_1, i_2}^M , there must exist a face in Nb_v that is hit by some path $\tilde{P}_{i_1, i_2}^{\overline{\mathcal{Q}}}$. By (B), the number of paths of the form P_{i_1, i_2}^M hitting v is bounded by

$$|\text{Nb}_v| \cdot C_1 (\log n)^6 \stackrel{(C)}{\leq} C_1 C_2 (\log n)^7. \quad (3.19)$$

□

Proof of condition 3 For all $v \in \mathcal{V}(M_n)$, we can check case by case that v and $\psi(\phi_n(v))$ are both on $\text{Nb}_{\tilde{\lambda}(\phi_n(v))}$. Therefore by (C) $\text{dist}(\psi(\phi_n(v)), v; M_n) \leq \max_{f \in \mathcal{F}(M_n)} |\text{Nb}_f| \leq C_2 \log n$.

Now we proceed to bound $\text{dist}(\phi(\psi(i)), i; \mathcal{G}_n)$ for $i \in [-n, n]_{\mathbb{Z}}$. If $\tilde{\lambda}(i) \in \mathcal{V}(M_n)$, then (C) implies that $\text{dist}(\tilde{\lambda}(\phi_n(\psi(i))), \tilde{\lambda}(i); \overline{\mathcal{Q}}_n) = \text{dist}(\tilde{\lambda}(\phi_n(\lambda(i))), \tilde{\lambda}(i); \overline{\mathcal{Q}}_n) \leq C_2 \log n$. If $\tilde{\lambda}(i) \in \mathcal{F}(M_n)$, then $\tilde{\lambda}(\phi_n(\psi(i)))$ and $\tilde{\lambda}(i)$ are both in $\text{Nb}_{\psi(i)} \cap \mathcal{F}(M_n)$. By Lemma 3.26 and (C), we have $\text{dist}(\tilde{\lambda}(\phi_n(\psi(i))), \tilde{\lambda}(i); \overline{\mathcal{Q}}_n) \leq C_2 \log n$. Now (A) implies that $\text{dist}(\phi_n(\psi(i)), i; \mathcal{G}_n) \leq C_2 \log n \cdot C_1 (\log n)^3 = C_1 C_2 (\log n)^4$. \square

3.3.4 Bipolar orientations with other weights and Schnyder woods

The only properties of $\mathcal{Z} = (\mathcal{L}, \mathcal{R})$ that we used in Sects. 3.3.2 and 3.3.3 are as follows:

$$\{\Delta \mathcal{Z}_i\}_{i \in \mathbb{Z}} \text{ are i.i.d with } \mathbb{E}[\Delta \mathcal{Z}_i] = (0, 0) \text{ and the correlation of } \mathcal{L} \text{ and } \mathcal{R} \text{ is in } (-1, 1) \quad (3.20)$$

The support of $\Delta \mathcal{Z}_1$ is contained in (3.7) and $\mathbb{E}[e^{b|\Delta \mathcal{Z}_1|}] < \infty$ for some $b > 0$. $\quad (3.21)$

In particular, as long as \mathcal{Z} satisfies (3.20)–(3.21), we can construct an infinite planar map associated with \mathcal{Z} , analogous to the UIBOM, via the sewing procedure. This gives a large family of measures on infinite bipolar-oriented maps. By [46, Theorem 6, Remark 2] and the argument for Proposition 3.16, this family of infinite bipolar-oriented maps includes the Benjamini–Schramm limits of the uniform measures on size n bipolar-oriented k -angulations⁸ for any $k \geq 3$. More generally, one can consider Boltzmann-type distributions where each face has a certain weight. We note that the above family includes maps encoded by walks with any correlation in $(-1, 0)$, which corresponds to $\kappa > 8$ and $\gamma < \sqrt{2}$.

The exact same argument as in Sect. 3.3.3 shows that Theorem 1.9 still holds for this family of maps (with M_n, ϕ_n , and ψ_n defined in exactly the same way and γ determined by $\text{corr}(\mathcal{L}_1, \mathcal{R}_1) = -\cos(\pi \gamma^2/4)$). We remark that in Sect. 3.3.3 we sometimes rotate the plane 180 degree and switch $\mathcal{Z}_t = (\mathcal{L}_t, \mathcal{R}_t)$ to $(\mathcal{R}_{-t}, \mathcal{L}_{-t})$. But we never need the fact that $(\mathcal{L}_t, \mathcal{R}_t) \stackrel{d}{=} (\mathcal{R}_{-t}, \mathcal{L}_{-t})$. All we need is that the walk $(\mathcal{R}_{-t}, \mathcal{L}_{-t})$ also satisfies (3.20)–(3.21), so we do not have to impose this kind of symmetry condition on $\Delta \mathcal{Z}$ to ensure our generalization of Sect. 3.3.3. Thus, we have established Theorem 1.9 in case 4.

As an example, suppose \mathcal{Z} is a bi-infinite random walk satisfies (3.20)–(3.21) with step distribution

$$\nu(-1, 1) = 2^{-1} \quad \text{and} \quad \nu(1, -j) = 2^{-j-2} \quad \text{for } j \geq 0. \quad (3.22)$$

Let $M^{\mathcal{Z}}$ be the infinite map produced by \mathcal{Z} via the sewing procedure. For any positive integer n we apply the finite sewing procedure to the walk $\{\mathcal{Z}_t\}_{0 \leq t \leq n-1}$. We condition on $\{\mathcal{Z}_t\}_{0 \leq t \leq n-1}$ starting at $(0, 0)$, ending at $(0, 1)$, and staying in $\mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}$. This produces a uniform sample of an n -edge bipolar-oriented map with the following property:

⁸ The infinite bipolar-oriented triangulation has been considered in [34], where the authors shows that the triangulation and its dual map jointly converge to two SLE₁₂’s coupled in imaginary geometry.

The west (resp. east) boundary length equals 1 (resp. 2),
and each bounded face has exactly 2 east edges. (3.23)

By the same argument as in Proposition 3.16, the Benjamini–Schramm limit as $n \rightarrow \infty$ is $M^{\mathcal{Z}}$.

We choose to elaborate on the example of $M^{\mathcal{Z}}$ because of its close relation to Schnyder wood-decorated random planar maps. Following the notations in [51], consider a simple plane triangulation (i.e. without self-loops or multiple edges) M with unbounded face $\triangle A_b A_r A_g$, where the *outer vertices* A_b , A_r , and A_g are arranged in counterclockwise order. We denote by $\overline{A_r A_g}$ the edge connecting A_r and A_g , and similarly for $\overline{A_g A_b}$ and $\overline{A_b A_r}$. Vertices, edges, and faces of M other than A_b , A_r , A_g , $\overline{A_r A_g}$, $\overline{A_g A_b}$, $\overline{A_b A_r}$ and $\triangle A_b A_r A_g$ are called *inner* vertices, edges, and faces respectively. A *3-orientation* on M is an orientation on inner edges where every inner vertex has out-degree 3. Given a 3-orientation \mathcal{O} on M , the following algorithm, gives a way of coloring the inner edges in blue, red and green:

1. Color A_b , A_r , A_g blue, red, and green.
2. For an inner edge e , set $e_1 = e$, and inductively for $k \geq 1$ do the following: if the head of e_k is an outer vertex, we stop. Otherwise, let e_{k+1} be the second outgoing edge encountered when clockwise (or equivalently, counterclockwise) rotating e_k about its head. This procedure always terminates in finitely many steps and yields a path $\mathcal{P} = [e_1, e_2, \dots, e_\ell]$.
4. Assign to e the color of the outer vertex at which \mathcal{P} terminates.

It is elementary to check (see e.g. [51, Introduction]) that the set of blue edges forms a tree T_b spanning all inner vertices plus A_b , and similar statements hold for red and green edges, where we denote the corresponding trees by T_r and T_g respectively. $\{T_b, T_r, T_g\}$ is called the *Schnyder wood* associated with the 3-orientation \mathcal{O} . In this paper we treat the notion of 3-orientation and Schnyder wood interchangeably and call (M, \mathcal{O}) a *Schnyder wood-decorated triangulation*, or simply a wooded triangulation. See Fig. 19 for an illustration.

The Schnyder wood was first introduced by Walter Schnyder [59] to prove a characterization of graph planarity. He later used Schnyder woods to describe an efficient algorithm for embedding planar graph in such a way that its edges are straight lines and its vertices lie on a grid [60]. Schnyder's celebrated construction has continued to play an important role in graph theory and enumerative combinatorics [11, 32]. It is proved in [51] that uniformly wooded triangulation converges to a triple of SLE₁₆ curves on Liouville quantum gravity with $\gamma = 1$ in the peanosphere topology. Although inspired by [51], arguments in this paper do not rely on any ingredient from [51], thus can be read independently.

Suppose M has n inner vertices. Let M' be the submap of M with all of the green edges removed. Reverse the orientation of each red edge, and give the outer edges the orientations $A_r A_b$, $A_r A_g$ and $A_g A_b$. It is explained in [29, Proposition 7.1] that this operation gives a bipolar orientation (M', \mathcal{O}') and is a bijection between the set of wooded triangulations with n inner vertices and the set of bipolar-oriented maps with $2n + 3$ edges satisfying (3.23). Moreover, to recover M from M' , one just needs to connect an edge between any pair vertices v, w where




Fig. 19 On the left is the output of the coloring algorithm for a 3-orientation on a simple triangulation, (M, \mathcal{O}) . On the right is the bipolar-oriented map (M', \mathcal{O}') obtained by deleting all green edges and reversing the direction of red edges. We see that (M', \mathcal{O}') satisfies (3.23) (color figure online)

$$\begin{aligned} v, w \text{ lie on opposite sides of a bounded face } f \text{ of } M' \\ \text{so that neither } v \text{ nor } w \text{ is a pole of } f. \end{aligned} \quad (3.24)$$

These edges form the green tree T_g that we deleted. Then we can reverse the lower right edge of each face to recover the red tree. In light of this bijection, given a sample of $M^{\mathcal{Z}}$, (deterministically) connecting all such pairs in $\mathcal{V}(M^{\mathcal{Z}})$ as in (3.24), we obtain a sample of the Benjamini–Schramm limit of a uniform wooded triangulation with n inner vertices, which we call the *uniform infinite wooded triangulation* or UIWT.⁹ This provides an alternative proof of the existence of UIWT other than [51, Proposition 2.9].

Proof of Theorem 1.9 and Lemma 1.12 for the UIWT Define $\psi_n, \phi_n, \tilde{\lambda}, \text{Nb}_\bullet, M_n^{\mathcal{Z}}$ in the same way as in Sects. 3.3.2–3.3.3 with M replaced by $M^{\mathcal{Z}}$. Since \mathcal{Z} satisfies (3.20)–(3.21), the same argument for case 3 of Lemma 1.12, Lemma 3.20, and Theorem 1.9 yields the analogs of these results for $M^{\mathcal{Z}}$. Now let M be a sample of UIWT generated by adding green edges of M to $M^{\mathcal{Z}}$ as in (3.24), with $M^{\mathcal{Z}}$ in place of M' . Let M_n be obtained from $M_n^{\mathcal{Z}}$ by adding edges as in (3.24) with $M_n^{\mathcal{Z}}$ in place of M' . Since $\mathcal{V}(M) = \mathcal{V}(M^{\mathcal{Z}})$, we see that ψ_n maps $[-n, n]_{\mathbb{Z}}$ to $\mathcal{V}(M_n)$ and ϕ_n maps $\mathcal{V}(M_n)$ to $[-n, n]_{\mathbb{Z}}$.

Since $\mathcal{V}(\partial M_n^{\mathcal{Z}}) = \mathcal{V}(\partial M_n)$ for each $n \in \mathbb{N}$, Lemma 1.12 in case 5 follows from the same statement for $M^{\mathcal{Z}}$. Moreover, conditions 2 and 3 in Theorem 1.9 hold for $M_n^{\mathcal{Z}}$ implies that the same holds for M_n . So it only remains to prove condition 1. In this case we work on the event that (A)–(C) in Sect. 3.3.3 hold, with $M^{\mathcal{Z}}$ in place of the UIBOM. In this case the only additional case to consider is when v_1, v_2 are connected by a green edge in M_n . By (3.24) we can find a face f incident to v_1, v_2

⁹ If we take a Benjamini–Schramm limit of a uniform wooded triangulation rooted at typical edge, then the rooted edge is colored blue, red or green with equal probability. But if we first sample $M^{\mathcal{Z}}$ then obtain UIWT by adding edges as in (3.24), the rooted edge must be blue or red. However, if we think of the UIWT as a vertex rooted map whose vertex is the tail of the root edge, since each vertex has a unique outgoing edge of each color, these two ways of generating UIWT are equivalent.

so that $|\tilde{\lambda}^{-1}(f)| \leq n$. Therefore the exact same proof of condition 1 for the UIBOM still works. \square

Proof of Theorems 1.6 and 1.7 in cases 3 through 5 This follows from Theorem 1.9, Lemma 1.11, and the results of [35], exactly as in case 1. \square

4 Open questions

The proofs of our main results extend to any random planar map which can be encoded by a peanosphere-type bijection for which the encoding random walk can be sufficiently closely approximated by Brownian motion. This leads us to ask whether there are any more random planar map models, aside from the ones listed in Sect. 1.4, for which this is the case.

Question 4.1 *Is there an analogue of the strong coupling theorem [47,64] for random walk and Brownian motion for the non-Markovian random walk appearing in Sheffield's hamburger-cheeseburger bijection [62] for $p \neq 0$, or for the non-Markovian random walk appearing in the generalized hamburger-cheeseburger model of [36]?*

An affirmative answer to Question 4.1 would enable us to generalize the results of this paper to critical Fortuin-Kasteleyn planar maps for $q \in (0, 4)$ which are encoded by Sheffield's hamburger cheeseburger bijection and belong to the γ -LQG universality class for $\gamma \in (\sqrt{2}, 2)$; and to the planar maps decorated by active spanning trees with bending energy which are encoded by the model of [36] and which belong to the γ -LQG universality class for $\gamma \in (0, \sqrt{2})$.

The results of the present paper can be easily generalized to other random planar map models for which the corresponding encoding walk has i.i.d. increments since the needed strong coupling result comes from [64].

Question 4.2 *What random planar map models, other than the ones listed in Sect. 1.4, can be encoded by a mating-of-trees bijection where the random walk has i.i.d. increments?*

It is natural to ask for a bijective encoding of a planar map decorated by a variant of the Ising model (or some other decorated random planar map model in the $\sqrt{3}$ -LQG universality class) by a random walk with i.i.d. increments, since this model is in some sense the simplest statistical mechanics model for which no such encoding is currently known. Note that the hamburger-cheeseburger bijection encodes a planar map decorated by an instance of the FK Ising model, but the walk in this case does not have i.i.d. increments.

A potential candidate for the encoding walk in such a hypothetical bijection is the *Gessel walk* [17,19,44], whose increment distribution is uniform on $\{(-1, 0), (1, 0), (1, 1), (-1, -1)\}$. This walk is exactly solvable and its coordinates have correlation $1/\sqrt{2} = -\cos(\pi(\sqrt{3})^2/4)$, which is the same as the correlation of the Brownian motion used to define the mated-CRT map when $\gamma = \sqrt{3}$.

We also note that the results of the present paper provide additional motivation to prove stronger estimates for distances in the mated-CRT map \mathcal{G} (improving on the results of [9, 24, 35, 41]) since these results can be immediately transferred to other random planar map models by means of Theorem 1.5.

Acknowledgements We thank an anonymous referee for helpful comments on an earlier version of this article. We thank Jason Miller for helpful discussions. E.G. was partially funded by NSF Grant DMS 1209044. N.H. was supported by a fellowship from the Norwegian Research Council. X.S. was supported by the Simons Foundation as a Junior Fellow at Simons Society of Fellows.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Ambjørn, J., Budd, T.G.: Geodesic distances in Liouville quantum gravity. *Nucl. Phys. B* **889**, 676–691 (2014). [arXiv:1405.3424](https://arxiv.org/abs/1405.3424)
2. Abrams, A., Kenyon, R.: Fixed-energy harmonic functions. ArXiv e-prints, May (2015) [arXiv:1505.05785](https://arxiv.org/abs/1505.05785)
3. Andres, S., Kajino, N.: Continuity and estimates of the Liouville heat kernel with applications to spectral dimensions. *Probab. Theory Relat. Fields* **166**(3–4), 713–752 (2016). [arXiv:1407.3240](https://arxiv.org/abs/1407.3240)
4. Aldous, D.: The continuum random tree. I. *Ann. Probab.* **19**(1), 1–28 (1991)
5. Aldous, D.: The continuum random tree. II. An overview. In: *Stochastic analysis* (Durham, 1990), volume 167 of London Math. Soc. Lecture Note Ser. (pp. 23–70). Cambridge University Press, Cambridge (1991)
6. Aldous, D.: The continuum random tree. III. *Ann. Probab.* **21**(1), 248–289 (1993)
7. Angel, O.: Growth and percolation on the uniform infinite planar triangulation. *Geom. Funct. Anal.* **13**(5), 935–974 (2003)
8. Angel, O.: Scaling of Percolation on Infinite Planar Maps, I. ArXiv Mathematics e-prints, December (2005) [arXiv:math/0501006](https://arxiv.org/abs/math/0501006)
9. Ang, M.: Comparison of discrete and continuum Liouville first passage percolation. ArXiv e-prints, Apr (2019) [arXiv:1904.09285](https://arxiv.org/abs/1904.09285)
10. Angel, O., Schramm, O.: Uniform infinite planar triangulations. *Commun. Math. Phys.* **241**(2–3), 191–213 (2003)
11. Bernardi, O., Bonichon, N.: Intervals in Catalan lattices and realizers of triangulations. *J. Combin. Theory Ser. A* **116**(1), 55–75 (2009)
12. Bonichon, N., Bousquet-Mélou, M., Fusy, É.: Baxter permutations and plane bipolar orientations. *Sém. Lothar. Combin.*, 61A:Art. B61Ah, 29, (2009/11), [arXiv:0805.4180](https://arxiv.org/abs/0805.4180)
13. Bouttier, J., Di Francesco, P., Guitter, E.: Planar maps as labeled mobiles. *Electron. J. Combin.* **11**(1):Research Paper 69, 27 (2004), [arXiv:math/0405099](https://arxiv.org/abs/math/0405099)
14. Bernardi, O.: Bijective counting of Kreweras walks and loopless triangulations. *J. Combin. Theory Ser. A* **114**(5), 931–956 (2007)
15. Bernardi, O.: Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Electron. J. Combin.*, **14**(1):Research Paper 9, 36 pp. (electronic) (2007) [arXiv:math/0601684](https://arxiv.org/abs/math/0601684)
16. Bernardi, O., Holden, N., Sun, X.: Percolation on triangulations: a bijective path to Liouville quantum gravity. ArXiv e-prints, July (2018) [arXiv:1807.01684](https://arxiv.org/abs/1807.01684)
17. Bostan, A., Kurkova, I., Raschel, K.: A human proof of Gessel's lattice path conjecture. *Trans. Am. Math. Soc.* **369**(2), 1365–1393 (2017). [arXiv:1309.1023](https://arxiv.org/abs/1309.1023)

18. Berestycki, N., Laslier, B., Ray, G.: Critical exponents on Fortuin–Kasteleyn weighted planar maps. *Commun. Math. Phys.* **355**(2), 427–462 (2017). [arXiv:1502.00450](https://arxiv.org/abs/1502.00450)
19. Bousquet-Mélou, M.: An elementary solution of Gessel’s walks in the quadrant. *Adv. Math.* **303**, 1171–1189 (2016). [arXiv:1503.08573](https://arxiv.org/abs/1503.08573)
20. Benjamini, I., Schramm, O.: Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.* **6**(23), 13 (electronic) (2001). [arXiv:math/0011019](https://arxiv.org/abs/math/0011019)
21. Chen, L.: Basic properties of the infinite critical-FK random map. *Ann. Inst. Henri Poincaré D* **4**(3), 245–271 (2017). [arXiv:1502.01013](https://arxiv.org/abs/1502.01013)
22. de Fraysseix, H., de Mendez, P.O., Rosenstiehl, P.: Bipolar orientations revisited. *Discrete Appl. Math.* **56**(2–3), 157–179 (1995). Special Issue: Fifth Franco-Japanese Days (Kyoto, 1992)
23. Ding, J., Goswami, S.: Upper bounds on Liouville first passage percolation and Watabiki’s prediction. *Commun. Pure Appl. Math.*, to appear (2016) [arXiv:1610.09998](https://arxiv.org/abs/1610.09998)
24. Ding, J., Gwynne, E.: The fractal dimension of Liouville quantum gravity: universality, monotonicity, and bounds. *Commun. Math. Phys.*, to appear (2018) [arXiv:1807.01072](https://arxiv.org/abs/1807.01072)
25. Duplantier, B., Miller, J., Sheffield, S.: Liouville quantum gravity as a mating of trees. ArXiv e-prints, September (2014) [arXiv:1409.7055](https://arxiv.org/abs/1409.7055)
26. Denisov, D., Wachtel, V.: Random walks in cones. *Ann. Probab.* **43**(3), 992–1044 (2015). [arXiv:1110.1254](https://arxiv.org/abs/1110.1254)
27. Duraj, J., Wachtel, V.: Invariance principles for random walks in cones. ArXiv e-prints, August (2015) [arXiv:1508.07966](https://arxiv.org/abs/1508.07966)
28. Ding, J., Zeitouni, O., Zhang, F.: Heat kernel for Liouville Brownian motion and Liouville graph distance. *Commun. Math. Phys.*, to appear (2018) [arXiv:1807.00422](https://arxiv.org/abs/1807.00422)
29. Felsner, S., Fusy, E., Noy, M., Orden, D.: Bijections for Baxter families and related objects. *J. Combin. Theory Ser. A* **118**(3), 993–1020 (2011)
30. Fortuin, C.M., Kasteleyn, P.W.: On the random-cluster model. I. Introduction and relation to other models. *Physica* **57**, 536–564 (1972)
31. Fusy, É., Poulalhon, D., Schaeffer, G.: Bijective counting of plane bipolar orientations and Schnyder woods. *Eur. J. Combin.* **30**(7), 1646–1658 (2009). [arXiv:0803.0400](https://arxiv.org/abs/0803.0400)
32. Felsner, S., Zickfeld, F.: Schnyder woods and orthogonal surfaces. *Discrete Comput. Geom.* **40**(1), 103–126 (2008)
33. Gwynne, E., Hutchcroft, T.: Anomalous diffusion of random walk on random planar maps. ArXiv e-prints July (2018) [arXiv:1807.01512](https://arxiv.org/abs/1807.01512)
34. Gwynne, E., Holden, N., Sun, X.: Joint scaling limit of a bipolar-oriented triangulation and its dual in the peanosphere sense. ArXiv e-prints, March (2016) [arXiv:1603.01194](https://arxiv.org/abs/1603.01194)
35. Gwynne, E., Holden, N., Sun, X.: A distance exponent for Liouville quantum gravity. *Probab. Theory Relat. Fields* **173**(3), 931–997 (2019). [arXiv:1606.01214](https://arxiv.org/abs/1606.01214)
36. Gwynne, E., Kassel, A., Miller, J., Wilson, D.B.: Active spanning trees with bending energy on planar maps and SLE-decorated liouville quantum gravity for $\kappa > 8$. *Commun. Math. Phys.* **358**(3), 1065–1115 (2018). [arXiv:1603.09722](https://arxiv.org/abs/1603.09722)
37. Gwynne, E., Miller, J.: Random walk on random planar maps: spectral dimension, resistance, and displacement. ArXiv e-prints, November (2017) [arXiv:1711.00836](https://arxiv.org/abs/1711.00836)
38. Gwynne, E., Miller, J.: Existence and uniqueness of the Liouville quantum gravity metric for $\gamma \in (0, 2)$. ArXiv e-prints, May (2019) [arXiv:1905.00383](https://arxiv.org/abs/1905.00383)
39. Gwynne, E., Miller, J., Sheffield, S.: The Tutte embedding of the mated-CRT map converges to Liouville quantum gravity. ArXiv e-prints, May (2017) [arXiv:1705.11161](https://arxiv.org/abs/1705.11161)
40. Gwynne, E., Mao, C., Sun, X.: Scaling limits for the critical Fortuin–Kasteleyn model on a random planar map I: cone times. *Ann. Inst. Henri Poincaré Probab. Stat.* **55**(1), 1–60 (2019). [arXiv:1502.00546](https://arxiv.org/abs/1502.00546)
41. Gwynne, E., Pfeffer, J.: Bounds for distances and geodesic dimension in Liouville first passage percolation. *Electron. Commun. Probab.* **24**(56), 12 (2019). [arXiv:1903.09561](https://arxiv.org/abs/1903.09561)
42. Gwynne, E., Pfeffer, J.: External diffusion limited aggregation on a spanning-tree-weighted random planar map. ArXiv e-prints, January (2019) [arXiv:1901.06860](https://arxiv.org/abs/1901.06860)
43. Gwynne, E., Pfeffer, J.: KPZ formulas for the Liouville quantum gravity metric. *Trans. Am. Math. Soc.* to appear (2019)
44. Kauers, M., Koutschan, C., Zeilberger, D.: Proof of Ira Gessel’s lattice path conjecture. *Proc. Natl. Acad. Sci.* **106**, 11502–11505 (2009). [arXiv:0806.4300](https://arxiv.org/abs/0806.4300)
45. Kenyon, R., Miller, J., Sheffield, S., Wilson, D.B.: The six-vertex model and Schramm–Loewner evolution. ArXiv e-prints, May (2016) [arXiv:1605.06471](https://arxiv.org/abs/1605.06471)

46. Kenyon, R., Miller, J., Sheffield, S., Wilson, D.B.: Bipolar orientations on planar maps and SLE₁₂. *Ann. Probab.* **47**(3), 1240–1269 (2019). [arXiv:1511.04068](https://arxiv.org/abs/1511.04068)
47. Komlós, J., Major, P., Tusnády, G.: An approximation of partial sums of independent RV's, and the sample DF. II. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete* **34**(1), 33–58 (1976)
48. Le Gall, J.-F.: Random trees and applications. *Probab. Surv.* **2**, 245–311 (2005). [arXiv:math/0511151](https://arxiv.org/abs/math/0511151)
49. Le Gall, J.-F.: Uniqueness and universality of the Brownian map. *Ann. Probab.* **41**(4), 2880–2960 (2013). [arXiv:1105.4842](https://arxiv.org/abs/1105.4842)
50. Lawler, G.F., Limic, V.: *Random Walk: A Modern Introduction*. Cambridge Studies in Advanced Mathematics, vol. 123. Cambridge University Press, Cambridge (2010)
51. Li, Y., Sun, X., Watson, S.S.: Schnyder woods, SLE(16), and Liouville quantum gravity. ArXiv e-prints, May (2017). [arXiv:1705.03573](https://arxiv.org/abs/1705.03573)
52. Miermont, G.: The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.* **210**(2), 319–401 (2013). [arXiv:1104.1606](https://arxiv.org/abs/1104.1606)
53. Maillard, P., Rhodes, R., Vargas, V., Zeitouni, O.: Liouville heat kernel: regularity and bounds. *Ann. Inst. Henri Poincaré Probab. Stat.* **52**(3), 1281–1320 (2016). [arXiv:1406.0491](https://arxiv.org/abs/1406.0491)
54. Miller, J., Sheffield, S.: Liouville quantum gravity and the Brownian map I: the QLE(8/3,0) metric. *Inventiones Mathematicae*, to appear (2015) [arXiv:1507.00719](https://arxiv.org/abs/1507.00719)
55. Miller, J., Sheffield, S.: Liouville quantum gravity and the Brownian map II: geodesics and continuity of the embedding. ArXiv e-prints, May (2016) [arXiv:1605.03563](https://arxiv.org/abs/1605.03563)
56. Miller, J., Sheffield, S.: Liouville quantum gravity and the Brownian map III: the conformal structure is determined. ArXiv e-prints, August (2016) [arXiv:1608.05391](https://arxiv.org/abs/1608.05391)
57. Miller, J., Sheffield, S.: Imaginary geometry IV: interior rays, whole-plane reversibility, and space-filling trees. *Probab. Theory Relat. Fields* **169**(3–4), 729–869 (2017). [arXiv:1302.4738](https://arxiv.org/abs/1302.4738)
58. Mullin, R.C.: On the enumeration of tree-rooted maps. *Can. J. Math.* **19**, 174–183 (1967)
59. Schnyder, W.: Planar graphs and poset dimension. *Order* **5**(4), 323–343 (1989)
60. Schnyder, W.: Embedding planar graphs on the grid. In: *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '90, pp. 138–148, Philadelphia, PA, USA (1990). Society for Industrial and Applied Mathematics
61. Schaeffer, G.: Bijective census and random generation of Eulerian planar maps with prescribed vertex degrees. *Electron. J. Combin.* **4**(1):Research Paper 20, 14 pp. (electronic) (1997)
62. Sheffield, S.: Quantum gravity and inventory accumulation. *Ann. Probab.* **44**(6), 3804–3848 (2016). [arXiv:1108.2241](https://arxiv.org/abs/1108.2241)
63. Watabiki, Y.: Analytic study of fractal structure of quantized surface in two-dimensional quantum gravity. *Prog. Theor. Phys. Suppl.* **114**, 1–17 (1993). Quantum gravity (Kyoto, 1992)
64. Zaitsev, A.Y.: Multidimensional version of the results of Komlós, Major and Tusnády for vectors with finite exponential moments. *ESAIM Probab. Stat.* **2**, 41–108 (1998)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.