

Scaling limit of the uniform infinite half-plane quadrangulation in the Gromov-Hausdorff-Prokhorov-uniform topology

Ewain Gwynne* Jason Miller†

Abstract

We prove that the uniform infinite half-plane quadrangulation (UIHPQ), with either general or simple boundary, equipped with its graph distance, its natural area measure, and the curve which traces its boundary, converges in the scaling limit to the Brownian half-plane. The topology of convergence is given by the so-called Gromov-Hausdorff-Prokhorov-uniform (GHPU) metric on curve-decorated metric measure spaces, which is a generalization of the Gromov-Hausdorff metric whereby two such spaces $(X_1, d_1, \mu_1, \eta_1)$ and $(X_2, d_2, \mu_2, \eta_2)$ are close if they can be isometrically embedded into a common metric space in such a way that the spaces X_1 and X_2 are close in the Hausdorff distance, the measures μ_1 and μ_2 are close in the Prokhorov distance, and the curves η_1 and η_2 are close in the uniform distance.

Keywords: random planar maps; Brownian map; Brownian half-plane; uniform infinite half-plane quadrangulation; Gromov-Hausdorff-Prokhorov-uniform topology.

AMS MSC 2010: 60D05; 60F17; 05C80.

Submitted to EJP on October 26, 2016, final version accepted on September 2, 2017.

Supersedes arXiv:1608.00954.

Contents

1 Introduction	2
1.1 Overview	2
1.2 Preliminary definitions	5
1.3 The Gromov-Hausdorff-Prokhorov-uniform metric	6
1.4 Basic definitions for quadrangulations	9
1.5 The Brownian half-plane	10
1.6 Theorem statements: scaling limit of the UIHPQ and UIHPQ _S	11
1.7 Outline	12

*Massachusetts Institute of Technology, United States. E-mail: ewain@mit.edu

†University of Cambridge, United Kingdom. E-mail: jpmiller@statslab.cam.ac.uk

2 Properties of the Gromov-Hausdorff-Prokhorov-uniform metric	13
2.1 Subsequential limits of isometries	13
2.2 Proofs for the GHPU metric	15
2.3 Proofs for the local GHPU metric	18
2.4 Conditions for GHPU convergence using the m -fold Gromov-Prokhorov topology	21
3 Schaeffer-type constructions	23
3.1 Encoding quadrangulations with boundary	23
3.2 Encoding the UIHPQ	26
3.3 Distance bounds for quadrangulations with boundary	27
3.4 Pruning the UIHPQ to get the UIHPQ _S	28
3.5 The Brownian disk	30
4 Scaling limit of the UIHPQ and UIHPQ_S	31
4.1 Convergence to the Brownian disk in the GHPU topology	31
4.2 Coupling the Brownian half-plane with the Brownian disk	33
4.3 Coupling the UIHPQ with a finite uniform quadrangulation with boundary	36
4.4 Proof of Theorem 1.11	40
4.5 Proof of Theorem 1.12	40
References	44

1 Introduction

1.1 Overview

There has been substantial interest in recent years in the scaling limits of random planar maps. Various uniform random planar maps (equipped with the graph distance) have been shown to converge in the Gromov-Hausdorff topology to *Brownian surfaces*, the best known of which is the Brownian map, which is the scaling limit of uniform random quadrangulations of the sphere [37, 41]. These results have been generalized in [1, 3, 9] to other ensembles of random maps on the sphere and in [18] (resp. [10]) to give the convergence of the uniform infinite plane quadrangulation (resp. uniformly random quadrangulations with boundary) toward the Brownian plane (resp. disk).

A planar map is naturally endowed with a measure μ (e.g., the one which assigns mass to each vertex equal to its degree). Many interesting random planar maps M are also equipped with a curve η . Examples include:

1. The path which visits the boundary ∂M in cyclic order of a planar map M with boundary.
2. A simple random walk or self-avoiding walk (SAW) on M .
3. The Peano curve associated with a distinguished spanning tree of M .
4. The exploration path associated with a percolation configuration on M .

Hence it is natural to consider scaling limits of random planar maps in a topology which describes not only their metric structure but also a distinguished measure and curve.

This article has two main aims. First, we will introduce such a topology, which arises from the *Gromov-Hausdorff-Prokhorov-uniform (GHPU) metric* on 4-tuples (X, d, μ, η) consisting of a metric space (X, d) , a measure μ on X , and a curve η in X . Two such 4-tuples $(X_1, d_1, \mu_1, \eta_1)$ and $(X_2, d_2, \mu_2, \eta_2)$ are close in the GHPU metric if they can be isometrically embedded into a common metric space (W, D) in such a way that X_1 and X_2

are close in the D -Hausdorff distance, μ_1 and μ_2 are close in the D -Prokhorov distance, and η_1 and η_2 are close in the D -uniform distance. We will consider a version of the GHPU metric for compact spaces as well as a local version for locally compact spaces. See Section 1.3 for a precise definition.

The definition of the GHPU metric is inspired by other metrics on types of metric spaces such as the Gromov-Hausdorff metric [14, 24], the Gromov-Prokhorov metric [23], and the Gromov-Hausdorff-Prokhorov metric [2, 40].

Second, we will prove scaling limit results for the uniform infinite half-plane quadrangulation (UIHPQ) in the local GHPU topology. The UIHPQ is the Benjamini-Schramm local limit [5] of uniform random quadrangulations with boundary as the total number of edges and then the perimeter tends to ∞ [15, 19], where the map is viewed from a root which is chosen uniformly at random from the boundary. There are two variants of the UIHPQ. The first is the UIHPQ with general boundary (which we will refer to as the UIHPQ), which may have boundary vertices with multiplicity greater than 1 in the external face; and the UIHPQ with simple boundary (UIHPQ_S), where we require that the boundary is simple (i.e., it is a path with no self-intersections). In this paper, we will prove that both the UIHPQ and the UIHPQ_S (equipped with the measure which assigns mass to each vertex equal to its degree and the curve which traces the boundary) converge in the scaling limit in the local GHPU topology to the Brownian half-plane, which we define in Section 1.5 below (see also [15, Section 5.3] for a different definition, which we expect is equivalent). Along the way, we will also improve the Gromov-Hausdorff scaling limit result for finite uniform quadrangulations with boundary toward the Brownian disk in [10] to a scaling limit result in the GHPU topology.

One particular reason to be interested in random quadrangulations with simple boundary (such as the UIHPQ_S) is that one can glue two such surfaces along their boundary to obtain a uniform random quadrangulation decorated by a SAW. See [8, Section 8.2] (which builds on [11, 13]) for the case of finite quadrangulations with simple boundary and [17, Part III], [16] for the case of the UIHPQ_S.

In [29], we will build upon the present work to prove, among other things, that the random planar map obtained by gluing a pair of independent UIHPQ_S's together along the boundary rays lying to the right of their respective root edges (i.e., the uniform infinite SAW-decorated half-plane) converges in the scaling limit in the GHPU topology, with the SAW playing the role of the distinguished curve, to a pair of independent Brownian half-planes glued together in the same way. We will also prove analogous scaling limit results for two independent UIHPQ_S's glued along their entire boundary and for a single UIHPQ_S with its positive and negative boundary rays glued together. The proofs of these results use the scaling limit statement for the UIHPQ_S proven in the present paper. See also [31, 32] for additional GHPU scaling limit results.

Remark 1.1. In an independent (and essentially simultaneous) work [4], Baur, Miermont, and Ray proved several scaling limit results for uniform quadrangulations with general boundary which include the statement that the UIHPQ with general boundary converges in the scaling limit to the Brownian half-plane in the Gromov-Hausdorff topology [4, Theorem 3.6]. The work [4] also includes a number of more general scaling limit statements for uniform random quadrangulations with boundary under different scaling regimes that we do not treat here. In the present paper we will deduce the scaling limit of the UIHPQ to the Brownian half-plane in a stronger topology than in [4] and also treat the case of the UIHPQ_S. Our proof is somewhat simpler than that in [4] since our coupling statement is less general.

We will now explain how the aforementioned results about scaling limits of glued UIHPQ_S's allow us to identify the scaling limit of the SAW on a random quadrangulation

Scaling limit of the UIHPQ in the GHPU topology

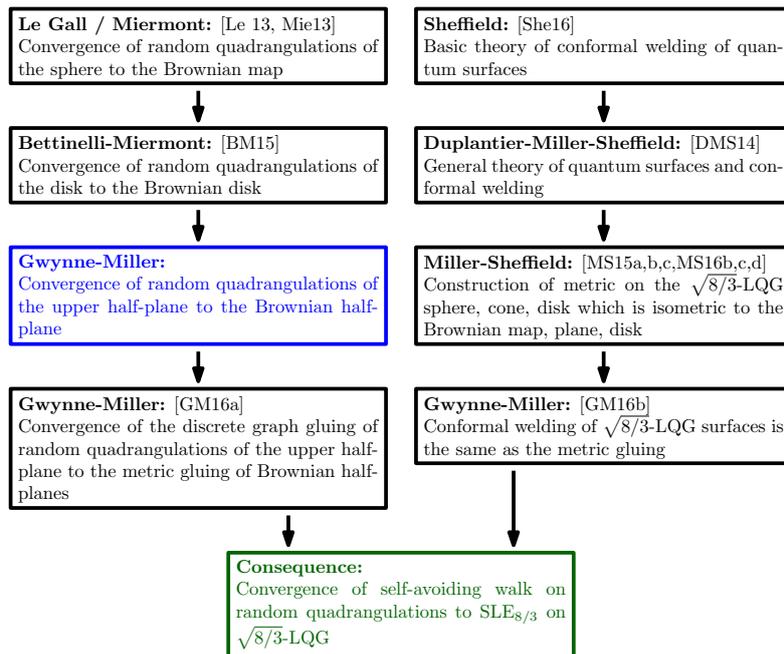


Figure 1: A chart of the different components which serve as input into the proof that self-avoiding walk on random quadrangulations converges to $SLE_{8/3}$ on $\sqrt{8/3}$ -LQG. The present article corresponds to the blue box and implies that a random quadrangulation of the upper half-plane converges in the GHPU topology to the Brownian half-plane. (See also [4] for another proof that the UIHPQ converges to the Brownian half-plane, and a more general treatment of scaling limits of quadrangulations with boundary.)

with $SLE_{8/3}$ on a $\sqrt{8/3}$ -Liouville quantum gravity (LQG) surface. Recently, it has been proven by Miller and Sheffield [42–46], building on [48], that Brownian surfaces are equivalent to $\sqrt{8/3}$ -LQG surfaces. Heuristically speaking, a γ -LQG surface for $\gamma \in (0, 2)$ is the random Riemann surface parameterized by a domain $D \subset \mathbb{C}$ whose Riemannian metric tensor is $e^{\gamma h} dx \otimes dy$, where $dx \otimes dy$ is the Euclidean metric tensor on D and h is some variant of the Gaussian free field (GFF) on D [21, 47, 50, 51]. This definition does not make rigorous sense since the GFF is a generalized function, not a function, so does not take values at points.

Miller and Sheffield showed that in the special case when $\gamma = \sqrt{8/3}$, one can make rigorous sense of a $\sqrt{8/3}$ -LQG surface as a metric space. Certain particular types of $\sqrt{8/3}$ -LQG surfaces introduced in [20], namely the quantum sphere, quantum disk, and weight-4/3 quantum cone, respectively, are isometric to the Brownian map, Brownian disk, and Brownian plane, respectively [45, Corollary 1.5]. In this paper we will extend this identification by proving that the Brownian half-plane is isometric to the weight-2 quantum wedge.

The results of [30] together with the identification between the Brownian half-plane and the weight-2 wedge proven in the present paper imply that the gluing of two Brownian half-planes along their positive boundary has the same law as a weight-4 quantum wedge (a particular type of $\sqrt{8/3}$ -LQG surface) decorated by an independent chordal $SLE_{8/3}$ curve [49], which is the gluing interface. Hence the scaling limit result of [29] discussed above yields the convergence of the SAW on a random quadrangulation to $SLE_{8/3}$ on a $\sqrt{8/3}$ -LQG surface. See Figure 1 for a diagram of how the different works fit together to establish this result.

LQG surfaces arise as the scaling limits of random planar maps for all values of $\gamma \in (0, 2)$, not just $\gamma = \sqrt{8/3}$. Values of γ other than $\sqrt{8/3}$ correspond to maps sampled with probability proportional to the partition function of some statistical mechanics model, rather than sampled uniformly. For general values of γ , certain random planar map models decorated by a space-filling curve, which is the Peano curve of a certain spanning tree, have been shown to converge to SLE-decorated LQG in the so-called *peanosphere topology*. This means that the joint law of the contour functions (or some variant thereof) of the spanning tree and its dual, appropriately re-scaled, converges to the law of the correlated Brownian motion which encodes a γ -LQG cone or sphere decorated by a space-filling SLE $_{16/\gamma^2}$ curve in [20, 44]. See [25, 27, 28, 33–35, 52] for results of this type.

Neither peanosphere convergence nor GHPU convergence implies the other. However, we expect that the curve-decorated planar maps which converge to SLE $_{16/\gamma^2}$ -decorated γ -LQG in the peanosphere topology also converge in the GHPU topology (this uses the γ -LQG metric space, which has so far only been constructed for $\gamma = \sqrt{8/3}$), and in fact converge in both topologies jointly. In the case of site percolation on a uniform triangulation (which corresponds to $\gamma = \sqrt{8/3}$), this joint GHPU/peanosphere convergence will be proven in the forthcoming work [26], building on [31] which shows GHPU convergence of a random planar map decorated by a single percolation interface. However, it remains open for other models.

Acknowledgements We thank two anonymous referees for helpful comments on an earlier version of this paper. E.G. was supported by the U.S. Department of Defense via an NDSEG fellowship. E.G. also thanks the hospitality of the Statistical Laboratory at the University of Cambridge, where this work was started. J.M. thanks Institut Henri Poincaré for support as a holder of the Poincaré chair, during which this work was completed.

1.2 Preliminary definitions

Before stating our main results, we set some standard notation and definitions which will be used throughout this paper.

1.2.1 Basic notation

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

For $a < b \in \mathbb{R}$, 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, we write $a \preceq b$ (resp. $a \succeq b$) if there is a constant C (independent of the parameters of interest) such that $a \leq Cb$ (resp. $a \geq Cb$). We write $a \asymp b$ if $a \preceq b$ and $a \succeq b$.

If f is a function, we write $a = o_b(f(b))$ if $a/f(b) \rightarrow 0$ as $b \rightarrow \infty$ or as $b \rightarrow 0$, depending on context. We write $a = O_b(f(b))$ if there is a constant $C > 0$, independent of the parameters of interest, such that $a \leq Cf(b)$.

1.2.2 Graphs

For a planar map G , we write $\mathcal{V}(G)$, $\mathcal{E}(G)$, and $\mathcal{F}(G)$, respectively, for the set of vertices, edges, and faces of G .

By a *path* in G , we mean a function $\lambda : I \rightarrow \mathcal{E}(G)$ for some (possibly infinite) discrete interval $I \subset \mathbb{Z}$, with the property that the edges $\{\lambda(i)\}_{i \in I}$ can be oriented in such a way that the terminal endpoint of $\lambda(i)$ coincides with the initial endpoint of $\lambda(i+1)$ for each $i \in I$ other than the right endpoint of I . We define the *length* of λ , denoted $|\lambda|$, to be the integer $\#I$. We say that λ is *simple* if the vertices hit by λ are all distinct.

For sets A_1, A_2 consisting of vertices and/or edges of G , we write $\text{dist}(A_1, A_2; G)$ for the graph distance from A_1 to A_2 in G , i.e. the minimum of the lengths of paths in G whose initial edge either has an endpoint which is a vertex in A_1 or shares an endpoint with an edge in A_1 ; and whose final edge satisfies the same condition with A_2 in place of A_1 .

For $r > 0$, we define the graph metric ball $B_r(A_1; G)$ to be the subgraph of G consisting of all vertices of G whose graph distance from A_1 is at most r and all edges of G whose endpoints both lie at graph distance at most r from A_1 . If $A_1 = \{x\}$ is a single vertex or edge, we write $B_r(\{x\}; G) = B_r(x; G)$.

1.2.3 Metric spaces

If (X, d) is a metric space, $A \subset X$, and $r > 0$, we write $B_r(A; d)$ for the set of $x \in X$ with $d(x, A) \leq r$. We emphasize that $B_r(A; d)$ is closed (this will be convenient when we work with the local GHPU topology). If $A = \{y\}$ is a singleton, we write $B_r(\{y\}; d) = B_r(y; d)$.

For a curve $\gamma : [a, b] \rightarrow X$, the d -length of γ is defined by

$$\text{len}(\gamma; d) := \sup_P \sum_{i=1}^{\#P} d(\gamma(t_i), \gamma(t_{i-1}))$$

where the supremum is over all partitions $P : a = t_0 < \dots < t_{\#P} = b$ of $[a, b]$. Note that the d -length of a curve may be infinite.

We say that (X, d) is a *length space* if for each $x, y \in X$ and each $\epsilon > 0$, there exists a curve of d -length at most $d(x, y) + \epsilon$ from x to y .

1.3 The Gromov-Hausdorff-Prokhorov-uniform metric

In this paper (and in [29]) we will consider scaling limits of metric measure spaces endowed with a distinguished continuous curve. A natural choice of topology for this convergence is the one induced by the *Gromov-Hausdorff-Prokhorov-uniform (GHPU) metric*, which we introduce in this subsection and study further in Section 2. This topology generalizes the Gromov-Hausdorff topology [14, 24], the Gromov-Prokhorov topology [23], and the Gromov-Hausdorff-Prokhorov topology [2, 40].

Let (X, d) be a metric space. The metric d gives rise to the d -Hausdorff metric \mathfrak{d}_d^H on compact subsets of X and the d -Prokhorov metric \mathfrak{d}_d^P on finite measures on X in the standard way.

The definition of the d -uniform metric on curves in X requires some discussion since we want to allow curves defined on an arbitrary interval in \mathbb{R} . Let $C_0(\mathbb{R}, X)$ be the set of continuous curves $\eta : \mathbb{R} \rightarrow X$ such that for each $\epsilon > 0$, there exists $T > 0$ such that $d(\eta(t), \eta(T)) \leq \epsilon$ and $d(\eta(-t), \eta(-T)) \leq \epsilon$ whenever $t \geq T$. If $\eta : [a, b] \rightarrow X$ is a curve defined on a compact interval, we identify η with the element of $C_0(\mathbb{R}, X)$ which agrees with η on $[a, b]$ and satisfies $\eta(t) = a$ for $t \leq a$ and $\eta(t) = b$ for $t \geq b$. We equip $C_0(\mathbb{R}, X)$ with the d -uniform metric, defined by

$$\mathfrak{d}_d^U(\eta_1, \eta_2) = \sup_{t \in \mathbb{R}} d(\eta_1(t), \eta_2(t)), \quad \forall \eta_1, \eta_2 \in C_0(\mathbb{R}, X). \tag{1.1}$$

Remark 1.2 (Graphs as connected metric spaces). In this paper we will often be interested in a graph G equipped with its graph distance d_G . In order to study continuous curves in G , we need to linearly interpolate G . We do this by identifying each edge of G with a copy of the unit interval $[0, 1]$. We extend the graph metric on G by requiring that this identification is an isometry.

If λ is a path in G , mapping some discrete interval $[a, b]_{\mathbb{Z}}$ to $\mathcal{E}(G)$, we extend λ from $[a, b]_{\mathbb{Z}}$ to $[a - 1, b]$ by linear interpolation, so that for $i \in [a, b]_{\mathbb{Z}}$, λ traces each edge $\lambda(i)$

at unit speed during the time interval $[i - 1, i]$. In particular, the Gromov-Hausdorff-Prokhorov-uniform metric and its local variant, to be defined below, make sense for graphs equipped with a measure and a curve.

1.3.1 Compact case

Let \mathbb{M}^{GHPU} be the set of 4-tuples $\mathfrak{X} = (X, d, \mu, \eta)$ where (X, d) is a compact metric space, d is a metric on X , μ is a finite Borel measure on X , and $\eta \in C_0(\mathbb{R}, X)$. We remark that an element of \mathbb{M}^{GHPU} has a natural marked point, namely $\eta(0)$.

Suppose that we are given elements $\mathfrak{X}_1 = (X_1, d_1, \mu_1, \eta_1)$ and $\mathfrak{X}_2 = (X_2, d_2, \mu_2, \eta_2)$ of \mathbb{M}^{GHPU} . For a compact metric space (W, D) and isometric embeddings $\iota_1 : X_1 \rightarrow W$ and $\iota_2 : X_2 \rightarrow W$, we define their *GHPU distortion* by

$$\text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_2}^{\text{GHPU}}(W, D, \iota_1, \iota_2) := \mathfrak{d}_D^{\text{H}}(\iota_1(X_1), \iota_2(X_2)) + \mathfrak{d}_D^{\text{P}}(((\iota_1)_*\mu_1, (\iota_2)_*\mu_2)) + \mathfrak{d}_D^{\text{U}}(\iota_1 \circ \eta_1, \iota_2 \circ \eta_2). \tag{1.2}$$

We define the *Gromov-Hausdorff-Prokhorov-uniform (GHPU) distance* by

$$\mathfrak{d}^{\text{GHPU}}(\mathfrak{X}_1, \mathfrak{X}_2) = \inf_{(W, D), \iota_1, \iota_2} \text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_2}^{\text{GHPU}}(W, D, \iota_1, \iota_2), \tag{1.3}$$

where the infimum is over all compact metric spaces (W, D) and isometric embeddings $\iota_1 : X_1 \rightarrow W$ and $\iota_2 : X_2 \rightarrow W$.

It will be proven in Lemma 2.4 below that $\mathfrak{d}^{\text{GHPU}}$ defines a pseudometric on \mathbb{M}^{GHPU} . It is not quite a metric since two elements $(X_1, d_1, \mu_1, \eta_1), (X_2, d_2, \mu_2, \eta_2) \in \mathbb{M}^{\text{GHPU}}$ lie at GHPU distance zero if there is a measure-preserving isometry from X_1 to X_2 which takes η_1 to η_2 . Let $\overline{\mathbb{M}}^{\text{GHPU}}$ be the set of equivalence classes of elements of \mathbb{M}^{GHPU} under the equivalence relation whereby $(X_1, d_1, \mu_1, \eta_1) \sim (X_2, d_2, \mu_2, \eta_2)$ if and only if there exists such an isometry $f : (X_1, d_1) \rightarrow (X_2, d_2)$ with $f_*\mu_1 = \mu_2$ and $f \circ \eta_1 = \eta_2$.

The following statement will be proven in Section 2.2.

Proposition 1.3. *The function $\mathfrak{d}^{\text{GHPU}}$ is a complete separable pseudometric on \mathbb{M}^{GHPU} and the quotient metric space $\mathbb{M}^{\text{GHPU}} / \{\mathfrak{d}^{\text{GHPU}} = 0\}$ is $\overline{\mathbb{M}}^{\text{GHPU}}$.*

Restricting $\mathfrak{d}^{\text{GHPU}}$ to elements of \mathbb{M}^{GHPU} for which the measure μ is identically equal to zero and/or the curve η is constant gives a natural metric on the space of compact metric spaces which are not equipped with a measure and/or a curve. In particular, the Gromov-Hausdorff and Gromov-Hausdorff-Prokhorov metrics are special cases of the GHPU metric and we also obtain a metric on curve-decorated compact metric spaces (which should be called the Gromov-uniform metric).

A particularly useful fact about the GHPU metric, which will be proven in Section 2.2 and used in the proof of Proposition 1.3, is that GHPU convergence is equivalent to Hausdorff, Prokhorov, and uniform convergence within a fixed compact metric space, in a sense which we will now make precise.

Definition 1.4 (HPU convergence). *Let (W, D) be a metric space. Let $\mathfrak{X} = (X, d, \mu, \eta)$ and $\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)$ for $n \in \mathbb{N}$ be elements of \mathbb{M}^{GHPU} such that $X, X^n \subset W$, $D|_X = d$, and $D|_{X^n} = d^n$. We say that $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the D -Hausdorff-Prokhorov-uniform (HPU) sense if $X^n \rightarrow X$ in the D -Hausdorff metric, $\mu^n \rightarrow \mu$ in the D -Prokhorov metric, and $\eta^n \rightarrow \eta$ in the D -uniform metric.*

Proposition 1.5. *Suppose $\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)$ for $n \in \mathbb{N}$ and $\mathfrak{X} = (X, d, \mu, \eta)$ are elements of \mathbb{M}^{GHPU} . Then $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the GHPU metric if and only if there exists a compact metric space (W, D) and isometric embeddings $X \rightarrow W$ and $X^n \rightarrow W$ for $n \in \mathbb{N}$ such that if we identify X and X^n with their images under these embeddings, then $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the D -HPU sense.*

Analogs of Proposition 1.5 for the Gromov-Hausdorff and Gromov-Prokhorov metrics are proven in [23, Lemmas 5.8 and A.1], respectively. The proof of Proposition 1.5 below will be similar to the proofs of these lemmas.

1.3.2 Non-compact case

In this paper we will also have occasion to consider non-compact curve-decorated metric measure spaces (such as the Brownian half-plane). In this subsection we consider a variant of the GHPU metric in this setting. We restrict attention to length spaces to avoid technical complications with convergence of metric balls. However, we expect that it is possible to relax this restriction with some modifications to the definition. See [14, Section 8.1] for a definition of the local Gromov-Hausdorff topology which does not require the length space condition.

Let $\mathbb{M}_\infty^{\text{GHPU}}$ be the set of 4-tuples $\mathfrak{X} = (X, d, \mu, \eta)$ where (X, d) is a locally compact length space, μ is a measure on X which assigns finite mass to each finite-radius metric ball in X , and $\eta : \mathbb{R} \rightarrow X$ is a curve in X . Note that \mathbb{M}^{GHPU} is not contained in $\mathbb{M}_\infty^{\text{GHPU}}$ since elements of the former are not required to be length spaces.

Let $\overline{\mathbb{M}}_\infty^{\text{GHPU}}$ be the set of equivalence classes of elements of $\mathbb{M}_\infty^{\text{GHPU}}$ under the equivalence relation whereby $(X_1, d_1, \mu_1, \eta_1) \sim (X_2, d_2, \mu_2, \eta_2)$ if and only if there is an isometry $f : X_1 \rightarrow X_2$ such that $f_*\mu_1 = \mu_2$ and $f \circ \eta_1 = \eta_2$.

We will define a local version of the GHPU metric on $\overline{\mathbb{M}}_\infty^{\text{GHPU}}$ by truncating \mathfrak{X} at the metric ball $B_r(\eta(0); d)$, then integrating the GHPU metric over all metric balls. The truncation is done in the following manner.

Definition 1.6. Let $\mathfrak{X} = (X, d, \mu, \eta)$ be an element of $\mathbb{M}_\infty^{\text{GHPU}}$. For $r > 0$, let

$$\underline{\tau}_r^\eta := (-r) \vee \sup\{t < 0 : d(\eta(0), \eta(t)) = r\} \quad \text{and} \quad \overline{\tau}_r^\eta := r \wedge \inf\{t > 0 : d(\eta(0), \eta(t)) = r\}. \tag{1.4}$$

The r -truncation of η is the curve $\mathfrak{B}_r\eta \in C_0(\mathbb{R}; X)$ defined by

$$\mathfrak{B}_r\eta(t) = \begin{cases} \eta(\underline{\tau}_r^\eta), & t \leq \underline{\tau}_r^\eta \\ \eta(t), & t \in (\underline{\tau}_r^\eta, \overline{\tau}_r^\eta) \\ \eta(\overline{\tau}_r^\eta), & t \geq \overline{\tau}_r^\eta. \end{cases}$$

The r -truncation of X is the curve-decorated metric measure space

$$\mathfrak{B}_r\mathfrak{X} = (B_r(\eta(0); d), d|_{B_r(\eta(0); d)}, \mu|_{B_r(\eta(0); d)}, \mathfrak{B}_r\eta).$$

If $\mathfrak{X} = (X, d, \mu, \eta) \in \mathbb{M}_\infty^{\text{GHPU}}$, then the Hopf-Rinow theorem [14, Theorem 2.5.28] implies that every closed metric ball in X is compact. Hence for $\mathfrak{X} \in \mathbb{M}_\infty^{\text{GHPU}}$, we have $\mathfrak{B}_r\mathfrak{X} \in \mathbb{M}^{\text{GHPU}}$ for each $r > 0$. Furthermore, for $R > r > 0$ we have $\mathfrak{B}_r\mathfrak{B}_R\mathfrak{X} = \mathfrak{B}_r\mathfrak{X}$.

The local GHPU metric on $\mathbb{M}_\infty^{\text{GHPU}}$ is the function on $\mathbb{M}_\infty^{\text{GHPU}} \times \mathbb{M}_\infty^{\text{GHPU}} \rightarrow [0, \infty)$ defined by

$$d_\infty^{\text{GHPU}}(\mathfrak{X}_1, \mathfrak{X}_2) = \int_0^\infty e^{-r} (1 \wedge d^{\text{GHPU}}(\mathfrak{B}_r\mathfrak{X}_1, \mathfrak{B}_r\mathfrak{X}_2)) dr \tag{1.5}$$

where d^{GHPU} is as in (1.3).

We let $\overline{\mathbb{M}}_\infty^{\text{GHPU}}$ be the set of equivalence classes of elements of $\mathbb{M}_\infty^{\text{GHPU}}$ under the equivalence relation whereby $(X_1, d_1, \mu_1, \eta_1) \sim (X_2, d_2, \mu_2, \eta_2)$ if and only if there is an isometry $f : (X_1, d_1) \rightarrow (X_2, d_2)$ such that $f_*\mu_1 = \mu_2$ and $f \circ \eta_1 = \eta_2$. The following is the analog of Proposition 1.3 for the local GHPU metric, and will be proven in Section 2.3.

Proposition 1.7. The function d_∞^{GHPU} is a complete separable pseudometric on $\mathbb{M}_\infty^{\text{GHPU}}$ and the quotient metric space $\mathbb{M}_\infty^{\text{GHPU}} / \{d_\infty^{\text{GHPU}} = 0\}$ is $\overline{\mathbb{M}}_\infty^{\text{GHPU}}$.

We next want to state an analog of Proposition 1.5 for the local GHPU metric. To this end, we make the following definition.

Definition 1.8 (Local HPU convergence). *Let (W, D) be a metric space. Let $\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)$ for $n \in \mathbb{N}$ and $\mathfrak{X} = (X, d, \mu, \eta)$ be elements of $\mathfrak{M}_\infty^{\text{GHPU}}$ such that X and each X^n is a subset of W satisfying $D|_X = d$ and $D|_{X^n} = d^n$. We say that $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the D -local Hausdorff-Prokhorov-uniform (HPU) sense if the following is true.*

- For each $r > 0$ we have $B_r(\eta^n(0); d^n) \rightarrow B_r(\eta(0); d)$ in the D -Hausdorff metric.
- For each $r > 0$ such that $\mu(\partial B_r(\eta(0); d)) = 0$, we have $\mu^n|_{B_r(\eta^n(0); d^n)} \rightarrow \mu|_{B_r(\eta(0); d)}$ in the D -Prokhorov metric.
- For each $a, b \in \mathbb{R}$ with $a < b$, we have $\eta^n|_{[a, b]} \rightarrow \eta|_{[a, b]}$ in the D -uniform metric.

The following is our analog of Proposition 1.5 for the local GHPU metric.

Proposition 1.9. *Let $\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)$ for $n \in \mathbb{N}$ and $\mathfrak{X} = (X, d, \mu, \eta)$ be elements of $\mathfrak{M}_\infty^{\text{GHPU}}$. Then $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the local GHPU topology if and only if there exists a boundedly compact (i.e., closed bounded sets are compact) metric space (W, D) and isometric embeddings $X^n \rightarrow W$ for $n \in \mathbb{N}$ and $X \rightarrow W$ such that the following is true. If we identify X^n and X with their embeddings into W , then $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the D -local HPU sense.*

1.4 Basic definitions for quadrangulations

The main results of this paper are scaling limit statements for quadrangulations with boundary in the GHPU topology. In this subsection we introduce notation to describe these objects.

A *quadrangulation with boundary* is a (finite or infinite) planar map Q with a distinguished face f_∞ , called the *exterior face*, such that every face of Q other than f_∞ has degree 4. The *boundary* of Q , denoted by ∂Q , is the smallest subgraph of Q which contains every edge of Q incident to f_∞ . The *perimeter* $\text{Perim}(Q)$ of Q is defined to be the degree of the exterior face, with edges counted with multiplicity (i.e., the number of half-edges on the boundary).

A *boundary path* of Q is a path λ from $[1, \text{Perim}(Q)]_{\mathbb{Z}}$ (if ∂Q is finite) or \mathbb{Z} (if ∂Q is infinite) to $\mathcal{E}(\partial Q)$ which traces the edge of ∂Q (counted with multiplicity) in cyclic order around the exterior face. Choosing a boundary path is equivalent to choosing an oriented root (half-)edge on the boundary. This root edge is $\lambda(\text{Perim}(Q))$, oriented toward $\lambda(1)$ in the finite case; or $\lambda(0)$, oriented toward $\lambda(1)$, in the infinite case (here we note that a quadrangulation cannot have any self-loops).

We say that ∂Q is *simple* if some (equivalently every) boundary path for Q hits each vertex exactly once.

For $n, l \in \mathbb{N}_0$, we write $\mathcal{Q}(n, l)$ for the set of quadrangulations with general boundary having n interior faces and $2l$ boundary edges (counted with multiplicity).

We write $\mathcal{Q}^\bullet(n, l)$ for the set of triples (Q, e_0, v_*) where $Q \in \mathcal{Q}(n, l)$, e_0 is a distinguished oriented half-edge of ∂Q (meaning that if e_0 has multiplicity 2, we need to specify a “side” of e_0), and $v_* \in \mathcal{V}(Q)$ is a distinguished vertex.

The *uniform infinite half-plane quadrangulation (UIHPQ)* is the infinite boundary-rooted quadrangulation (Q_∞, e_∞) which is the limit in law with respect to the Benjamini-Schramm topology [5] of a uniform sample from $\mathcal{Q}(n, l)$ (rooted at a uniformly random boundary edge) if we first send $n \rightarrow \infty$ and then $l \rightarrow \infty$ [15, 19].

The *uniform infinite planar quadrangulation with simple boundary (UIHPQ_S)* is the infinite boundary-rooted quadrangulation (Q_S, e_S) with simple boundary which is the limit in law with respect to the Benjamini-Schramm topology [5] of a uniformly random quadrangulation with simple boundary (rooted at a uniformly random boundary edge)

with n interior faces and $2l$ boundary edges if we first send $n \rightarrow \infty$ and then $l \rightarrow \infty$. The UIHPQ_S can be recovered from the UIHPQ by “pruning” the dangling quadrangulations which are disconnected from ∞ by a single vertex [15, 19]; see Section 3.4 for a review of this procedure.

1.5 The Brownian half-plane

The limiting object in the main scaling limit results of this paper is the Brownian half-plane, which we define in this section. The construction given here is of the “unconstrained” type (corresponding to the version of the Schaeffer bijection in which labels are not required to be positive). There is also a constrained construction of the Brownian half-plane in [15, Section 5.3]. We expect (but do not prove) that this construction is equivalent to the one we give here. Our construction is a continuum analog of the Schaeffer-type construction of the UIHPQ_S found in [19] (c.f. [15]), which we review in Section 3.2.

Let $X_\infty : \mathbb{R} \rightarrow [0, \infty)$ be the process such that $\{X_\infty(t)\}_{t \geq 0}$ is a standard linear Brownian motion and $\{X_\infty(-t)\}_{t \geq 0}$ is an independent Brownian motion conditioned to stay positive (i.e., a 3-dimensional Bessel process). For $r \in \mathbb{R}$, let

$$T_\infty(r) := \inf\{t \in \mathbb{R} : X_\infty(t) = -r\},$$

so that $r \mapsto T_\infty(r)$ is non-decreasing and for each $r \in \mathbb{R}$,

$$\{X_\infty(T_\infty(r) + t) + r\}_{t \in \mathbb{R}} \stackrel{d}{=} \{X_\infty(t)\}_{t \in \mathbb{R}}. \tag{1.6}$$

Also let $T_\infty^{-1} : \mathbb{R} \rightarrow \mathbb{R}$ be the right-continuous inverse of T , so that

$$T_\infty^{-1}(t) = -\inf\{X_\infty(s) : s \leq t\}. \tag{1.7}$$

For $s, t \in \mathbb{R}$, let

$$d_{X_\infty}(s, t) := X_\infty(s) + X_\infty(t) - 2 \inf_{u \in [s \wedge t, s \vee t]} X_\infty(u). \tag{1.8}$$

Then d_{X_∞} defines a pseudometric on \mathbb{R} and the quotient metric space $\mathbb{R}/\{d_{X_\infty} = 0\}$ is a forest of continuum random trees, indexed by the excursions of X_∞ away from its running infimum.

Conditioned on X_∞ , let Z_∞^0 be the centered Gaussian process with

$$\text{Cov}(Z_\infty^0(s), Z_\infty^0(t)) = \inf_{u \in [s \wedge t, s \vee t]} \left(X_\infty(u) - \inf_{v \leq u} X_\infty(v) \right), \quad s, t \in \mathbb{R}. \tag{1.9}$$

By the Kolmogorov continuity criterion, Z_∞^0 a.s. admits a continuous modification which is locally α -Hölder continuous for each $\alpha < 1/4$. For this modification we have $Z_\infty^0(s) = Z_\infty^0(t)$ whenever $d_{X_\infty}(s, t) = 0$, so Z_∞^0 defines a function on the continuum random forest $\mathbb{R}/\{d_{X_\infty} = 0\}$.

Let $\mathfrak{b}_\infty : \mathbb{R} \rightarrow \mathbb{R}$ be $\sqrt{3}$ times a two-sided standard linear Brownian motion. For $t \in \mathbb{R}$, define

$$Z_\infty(t) := Z_\infty^0(t) + \mathfrak{b}_\infty(T_\infty^{-1}(t)),$$

with T_∞^{-1} as in (1.7).

For $s, t \in \mathbb{R}$, define

$$d_{Z_\infty}(s, t) = Z_\infty(s) + Z_\infty(t) - 2 \inf_{u \in [s \wedge t, s \vee t]} Z_\infty(u). \tag{1.10}$$

Also define the pseudometric

$$d_\infty^0(s, t) = \inf \sum_{i=1}^k d_{Z_\infty}(s_i, t_i) \quad (1.11)$$

where the infimum is over all $k \in \mathbb{N}$ and all $2k+2$ -tuples $(t_0, s_1, t_1, \dots, s_k, t_k, s_{k+1}) \in \mathbb{R}^{2k+2}$ with $t_0 = s$, $s_{k+1} = t$, and $d_{X_\infty}(t_{i-1}, s_i) = 0$ for each $i \in [1, k+1]_{\mathbb{Z}}$. In other words, d_∞^0 is the largest pseudometric on \mathbb{R} which is at most d_{Z_∞} and is zero whenever d_{X_∞} is 0.

The *Brownian half-plane* is the quotient space $H_\infty = \mathbb{R}/\{d_\infty^0 = 0\}$ equipped with the quotient metric, which we call d_∞ . We write $\mathbb{p}_\infty : \mathbb{R} \rightarrow H_\infty$ for the quotient map. It follows from [8, Theorem 2] (which says that the Brownian disk has the topology of the closed disk) and Proposition 4.2 below that H_∞ has the topology of the closed half-plane.

The *boundary* of H_∞ is the set $\partial H_\infty = \mathbb{p}_\infty(\{T_\infty(r) : r \in \mathbb{R}\})$. It follows from the proof of Proposition 4.2 below and the analogous property of the Brownian disk [8, Proposition 21] that ∂H_∞ is in fact the boundary of H_∞ in the topological sense (i.e., the set of points which do not have a neighborhood which is homeomorphic to the disk).

The *area measure* of H_∞ is the pushforward of Lebesgue measure on \mathbb{R} under \mathbb{p}_∞ , and is denoted by μ_∞ . The *boundary measure* of H_∞ is the pushforward of Lebesgue measure on \mathbb{R} under the map $r \mapsto \mathbb{p}_\infty(T_\infty(r))$. The *boundary path* of H_∞ is the path $\eta_\infty : \mathbb{R} \rightarrow \mathbb{R}$ defined by $\eta_\infty(r) = \mathbb{p}_\infty(T_\infty(r))$. Note that η_∞ travels one unit of boundary length in one unit of time.

Although it is not needed for the statement or proof of our main results, we record for reference the $\sqrt{8/3}$ -LQG description of the Brownian half-plane, which will be proven in this paper.

Proposition 1.10. *Let $(\mathbb{H}, h, 0, \infty)$ be a $\sqrt{8/3}$ -quantum gravity wedge (i.e., a quantum wedge of weight equal to 2) with LQG parameter $\gamma = \sqrt{8/3}$ [20]. Let μ_h and ν_h , respectively, be the $\sqrt{8/3}$ -LQG area and boundary length measures induced by h [21]. Also let \mathfrak{d}_h be the $\sqrt{8/3}$ -LQG metric induced by h [43, 45, 46]. Let $\eta_h : \mathbb{R} \rightarrow \mathbb{R}$ be the curve which parameterizes \mathbb{R} according to $\sqrt{8/3}$ -LQG length and satisfies $\eta_h(0) = 0$. Then $(\mathbb{H}, \mathfrak{d}_h, \mu_h, \eta_h)$ and $(H_\infty, d_\infty, \mu_\infty, \eta_\infty)$ (as defined just above) agree as elements of $\mathbb{M}_\infty^{\text{GHPU}}$, i.e. there exists an isometry $f : (\mathbb{H}, \mathfrak{d}_h) \rightarrow (H_\infty, d_\infty)$ satisfying $f_*\mu_h = \mu_\infty$ and $f \circ \eta_h = \eta_\infty$.*

Proposition 1.10 follows from the results of this paper together with the same argument to prove the analogous $\sqrt{8/3}$ -LQG description of the Brownian plane in [45, Corollary 1.5]. Indeed, Proposition 4.2 below tells us that the Brownian half-plane is the local limit of Brownian disks when we zoom in near a boundary point sampled uniformly from the boundary measure. The $\sqrt{8/3}$ -quantum wedge is the local limit of quantum disks when we zoom in near a boundary point [20]. We already know from [45, Corollary 1.5] that Brownian disks coincide with quantum disks in the sense of Proposition 1.10, so the proposition follows.

We remark that the weight-2 quantum wedge mentioned in Proposition 1.10 comes with some additional structure, namely an embedding into \mathbb{H} with the two marked points respectively sent to 0 and ∞ . It follows from the main result of [46] that this embedding is a.s. determined by the quantum wedge, viewed as a random variable taking values in $\mathbb{M}_\infty^{\text{GHPU}}$. This in particular implies that the Brownian half-plane a.s. determines its embedding into $\sqrt{8/3}$ -LQG.

1.6 Theorem statements: scaling limit of the UIHPQ and UIHPQ_S

In this subsection we state scaling limit results for the UIHPQ with general and simple boundary in the local GHPU topology.

Let (H_∞, d_∞) be an instance of the Brownian half-plane, as in Section 1.5. Let μ_∞ and $\eta_\infty : \mathbb{R} \rightarrow \partial H_\infty$, respectively, be its area measure and natural boundary path. Let

$$\mathfrak{H}_\infty := (H_\infty, d_\infty, \mu_\infty, \eta_\infty), \tag{1.12}$$

so that \mathfrak{H}_∞ is an element of $\mathbb{M}_\infty^{\text{GHPU}}$, defined as in Section 1.3.2.

Let $(Q_\infty, \mathfrak{e}_\infty)$ be a UIHPQ (with general boundary). We view Q_∞ as a connected metric space by replacing each edge with an isometric copy of the unit interval, as in Remark 1.2. For $n \in \mathbb{N}$, let d_∞^n be the graph distance on Q_∞ , re-scaled by $(9/8)^{1/4}n^{-1/4}$. Let μ_∞^n be the measure on $\mathcal{V}(Q_\infty)$ which assigns a mass to each vertex equal to $(4n)^{-1}$ times its degree (in the scaling limit, this is equivalent to assigning each face mass n^{-1}). Let $\lambda_\infty : \mathbb{R} \rightarrow \partial Q_\infty$ be the boundary path of Q_∞ started from \mathfrak{e}_∞ and extended by linear interpolation. Let $\eta_\infty^n(t) := \lambda_\infty(2^{3/2}n^{1/2}t)$ for $t \in \mathbb{R}$. For $n \in \mathbb{N}$, let

$$\mathfrak{Q}_\infty^n := (Q_\infty, d_\infty^n, \mu_\infty^n, \eta_\infty^n) \tag{1.13}$$

so that \mathfrak{Q}_∞^n is an element of $\mathbb{M}_\infty^{\text{GHPU}}$.

Theorem 1.11. *In the setting described just above, we have $\mathfrak{Q}_\infty^n \rightarrow \mathfrak{H}_\infty$ in law in the local GHPU topology, i.e. the UIHPQ converges in law in the scaling limit to the Brownian half-plane in the local GHPU topology.*

Next we state an analog of Theorem 1.11 for the UIHPQ with simple boundary. Let (Q_S, \mathfrak{e}_S) be a UIHPQ_S. We will define for each $n \in \mathbb{N}$ an element of $\mathbb{M}_\infty^{\text{GHPU}}$ associated with (Q_S, \mathfrak{e}_S) in the same manner as in the case of the UIHPQ, except that the time scaling for the boundary path is different.

As above we view Q_S as a connected metric space in the manner of Remark 1.2. For $n \in \mathbb{N}$, let d_S^n be the graph distance on Q_S^n , re-scaled by $(9/8)^{1/4}n^{-1/4}$ and let μ_S^n be the measure on $\mathcal{V}(Q_S)$ which assigns a mass to each vertex equal to $(4n)^{-1}$ times its degree. Let $\lambda_S : \mathbb{R} \rightarrow \partial Q_S$ be the boundary path of Q_S , started from \mathfrak{e}_S and extended by linear interpolation. Let $\eta_S^n(t) := \lambda_S\left(\frac{2^{3/2}}{3}n^{1/2}t\right)$ for $t \in \mathbb{R}$ (note that $\frac{2^{3/2}}{3}$ is replaced by $2^{3/2}$ in the UIHPQ case). For $n \in \mathbb{N}$, let

$$\mathfrak{Q}_S^n := (Q_S, d_S^n, \mu_S^n, \eta_S^n). \tag{1.14}$$

Theorem 1.12. *With \mathfrak{H}_∞ as in (1.12), we have $\mathfrak{Q}_S^n \rightarrow \mathfrak{H}_\infty$ in law in the local GHPU topology, i.e. the UIHPQ_S converges in law in the scaling limit to the Brownian half-plane in the local GHPU topology.*

We will also prove in Section 4.1 below a scaling limit result for finite quadrangulations with boundary toward the Brownian disk in the GHPU topology. We do not state this result here, however, as its proof is a straightforward extension of the proof of the analogous convergence statement in the Gromov-Hausdorff topology from [10].

1.7 Outline

The remainder of this article is structured as follows. In Section 2, we prove the statements about the Gromov-Hausdorff-Prokhorov-uniform metric described in Section 1.3 plus some additional properties, including compactness criteria and a measure-theoretic condition for GHPU convergence.

In Section 3, we review some facts about random planar maps in preparation for our proofs of Theorems 1.11 and 1.12, including the Schaeffer-type constructions of uniform quadrangulations with simple boundary and the UIHPQ, the relationship between the UIHPQ and the UIHPQ_S via the pruning procedure, and the definition of the Brownian disk.

In Section 4, we prove Theorems 1.11 and 1.12. The proof of Theorem 1.11 is similar in spirit to the proof of the scaling limit result for the Brownian plane in [18]. It proceeds by showing that the Brownian half-plane (resp. UIHPQ) can be closely approximated by a Brownian disk (resp. uniform quadrangulation with boundary) and applying a strengthened version of the scaling limit result for uniform quadrangulations with boundary from [10]. Theorem 1.12 is deduced from Theorem 1.11 and the pruning procedure.

2 Properties of the Gromov-Hausdorff-Prokhorov-uniform metric

In this section we will establish the important properties of the GHPU and local GHPU metrics, defined in Section 1.3, and in particular prove Propositions 1.3, 1.5, 1.7, and 1.9. We start in Section 2.1 by proving some elementary topological lemmas which give conditions for a sequence of 1-Lipschitz maps or isometries defined on a sequence of metric spaces to have a subsequential limit. These lemmas will be used several times in this section and in [29]. In Section 2.2, we establish the basic properties of the GHPU metric on compact curve-decorated metric measure spaces. In Section 2.3, we establish the basic properties of the local GHPU metric on non-compact curve-decorated metric measure spaces. In Section 2.4, we introduce a generalization of Gromov-Prokhorov convergence and use it to give a criterion for GHPU convergence which will be used for the proof of our scaling limit results in [29].

2.1 Subsequential limits of isometries

In this subsection we record two elementary topological lemmas which will be useful for our study of the GHPU metric.

Lemma 2.1. *Let (W, D, w) be a separable pointed metric space and let $(\widehat{W}, \widehat{D})$ be any metric space. Let $\{X^n\}_{n \in \mathbb{N}}$ and X be closed subsets of W and for $n \in \mathbb{N}$, let $f^n : X^n \rightarrow \widehat{W}$ be a 1-Lipschitz map. Suppose that the following are true.*

1. *For each $r > 0$, $B_r(w; D) \cap X^n \rightarrow B_r(w; D) \cap X$ in the D -Hausdorff metric.*
2. *For each $r > 0$, there exists a compact set $\widehat{W}_r \subset \widehat{W}$ such that $f^n(B_r(w; D) \cap X^n) \subset \widehat{W}_r$ for each $n \in \mathbb{N}$.*

Then there is a sequence \mathcal{N} of positive integers tending to ∞ and a 1-Lipschitz map $f : X \rightarrow \widehat{W}$ such that $f^n \rightarrow f$ as $\mathcal{N} \ni n \rightarrow \infty$ in the following sense. For any $x \in X$, any subsequence \mathcal{N}' of \mathcal{N} , and any sequence of points $x^n \in X^n$ for $n \in \mathcal{N}'$ with $x^n \rightarrow x$, we have $f^n(x^n) \rightarrow f(x)$ as $\mathcal{N}' \ni n \rightarrow \infty$. Moreover, if each f^n is an isometry onto its image, then f is also an isometry onto its image.

Proof. Let $\{x_j\}_{j \in \mathbb{N}}$ be a countable dense subset of X (which exists since X is separable). By assumption 1, $B_r(w; D) \cap X^n \rightarrow B_r(w; D) \cap X$ in the D -Hausdorff topology for each $r > 0$, so for each $j \in \mathbb{N}$ we can choose points $x_j^n \in X^n$ such that $D(x_j^n, x_j) \rightarrow 0$. By condition 2, each of the sequences $\{x_j^n\}_{n \in \mathbb{N}}$ is contained in a compact subset of \widehat{W} .

By a diagonalization argument we can find a sequence \mathcal{N} of positive integers tending to ∞ and points $\{\widehat{x}_j\}_{j \in \mathbb{N}}$ in \widehat{W} such that $f^n(x_j^n) \rightarrow \widehat{x}_j$ for each $j \in \mathbb{N}$ as $\mathcal{N} \ni n \rightarrow \infty$. Let $f(x_j) := \widehat{x}_j$ for $j \in \mathbb{N}$. Then for $j_1, j_2 \in \mathbb{N}$,

$$\widehat{D}(f(x_{j_1}), f(x_{j_2})) = \lim_{\mathcal{N} \ni n \rightarrow \infty} \widehat{D}(f^n(x_{j_1}^n), f^n(x_{j_2}^n)) \leq \lim_{\mathcal{N} \ni n \rightarrow \infty} D(x_{j_1}^n, x_{j_2}^n) = D(x_{j_1}, x_{j_2}),$$

with equality throughout if in fact each f^n is an isometry onto its image. Since $\{x_j\}_{j \in \mathbb{N}}$ is dense in X , the map f extends by continuity to a 1-Lipschitz map $X \rightarrow \widehat{W}$, which preserves distances in the case when each f^n is an isometry.

It remains to check that $f^n \rightarrow f$ in the sense described in the lemma. Suppose that we are given a subsequence \mathcal{N}' of \mathcal{N} and a sequence of points $x^n \in X^n$ with $x^n \rightarrow x \in X$. Fix $\epsilon > 0$ and choose $j \in \mathbb{N}$ with $x_j \in B_\epsilon(x; D)$. Then for large enough $n \in \mathcal{N}'$,

$$\widehat{D}(f^n(x_j^n), f^n(x^n)) \leq D(x_j^n, x^n) \leq D(x_j^n, x_j) + D(x_j, x) + D(x, x^n) \leq \epsilon + o_n(1).$$

Since $f^n(x_j^n) \rightarrow f(x_j)$, $f^n(x^n)$ lies within \widehat{D} -distance 2ϵ of $f(x_j)$ for large enough $n \in \mathcal{N}'$. On the other hand, $\widehat{D}(f(x_j), f(x)) \leq \epsilon$ (since f is 1-Lipschitz). Therefore $f^n(x^n) \rightarrow f(x)$ along the subsequence \mathcal{N}' . \square

In the case when the f^n 's are isometries onto their images and we assume convergence in the HPU sense, we obtain existence of an isometry f satisfying additional properties.

Lemma 2.2. *Let $\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)$ for $n \in \mathbb{N}$ and $\mathfrak{X} = (X, d, \mu, \eta)$ be elements of \mathbb{M}^{GHPU} (resp. $\mathbb{M}_\infty^{\text{GHPU}}$) such that X^n and X are subsets of a common boundedly compact (i.e., closed bounded subsets are compact) metric space (W, D) satisfying $d^n = D|_{X^n}$ and $d = D|_X$. Let $(\widehat{W}, \widehat{D})$ be another boundedly compact metric space and let $\widehat{\mathfrak{X}} = (\widehat{X}, \widehat{d}, \widehat{\mu}, \widehat{\eta})$ be an element of \mathbb{M}^{GHPU} (resp. $\mathbb{M}_\infty^{\text{GHPU}}$) such that $\widehat{X} \subset \widehat{D}$ and $\widehat{D}|_{\widehat{X}} = \widehat{d}$.*

Suppose that we are given distance-preserving maps $f^n : X^n \rightarrow \widehat{W}$ for each $n \in \mathbb{N}$ such that the following are true (using the terminology as in Definitions 1.4 and 1.8).

1. $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the D -(local) HPU sense.
2. $(f^n(X^n), \widehat{D}|_{f^n(X^n)}, f_*\mu^n, f^n \circ \eta^n) \rightarrow \widehat{\mathfrak{X}}$ in the \widehat{D} -(local) HPU sense.

There is a sequence \mathcal{N} of positive integers tending to ∞ and an isometry $f : X \rightarrow \widehat{X}$ with $f \circ \eta = \widehat{\eta}$ and $f_*\mu = \widehat{\mu}$ such that $f^n \rightarrow f$ as $\mathcal{N} \ni n \rightarrow \infty$ in the following sense. For any $x \in X$, any subsequence \mathcal{N}' of \mathcal{N} , and any sequence of points $x^n \in X^n$ for $n \in \mathcal{N}'$ with $x^n \rightarrow x$, we have $f^n(x^n) \rightarrow f(x)$ as $\mathcal{N}' \ni n \rightarrow \infty$.

Proof. By Lemma 2.1 and since $(\widehat{W}, \widehat{d})$ is boundedly compact, there is a sequence \mathcal{N} of positive integers tending to ∞ and a distance-preserving map $f : X \rightarrow \widehat{X}$ such $f^n \rightarrow f$ in the sense described in the statement of Lemma 2.1. It remains to check that f is surjective and satisfies $f \circ \eta = \widehat{\eta}$ and $f_*\mu = \widehat{\mu}$.

We start with surjectivity. Suppose given $\widehat{x} \in \widehat{X}$. Since $f^n(X^n) \rightarrow \widehat{X}$ in the \widehat{D} -local Hausdorff metric, we can find a sequence $\{x^n\}_{n \in \mathcal{N}}$ of points in X^n such that $f^n(x^n) \rightarrow \widehat{x}$. There is an $r > 0$ such that $\widehat{D}(f^n(x), f^n(\eta^n(0))) \leq r$ for each $n \in \mathcal{N}$, so since each f^n is an isometry we also have $d^n(x, \eta^n(0)) \leq r$ for each $n \in \mathcal{N}$. Since W is boundedly compact, by possibly passing to a subsequence, we can arrange so that $x^n \rightarrow x \in X$. Then the convergence of f^n to f implies that $f(x) = \widehat{x}$.

Next we check that $f \circ \eta = \widehat{\eta}$. For each $t \geq 0$ we have $\eta^n(t) \rightarrow \eta(t)$ and $f^n(\eta^n(t)) \rightarrow \widehat{\eta}(t)$. Therefore our convergence condition for f^n toward f implies that $f(\eta(t)) = \widehat{\eta}(t)$.

Finally, we show that $f_*\mu = \widehat{\mu}$. We do this in the case of $\mathbb{M}_\infty^{\text{GHPU}}$; the case of \mathbb{M}^{GHPU} is similar (but in fact simpler). Let $r > 0$ be a radius for which both

$$\mu(\partial B_r(\eta(0); d)) = 0 \quad \text{and} \quad \widehat{\mu}(\partial B_r(\widehat{\eta}(0); \widehat{d})) = 0. \tag{2.1}$$

For $n \in \mathcal{N}$ let x^n be sampled uniformly from $\mu^n|_{B_r(\eta^n(0); d^n)}$. By assumption 1, $\mu^n|_{B_r(\eta^n(0); d^n)} \rightarrow \mu|_{B_r(\eta(0); d)}$ in the D -Prokhorov metric. Therefore, $x^n \rightarrow x$ in law, where x is sampled uniformly from $\mu|_{B_r(\eta(0); d)}$. By the Skorokhod representation theorem, we can couple $\{x^n\}_{n \in \mathcal{N}}$ with x in such a way that $x^n \rightarrow x$ a.s. Then our convergence condition for f^n implies that $f^n(x^n) \rightarrow f(x)$ a.s. Since each f^n is an isometry, the random variable $f^n(x^n)$ is sampled uniformly from $f_*\mu^n|_{B_r(f^n(\eta^n(0)); \widehat{D}) \cap f^n(X^n)}$. By (2.1) and

assumption 2, we find that $f^n(x^n)$ converges in law to \hat{x} , where \hat{x} is sampled uniformly from $\mu|_{B_r(\hat{\eta}(0); \hat{d})}$. Hence the law of $f(x)$ is that of a uniform sample from $\mu|_{B_r(\hat{\eta}(0); \hat{d})}$. Since $\mu^n(B_r(\eta^n(0); d^n))$ converges to both $\mu(B_r(\eta(0); d))$ and $\hat{\mu}(\hat{\eta}(0); \hat{d})$, we find that

$$f_*(\mu|_{B_r(\eta(0); d)}) = \hat{\mu}|_{B_r(\hat{\eta}(0); \hat{d})}$$

for all but countably many $r > 0$. Therefore $f_*\mu = \hat{\mu}$. □

2.2 Proofs for the GHPU metric

In this subsection we prove Propositions 1.3 and 1.5. Most of the arguments in this subsection are straightforward adaptations of standard proofs for the Gromov-Hausdorff, Gromov-Prokhorov, and Gromov-Hausdorff-Prokhorov metrics; see [2, 14, 23, 24, 40], but we give full proofs here for the sake of completeness.

The following lemma tells us that the definition of $\mathfrak{d}^{\text{GHPU}}$ in (1.3) is unaffected if when taking the infimum we require that $W = X_1 \sqcup X_2$ and ι_1 and ι_2 are the natural inclusions.

Lemma 2.3. *Let $\mathfrak{X}_1 = (X_1, d_1, \mu_1, \eta_1)$ and $\mathfrak{X}_2 = (X_2, d_2, \mu_2, \eta_2)$ be in $\mathfrak{M}^{\text{GHPU}}$ and identify X_1 and X_2 with their inclusions into the disjoint union $X_1 \sqcup X_2$. For a metric d_\sqcup on $X_1 \sqcup X_2$ with $d_\sqcup|_{X_1} = d_1$ and $d_\sqcup|_{X_2} = d_2$, we define*

$$\text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_2}^{\text{GHPU}, \sqcup}(d_\sqcup) = \mathfrak{d}_{d_\sqcup}^{\text{H}}(X_1, X_2) + \mathfrak{d}_{d_\sqcup}^{\text{P}}(\mu_1, \mu_2) + \mathfrak{d}_{d_\sqcup}^{\text{U}}(\eta_1, \eta_2). \tag{2.2}$$

Then

$$\mathfrak{d}^{\text{GHPU}}(\mathfrak{X}_1, \mathfrak{X}_2) = \inf \text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_2}^{\text{GHPU}}(d_\sqcup) \tag{2.3}$$

where the infimum is over all metrics on $X_1 \sqcup X_2$ with $d_\sqcup|_{X_1} = d_1$ and $d_\sqcup|_{X_2} = d_2$.

Proof. It is clear that the infimum in (2.3) is at most the infimum in (1.3), so we only need to prove the reverse inequality. Suppose given a compact metric space (W, D) and isometric embeddings $\iota_1 : X_1 \rightarrow W$ and $\iota_2 : X_2 \rightarrow W$. Given $\epsilon > 0$, we define a metric on $X_1 \sqcup X_2$ by

$$d_\sqcup^\epsilon(x, y) = \begin{cases} d_1(x, y), & x, y \in X_1 \\ d_2(x, y), & x, y \in X_2 \\ D(\iota_1(x), \iota_2(y)) + \epsilon, & x \in X_1, y \in X_2 \\ D(\iota_2(x), \iota_1(y)) + \epsilon, & x \in X_2, y \in X_1. \end{cases}$$

It is easy to see that d_\sqcup^ϵ defines a metric on $X_1 \sqcup X_2$. Furthermore, since ι_1 and ι_2 are isometric embeddings, it follows that $d_\sqcup^\epsilon(x, y)$ differs from $D(x, y)$ by at most ϵ . Hence

$$\text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_2}^{\text{GHPU}, \sqcup}(d_\sqcup) \leq \text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_2}^{\text{GHPU}}(W, D, \iota_1, \iota_2) + 3\epsilon.$$

Since $\epsilon > 0$ is arbitrary the statement of the lemma follows. □

We now verify the triangle inequality for $\mathfrak{d}^{\text{GHPU}}$, which in particular implies that $\mathfrak{d}^{\text{GHPU}}$ is a pseudometric.

Lemma 2.4. *The function $\mathfrak{d}^{\text{GHPU}}$ satisfies the triangle inequality, i.e. for $\mathfrak{X}_1, \mathfrak{X}_2, \mathfrak{X}_3 \in \mathfrak{M}^{\text{GHPU}}$ we have*

$$\mathfrak{d}^{\text{GHPU}}(\mathfrak{X}_1, \mathfrak{X}_3) \leq \mathfrak{d}^{\text{GHPU}}(\mathfrak{X}_1, \mathfrak{X}_2) + \mathfrak{d}^{\text{GHPU}}(\mathfrak{X}_2, \mathfrak{X}_3).$$

Proof. Write $\mathfrak{X}_i = (X_i, d_i, \mu_i, \eta_i)$ for $i \in \{1, 2, 3\}$ and fix $\epsilon > 0$. By Lemma 2.3 we can find metrics d_{12} and d_{23} on $X_1 \sqcup X_2$ and $X_2 \sqcup X_3$, respectively, which restrict to the given metrics on each factor such that

$$\text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_2}^{\text{GHPU}, \sqcup}(d_{12}) \leq \mathfrak{d}^{\text{GHPU}}(\mathfrak{X}_1, \mathfrak{X}_2) + \epsilon \quad \text{and} \quad \text{Dis}_{\mathfrak{X}_2, \mathfrak{X}_3}^{\text{GHPU}, \sqcup}(d_{23}) \leq \mathfrak{d}^{\text{GHPU}}(\mathfrak{X}_2, \mathfrak{X}_3) + \epsilon.$$

We define a metric on $X_1 \sqcup X_2 \sqcup X_3$ as follows. If $x, y \in X_1 \sqcup X_2 \sqcup X_3$ and both x and y belong to $X_1 \sqcup X_2$ or $X_2 \sqcup X_3$ we set $d_{13}(x, y) = d_{12}(x, y)$ or $d_{23}(x, y)$, respectively. For $x_1 \in X_1$ and $x_3 \in X_3$, we set

$$d_{13}(x_1, x_3) = d_{13}(x_3, x_1) = \inf_{x_2 \in X_2} (d_{12}(x_1, x_2) + d_{23}(x_2, x_3)).$$

It is easily checked that d_{13} is a metric on $X_1 \sqcup X_2 \sqcup X_3$, so restricts to a metric on $X_1 \sqcup X_3$ which in turn restricts to d_1 on X_1 and d_3 on X_3 . Furthermore, by the triangle inequalities for the d_{13} -Hausdorff, Prokhorov, and uniform metrics, we have

$$\text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_3}^{\text{GHPU}, \sqcup}(d_{13}) \leq \text{Dis}_{\mathfrak{X}_1, \mathfrak{X}_2}^{\text{GHPU}, \sqcup}(d_{12}) + \text{Dis}_{\mathfrak{X}_2, \mathfrak{X}_3}^{\text{GHPU}, \sqcup}(d_{23}).$$

The statement of the lemma follows. □

Now we can prove Proposition 1.5, using a similar argument to that used to prove [23, Lemma A.1].

Proof of Proposition 1.5. By Lemma 2.3, for each $n \in \mathbb{N}$ there exists a metric d_{\sqcup}^n on $X \sqcup X^n$ such that $\text{Dis}_{\mathfrak{X}, \mathfrak{X}^n}^{\text{GHPU}, \sqcup}(d_{\sqcup}^n) \rightarrow 0$. Let $W := X \sqcup \bigsqcup_{n=1}^{\infty} X^n$ and identify X and each X^n with its natural inclusion into W . We define a metric D on W as follows. If $x, y \in W$ such that $x, y \in X \sqcup X^n$ for some $n \in \mathbb{N}$, we set $D(x, y) = d_{\sqcup}^n(x, y)$. If $x \in X^n$ and $y \in X^m$ for some $n, m \in \mathbb{N}$, we set

$$D(x, y) = \inf_{u \in X} (d_{\sqcup}^n(x, u) + d_{\sqcup}^m(u, y)).$$

As in the proof of Lemma 2.4, D is a metric on W and by definition this metric restricts to d^n on each X^n and to d on X . Furthermore we have $\text{Dis}_{\mathfrak{X}, \mathfrak{X}^n}^{\text{GHPU}, \sqcup}(d_{\sqcup}^n) = \text{Dis}_{\mathfrak{X}, \mathfrak{X}^n}^{\text{GHPU}, \sqcup}(D) \rightarrow 0$ as $n \rightarrow \infty$, which implies that $X^n \rightarrow X$ in the D -Hausdorff metric, $\mu^n \rightarrow \mu$ in the D -Prokhorov metric, and $\eta^n \rightarrow \eta$ in the D -uniform metric.

By replacing W with its metric completion, we can assume that W is complete. Since X is totally bounded, for each $\epsilon > 0$ we can find $N \in \mathbb{N}$ and x_1, \dots, x_N such that $X \subset \bigcup_{i=1}^N B_{\epsilon}(x_i; d)$. Since $X^n \rightarrow X$ in the D -Prokhorov metric and $D|_X = d$, it follows that there exists $n_0 \in \mathbb{N}$ such that $X^n \subset \bigcup_{i=1}^N B_{\epsilon}(x_i; D)$ for $n \geq n_0$. Since each X^n for $n \leq n_0$ is totally bounded, we infer that W is totally bounded, hence compact. □

Lemma 2.5 (Positive definiteness). *Let $\mathfrak{X}_1 = (X_1, d_1, \mu_1, \eta_1)$ and $\mathfrak{X}_2 = (X_2, d_2, \mu_2, \eta_2)$ be in \mathbb{M}^{GHPU} . If $d^{\text{GHPU}}(\mathfrak{X}_1, \mathfrak{X}_2) = 0$, then there is an isometry $f : X_1 \rightarrow X_2$ with $f_*\mu_1 = \mu_2$ and $f \circ \eta_1 = \eta_2$.*

Proof. By Lemma 2.3, we can find a sequence of metrics d_{\sqcup}^n on $X_1 \sqcup X_2$ whose GHPU distortion tends to 0. By Proposition 1.5 there is a compact metric space (W, D) and isometric embeddings $\iota : X_1 \rightarrow W$ and $\iota^n : X_2 \rightarrow W$ for $n \in \mathbb{N}$ such that

$$(\iota^n(X_2), D|_{\iota^n(X_2)}, \iota_*^n \mu_2, \iota^n \circ \eta) \rightarrow (X_1, d_1, \mu_1, \eta_1)$$

in the D -HPU topology. By Lemma 2.2 (applied with $\mathfrak{X}^n = \mathfrak{X}_2$ for each $n \in \mathbb{N}$) we can find a subsequence along which the embeddings ι^n converge to an isometry $g : X_2 \rightarrow \iota(X_1)$ with $g_*\mu_2 = \iota_*\mu_2$ and $g \circ \eta_2 = \iota \circ \eta_1$. The statement of the lemma follows by taking $f = g^{-1} \circ \iota$. □

For our proof of completeness, we will use the following compactness criterion for \mathbb{M}^{GHPU} , which is also of independent interest.

Lemma 2.6 (Compactness criterion). *Let \mathcal{K} be a subset of \mathbb{M}^{GHPU} and suppose the following conditions are satisfied.*

1. \mathcal{K} is uniformly totally bounded, i.e. for each $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that for each $(X, d, \mu, \eta) \in \mathcal{K}$, the set X can be covered by at most N d -balls of radius at most ϵ .
2. There is a constant $C > 0$ such that for each $(X, d, \mu, \eta) \in \mathcal{K}$, we have $\mu(X) \leq C$.
3. \mathcal{K} is equicontinuous, i.e. for each $\epsilon > 0$, there exists $\delta > 0$ such that for each (X, d, μ, η) in \mathcal{K} and each $s, t \in \mathbb{R}$ with $|s - t| \leq \delta$, we have $d(\eta(s), \eta(t)) \leq \epsilon$; and for each $t \in \mathbb{R}$ with $|t| \geq \delta^{-1}$, we have $d(\eta(s), \eta(\pm\delta^{-1})) \leq \epsilon$, where \pm is the sign of t .

Then every sequence in \mathcal{K} has a subsequence which converges with respect to the GHPU topology.

Proof. Let $\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)$ for $n \in \mathbb{N}$ be elements of \mathcal{K} . By condition 1 and the Gromov compactness criterion [14, Theorem 7.4.15], we can find a sequence $n_k \rightarrow \infty$ and a compact metric space (X, d) such that $(X^{n_k}, d^{n_k}) \rightarrow (X, d)$ in the Gromov-Hausdorff topology. By [23, Lemma A.1] (or Proposition 1.5 applied with $\mu = 0$ and η a constant curve) we can find a compact metric space (W, D) and isometric embeddings $X^{n_k} \rightarrow W$ for $k \in \mathbb{N}$ and $X \rightarrow W$ such that if we identify X^{n_k} and X with their embeddings, we have $X^{n_k} \rightarrow X$ in the D -Hausdorff metric.

By conditions 2 and 3, the Prokhorov theorem, and the Arzela-Ascoli theorem, after possibly passing to a further subsequence we can find a finite Borel measure μ on X and a curve η in X such that $\mu^{n_k} \rightarrow \mu$ in the D -Prokhorov metric and $\eta^{n_k} \rightarrow \eta$ in the D -uniform metric as $k \rightarrow \infty$. Therefore $\mathfrak{X}^{n_k} \rightarrow \mathfrak{X}$ in the GHPU metric. \square

Lemma 2.7 (Completeness). *The pseudometric space $(\mathbb{M}^{\text{GHPU}}, d^{\text{GHPU}})$ is complete.*

Proof. Let $\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)$ for $n \in \mathbb{N}$ be a Cauchy sequence with respect to d^{GHPU} . It is clear that $\mathcal{K} = \{\mathfrak{X}^n\}_{n \in \mathbb{N}}$ satisfies the conditions of Lemma 2.6, so has a convergent subsequence. The Cauchy condition implies that in fact the whole sequence converges. \square

Next we check separability.

Lemma 2.8 (Separability). *The space \mathbb{M}^{GHPU} with the metric d^{GHPU} is separable.*

The proof of Lemma 2.8 is slightly more difficult than one might expect. The reason for this is that we cannot simply approximate elements of \mathbb{M}^{GHPU} by finite metric spaces, since such spaces do not admit non-trivial continuous paths. We get around this as follows. Given $\mathfrak{X} = (X, d, \mu, \eta) \in \mathbb{M}^{\text{GHPU}}$, we first find a finite subset A of X equipped with a measure μ_0 such that $(A, d|_A, \mu_0)$ closely approximates (X, d, μ) in the Hausdorff and Prokhorov metrics. We then isometrically embed $(A, d|_A)$ into (a very high dimensional) Euclidean space equipped with the L^∞ distance and draw in a piecewise linear path which approximates η .

Proof of Lemma 2.8. Let \mathbb{F} be the set of $(X, d, \mu, \eta) \in \mathbb{M}^{\text{GHPU}}$ for which the following is true.

- X is a subset of \mathbb{R}^N for some $N \in \mathbb{N}$ and d is the restriction of the L^∞ metric d^∞ on \mathbb{R}^N (i.e. $d^\infty(z, w) = \max_{i \in [1, N]} |z_i - w_i|$).
- X is the union of finitely many points in \mathbb{Q}^N with and finitely many linear segments with endpoints in \mathbb{Q}^N .
- The measure μ is supported on $X \cap \mathbb{Q}^N$, and $\mu(x) \in \mathbb{Q} \cap [0, \infty)$ for each $x \in X \cap \mathbb{Q}^N$.
- The curve η is the concatenation of finitely many (possibly degenerate) linear segments with endpoints in \mathbb{Q}^N , each traversed at a constant d^∞ -speed which belongs to $\mathbb{Q} \cap [0, \infty)$.

It is clear that \mathbb{F} is countable. We claim that \mathbb{F} is dense in \mathbb{M}^{GHPU} . It is clear that \mathbb{F} is dense in the set $\widehat{\mathbb{F}} \subset \mathbb{M}^{\text{GHPU}}$ which is defined in the same manner as \mathbb{F} but with every instance of \mathbb{Q} replaced by \mathbb{R} . Hence it suffices to show that $\widehat{\mathbb{F}}$ is dense in \mathbb{M}^{GHPU} .

Suppose that we are given $\mathfrak{X} = (X, d, \mu, \eta) \in \mathbb{M}^{\text{GHPU}}$ and $\epsilon > 0$. We will construct $\widehat{\mathfrak{X}} = (\widehat{X}, \widehat{d}, \widehat{\mu}, \widehat{\eta}) \in \widehat{\mathbb{F}}$ which approximates X in the GHPU sense. We first define a finite subset A of X as follows.

- Let A_1 be a finite ϵ -dense subset of X .
- Let μ_0 be a finitely supported measure on X with $\phi_d^{\text{P}}(\mu, \mu_0) \leq \epsilon$ (which can be obtained, e.g., by sampling M i.i.d. points from μ for M large and defining the mass of each to be $\mu(X)/M$). Let A_2 be the support of μ_0 .
- Let $K \in \mathbb{N}$ be chosen so that $d(\eta(s), \eta(t)) \leq \epsilon$ whenever $s, t \in \mathbb{R}$ with $|t - s| \leq 1/K$ and $d(\eta(\pm t), \eta(\pm \delta^{-1})) \leq \epsilon$ whenever $t \geq K$. For $k \in [-K^2, K^2]_{\mathbb{Z}}$ let $y_k := \eta(k/K)$. Let

$$A_3 := \{y_{-K^2}, \dots, y_{K^2}\}.$$

- Let $A := A_1 \cup A_2 \cup A_3$.

It is not hard to see (and is proven, e.g., in [39, Theorem 15.7.1]) that there is an isometric embedding ι of the metric space $(A, d|_A)$ into (\mathbb{R}^N, d^∞) for $N = \#A$ (here we recall that d^∞ is the L^∞ metric on \mathbb{R}^N). For $k \in [-K^2, K^2]_{\mathbb{Z}}$, let $\widehat{\eta}_k$ be the straight line path in \mathbb{R}^N from $\iota(y_{k-1})$ to $\iota(y_k)$ with constant d^∞ -speed which is traversed in $1/K$ units of time. By our choice of the y_k 's, $\widehat{\eta}_k \in B_\epsilon(\iota(y_k); d^\infty)$. Let $\widehat{\eta}$ be the concatenation of the paths $\widehat{\eta}_k$.

Let $\widehat{X} := \iota(A) \cup \widehat{\eta}$. Let $\widehat{d} := d^\infty|_{\widehat{X}}$. Let $\widehat{\mu}$ be the measure on \widehat{X} which is the pushforward of μ_0 under ι . Then $\widehat{\mathfrak{X}} := (\widehat{X}, \widehat{d}, \widehat{\mu}, \widehat{\eta}) \in \widehat{\mathbb{F}}$.

It remains only to compare \mathfrak{X} with $\widehat{\mathfrak{X}}$. Let W be the set obtained from $X \sqcup \widehat{X}$ by identifying $A \subset X$ with $\iota(A) \subset \widehat{X}$. Let D be the metric on W which restricts to d (resp. \widehat{d}) on X (resp. \widehat{X}) and which is defined for $x \in X$ and $\widehat{x} \in \widehat{X}$ by

$$d_Z(x, \widehat{x}) = d_Z(\widehat{x}, x) = \inf_{a \in A} \left(d(x, a) + \widehat{d}(\iota(a), \widehat{x}) \right).$$

Note that this is well defined and satisfies the triangle inequality since ι is an isometry. By our choices of A , μ_0 , K , and the path $\widehat{\eta}$, it follows that the GHPU distortion of (W, D) and the natural inclusions of X and \widehat{X} into W is at most 4ϵ , so since $\epsilon > 0$ is arbitrary we obtain the desired separability. \square

Proof of Proposition 1.3. This follows by combining Lemmas 2.4, 2.5, 2.7, and 2.8. \square

2.3 Proofs for the local GHPU metric

In this subsection we will prove Propositions 1.7 and 1.9. These statements will, for the most part, be deduced from the analogous statements in the compact case proven in Section 2.2. We first state a lemma which will enable us to relate GHPU convergence and local GHPU convergence. For the statement of this lemma and in what follows, it will be convenient to have the following definition.

Definition 2.9 (Good radius). *Let $\mathfrak{X} = (X, d, \mu, \eta) \in \mathbb{M}_\infty^{\text{GHPU}}$ and $r > 0$. We say that r is a good radius for \mathfrak{X} if*

$$\mu(\partial B_r(\eta(0); d)) = 0 \quad \text{and} \quad |\eta^{-1}(\partial B_r(\eta(0); d))| = 0, \tag{2.4}$$

where $|\eta^{-1}(\partial B_r(\eta(0); d))|$ is the Lebesgue measure of $\eta^{-1}(\partial B_r(\eta(0); d))$.

Since the sets $\partial B_r(\eta(0); d)$ for $r > 0$ are disjoint, it follows that all but countably many radii are good.

Lemma 2.10. *Let $\mathfrak{X} = (X, d, \mu, \eta) \in \mathbb{M}^{\text{GHPU}}$ and suppose that for each $\zeta > 0$, each point $x \in X$ can be joined to $\eta(0)$ by a path of d -length at most $d(\eta(0), x) + \zeta$. Let $R > r > 0$ and suppose that the radius r is good, in the sense of (2.4). For each $\epsilon > 0$, there exists $\delta > 0$ depending only on ϵ and the R -truncation $\mathfrak{B}_R \mathfrak{X}$ (Definition 1.6) such that the following is true. Let $\tilde{\mathfrak{X}} = (\tilde{X}, \tilde{d}, \tilde{\mu}, \tilde{\eta}) \in \mathbb{M}^{\text{GHPU}}$ and suppose that d_{\sqcup} is a metric on $X \sqcup \tilde{X}$ which restricts to d on X and \tilde{d} on \tilde{X} and whose GHPU distortion (recall (2.2)) satisfies $\text{Dis}_{\mathfrak{X}, \tilde{\mathfrak{X}}}^{\text{GHPU}, \sqcup}(d_{\sqcup}) \leq \delta$. Then*

$$\text{Dis}_{\mathfrak{B}_r \mathfrak{X}, \mathfrak{B}_r \tilde{\mathfrak{X}}}^{\text{GHPU}, \sqcup}(d_{\sqcup}) \leq \epsilon. \tag{2.5}$$

The proof of Lemma 2.10 is a straightforward but tedious application of the definitions together with the triangle inequality, so we omit it. In light of Lemma 2.3, Lemma 2.10 implies that if $d^{\text{GHPU}}(\mathfrak{X}, \tilde{\mathfrak{X}}) < \delta$, then $d^{\text{GHPU}}(\mathfrak{B}_r \mathfrak{X}, \mathfrak{B}_r \tilde{\mathfrak{X}}) \leq \epsilon$.

As a consequence of Lemma 2.10, we see that local GHPU convergence is really just GHPU convergence of curve-decorated metric measure spaces truncated at an appropriate sequence of metric balls.

Lemma 2.11. *Let \mathfrak{X} and $\{\mathfrak{X}^n\}_{n \in \mathbb{N}}$ be elements of $\mathbb{M}_{\infty}^{\text{GHPU}}$. If $\{r_k\}_{k \in \mathbb{N}}$ is a sequence of positive real numbers tending to ∞ and the r_k -truncations (Definition 1.6) satisfy $\mathfrak{B}_{r_k} \mathfrak{X}^n \rightarrow \mathfrak{B}_{r_k} \mathfrak{X}$ in the GHPU topology for each $k \in \mathbb{N}$, then $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the local GHPU topology. Conversely, if $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the local GHPU topology, then for each good radius r (Definition 2.9), we have $\mathfrak{B}_r \mathfrak{X}^n \rightarrow \mathfrak{B}_r \mathfrak{X}$ in the GHPU topology.*

Proof. Let \mathcal{R} be the set of good radii $r > 0$ for \mathfrak{X} . Recall that $(0, \infty) \setminus \mathcal{R}$ is countable.

Suppose first that we are given a sequence $r_k \rightarrow \infty$ such that $\mathfrak{B}_{r_k} \mathfrak{X}^n \rightarrow \mathfrak{B}_{r_k} \mathfrak{X}$ in the GHPU topology for each $k \in \mathbb{N}$. By Lemma 2.10 applied with $\mathfrak{B}_{r_k} \mathfrak{X}$ for $r_k > r$ in place of \mathfrak{X} , we find that for each $r \in \mathcal{R}$, $d^{\text{GHPU}}(\mathfrak{B}_r \mathfrak{X}^n, \mathfrak{B}_r \mathfrak{X}) \rightarrow 0$. By the dominated convergence theorem applied to the formula (1.5), it follows that $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the local GHPU topology.

Conversely, suppose $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the local GHPU topology and let $r \in \mathcal{R}$. By Lemma 2.10 applied with $R = r + 1$, and $\mathfrak{B}_R \mathfrak{X}$ for $R > r + 1$ in place of \mathfrak{X} , for each $k \in \mathbb{N}$ and each $\epsilon > 0$, there exists $\delta = \delta(k, \epsilon) > 0$ such that whenever $R \geq r + 1$ and $d^{\text{GHPU}}(\mathfrak{B}_R \mathfrak{X}^n, \mathfrak{B}_R \mathfrak{X}) \leq \delta$, we have $d^{\text{GHPU}}(\mathfrak{B}_r \mathfrak{X}^n, \mathfrak{B}_r \mathfrak{X}) \leq \epsilon$. Local GHPU convergence implies that for large enough $n \in \mathbb{N}$, there exists $R \geq r + 1$ with $d^{\text{GHPU}}(\mathfrak{B}_R \mathfrak{X}^n, \mathfrak{B}_R \mathfrak{X}) \leq \delta$. Hence $\mathfrak{B}_r \mathfrak{X}^n \rightarrow \mathfrak{B}_r \mathfrak{X}$. \square

Proof of Proposition 1.9. It is clear from Lemma 2.11 that the existence of (W, D) and isometric embeddings as in the statement of the lemma implies that $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the local GHPU topology.

Conversely, suppose $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the local GHPU topology. The proof of this direction is a generalization of that of Proposition 1.5. To lighten notation, for $r > 0$ and $n \in \mathbb{N}$ we write

$$B_r^n := B_r(\eta^n(0); d^n) \quad \text{and} \quad B_r := B_r(\eta(0); d).$$

Choose a sequence of good radii $r_k \rightarrow \infty$ for \mathfrak{X} . By Lemma 2.11, we have, in the notation of Definition 1.6, $\mathfrak{B}_{r_k} \mathfrak{X}^n \rightarrow \mathfrak{B}_{r_k} \mathfrak{X}$ for each $n \in \mathbb{N}$. For each $k \in \mathbb{N}$, choose $N_k \in \mathbb{N}$ such that for $n \geq N_k$, we have $d_{\infty}^{\text{GHPU}}(\mathfrak{B}_{r_k} \mathfrak{X}^n, \mathfrak{B}_{r_k} \mathfrak{X}) \leq 1/k$. For $n \in \mathbb{N}$, let k_n be the largest $k \in \mathbb{N}$ such that $N_k \leq n$, and note that $k_n \rightarrow \infty$ as $n \rightarrow \infty$.

By Lemma 2.3, for each $n \in \mathbb{N}$ there exists a metric \tilde{d}_{\sqcup}^n on $B_{r_{k_n}} \sqcup B_{r_{k_n}}^n$ which restricts to d^n on $B_{r_{k_n}}^n$ and d on $B_{r_{k_n}}$ and whose GHPU distortion is at most $1/k_n$. We now extend

\tilde{d}_\square^n to a metric d_\square^n on $X \sqcup X$ by the formula

$$d_\square^n(x, y) := \begin{cases} d(x, y) & x, y \in X \\ d^n(x, y) & x, y \in X^n \\ \inf_{(u,v) \in \tilde{B}_{r_{k_n}} \times B_{r_{k_n}}^n} \left(d(x, u) + d^n(y, v) + \tilde{d}_\square^n(u, v) \right) & x \in X, y \in X^n \end{cases}$$

together with the requirement that $d_\square^n(y, x) = d_\square^n(x, y)$ when $y \in X$ and $x \in X^n$. It is easily verified that d_\square^n is a metric on $X \sqcup X^n$ and that $d_\square^n|_{B_{r_{k_n}} \sqcup B_{r_{k_n}}^n} = \tilde{d}_\square^n$.

As in the proof of Proposition 1.5, let $W := X \sqcup \bigsqcup_{n=1}^\infty X^n$ and identify X and each X^n with its natural inclusion into W . We define a metric D on W as follows. If $x, y \in W$ such that $x, y \in X \sqcup X^n$ for some $n \in \mathbb{N}$, we set $D(x, y) = d_\square^n(x, y)$. If $x \in X^n$ and $y \in X^m$ for some $n, m \in \mathbb{N}$, we set

$$D(x, y) = \inf_{u \in X} (d_\square^n(x, u) + d_\square^m(u, y)).$$

Then D is a metric on W which restricts to d^n on each X^n and to d on X .

For each $n \in \mathbb{N}$, the restriction of D to $B_{r_{k_n}} \sqcup B_{r_{k_n}}^n$ agrees with the corresponding restriction of d_\square^n , which agrees with \tilde{d}_\square^n . Since the GHPU distortion of \tilde{d}_\square^n is at most $1/k_n \rightarrow 0$ and $r_{k_n} \rightarrow \infty$ as $n \rightarrow \infty$, it follows from Lemma 2.10 that $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the D -local HPU topology.

By possibly replacing W with its metric completion, we can take W to be complete. Since (X, d) and each (X^n, d^n) is boundedly compact (since they are locally compact length spaces and by the Hopf-Rinow Theorem), it follows easily from D -Hausdorff metric convergence of the metric balls $B_r(\eta^n(0); d^n)$ that each metric ball in X with finite radius is totally bounded, hence compact. \square

Next we record a compactness criterion for the local GHPU metric which will be used to prove completeness.

Lemma 2.12 (Compactness criterion). *Let \mathcal{K} be a subset of $\mathbb{M}_\infty^{\text{GHPU}}$ and suppose that there is a sequence $r_k \rightarrow \infty$ such that for each $k \in \mathbb{N}$, the set of r_k -truncations $\mathfrak{B}_{r_k} \mathcal{K} = \{\mathfrak{B}_{r_k} \mathfrak{X} : \mathfrak{X} \in \mathcal{K}\}$ (Definition 1.6) satisfies the conditions of Lemma 2.6 with $\mathfrak{B}_{r_k} \mathfrak{X}$ in place of \mathfrak{X} , i.e. $\mathfrak{B}_{r_k} \mathcal{K}$ is totally bounded, has bounded total mass, and is equicontinuous. Then every sequence in \mathcal{K} has a subsequence which converges with respect to the local GHPU metric.*

Proof. By the compactness criterion for the local Gromov-Hausdorff-Prokhorov topology [2, Theorem 2.9], the set of metric measure spaces $\{(X, d, \mu, \eta) : (X, d, \mu, \eta) \in \mathcal{K}\}$ is pre-compact with respect to the pointed local Gromov-Hausdorff-Prokhorov topology. Hence for any sequence $\{\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)\}_{n \in \mathbb{N}}$ of elements of \mathcal{K} , there exists a sequence $n_k \rightarrow \infty$ and a locally compact pointed length space equipped with a finite measure (X, d, μ, x) such that $(X^{n_k}, d^{n_k}, \mu^{n_k}, \eta^{n_k}(0)) \rightarrow (X, d, \mu, x)$ in the pointed local Gromov-Hausdorff-Prokhorov topology.

By the analog of Proposition 1.9 for local Gromov-Hausdorff-Prokhorov convergence (which follows from Proposition 1.5 by taking η to be a constant curve) we can find a boundedly compact metric space (W, D) and isometric embeddings of (X^{n_k}, d^{n_k}) for $k \in \mathbb{N}$ and (X, d) into (W, D) such that if we identify these spaces with their embeddings, then $\eta^{n_k}(0) \rightarrow x$, $B_r(\eta^{n_k}(0); d^{n_k}) \rightarrow B_r(x; d)$ in the D -Hausdorff metric for each $r > 0$, and $\mu^{n_k}|_{B_r(\eta^{n_k}(0); d^{n_k})} \rightarrow \mu|_{B_r(x; d)}$ in the D -Prokhorov metric for each $r > 0$ such that $\mu(\partial B_r(x; d)) = 0$. By the Arzela-Ascoli theorem and a diagonalization argument, after possibly passing to a further subsequence we can find a curve η in X such that $(X^{n_k}, d^{n_k}, \mu^{n_k}, \eta^{n_k}) \rightarrow (X, d, \mu, \eta)$ in the D -HPU sense (Definition 1.8), so by Proposition 1.9 $\mathfrak{X}^{n_k} \rightarrow (X, d, \mu, \eta)$ in the local GHPU metric. \square

Proof of Proposition 1.7. Symmetry and the triangle inequality are immediate from the formula (1.5) and the analogous properties for the GHPU metric, so d_∞^{GHPU} is a pseudo-metric.

The fact that $d_\infty^{\text{GHPU}}(\mathfrak{X}_1, \mathfrak{X}_2) = 0$ implies that \mathfrak{X}_1 and \mathfrak{X}_2 agree as elements of $\overline{\mathbb{M}}_\infty^{\text{GHPU}}$ follows from Proposition 1.9, Lemma 2.2, and the same argument used to prove Lemma 2.5.

Completeness follows from Lemma 2.12 (c.f. the proof of Lemma 2.7).

Finally, we check separability. We know that \mathbb{M}^{GHPU} is separable, so the set $\mathbb{M}_\infty^{\text{GHPU}} \cap \overline{\mathbb{M}}_\infty^{\text{GHPU}}$ consisting of length spaces in \mathbb{M}^{GHPU} is separable with respect to the GHPU metric. By Lemma 2.11, $\mathbb{M}^{\text{GHPU}} \cap \overline{\mathbb{M}}_\infty^{\text{GHPU}}$ is also separable with respect to the local GHPU metric. So, it suffices to show that $\mathbb{M}^{\text{GHPU}} \cap \overline{\mathbb{M}}_\infty^{\text{GHPU}}$ is dense in $\overline{\mathbb{M}}_\infty^{\text{GHPU}}$.

We cannot approximate an element of $\mathfrak{X} = (X, d, \mu, \eta) \in \overline{\mathbb{M}}_\infty^{\text{GHPU}}$ by the r -truncation $\mathfrak{B}_r \mathfrak{X}$ since $B_r(\eta(0); d)$ may not be a length space with the restricted metric. We instead construct a modified version of $\mathfrak{B}_r \mathfrak{X}$ which is a length space. Given $r > 0$, let X_r be the quotient space of $B_r(\eta(0); d)$ under the equivalence relation which identifies $\partial B_r(\eta(0); d)$ to a point. Let \tilde{d}_r be the quotient metric on X_r and let $p_r : B_r(\eta(0); d) \rightarrow X_r$ be the quotient map.

We note that \tilde{d}_r is a metric, not a pseudometric since every point of $B_r(\eta(0); d) \setminus \partial B_r(\eta(0); d)$ lies at positive distance from $\partial B_r(\eta(0); d)$. Furthermore, the triangle inequality implies that the restriction of \tilde{d}_r to $p_r(B_{r/3}(\eta(0); d))$ coincides with the corresponding restriction of d , pushed forward under p_r .

Let d_r be the smallest length metric on X_r which is greater than or equal to \tilde{d}_r . Equivalently, for $x, y \in X_r$, $d_r(x, y)$ is the infimum of the \tilde{d}_r -lengths of paths from x to y contained in X_r . Then X_r is complete and totally bounded with respect to d_r , so is compact with respect to d_r .

Since d is a length metric, we find that the restriction of d_r to $p_r(B_{r/3}(\eta(0); d))$ coincides with the corresponding restriction of d , pushed forward under p_r .

Let

$$\mathfrak{X}_r := (X_r, d_r, (p_r)_* \mu, p_r \circ \mathfrak{B}_r \eta),$$

with $\mathfrak{B}_r \eta$ as in Definition 1.6. Then $\mathfrak{X}_r \in \mathbb{M}^{\text{GHPU}} \cap \overline{\mathbb{M}}_\infty^{\text{GHPU}}$ and $\mathfrak{B}_{r/3} \mathfrak{X}_r$ agrees with $\mathfrak{B}_{r/3} \mathfrak{X}$ as elements of $\overline{\mathbb{M}}_\infty^{\text{GHPU}}$. Hence $d_\infty^{\text{GHPU}}(\mathfrak{X}, \mathfrak{X}_r) \leq e^{-r/3}$. Since $r > 0$ is arbitrary, we obtain the desired density. \square

2.4 Conditions for GHPU convergence using the m -fold Gromov-Prokhorov topology

In this subsection we prove a lemma giving conditions for GHPU convergence which will be used in subsequent sections to prove convergence of uniform quadrangulations with boundary to the Brownian disk in the GHPU topology. To state the lemma we first need to consider a variant of the Gromov-Prokhorov topology [23], which we now define.

For $m \in \mathbb{N}$ let \mathbb{M}_m^{GP} be the space of $(m + 2)$ -tuples $(X, d, \mu_1, \dots, \mu_m)$ where (X, d) is a separable metric space and $\{\mu_r\}_{r \in [1, m]_{\mathbb{Z}}}$ are finite Borel measures on X . Given $k \in \mathbb{N}$ and $r \in [1, m]_{\mathbb{Z}}$ let $\{x^{j+(r-1)k}\}_{j \in [1, k]_{\mathbb{Z}}}$ be i.i.d. samples from μ_r and let $M_k(X, d, \mu_1, \dots, \mu_m)$ be the $km \times km$ matrix whose i, j th entry is $d(x^i, x^j)$ for $i, j \in [1, km]_{\mathbb{Z}}$. We define the m -fold Gromov-Prokhorov (GP) topology on \mathbb{M}^{GP} to be the weakest one for which the functional

$$(X, d, \mu_1, \dots, \mu_m) \mapsto \mathbb{E}[\phi(M_k(X, d, \mu_1, \dots, \mu_m), \mu_1(X), \dots, \mu_m(X))] \tag{2.6}$$

is continuous for each bounded continuous function $\phi : \mathbb{R}^{(km)^2} \times \mathbb{R}^m \rightarrow \mathbb{R}$. Note that convergence in the m -fold GP topology is equivalent to convergence of each of these functionals.

If $x \in X$ is a marked point, we set $x^0 = x$ and define $M_k^\bullet(X, d, \mu_1, \dots, \mu_m, x)$ to be the random $(km + 1) \times (km + 1)$ matrix whose (i, j) th entry is $d(x^i, x^j)$ for $i, j \in [0, km]_{\mathbb{Z}}$, where x^i for $i \in [1, km]_{\mathbb{Z}}$ is defined as above. We define the m -fold Gromov-Prokhorov (GP) topology on pointed metric spaces with m finite Borel measures to be the weakest one for which the functional

$$(X, d, \mu_1, \dots, \mu_m, x) \mapsto \mathbb{E}[\phi(M_k^\bullet(X, d, \mu_1, \dots, \mu_m, x), \mu_1(X), \dots, \mu_m(X))] \tag{2.7}$$

is continuous for each bounded continuous function $\phi : \mathbb{R}^{(km+1)^2} \times \mathbb{R}^m \rightarrow \mathbb{R}$.

Like the Gromov-Prokhorov topology, the m -fold Gromov-Prokhorov topology also separates points.

Lemma 2.13. *Let $(X, d, \mu_1, \dots, \mu_m)$ and $(\tilde{X}, \tilde{d}, \tilde{\mu}_1, \dots, \tilde{\mu}_m)$ be elements of \mathbb{M}_m^{GP} . Suppose that the functionals (2.6) agree on $(X, d, \mu_1, \dots, \mu_m)$ and $(\tilde{X}, \tilde{d}, \tilde{\mu}_1, \dots, \tilde{\mu}_m)$ for each $k \in \mathbb{N}$. Suppose further that the union of the closed supports of μ_r (resp. $\tilde{\mu}_r$) for $r \in [1, m]_{\mathbb{Z}}$ is all of X (resp. \tilde{X}). Then there is an isometry $f : X \rightarrow \tilde{X}$ with $f_*\mu_r = \tilde{\mu}_r$ for each $r \in [1, m]_{\mathbb{Z}}$. If X and \tilde{X} are endowed with marked points x and \tilde{x} , respectively, for which the functionals (2.7) agree for each $k \in \mathbb{N}$, we can also take f to satisfy $f(x) = \tilde{x}$.*

Proof. We treat the unpointed case; the pointed case is treated in an identical manner. For $r \in [1, m]_{\mathbb{Z}}$, let $\{x_r^j\}_{j \in \mathbb{N}}$ and $\{\tilde{x}_r^j\}_{j \in \mathbb{N}}$ be i.i.d. samples from μ_r and $\tilde{\mu}_r$, respectively. By our assumption about the supports of the measures μ_r and $\tilde{\mu}_r$, the sets $\{x_r^j : (r, j) \in [1, m]_{\mathbb{Z}} \times \mathbb{N}\}$ and $\{\tilde{x}_r^j : (r, j) \in [1, m]_{\mathbb{Z}} \times \mathbb{N}\}$ are a.s. dense in X and \tilde{X} , respectively. The agreement of the functionals (2.6) implies that $\mu_r(X) = \tilde{\mu}_r(\tilde{X})$ for each $r \in [1, m]_{\mathbb{Z}}$. Furthermore, the collections of distances $\{d(x_r^j, x_{r'}^{j'}) : (r, j), (r', j') \in [1, m]_{\mathbb{Z}} \times \mathbb{N}\}$ and $\{\tilde{d}(\tilde{x}_r^j, \tilde{x}_{r'}^{j'}) : (r, j), (r', j') \in [1, m]_{\mathbb{Z}} \times \mathbb{N}\}$ agree in law, so we can couple everything in such a way that these collections of distances agree a.s. Let $f : \{x_r^j : (r, j) \in [1, m]_{\mathbb{Z}} \times \mathbb{N}\} \rightarrow \tilde{X}$ be the function which sends each x_r^j to \tilde{x}_r^j . By our choice of coupling f is distance preserving, and since the domain and image of f are a.s. dense in X and \tilde{X} , respectively, f extends by continuity to an isometry $X \rightarrow \tilde{X}$.

For each $N \in \mathbb{N}$, the set $\{x_r^j : (r, j) \in [1, m]_{\mathbb{Z}} \times [N, \infty]_{\mathbb{Z}}\}$ is still a.s. dense in X , so the restriction of f to this set a.s. determines f . By the Kolmogorov zero-one law, f is a.s. equal to a certain deterministic map $X \rightarrow \tilde{X}$. In particular, for each $r \in [1, m]_{\mathbb{Z}}$ and each $A \subset X$, we have

$$\mu_r(A) = \mu_r(X)\mathbb{P}[x_r^1 \in A] = \tilde{\mu}_r(X)\mathbb{P}[\tilde{x}_r^1 \in f(A)] = \tilde{\mu}_r(f(A)),$$

so $f_*\mu_r = \tilde{\mu}_r$. □

We now state our condition for GHPU convergence.

Lemma 2.14. *Let $\mathfrak{X}^n = (X^n, d^n, \mu^n, \eta^n)$ for $n \in \mathbb{N}$ and $\mathfrak{X} = (X, d, \mu, \eta)$ be elements of \mathbb{M}^{GHPU} . Suppose that the curves η^n for $n \in \mathbb{N}$ (resp. η) are each constant outside some bounded interval $[0, T^n]$ (resp. $[0, T]$). Let ν^n (resp. ν) be the pushforward of Lebesgue measure on $[0, T^n]$ (resp. $[0, T]$) under η^n (resp. η). Suppose the following conditions are satisfied.*

1. $(X^n, d^n, \eta(0)) \rightarrow (X, d, \eta(0))$ in the pointed Gromov-Hausdorff metric.
2. $(X^n, d^n, \mu^n, \nu^n, \eta^n(0)) \rightarrow (X, d, \mu, \nu, \eta(0))$ in the 2-fold GP topology.
3. The curves $\{\eta^n\}_{n \in \mathbb{N}}$ are equicontinuous, i.e. for each $\epsilon > 0$ there exists $\delta > 0$ such that for each $n \in \mathbb{N}$, we have $d^n(\eta^n(s), \eta^n(t)) \leq \epsilon$ whenever $s, t \in [0, T^n]$ with $|s - t| \leq \delta$.
4. The closed support of μ is all of X .

5. η is a simple curve.

Then $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the GHPU topology.

Proof. By Lemma 2.6 and condition 1, we can find a sequence $\mathcal{N} \subset \mathbb{N}$ and $\tilde{\mathfrak{X}} = (\tilde{X}, \tilde{d}, \tilde{\mu}, \tilde{\eta}) \in \mathbb{M}^{\text{GHPU}}$ such that $(X, d, \eta(0))$ and $(\tilde{X}, \tilde{d}, \tilde{\eta}(0))$ are isometric as pointed metric spaces and $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the GHPU topology. By Proposition 1.5, there is a compact metric space (W, D) , and isometric embeddings of X^n for $n \in \mathcal{N}$ and X into (W, D) such that if we identify X^n and X with their images under these embeddings, we have $\mathfrak{X}^n \rightarrow \mathfrak{X}$ in the D -HPU sense as $\mathcal{N} \ni n \rightarrow \infty$.

From the convergences $\nu^n \rightarrow \tilde{\nu}$, $\mu^n \rightarrow \tilde{\mu}$, and $\eta^n(0) \rightarrow \tilde{\eta}(0)$, we infer that $(X^n, d^n, \mu^n, \nu^n, \eta^n(0)) \rightarrow (\tilde{X}, \tilde{d}, \tilde{\mu}, \tilde{\nu}, \tilde{\eta}(0))$ in the 2-fold GP topology. Hence condition 4 and Lemma 2.13 imply that we can find a distance-preserving map $f : X \rightarrow \tilde{X}$ (whose range is the union of the closed supports of $\tilde{\nu}$ and $\tilde{\mu}$) such that $f_*\tilde{\mu} = \mu$, $f_*\tilde{\nu} = \nu$, and $f(\tilde{\eta}(0)) = \eta(0)$. Since (X, d) and (\tilde{X}, \tilde{d}) are isometric as metric spaces and a compact metric space cannot be isometric to a proper subset of itself, we find that f is in fact surjective, so is an isometry from X to \tilde{X} .

We claim that also $f \circ \tilde{\eta} = \eta$, which will imply that $(X, d, \tilde{\mu}, \tilde{\eta}) = (X, d, \mu, \eta)$ as elements of $\overline{\mathbb{M}}^{\text{GHPU}}$.

Since each η^n is constant outside of $[0, T^n]$ and $T^n \rightarrow T$, we find that $\tilde{\eta}$ is constant outside of $[0, T]$. It is clear that $\tilde{\nu}$ is supported on $\tilde{\eta}$. For $0 \leq a \leq b \leq T$ we have $\eta^n([a, b]) \rightarrow \tilde{\eta}([a, b])$ in the D -Hausdorff distance, so since $\nu^n \rightarrow \tilde{\nu}$ it follows that $\tilde{\nu}(\tilde{\eta}([a, b])) \geq b - a$. By condition 5 and the existence of the isometry f above, the measure $\tilde{\nu}$ has no point masses so for each $\epsilon > 0$ we can find $\delta > 0$ such that $\tilde{\nu}(B_\delta(\tilde{\eta}([a, b]))) \leq b - a + \epsilon$. Hence in fact $\tilde{\nu}(\tilde{\eta}([a, b])) = b - a$. Therefore $\tilde{\nu}$ is the pushforward under $\tilde{\eta}$ of Lebesgue measure on $[0, T]$.

The map f takes the closed support of $\tilde{\nu}$ to the closed support of ν , so takes the range of $\tilde{\eta}$ to the range of η . Since $f_*\tilde{\nu} = \nu$ and $f(\tilde{\eta}(0)) = \eta(0)$, for each $t \in [0, T]$ the set $f(\tilde{\eta}([0, t]))$ is a connected subset of $\tilde{\eta}$ containing 0, with $\tilde{\nu}$ -mass equal to t . Therefore $\tilde{\eta}^{-1}(f(\tilde{\eta}([0, t]))) = [0, t]$. In other words, the map $\tilde{\eta}^{-1} \circ f \circ \eta$ is the identity, so $f \circ \eta = \tilde{\eta}$. \square

3 Schaeffer-type constructions

In this section we will review constructions from the planar map literature which are needed for the proofs of our scaling limit results. In Sections 3.1 and 3.2, respectively, we review the Schaeffer-type constructions of uniform quadrangulations with boundary and the UIHPQ_S. In Section 3.3, we record some distance estimates in terms of the encoding functions in these constructions. In Section 3.4, we recall how to “prune” the UIHPQ to get an instance of the UIHPQ_S. In Section 3.5, we recall the definition of the Brownian disk from [10].

3.1 Encoding quadrangulations with boundary

Recall from Section 1.4 the set $\mathcal{Q}^\bullet(n, l)$ of boundary-rooted, pointed quadrangulations with n internal faces and $2l$ boundary edges. In this subsection we review a variant of the Schaeffer bijection for elements of $\mathcal{Q}^\bullet(n, l)$ which is really a special case of the Bouttier-Di Francesco-Guitter bijection [12]. Our presentation is similar to that in [19, Section 3.3] and [10, Section 3.3]. See Figure 2 for an illustration.

For $l \in \mathbb{N}$, a *bridge* of length $2l$ is a function $b^0 : [0, 2l]_{\mathbb{Z}} \rightarrow \mathbb{Z}$ such that $b^0(j+1) - b^0(j) \in \{-1, 1\}$ for each $j \in [0, 2l-1]_{\mathbb{Z}}$ and $b^0(0) = b^0(2l) = 0$. For a bridge b^0 , we associate a function $b : [0, l]_{\mathbb{Z}} \rightarrow \mathbb{Z}$ as follows. We set $j_0 = 0$ and for $k \in [1, l]_{\mathbb{Z}}$ we let j_k be the k th smallest $j \in [0, 2l-1]_{\mathbb{Z}}$ for which $b^0(j+1) - b^0(j) = -1$. We then let $b(k) := b^0(j_k)$.

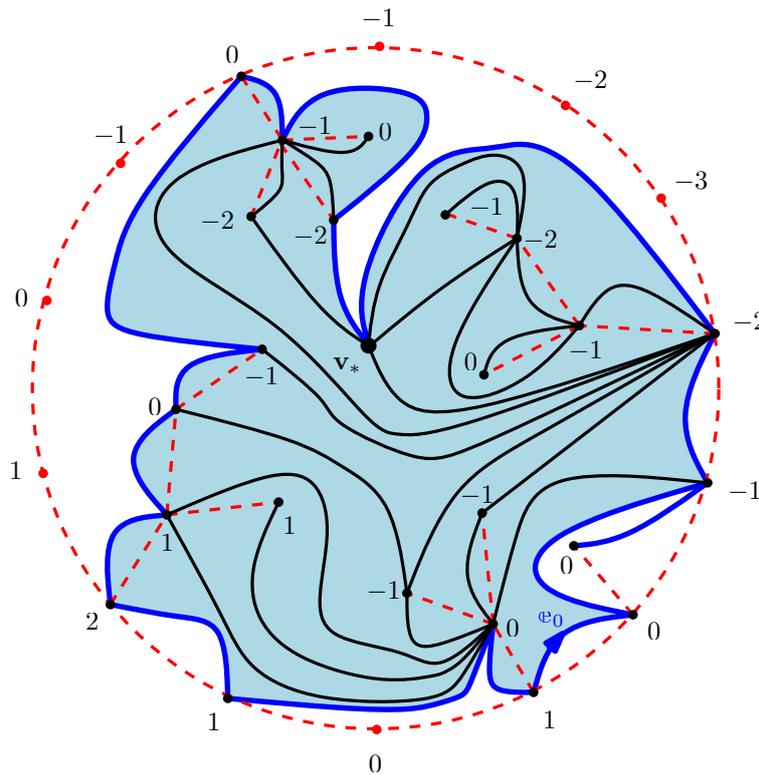


Figure 2: Illustration of the Schaeffer-type encoding of a boundary-rooted, pointed quadrangulation with boundary (Q, e_0, v_*) . The region enclosed by Q is shown in light blue. The graph F containing the trees t_k is the union of the dashed red lines and the black vertices. The red vertices correspond to upward steps of b^0 , and do not belong to F or Q . Each vertex is labeled with either its label $L(v)$ or the corresponding value of b^0 . The black edges are part of the interior of Q and the blue edges are part of ∂Q . The oriented boundary root edge e_0 is indicated with an arrow.

For $n, l \in \mathbb{N}$, a *treed bridge* of area n and boundary length $2l$ is an $(l + 1)$ -tuple $(b^0; (t_0, v_0, L_0), \dots, (t_{l-1}, v_{l-1}, L_{l-1}))$ where b^0 is a bridge of length $2l$ and (t_k, v_k, L_k) for $k \in [0, l - 1]_{\mathbb{Z}}$ is a rooted plane tree with a label function $L_k : \mathcal{V}(t_k) \rightarrow \mathbb{Z}$ satisfying $L_k(v) - L_k(v') \in \{-1, 0, 1\}$ whenever v and v' are joined by an edge and $L_k(v_k) = b(k)$, where b is constructed from b^0 as above, such that the total number of edges in the trees t_k is n . Let $\mathcal{T}^\bullet(n, l)$ be the set of treed bridges of area n and boundary length $2l$, together with a sign $\xi \in \{-, +\}$ (which will be used to determine the orientation of the root edge).

We associate a treed bridge with a rooted, labeled planar map (F, e_0, L) with two faces as follows. Draw an edge from v_k to v_{k+1} for each $k \in [0, l - 2]_{\mathbb{Z}}$ and an edge from v_{l-1} to v_0 . This gives us a cycle which we embed into \mathbb{C} in such a way that the vertices v_k all lie on the unit circle. We extend this embedding to the trees t_k in such a way that each is mapped into the unit disk. This gives us a planar map F with an inner face of degree $2n + l$ (containing all of the trees t_k) and an outer face of degree l . Let e_0 be the oriented edge of F from v_{l-1} to v_0 . Let L be the label function on vertices of F inherited from the label functions L_k for $k \in [0, l - 1]_{\mathbb{Z}}$.

We now associate a rooted, pointed quadrangulation with boundary to (F, e_0, L) and the sign ξ via a variant of the Schaeffer bijection. Let $\mathbb{p} : [0, 2n + l]_{\mathbb{Z}} \rightarrow \mathcal{V}(F)$ be the contour exploration of the inner face of F started from v_1 , i.e. the concatenation of the

contour explorations of the trees t_0, \dots, t_{l-1} . Also define (by a slight abuse of notation) $L(i) = L(\mathbb{p}(i))$. Note that each $\mathbb{p}(i)$ for $i \in [0, 2n + l]_{\mathbb{Z}}$ is associated with a unique corner of the inner face of F (i.e. a connected component of $B_\epsilon(\mathbb{p}(i)) \setminus F$ for small $\epsilon > 0$). Let v_* be an extra vertex not connected to any vertex of F , lying in the interior face of F . For $i \in [0, 2n + l]_{\mathbb{Z}}$, define the successor $s(i)$ of i to be the smallest $i' \geq i$ (with elements of $[0, 2n + l]_{\mathbb{Z}}$ viewed modulo $2n + l$) such that $L(i') = L(i) - 1$, or let $s(i) = \infty$ if no such i' exists. For $i \in [0, 2n + l]_{\mathbb{Z}}$, draw an edge from (the corner associated with) $\mathbb{p}(i)$ to (the corner associated with) $\mathbb{p}(s(i))$, or an edge from $\mathbb{p}(i)$ to v_* if $s(i) = \infty$. Then, delete all of the edges of F to obtain a map Q . We take Q to be rooted at the oriented edge $e_0 \in \mathcal{E}(\partial Q)$ from v_0 to $\mathbb{p}(s(0))$ (if $\xi = -$) or from $\mathbb{p}(s(0))$ to v_0 (if $\xi = +$), viewed as a half-edge on the boundary of the external face.

As explained in, e.g., [19, Section 3.2] and [10, Section 3.3], this construction defines a bijection from $\mathcal{T}^\bullet(n, l)$ to $\mathcal{Q}^\bullet(n, l)$.

Remark 3.1. It is explained in [19, Section 3.3.1] that there is a canonical boundary path $\lambda : [1, 2l]_{\mathbb{Z}} \rightarrow \mathcal{E}(\partial Q)$ starting and ending at the terminal vertex v_0 of the root edge e_0 which traces all of the edges in ∂Q in cyclic order, defined as follows. Recall the definition of the times j_k for $k \in \mathbb{Z}$ at which the walk b^0 has a downward step. For each $k \in [1, l]_{\mathbb{Z}}$, there is a unique connected component of the complement of Q in the inner face of the map F which contains the edge from v_k to v_{k+1} (or from v_l to v_1 if $k = l$) on its boundary. It is easy to see from the Schaeffer bijection that there are precisely $j_{k+1} - j_k + 1$ vertices and $j_{k+1} - j_k$ edges of ∂Q on the boundary of this component, counted with multiplicity. The ordered sequence of labels of the vertices coincides with the ordered sequence of values of $b^0(j)$ for $j \in [j_k, j_{k+1}]_{\mathbb{Z}}$ (which by definition of j_k is the same as $b(j_k), b(j_k) - 1, b(j_k), b(j_k + 1), \dots, b(j_k) + j_{k+1} - j_k - 2 = b(j_{k+1})$). For $j \in [j_k, j_{k+1} - 1]_{\mathbb{Z}}$, we let $\lambda(j)$ be the j th edge along the boundary of this component, in order started from v_k and counted with multiplicity. Then λ is a bijection from $[1, 2l]_{\mathbb{Z}}$ to the set of edges of ∂Q if we count the latter according to their multiplicity in the external face.

As in the case of the ordinary Schaeffer bijection, the above construction can also be phrased in terms of walks. For $i \in [0, 2n + l]_{\mathbb{Z}}$, let $k_i \in [0, l - 1]_{\mathbb{Z}}$ be chosen so that the vertex $\mathbb{p}(i)$ belongs to the tree t_{k_i} and let

$$C(i) := \text{dist}(\mathbb{p}(i), v_{k_i}; t_{k_i}) - k_i, \quad \forall i \in [0, 2n + l - 1]_{\mathbb{Z}} \quad \text{and} \quad C(2n + l) = -l. \quad (3.1)$$

Then C is the concatenation of the contour functions of the trees t_k , but with an extra downward step whenever it moves to a new tree. Let

$$I(k) := \min\{i \in [0, 2n + l]_{\mathbb{Z}} : C(i) = -k\}, \quad \forall k \in [0, l]_{\mathbb{Z}}. \quad (3.2)$$

Then for $k \in [0, l - 1]_{\mathbb{Z}}$, $I(k)$ is the first time i for which $\mathbb{p}(i) \in t_k$ and the range of I is precisely the set of vertices lying on the outer boundary of the graph F . Also let $L^0(i) := L(i) - b(k_i)$. To describe the law of the pair (C, L^0) we need the following definition.

Definition 3.2. Let $[a, b]_{\mathbb{Z}}$ be a (possibly infinite) discrete interval and let $S : [a, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}$ be a (deterministic or random) path with $S(i) - S(i - 1) \in \{-1, 0, 1\}$ for each $i \in [a + 1, b]_{\mathbb{Z}}$. The head of the discrete snake driven by S is the function $H : [a, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}$ whose conditional law given S is described as follows. We set $H(a) = 0$. Inductively, suppose $i \in [a + 1, b]_{\mathbb{Z}}$ and $H(i)$ has been defined for $j \in [a, i - 1]_{\mathbb{Z}}$. If $S(i) - S(i - 1) \in \{-1, 0\}$, let i' be the largest $j \in [a, i - 1]_{\mathbb{Z}}$ for which $H(i) = H(i')$; or $i' = -\infty$. If $i' \neq -\infty$, we set $H(i) = H(i')$. Otherwise, we sample $H(i) - H(i - 1)$ uniformly from $\{-1, 0, 1\}$.

The following lemma is immediate from the definitions and the fact that the above construction is a bijection.

Lemma 3.3. *If we sample (Q, e_0, v_*) uniformly from $\mathcal{Q}^\bullet(n, l)$, then the law of C is that of a simple random walk started from 0 and conditioned to reach $-l$ for the first time at time $2n+l$. The process L^0 is the head of the discrete snake driven by $i \mapsto C(i) - \min_{j \in [1, i]_{\mathbb{Z}}} C(j)$. The pair (C, L^0) is independent from b^0 .*

3.2 Encoding the UIHPQ

In this subsection we describe an infinite-volume analog of the bijection of Section 3.1 which encodes a UIHPQ which is alluded to but not described explicitly in [19, Section 6.1] (see also [15] for a different encoding). See Figure 3 for an illustration.

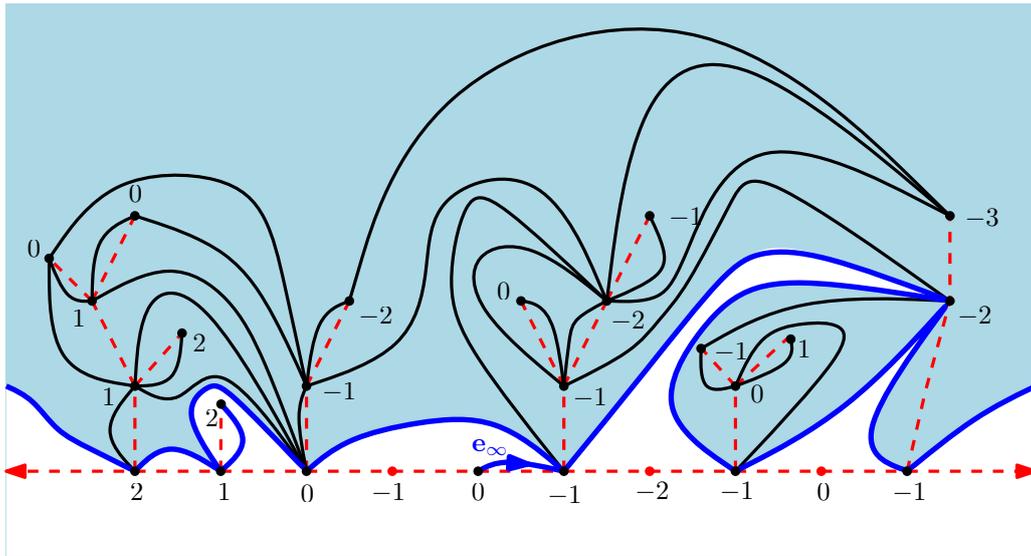


Figure 3: Illustration of the Schaeffer-type encoding of the UIHPQ. The graph F_∞ containing the trees $t_{\infty, k}$ is the union of the dashed red lines and the black vertices. The red vertices correspond to upward steps of b_∞^0 , and do not belong to F_∞ or Q_∞ . Each vertex is labelled with either its label $L_\infty(v)$ or the corresponding value of b_∞^0 . The black edges are part of the interior of the UIHPQ, and the blue edges are part of its boundary. Note that ∂Q_∞ is not a simple path. The oriented boundary root edge e_∞ is indicated with an arrow.

Let $b_\infty^0 : \mathbb{Z} \rightarrow \mathbb{N}_0$ be a two-sided simple random walk reflected at 0 (with increments uniform in $\{-1, 1\}$). Let $\{j_k\}_{k \in \mathbb{Z}}$ be the ordered set of times for which $b_\infty^0(j+1) - b_\infty^0(j) = -1$, shifted so that j_1 is the smallest $j \geq 0$ for which $b_\infty^0(j+1) - b_\infty^0(j) = -1$. Let $b_\infty^0(k) := b_\infty^0(j_k)$.

Conditional on b_∞^0 , let $\{(t_{\infty, k}, v_{\infty, k}, L_{\infty, k})\}_{k \in \mathbb{Z}}$ be a bi-infinite sequence of independent triples where each $(t_{\infty, k}, v_{\infty, k})$ is a rooted Galton-Watson tree whose offspring distribution is geometric with parameter $1/2$ and, conditional on $t_{\infty, k}$, each $L_{\infty, k}$ is uniformly distributed on the set of label functions $\mathcal{V}(t_{\infty, k}) \rightarrow \mathbb{Z}$ which satisfy $L_{\infty, k}(v_{\infty, k}) = b_\infty^0(k)$ and $|L_{\infty, k}(u) - L_{\infty, k}(v)| \leq 1$ whenever $u, v \in \mathcal{V}(t_{\infty, k})$ are connected by an edge.

We define a graph F_∞ as follows. Equip \mathbb{Z} with the standard nearest-neighbor graph structure and embed it as the real line in \mathbb{C} . For $k \in \mathbb{Z}$, embed the tree $t_{\infty, k}$ into the upper half-plane in such a way that the vertex $v_{\infty, k}$ is identified with $k \in \mathbb{Z}$ and none of the trees $t_{\infty, k}$ intersect each other or intersect \mathbb{R} except at their root vertices. The graph F_∞^0 is the union of \mathbb{Z} and the trees $t_{\infty, k}$ for $k \in \mathbb{Z}$ with this graph structure, with each integer j_k identified with the corresponding root vertex $v_{\infty, k}$.

We define a label function L_∞ on the vertices of F_∞ by setting $L_\infty|_{\mathcal{V}(t_{\infty,k})} = L_{\infty,k}$ for each $k \in \mathbb{Z}$. We let $\mathbb{p}_\infty : \mathbb{Z} \rightarrow \mathcal{V}(F_\infty)$ be the contour exploration of the upper face of F_∞ , shifted so that \mathbb{p}_∞ starts exploring the tree $t_{\infty,1}$ at time 0. We then define the successor $s_\infty(i)$ of each time $i \in \mathbb{Z}$ exactly as in the Schaeffer bijection, except that there is no need to add an extra vertex since a.s. $\liminf_{i \rightarrow \infty} L_\infty(i) = -\infty$. We then draw an edge connecting each vertex $\mathbb{p}_\infty(i)$ to $\mathbb{p}_\infty(s_\infty(i))$ for each $i \in \mathbb{Z}$ to obtain an infinite quadrangulation with boundary Q_∞ , which we take to be rooted at the oriented edge e_∞ which goes from $v_{\infty,0}$ to $\mathbb{p}_\infty(s_\infty(0))$. Then (Q_∞, e_∞) is an instance of the UIHPQ. Furthermore, our construction of F_∞ gives rise to an embedding of Q_∞ into \mathbb{H} .

We note that the obvious analog of Remark 3.1 holds in this setting.

Remark 3.4. There is a canonical choice of boundary path $\lambda_\infty : \mathbb{Z} \rightarrow \mathcal{E}(\partial Q)$ which hits the terminal vertex v_∞ of the root edge e_∞ at time 0 and which traces all of the edges in ∂Q in cyclic order, defined as follows. Recall the definition of the times j_k for $k \in \mathbb{Z}$ at which the walk b_∞^0 has a downward step. For each $k \in \mathbb{Z}$, there is a unique connected component of the complement of Q_∞ in the upper face of the map F_∞ which contains the edge from $v_{\infty,k}$ to $v_{\infty,k+1}$ on its boundary. There are precisely $j_{k+1} - j_k + 1$ vertices and $j_{k+1} - j_k$ edges of ∂Q_∞ on the boundary of this component, counted with multiplicity. For $j \in [j_k, j_{k+1} - 1]_{\mathbb{Z}}$, we let $\lambda_\infty(j)$ be the j th edge along the boundary of this component, in order started from $v_{\infty,k}$ and counted with multiplicity. Then λ_∞ is a bijection from \mathbb{Z} to the set of edges of ∂Q_∞ if we count the latter according to their multiplicity in the external face.

As in Section 3.1, we now define random paths which encode (Q_∞, e_∞) . For $i \in \mathbb{Z}$, let k_i be chosen so that the vertex $\mathbb{p}_\infty(i)$ belongs to the tree t_{∞,k_i} and let

$$C_\infty(i) := \text{dist}(\mathbb{p}_\infty(i), v_{\infty,k_i}; t_{\infty,k_i}) - k_i, \quad \forall i \in \mathbb{Z}. \tag{3.3}$$

Let

$$I_\infty(k) := \min\{i \in \mathbb{Z} : C_\infty(i) = -k\}, \quad \forall k \in \mathbb{Z}. \tag{3.4}$$

Then $I_\infty(k)$ is the first time i for which $\mathbb{p}_\infty(i) \in t_{\infty,k}$ and the range of I_∞ is the set of vertices lying on the outer boundary of the graph F_∞ . We let $L_\infty(i) := L_\infty(\mathbb{p}_\infty(i))$ and $L_\infty^0(i) := L_\infty(i) - b_\infty(k_i)$.

Lemma 3.5. *The pair (C_∞, L_∞^0) is independent from b_∞^0 and its law can be described as follows. The law of $C_\infty|_{\mathbb{N}_0}$ is that of a simple random walk started from 0 and the law of $C_\infty(\cdot)|_{\mathbb{N}_0}$ is that of a simple random walk started from 0 and conditioned to stay positive for all time (see, e.g., [6] for a definition of this conditioning for a large class of random walks). Furthermore, L_∞^0 is the head of the discrete snake driven by $i \mapsto C_\infty(i) - \min_{j \in (-\infty, i]_{\mathbb{Z}}} C_\infty(j)$ (Definition 3.2).*

Proof. The process C_∞ is the concatenation of the contour functions of countably many i.i.d. Galton-Watson trees with offspring distribution given by a geometric random variable with parameter 1/2, separated by downward steps. Each of these contour functions has the law of a simple random walk run until the first time it hits 0. Therefore, C_∞ has the law described in the statement of the lemma. Since each $L_\infty^0|_{[I(k), I(k+1)-1]_{\mathbb{Z}}}$ is the head of the discrete snake driven by $C_\infty|_{[I(k), I(k+1)-1]_{\mathbb{Z}}}$, we see that L_∞^0 is the head of the discrete snake driven by C_∞ . The pair (C_∞, L_∞^0) is determined by the labeled trees $\{(t_{\infty,k}, v_{\infty,k}, L_{\infty,k})\}_{k \in \mathbb{Z}}$, so is independent from b_∞^0 . \square

3.3 Distance bounds for quadrangulations with boundary

In this subsection we record elementary upper and lower bounds for distances in quadrangulations with boundary in terms of the encoding processes in the Schaeffer bijection. We start with an upper bound.

Lemma 3.6. *Suppose we are in the setting of Section 3.1. In particular, let $(Q, \mathfrak{e}_0, \mathfrak{v}_*) \in \mathcal{Q}^\bullet(n, l)$, let (C, L, b) be its Schaeffer encoding process, and let $\mathbb{p} : [0, 2n + l]_{\mathbb{Z}} \rightarrow \mathcal{V}(Q)$ be the projection map. For $i_1, i_2 \in \mathbb{Z}$, we have*

$$\text{dist}(\mathbb{p}(i_1), \mathbb{p}(i_2); Q) \leq L(i_1) + L(i_2) - 2 \min_{j \in [i_1 \wedge i_2, i_1 \vee i_2]_{\mathbb{Z}}} L(j) + 2$$

Proof. In the finite-volume case, this follows from [40, Lemma 3] (which gives the analogous estimate in a more general setting). The infinite-volume version follows from exactly the same argument. See also [36, Lemma 3.1] (which is the analogous estimate for quadrangulations without boundary, and is proven in the same manner) and/or [8, Equation (3)] (which states the precise estimate given in the present lemma in the finite-volume case). \square

We also have a lower bound for distances, which is a variant of the so-called cactus bound for the Brownian map (see, e.g., [38, Proposition 5.9]).

Lemma 3.7. *Suppose we are in the setting of Section 3.1. In particular, let $(Q, \mathfrak{e}_0, \mathfrak{v}_*) \in \mathcal{Q}^\bullet(n, l)$, let (C, L) be its Schaeffer encoding process, let $I(m)$ for $m \in [0, l]_{\mathbb{Z}}$ be as in (3.2), and let $\mathbb{p} : [0, 2n + l]_{\mathbb{Z}} \rightarrow \mathcal{V}(Q) \setminus \{\mathfrak{v}_*\}$ be the contour exploration. For $i_1, i_2 \in [1, 2n + l]_{\mathbb{Z}}$ with $i_1 < i_2$, let*

$$J(i_1, i_2) := (I([0, l]_{\mathbb{Z}}) \cap [i_1, i_2]_{\mathbb{Z}}) \cup \{i_1, i_2\} \quad \text{and} \quad J'(i_1, i_2) := (I([0, l]_{\mathbb{Z}}) \setminus [i_1, i_2]_{\mathbb{Z}}) \cup \{i_1, i_2\}, \tag{3.5}$$

so that $\mathbb{p}(J(i_1, i_2))$ (resp. $\mathbb{p}(J'(i_1, i_2))$) consists of $\mathbb{p}(i_1), \mathbb{p}(i_2)$, and the set of vertices of ∂Q which are contained in the image of $\mathbb{p} \circ I$ and which are (resp. are not) contained in $\mathbb{p}([i_1, i_2]_{\mathbb{Z}})$. Then

$$\text{dist}(\mathbb{p}(i_1), \mathbb{p}(i_2); Q) \geq L(i_1) + L(i_2) - 2 \max \left\{ \min_{j \in J(i_1, i_2)} L(j), \min_{j \in J'(i_1, i_2)} L(j) \right\}. \tag{3.6}$$

The analogous estimate also holds in the setting of the UIHPQ (but in this case the minimum over $J'(i_1, i_2)$ in (3.6) is a.s. equal to $-\infty$, so only the first term in the maximum is present).

Proof. This follows from essentially the same proof as the ordinary cactus bound for quadrangulations without boundary (see [38, Proposition 5.9(ii)]) but in the finite-volume case one has to consider two paths in the graph F associated with the treed bridge from Section 3.1 since F has a single cycle. See also the proof of [8, Theorem 5] for a lower bound which immediately implies the one in the statement of the lemma in the finite-volume case. \square

3.4 Pruning the UIHPQ to get the UIHPQ_S

Recall from Section 1.4 that the UIHPQ_S is the Benjamini-Schramm limit of uniformly random quadrangulations with simple boundary, as viewed from a uniformly random vertex on the boundary, as the area and then the perimeter tend to ∞ . In this subsection we explain how to *prune* an instance of the UIHPQ to obtain an instance of the UIHPQ_S.

Suppose $(Q_\infty, \mathfrak{e}_\infty)$ is a UIHPQ. There are infinitely many vertices $v \in \mathcal{V}(\partial Q_\infty)$ which have multiplicity at least 2 in the external face, so are hit twice by the boundary path $\lambda_\infty : \mathbb{Z} \rightarrow \mathcal{E}(\partial Q_\infty)$ of Remark 3.4. Attached to each of these vertices is a finite dangling quadrangulation which is disconnected from ∞ in Q_∞ by removing a single boundary vertex.

Let Q_S be the largest subgraph of Q_∞ with the property that none of its vertices or edges can be disconnected from ∞ in Q_∞ by removing a single boundary vertex. In other words, Q_S is obtained by removing all of the “dangling quadrangulations” of Q_∞ .

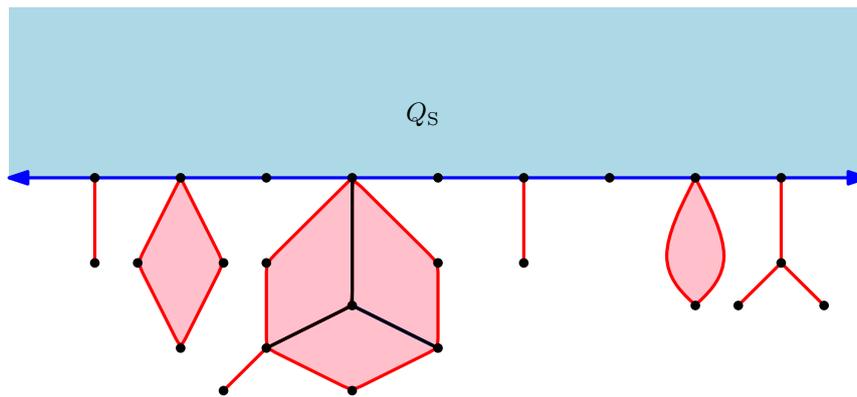


Figure 4: Pruning the red dangling quadrangulations from the UIHPQ Q_∞ (light blue and pink regions with blue and red boundary) produces a UIHPQ $_S$ Q_S (light blue region with blue boundary).

which are joined to ∞ by a single vertex. Also let e_S be the edge immediately to the left of the vertex which can be removed to disconnect e_∞ from ∞ (if such a vertex exists) or let $e_S = e_\infty$ if e_∞ belongs to ∂Q_S . Then (Q_S, e_S) is an instance of the UIHPQ $_S$.

One can recover a canonical boundary path $\lambda_S : \mathbb{Z} \rightarrow \mathcal{E}(\partial Q_S)$ which traces the edges of ∂Q_∞ and hits the terminal endpoint of e_S at time 0 from the analogous boundary path λ_∞ of the UIHPQ by skipping all of the intervals of time during which λ_S is tracing a dangling quadrangulation.

It is shown in [19, Section 6.1.2] and explained more explicitly in [15, Section 6] that if we start with an instance (Q_S, e_S) of the UIHPQ $_S$, we can construct a UIHPQ (Q_∞, e_∞) which can be pruned as above to recover (Q_S, e_S) via an explicit sampling procedure. Conditional on Q_S , let $\{(q_v, e_v)\}_{v \in \mathcal{V}(\partial Q_S)}$ be an independent sequence of random finite quadrangulations with general boundary with an oriented boundary root edge, with distributions described as follows. Let v_0 be the right endpoint of the root edge e_S . Each q_v for $v \neq v_0$ is distributed according to the so-called *free Boltzmann distribution on quadrangulations with general boundary*, which is given by

$$\mathbb{P}[(q_v, e_v) = (\mathfrak{q}, \epsilon)] = C^{-1} \left(\frac{1}{12}\right)^n \left(\frac{1}{8}\right)^l \tag{3.7}$$

for any quadrangulation \mathfrak{q} with n interior faces and $2l$ boundary edges (counted with multiplicity) with a distinguished oriented root edge $\epsilon \in \partial \mathfrak{q}$, where here $C > 0$ is a normalizing constant. The quadrangulation q_{v_0} is instead distributed according to

$$\mathbb{P}[(q_{v_0}, e_{v_0}) = (\mathfrak{q}, \epsilon)] = \tilde{C}^{-1} (2l + 1) \left(\frac{1}{12}\right)^n \left(\frac{1}{8}\right)^l \tag{3.8}$$

for a different normalizing constant \tilde{C} . (Intuitively, the reason for the extra factor $2l + 1$ in (3.8) is that a dangling quadrilateral with longer boundary length is more likely to contain the root edge.)

For each $v \in \mathcal{V}(\partial Q_S)$, identify the terminal endpoint of e_v with v . This gives us a new quadrangulation Q_∞ with general boundary. We choose an oriented root edge e_∞ for Q_∞ by uniformly sampling one of the oriented edges of $\mathcal{E}(\partial q_{v_0}) \cup \{e_S\}$. Then (Q_∞, e_∞) is a UIHPQ with general boundary which can be pruned to recover (Q_S, e_S) .

3.5 The Brownian disk

In this subsection we will review the definition of the Brownian disk from [10], which is the scaling limit of finite uniform quadrangulations with boundary [10]. The construction is a finite-volume analog of the construction of the Brownian half-plane in Section 1.5 and a continuum analog of the Schaeffer-type construction in Section 3.1. We follow closely the exposition given in [30, Section 3.1].

Fix an area parameter $a > 0$ and a boundary length parameter $\ell > 0$. Let $X : [0, a] \rightarrow [0, \infty)$ be a standard Brownian motion started from ℓ and conditioned to hit 0 for the first time at time a (such a Brownian motion is defined rigorously in, e.g., [10, Section 2.1]). For $s, t \in [0, 1]$, set

$$d_X(s, t) := X(s) + X(t) - 2 \inf_{u \in [s \wedge t, s \vee t]} X(u). \tag{3.9}$$

Conditioned on X , let Z^0 be the centered Gaussian process with

$$\text{Cov}(Z^0(s), Z^0(t)) = \inf_{u \in [s \wedge t, s \vee t]} \left(X(u) - \inf_{v \in [0, u]} X(v) \right), \quad s, t \in [0, a]. \tag{3.10}$$

Using the Kolmogorov continuity criterion, one can check that Z^0 a.s. admits a continuous modification which is α -Hölder continuous for each $\alpha < 1/4$. For this modification we have $Z_s^0 = Z_t^0$ whenever $d_X(s, t) = 0$.

Let \mathfrak{b} be $\sqrt{3}$ times a Brownian bridge from 0 to 0 independent from (X, Z) with time duration ℓ . For $r \in [0, \ell]$, let

$$T(r) := \inf\{t \geq 0 : X_t = \ell - r\} \tag{3.11}$$

and for $t \in [0, a]$, let $T^{-1}(t) := \sup\{r \in [0, \ell] : T(r) \leq t\}$. Set

$$Z(t) := Z^0(t) + \mathfrak{b}(T^{-1}(t)).$$

We view $[0, a]$ as a circle by identifying 0 with a and for $s, t \in [0, a]$ we define $\underline{Z}(s, t)$ to be the minimal value of Z on the counterclockwise arc of $[0, a]$ from s to t . For $s, t \in [0, a]$, define

$$d_Z(s, t) = Z(s) + Z(t) - 2(\underline{Z}(s, t) \vee \underline{Z}(t, s)) \tag{3.12}$$

and

$$d^0(s, t) = \inf \sum_{i=1}^k d_Z(s_i, t_i) \tag{3.13}$$

where the infimum is over all $k \in \mathbb{N}$ and all $2k + 2$ -tuples $(t_0, s_1, t_1, \dots, s_k, t_k, s_{k+1}) \in [0, a]^{2k+2}$ with $t_0 = s, s_{k+1} = t$, and $d_X(t_{i-1}, s_i) = 0$ for each $i \in [1, k + 1]_{\mathbb{Z}}$. Equivalently, d^0 is the largest pseudometric on $[0, a]$ which is at most d_Z and is zero whenever d_X is 0.

The *Brownian disk* with area a and perimeter ℓ is the quotient space $H = [0, a] / \{d^0 = 0\}$ equipped with the quotient metric, which we call d . It is shown in [10] that (H, d) is a.s. homeomorphic to the closed disk.

Let $\mathbb{p} : [0, a] \rightarrow H$ for the quotient map. The *area measure* of H is the pushforward μ of Lebesgue measure on $[0, a]$ under \mathbb{p} . The *boundary* of H is the set $\partial H = \mathbb{p}(\{T_r : r \in [0, \ell]\})$ (this set is the topological boundary of H by [8, Proposition 21], and is homeomorphic to the circle). We note that ∂H has a natural orientation, obtained by declaring that the path $t \mapsto \mathbb{p}(t)$ traces ∂H in the counterclockwise direction. The *boundary measure* of H is the pushforward ν of Lebesgue measure on $[0, \ell]$ under $r \mapsto \mathbb{p}(T(r))$. The *boundary path* of H is the curve $\eta : [0, \ell] \rightarrow \partial H$ defined by $\eta(r) = \mathbb{p}(T(r))$.

By [45, Corollary 1.5], the law of the metric measure space (H, d, μ, ν) is the same as that of the $\sqrt{8/3}$ -LQG disk with area a and boundary length ℓ , equipped with its $\sqrt{8/3}$ -LQG area measure and boundary length measure.

4 Scaling limit of the UIHPQ and UIHPQ_S

In this section we will prove Theorems 1.11 and 1.12. We will extract Theorem 1.11 from [10] in a manner which is similar to that in which the scaling limit result for the uniform infinite planar quadrangulation without boundary (UIPQ) in [18] was extracted from [37, 41]. We start in Section 4.1 by showing that one can improve the Gromov-Hausdorff scaling limit result for finite uniformly random quadrangulations with boundary toward the Brownian disk [10] to get convergence in the GHPU topology. In Section 4.2, we will show that one can couple an instance of the Brownian half-plane with a Brownian disk in such a way that metric balls of a certain radius centered at the root point coincide with high probability. In Section 4.3, we prove an analogous coupling result for the UIHPQ with a finite uniformly random quadrangulation with boundary. In Section 4.4, we will deduce Theorem 1.11 from these coupling results and the scaling limit result of Section 4.1.

In Section 4.5, we will deduce Theorem 1.12 from Theorem 1.11 using the pruning procedure discussed in Section 3.4.

4.1 Convergence to the Brownian disk in the GHPU topology

It is proven in [10] that uniformly random quadrangulations with boundary converge in the scaling limit to the Brownian disk in the Gromov-Hausdorff topology, and it is not hard to see from the proof in [10] that one has convergence in the stronger GHPU topology as well. We will explain why this is the case just below.

Fix $\ell > 0$ and let (H, d) be a unit area Brownian disk with boundary length ℓ . Let μ (resp. η) be the natural area measure (resp. boundary path) on H , as in Section 3.5. Let

$$\mathfrak{H} := (H, d, \mu, \eta) \quad (4.1)$$

so that \mathfrak{H} is an element of $\mathcal{M}^{\text{GHPU}}$.

Let $\{l^n\}_{n \in \mathbb{N}}$ be a sequence of positive integers with $(2n)^{-1/2}l^n \rightarrow \ell$. For $n \in \mathbb{N}$, let $(Q^n, \mathfrak{e}_0^n, \mathfrak{v}_*^n)$ be sampled uniformly from $\mathcal{Q}^\bullet(n, l^n)$ (Section 1.4). We view Q^n as a connected metric space in the manner of Remark 1.2. Let d^n be the graph distance on Q^n , re-scaled by $(9/8)^{1/4}n^{-1/4}$. Let μ^n be the measure on $\mathcal{V}(Q^n)$ which assigns a mass to each vertex equal to $(4n)^{-1}$ times its degree. Let $\lambda^n : [0, 2l^n] \rightarrow \partial Q^n$ be the boundary path of Q^n started from \mathfrak{e}_0^n , extended by linear interpolation. Let $\eta^n(t) := \lambda^n(2^{3/2}n^{1/2}t)$ for $t \in [0, 2^{-1/2}n^{-1/2}l^n]$. Let

$$\mathfrak{Q}^n := (Q^n, d^n, \mu^n, \eta^n). \quad (4.2)$$

Theorem 4.1. *In the setting described just above, we have $\mathfrak{Q}^n \rightarrow \mathfrak{H}$ in law in the GHPU topology.*

Proof. We will deduce the theorem statement from the scaling limit result in [10] together with Lemma 2.14. The key point is that the encoding processes for Q^n from Section 3.1 converge jointly with the metric spaces (Q^n, d^n) to the encoding processes for (H, d) in Section 3.5 and the metric space (H, d) , in the uniform topology and the Gromov-Hausdorff topology, respectively; and the measures and curves defined above are determined by the encoding processes in a relatively simple way. This allows us to check the conditions of Lemma 2.14.

Let (X, Z, \mathfrak{b}) be the encoding process for (H, d) and for $r \in [0, \ell]$ let $T(r)$ be the first time that $X(t) = \ell - r$, as in (3.11). Let $\mathfrak{p} : [0, 1] \rightarrow H$ be the quotient map and for $s, t \in [0, 1]$ write

$$\tilde{d}(s, t) := d(\mathfrak{p}(s), \mathfrak{p}(t))$$

so that \tilde{d} is a pseudometric on $[0, 1]$. Recall that μ is the pushforward of Lebesgue measure on $[0, a]$ under \mathbb{p} and $\eta(r) = \mathbb{p}(T(r))$ for $r \in [0, \ell]$. Let ν be the pushforward of Lebesgue measure on $[0, \ell]$ under η , i.e. the boundary length measure of H .

For $n \in \mathbb{N}$, let $(C^n, L^n, b^{0,n})$ be the Schaeffer encoding triple of $(Q^n, \mathfrak{e}_0^n, \mathfrak{v}_*^n)$ as in Section 3.1. For $m \in [1, l^n]_{\mathbb{Z}}$, define $I^n(m) := \min\{i \in [1, 2n + l^n]_{\mathbb{Z}} : C^n(i) = -m\}$, as in (3.2). Extend C^n and L^n to $[0, 2n]$ by linear interpolation and for $t \in [0, 1]$, define

$$X^n(t) := (2n)^{-1/2}(C^n((2n + l^n)t) + l^n) \quad \text{and} \quad Z^n(t) := \left(\frac{9}{8n}\right)^{1/4} L^n((2n + l^n)t).$$

For $r \in [0, (2n)^{-1/2}l^n]$, also define (in analogy with (3.11))

$$T^n(r) := (2n)^{-1}I^n(\lfloor(2n)^{1/2}r\rfloor). \tag{4.3}$$

Let $\mathbb{p}^n : [0, 2n + l^n]_{\mathbb{Z}} \rightarrow \mathcal{V}(Q^n)$ be the contour exploration as in Section 3.1. For $s, t \in [0, 1]$ with $(2n + l^n)s, (2n + l^n)t \in \mathbb{N}_0$, let

$$\tilde{d}^n(s, t) := d^n(\mathbb{p}^n((2n + l^n)s), \mathbb{p}^n((2n + l^n)t))$$

and extend \tilde{d}^n to $[0, 1]^2$ by linear interpolation.

In the variant of the Schaeffer bijection described in Section 3.1, we add one edge from the vertex $\mathbb{p}^n(i)$ to the successor vertex (which has a smaller label) for each $i \in [0, 2n + l^n]_{\mathbb{Z}}$. Consequently, if we let $\hat{\mu}^n$ be the pushforward under \mathbb{p}^n of $(2n)^{-1}$ times counting measure on $[0, 2n + l^n]_{\mathbb{Z}}$, then for $v \in \mathcal{V}(Q^n)$ it holds that $\hat{\mu}^n(v)$ equals $(2n)^{-1}$ times the number of edges of Q^n which connect v to a vertex with a smaller label. Since μ^n assigns mass to each vertex equal to $(4n)^{-1}$ times its degree, the d^n -Prokhorov distance between $\hat{\mu}^n$ and μ^n is at most a universal constant times $n^{-1/4}$ (note that the scaling factors for $\hat{\mu}^n$ and μ^n differ by a factor of 2 since each edge is counted twice—once for each of its endpoints—when considering the measure μ^n).

Let ν^n be the pushforward of $2^{-3/2}n^{-1/2}$ times counting measure on $[1, 2l^n]_{\mathbb{Z}}$ under the (linearly interpolated) boundary path λ^n . Equivalently ν^n is the counting measure on $\mathcal{V}(\partial Q^n)$ (with vertices counted with multiplicity), rescaled by $2^{-3/2}n^{-1/2}$. The measure ν^n does not admit a simple description in terms of $(C^n, L^n, b^{0,n})$ but we can describe a closely related measure as follows.

Let $\hat{\nu}^n$ be the pushforward under $\mathbb{p}^n \circ I^n$ of $2^{-1/2}n^{-1/2}$ times the counting measure on $[1, l^n]_{\mathbb{Z}}$ (with I^n as above). Note that the scaling factor here is off by a factor of 2 as compared to the scaling factor of ν^n so that $\hat{\nu}^n$ has the same total mass as ν^n . The measure $\hat{\nu}^n$ can equivalently be described as follows. Let $j_0^n = 0$ and for $k \in [1, l^n]_{\mathbb{Z}}$, let j_k^n be the k th smallest downward step of the random walk bridge $b^{0,n}$. Then $\hat{\nu}^n$ is the measure on $\mathcal{V}(\partial Q^n)$ which assigns mass $2^{-1/2}n^{-1/2}$ to the j_k^n th vertex of ∂Q^n in counterclockwise cyclic order started from the root vertex. Since $b^{0,n}$ is a simple random walk bridge, we can use Hoeffding’s concentration inequality for binomial random variables to find that except on an event of probability decaying faster than any power of n ,

$$\max_{k \in [1, l^n]_{\mathbb{Z}}} |j_k^n - 2k| = n^{o_n(1)} \tag{4.4}$$

which implies that $\lambda^n(2k)$ differs from $\mathbb{p}^n(I^n(k))$ for $k \in \mathbb{N}$ by at most $n^{o_n(1)}$ units of boundary length and

$$|\nu^n(A) - \hat{\nu}^n(A)| \leq n^{-1/2+o_n(1)}, \quad \forall A \subset \partial Q^n.$$

It is shown in [10, Section 5] that, in the notation above,

$$(X^n, Z^n, \tilde{d}^n) \rightarrow (X, Z, \tilde{d}) \quad \text{and} \quad (Q^n, d_{Q^n}, \eta^n(0)) \rightarrow (H, d, \eta(0)) \tag{4.5}$$

in law in the uniform topology and the pointed GH topology, respectively. With T and T^n the hitting time processes as in (3.11) and (4.3), respectively, we have $(X^n, Z^n, T^n) \rightarrow (X, Z, T)$ in law in the uniform topology in the first two coordinates and the Skorokhod topology in the third coordinate. Since (X, Z) a.s. determines T , \tilde{d} , and $(H, d, \eta(0))$, the convergence (4.5) in law occurs jointly with the convergence $T^n \rightarrow T$ in law.

By the Skorokhod representation theorem, we can couple $\{(Q^n, e_0^n, \nu_*^n)\}_{n \in \mathbb{N}}$ with (BD, d, x) in such a way that the convergence (4.5) occurs a.s. and also $T^n \rightarrow T$ a.s. in the Skorokhod topology. Henceforth fix a coupling for which this is the case. Note that the Borel-Cantelli lemma implies that in any such coupling, (4.4) holds for large enough n .

By [30, Lemma 3.2], re-phrased in terms of the pseudometric \tilde{d} on $[0, 1]$, there a.s. exists $C > 0$ such that for each $r_1, r_2 \in [0, \ell]$ with $r_1 < r_2$,

$$\tilde{d}(T(r_1), T(r_2)) \leq C(r_2 - r_1)^{1/2}(|\log(r_2 - r_1)| + 1)^2. \tag{4.6}$$

The Skorokhod convergence $T^n \rightarrow T$ together with the uniform convergence $\tilde{d}^n \rightarrow \tilde{d}$ therefore implies that for each $\epsilon > 0$, there a.s. exists $\delta > 0$ such that for each $n \in \mathbb{N}$ and each $r_1, r_2 \in [0, \ell]$ with $|r_1 - r_2| \leq \delta$, we have $\tilde{d}^n(T^n(r_1), T^n(r_2)) \leq \epsilon$. From this together with (4.4) and the discussion just after, we infer that the re-scaled boundary paths $\{\eta^n\}_{n \in \mathbb{N}}$ are equicontinuous.

We will now apply Lemma 2.14 to deduce our desired GHPU convergence. The lemma does not apply directly in our setting since the curve η is not simple (it satisfies $\eta(0) = \eta(\ell)$), so a straightforward truncation argument is needed. For $r \in [0, (2n)^{-1/2}\ell^n]$, let $\hat{\nu}_r^n$ be the restriction of $\hat{\nu}^n$ to the counterclockwise arc of ∂Q^n started from $\eta^n(0)$ with $\hat{\nu}^n$ -length r . Also let ν_r^n be the pushforward of Lebesgue measure on $[0, r]$ under η^n . For $r \in [0, \ell]_{\mathbb{Z}}$, let ν_r be the pushforward of Lebesgue measure on $[0, r]$ under η .

Our choice of coupling and our above description of the relationship between the measures μ^n and $\hat{\mu}^n$,

$$(Q^n, d^n, \mu^n, \hat{\nu}_r^n, \eta^n(0)) \rightarrow (H, d, \mu, \nu_r, \eta(0))$$

in the 2-fold pointed Gromov-Prokhorov topology for each $r \in (0, \ell)$. By (4.4), we also have this convergence with ν_r^n in place of $\hat{\nu}_r^n$. By Lemma 2.14, we infer that

$$(Q^n, d^n, \mu^n, \eta^n|_{[0, r]}) \rightarrow (H, d, \mu, \eta|_{[0, r]})$$

in the GHPU topology. Since r can be made arbitrarily close to ℓ , by combining this with equicontinuity of the curves η^n we see that in fact $\Omega^n \rightarrow \mathfrak{H}$ in the GHPU topology. \square

4.2 Coupling the Brownian half-plane with the Brownian disk

In this section, we will prove the following coupling statement, which is an analog of [18, Proposition 4] for Brownian surfaces with boundary. This coupling statement is needed for the proof of Theorem 1.11 and also implies Proposition 1.10, as explained just below the proposition statement.

Proposition 4.2. *Fix $a, \ell > 0$ and let (H, d) be a Brownian disk with area a and boundary length ℓ . Also let (H_∞, d_∞) be a Brownian half-plane. For each $\epsilon \in (0, 1)$ there exists $\alpha > 0$ and a coupling of (H, d) with (H_∞, d_∞) such that the following is true. Let μ (resp. μ_∞) and η (resp. η_∞) be the area measure and boundary path of H (resp. H_∞). In the notation of Definition 1.6, it holds with probability at least $1 - \epsilon$ that*

$$\mathfrak{B}_\alpha(H, d, \mu, \eta) \quad \text{and} \quad \mathfrak{B}_\alpha(H_\infty, d_\infty, \mu_\infty, \eta_\infty).$$

agree as elements of $\overline{\mathfrak{M}}^{\text{GHPU}}$.

Note that scaling distances in a Brownian surface by $r > 0$ corresponds to scaling boundary lengths by r^2 and areas by r^4 . Indeed, this follows by the scaling properties of Brownian motion. Hence Proposition 4.2 (applied with $a = \ell = 1$) implies that for any $r > 0$ and any $\epsilon \in (0, 1)$, there exists $R > 0$ such that if (H, d, μ, η) is as in Proposition 4.2 for $a = R$ and $\ell = R^{1/2}$, then we can couple (H, d) with (H_∞, d_∞) in such a way that $\mathfrak{B}_r(H, d, \mu, \eta)$ and $\mathfrak{B}_r(H_\infty, d_\infty, \mu_\infty, \eta_\infty)$ coincide as elements of $\overline{\mathfrak{M}}^{\text{GHPU}}$ with probability at least $1 - \epsilon$.

The proof of Proposition 4.2 requires two lemmas. The first is a continuum analog of Lemma 3.7.

Lemma 4.3. *Suppose we are in the setting of Section 3.5. In particular, fix $a, \ell > 0$, let (H, d) be a Brownian disk with area a and boundary length ℓ , let $\mathbb{p} : [0, a] \rightarrow H$ be the quotient map, let (X, Z) be the encoding process, and let $T(r)$ for $r \in [0, \ell]$ be as in (3.11). Almost surely, the following is true. For $s, t \in [0, a]$ with $s \leq t$, let*

$$S(s, t) := (T([0, \ell]) \cap [s, t]) \cup \{s, t\} \quad \text{and} \quad S'(s, t) := (T([0, \ell]) \setminus [s, t]) \cup \{s, t\}$$

so that $\mathbb{p}(S(s, t))$ (resp. $\mathbb{p}(S'(s, t))$) consists of $\mathbb{p}(s)$, $\mathbb{p}(t)$, and the set of points in ∂H which are (resp. are not) contained in $\mathbb{p}([s, t])$. Then

$$d(\mathbb{p}(s), \mathbb{p}(t)) \geq Z(s) + Z(t) - 2 \max \left\{ \inf_{u \in S(s, t)} Z(u), \inf_{u \in S'(s, t)} Z(u) \right\}. \tag{4.7}$$

The analogous estimate also holds in the setting of the Brownian half-plane (note that in this case the infimum over $S'(s, t)$ in (4.7) is a.s. equal to $-\infty$).

Proof. This follows from essentially the same argument used to prove the cactus bound in the case of the Brownian map (see, e.g., [18, Equation (4)]). See also [8, Lemma 22] for an analogous estimate in the case when $\mathbb{p}(s), \mathbb{p}(t) \in \partial H$. □

Next we prove a coupling statement for the encoding processes.

Lemma 4.4. *Fix $a, \ell > 0$ and let (X, Z) be the encoding process for the Brownian disk with area a and boundary length ℓ and let T be as in (3.11). Also let (X_∞, Z_∞) be the encoding process for the Brownian half-plane and for $r \in \mathbb{R}$ let $T_\infty(r) = \inf\{t \in [0, \ell] : X_\infty(t) = -r\}$, as in Section 1.5. For each $\epsilon, \delta \in (0, 1)$, there exists $0 < \delta_0 < \delta_1 \leq \delta$ and a coupling of (X, Z) with (X_∞, Z_∞) such that with probability at least $1 - \epsilon$, the following is true. We have $T(r) = T_\infty(r)$ for each $r \in [0, \delta_1]$, $X(t) - \ell = X_\infty(t)$ for each $t \in [0, T(\delta_1)]$, and for each $t \in [T_\infty(\delta_0), T_\infty(\delta_1)]$,*

$$Z(t) - Z(T(\delta_0)) = Z_\infty(t) - Z_\infty(T_\infty(\delta_0)).$$

Proof. By [10, Proposition 9], for each $r \in [0, \ell]$ the Radon-Nikodym derivative of the law of $(X - \ell)|_{[0, T(r)]}$ with respect to the law of $X_\infty|_{[0, T_\infty(r)]}$ is given by $f_r(T_\infty(r))$, where

$$f_r(t) := \mathbb{1}_{(t < a)} \frac{(\ell - r)a^{3/2}}{\ell(a - t)^{3/2}} \exp\left(\frac{\ell^2}{2r} - \frac{(\ell - r)^2}{2(a - t)}\right).$$

Since $f_r(t)$ is continuous in both variables and $f_0(0) = 1$, if we are given $\epsilon \in (0, 1)$, then we can find $\zeta \in [0, \ell]$ such that $|f_r(t) - 1| \leq \epsilon/4$ for $t, r \in [0, \zeta]$. Choose $\delta_1 \in (0, \delta \wedge \zeta]$ such that $\mathbb{P}[T_\infty(\delta_1) \leq \zeta] \geq 1 - \epsilon/4$. Then except on an event of probability at most $\epsilon/4$, the Radon-Nikodym derivative of the law of $(X - \ell)|_{[0, T(\delta_1)]}$ with respect to the law of $X_\infty|_{[0, T_\infty(\delta_1)]}$ is between $1 - \epsilon/4$ and $1 + \epsilon/4$, so we can couple these two restricted processes so that they agree with probability at least $1 - \epsilon/2$. Henceforth assume we have chosen such a coupling.

Recalling (3.10) and (1.9), we see that for any fixed realization x of $(X - \ell)|_{[0, T(\delta_1)]}$, the regular conditional law of $Z^0|_{[0, T(\delta_1)]}$ given $\{(X - \ell)|_{[0, T(\delta_1)]} = x\}$ coincides with the regular conditional law of $Z_\infty|_{[0, T_\infty(\delta_1)]}$ given $\{X_\infty|_{[0, T_\infty(\delta_1)]} = x\}$. Hence we can extend our coupling so that in fact

$$T(\delta_1) = T_\infty(\delta_1) \quad \text{and} \quad (X - \ell, Z^0)|_{[0, T(\delta_1)]} = (X_\infty, Z_\infty)|_{[0, T_\infty(\delta_1)]} \tag{4.8}$$

with probability at least $1 - \epsilon/2$.

Recall that Z (resp. Z_∞) is obtained from Z^0 (resp. Z_∞^0) by adding the composition of \mathfrak{b} (resp. \mathfrak{b}_∞) with its running infimum process, where \mathfrak{b} is $\sqrt{3}$ times a Brownian bridge and \mathfrak{b}_∞ is $\sqrt{3}$ times a standard linear Brownian motion.

For $0 < \delta_0 < \delta_1$, the law of $(\mathfrak{b} - \mathfrak{b}(\delta_0))|_{[\delta_0, \delta_1]}$ is absolutely continuous with respect to the law of $(\mathfrak{b}_\infty - \mathfrak{b}_\infty(\delta_0))|_{[\delta_0, \delta_1]}$, with Radon-Nikodym derivative tending to 1 in probability as δ_0 increases to δ_1 . We can find $\delta_0 \in (0, \delta_1)$ for which this Radon-Nikodym derivative lies in $[1 - \epsilon/4, 1 + \epsilon/4]$ with probability at least $1 - \epsilon/4$, so we can couple $(\mathfrak{b} - \mathfrak{b}(\delta_0))|_{[\delta_0, \delta_1]}$ with $(\mathfrak{b}_\infty - \mathfrak{b}_\infty(\delta_0))|_{[\delta_0, \delta_1]}$ in such a way that

$$(\mathfrak{b} - \mathfrak{b}(\delta_0))|_{[\delta_0, \delta_1]} = (\mathfrak{b}_\infty - \mathfrak{b}_\infty(\delta_0))|_{[\delta_0, \delta_1]} \tag{4.9}$$

with probability at least $1 - \epsilon/2$.

Since \mathfrak{b} (resp. \mathfrak{b}_∞) is independent from (X, Z^0) (resp. (X_∞, Z_∞^0)), we can couple (X, Z, \mathfrak{b}) and $(X_\infty, Z_\infty, \mathfrak{b}_\infty)$ so that (4.8) and (4.9) hold simultaneously with probability at least $1 - 4\epsilon$. For such a coupling, the conditions in the statement of the lemma are satisfied. \square

Proof of Proposition 4.2. Let (X, Z) and (X_∞, Z_∞) be the encoding processes, as in Proposition 4.2. By Proposition 4.2, for each $\epsilon, \delta \in (0, 1)$ there exists $0 < \delta_0 < \delta_1 \leq \delta$ and a coupling of (X, Z) with (X_∞, Z_∞) such that with probability at least $1 - \epsilon/4$, we have $T(r) = T_\infty(r)$ for each $r \in [0, \delta_1]$, $X(t) - \ell = X_\infty(t)$ for each $t \in [0, T(\delta_1)]$, and for each $t \in [T_\infty(\delta_0), T_\infty(\delta_1)]$,

$$Z(t) - Z(T(\delta_0)) = Z_\infty(t) - Z_\infty(T_\infty(\delta_0)).$$

Note that with \mathfrak{b}_∞ (resp. \mathfrak{b}) the Brownian motion (resp. Brownian bridge) from Section 1.5 (resp. Section 3.5), we have $Z_\infty \circ T = \mathfrak{b}_\infty$ (resp. $Z \circ T = \mathfrak{b}$), so (4.2) implies that also $\mathfrak{b}(r) - \mathfrak{b}(\delta_0) = \mathfrak{b}_\infty(r) - \mathfrak{b}_\infty(\delta_0)$ for each $r \in [\delta_0, \delta_1]$.

If we choose δ_0 and δ_1 sufficiently small, then it is unlikely that the infimum of \mathfrak{b} over $[0, \delta_0]$ is smaller than the infimum of \mathfrak{b} over $[\delta_1, \ell]$. Since \mathfrak{b} is a constant times a Brownian bridge, $\mathfrak{b}(\delta_0)$ is a.s. strictly larger than the minimum value of \mathfrak{b} on each of $[0, \delta_0]$ and $[\delta_1, \ell]$. Furthermore, if we choose δ_1 sufficiently close to δ_0 (leaving δ_0 fixed), then by continuity it is likely that $|Z(s) - \mathfrak{b}(\delta_0)|$ is close to 0 for each $s \in [T(\delta_0), T(\delta_1)]$. Hence by choosing δ_0 sufficiently small and then choosing δ_1 sufficiently close to δ_0 , we can arrange that except on an event of probability at most $\epsilon/4$,

$$\max \left\{ \inf_{r \in [0, \delta_0]} \mathfrak{b}(r), \inf_{r \in [\delta_1, \ell]} \mathfrak{b}(r) \right\} < \inf_{s \in [T(\delta_0), T(\delta_1)]} Z(s). \tag{4.10}$$

Let $\delta_* := \frac{1}{2}(\delta_0 + \delta_1)$. Also define

$$\rho := - \left(\inf_{r \in [\delta_0, \delta_*]_{\mathbb{Z}}} (\mathfrak{b}_\infty(r) - \mathfrak{b}_\infty(\delta_0)) \vee \inf_{r \in [\delta_*, \delta_1]_{\mathbb{Z}}} (\mathfrak{b}_\infty(r) - \mathfrak{b}_\infty(\delta_0)) \right) \geq 0$$

and

$$\rho' := \mathfrak{b}_\infty(\delta_*) - \mathfrak{b}_\infty(\delta_0) + \rho.$$

Since \mathfrak{b}_∞ is a constant times a standard Brownian motion, there exists $\alpha > 0$ such that with probability at least $1 - \epsilon/2$ we have $\frac{1}{4}\rho' \geq \alpha$.

By definition, the quotient map $\mathbb{p}_\infty : \mathbb{R} \rightarrow H$ satisfies $\mathbb{p}_\infty(T_\infty(\delta_*)) = \eta_\infty(\delta_*)$. By invariance of the law of the Brownian half-plane under re-rooting along the boundary, which follows from (1.6),

$$(H_\infty, d_\infty, \mu_\infty, \eta_\infty(\delta_* + \cdot)) \stackrel{d}{=} (H_\infty, d_\infty, \mu_\infty, \eta_\infty).$$

It is immediate from Theorem 4.1 that the analogous property also holds for the Brownian disk. Hence it suffices to show that whenever (4.2) and (4.10) occur,

$$\mathfrak{B}_{\frac{1}{4}\rho'}(H, d, \mu, \eta(\delta_* + \cdot)) \quad \text{and} \quad \mathfrak{B}_{\frac{1}{4}\rho'}(H_\infty, d_\infty, \mu_\infty, \eta_\infty(\delta_* + \cdot)) \tag{4.11}$$

agree as elements of $\overline{\mathbb{M}}^{\text{GHPU}}$.

Henceforth assume that (4.2) and (4.10) occur (which happens with probability at least $1 - \epsilon/2$). Let $\mathbb{p} : [0, a] \rightarrow H$ and $\mathbb{p}_\infty : \mathbb{R} \rightarrow H_\infty$ be the quotient maps. Let $r_0 \in [\delta_0, \delta_*]_{\mathbb{Z}}$ and $r_1 \in [\delta_*, \delta_1]_{\mathbb{Z}}$ be chosen so that

$$\mathfrak{b}_\infty(r_0) - \mathfrak{b}_\infty(r_1) \leq -\rho \quad \text{and} \quad \mathfrak{b}_\infty(r_1) - \mathfrak{b}_\infty(\delta_0) \leq -\rho.$$

By Lemma 4.3 together with (4.2) and (4.10) (the latter equation is used to deal with the second term in the maximum in (4.7)), if $t \geq T(r_1)$ then

$$\begin{aligned} d(\mathbb{p}(t), \eta(\delta_*)) &\geq Z(t) + \mathfrak{b}(\delta_*) - 2 \max \left\{ \inf_{u \in S(T(\delta_*), t)} Z(u), \inf_{u \in S'(T(\delta_*), t)} Z(u) \right\} \\ &\geq Z(t) + \mathfrak{b}(\delta_*) - Z(t) + \rho - \mathfrak{b}(\delta_*) \geq \rho'. \end{aligned}$$

We have a similar estimate if $t \leq T(r_0)$. Again using Lemma 4.3, we similarly obtain that if $t \in [0, a] \setminus [T(r_0), T(r_1)]$ then $d_\infty(\mathbb{p}_\infty(t), \eta_\infty(\delta_*)) \geq \rho'$.

Hence

$$B_{\rho'}(\eta(\delta_*); d) \subset \mathbb{p}([T(r_0), T(r_1)]) \quad \text{and} \quad B_{\rho'}(\eta_\infty(\delta_*); d_\infty) \subset \mathbb{p}_\infty([T_\infty(r_0), T_\infty(r_1)]).$$

From the definitions (1.11) and (3.13) of the pseudometrics d^0 and d_∞^0 , respectively, and the triangle inequality, we find that the sets $\mathbb{p}^{-1}(B_{\frac{1}{4}\rho'}(\eta(\delta_*); d))$ and $\mathbb{p}_\infty^{-1}(B_{\frac{1}{4}\rho'}(\eta_\infty(\delta_*); d_\infty))$ (resp. the distances $d^0(s, t)$ and $d_\infty^0(s, t)$ for s and t in these sets) are given by the same deterministic functionals of $X|_{[0, T(\delta_1)]}$ and $(Z - Z(T(\delta_0)))|_{[T(\delta_0), T(\delta_1)]}$ and $(Z - Z(T(\delta_0)))|_{[T(\delta_0), T(\delta_1)]}$. Note that we use (4.10) to resolve the discrepancy between the definitions (1.10) and (3.12). The measures μ and μ_∞ and the paths η and η_∞ are determined by the same local functionals of the Schaeffer encoding functions. From these considerations, we see that the map $\mathbb{p}(t) \mapsto \mathbb{p}_\infty(t)$ for $t \in [T_\infty(r_0), T_\infty(r_1)]$ is well defined and restricts to a measure- and curve-preserving isometry from $B_{\frac{1}{4}\rho'}(\eta(\delta_*); d)$ to $B_{\frac{1}{4}\rho'}(\eta_\infty(\delta_*); d_\infty)$. Hence (4.11) holds. \square

4.3 Coupling the UIHPQ with a finite uniform quadrangulation with boundary

In this section we prove a discrete analog of Proposition 4.2, which is also an analog of [18, Proposition 9] for maps with boundary. Throughout this section, we let $(Q_\infty, \mathfrak{e}_\infty)$ be a UIHPQ. We also fix $\ell > 0$, a sequence of positive integers $\{l^n\}_{n \in \mathbb{N}}$ with $(2n)^{-1/2}l^n \rightarrow \ell$, and for $n \in \mathbb{N}$ we let $(Q^n, \mathfrak{e}_0^n, \mathfrak{v}_*^n)$ be sampled uniformly from $\mathcal{Q}^\bullet(n, l^n)$ (defined in Section 1.4). Also let \mathfrak{v}_∞ (resp. \mathfrak{v}_0^n) be the initial endpoint of \mathfrak{e}_∞ (resp. \mathfrak{e}_0^n), so that $\mathfrak{v}_\infty = \lambda_\infty(0)$, with λ_∞ the linearly interpolated boundary path of Q_∞ .

Proposition 4.5. *For each $\epsilon \in (0, 1)$ there exists $\alpha > 0$ and $n_* \in \mathbb{N}$ such that for $n \geq n_*$, there is a coupling of $(Q_\infty, \mathfrak{e}_\infty)$ with $(Q^n, \mathfrak{e}_0^n, \mathfrak{v}_*^n)$ with the following property. With probability at least $1 - \epsilon$, the graph metric balls $B_{\alpha n^{1/4}}(\mathfrak{v}_0^n; Q^n)$ and $B_{\alpha n^{1/4}}(\mathfrak{v}_\infty; Q_\infty)$*

equipped with the graph structures they inherit from Q^n and Q_∞ , respectively, are isomorphic (as graphs) via an isomorphism which takes \mathfrak{e}_0^n to \mathfrak{e}_∞ and $\partial Q^n \cap B_{\alpha n^{1/4}}(\mathfrak{v}_0^n; Q^n)$ to $\partial Q_\infty \cap B_{\alpha n^{1/4}}(\mathfrak{v}_\infty; Q_\infty)$. Furthermore, we can arrange that this isomorphism is an isometry for the metrics on $B_{\alpha n^{1/4}}(\mathfrak{v}_0^n; Q^n)$ and $B_{\alpha n^{1/4}}(\mathfrak{v}_\infty; Q_\infty)$ which they inherit from Q^n and Q_∞ , respectively.

The proof of Proposition 4.5 is similar to that of Proposition 4.2. We will first construct a coupling of the encoding functions, then transfer to a coupling of the maps using Lemma 3.7. The following lemma is needed to bound the Radon-Nikodym derivative of the encoding functions.

Lemma 4.6. *Let $\{S(i)\}_{i \in \mathbb{N}_0}$ be a simple random walk on \mathbb{Z} started from 0, with steps distributed uniformly in $\{-1, 1\}$. Let $\{\mathcal{F}_i\}_{i \in \mathbb{N}_0}$ be the filtration generated by S . For $m \in \mathbb{N}_0$, let $I(m) := \inf\{i \in \mathbb{N}_0 : S(i) = -m\}$. Let N be a stopping time for $\{\mathcal{F}_i\}_{i \in \mathbb{N}_0}$. Then for $m, n \in \mathbb{N}$, it holds on the event $\{N < I(m) \wedge n\}$ that*

$$\mathbb{P}[I(m) = n \mid \mathcal{F}_N] = \frac{1}{\sqrt{2\pi}}(S(N) + m)(n - N)^{-3/2} \exp\left(-\frac{(S(N) + m)^2}{2(n - N)}\right) + o((S(N) + m)^{-2})$$

with the rate of the $o((S_N + m)^{-2})$ error universal and deterministic.

Proof. By Donsker's theorem, we have $m^2 I(m) \rightarrow T$ in law, where T is the first time a standard linear Brownian motion hits -1 . By the local limit theorem for stable laws (see, e.g., [22, Section 50]), we find that for $n \in \mathbb{N}$,

$$\limsup_{n \rightarrow \infty} \sup_{n \in \mathbb{N}} \left| m^2 \mathbb{P}[I(m) = n] - g\left(\frac{n}{m^2}\right) \right| = 0,$$

where

$$g(t) = \frac{1}{\sqrt{2\pi}t^{3/2}} e^{-\frac{1}{2t}}$$

is the density of T . By the strong Markov property, on the event $\{N < I(m) \wedge n\}$ we have

$$\mathbb{P}[I(m) = n \mid \mathcal{F}_N] = \mathbb{P}[I(-s + m) = n - N]_{|s=S_N}.$$

The statement of the lemma follows. □

In what follows, we define the encoding paths $C_\infty, L_\infty^0, L_\infty, b_\infty$, and b_∞^0 as in Section 3.2 and the analogous finite-volume objects $C^n, L_\infty^{0,n}, L^n, b^n$, and $b^{0,n}$ as in Section 3.1 for the quadrangulation Q^n , but with an additional superscript n . For $m \in [0, l^n]$ (resp. $m \in \mathbb{Z}$) we let $I^n(m)$ (resp. $I_\infty(m)$) be the smallest $i \in [0, 2n + l^n]_{\mathbb{Z}}$ (resp. $i \in \mathbb{Z}$) for which $C^n(i) = -m$ (resp. $C_\infty(i) = -m$), as in (3.2) (resp. (3.4)). We extend these functions to $[0, 2n + l]$ (resp. \mathbb{R}) by setting $I^n(s) = I^n(\lfloor s \rfloor)$ (resp. $I_\infty(s) = I_\infty(\lfloor s \rfloor)$). The following lemma gives us a coupling of the Schaeffer encoding functions of the maps in Proposition 4.5.

Lemma 4.7. *For each $\epsilon, \delta \in (0, 1)$, there exists $0 < \delta_0 < \delta_1 \leq \delta$ and $n_* \in \mathbb{N}$ such that for $n \geq n_*$, there exists a coupling of the encoding pairs (C^n, L^n) and (C_∞, L_∞) such that with probability at least $1 - \epsilon$, the following is true. We have $I^n(m) = I_\infty(m)$ for each $m \in [0, \delta_1 n^{1/2}]$. Furthermore, for each $i \in [0, I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}$ we have $C^n(i) = C_\infty(i)$ and for each $i \in [I^n(\delta_0 n^{1/2}), I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}$ we have*

$$L^n(i) - L^n(I^n(\delta_0 n^{1/2})) = L_\infty(i) - L_\infty(I_\infty(\delta_0 n^{1/2})).$$

Proof. Recall that the law of C^n is that of a simple random walk conditioned to first hit $-l^n$ at time $2n + l^n$ and the law of $C_\infty|_{\mathbb{N}_0}$ is that of an unconditioned simple random walk (Lemmas 3.3 and 3.5). By Lemma 4.6 and Bayes' rule, for $\delta \in (0, \ell)$ and $n \in \mathbb{N}$, then

the Radon-Nikodym derivative of the law of $C^n|_{[0, I^n(\delta n^{1/2})]_{\mathbb{Z}}}$ with respect to the law of $C_\infty|_{[0, I_\infty(\delta n^{1/2})]_{\mathbb{Z}}}$ is given by $f_\delta^n(I_\infty(\delta n^{1/2}))$ where for $k \in [0, 2n + l^n]_{\mathbb{Z}}$,

$$f_\delta^n(k) = \frac{(l^n - \lfloor \delta n^{1/2} \rfloor)(2n + l^n - k)^{-3/2} \exp\left(-\frac{(l^n - \lfloor \delta n^{1/2} \rfloor)^2}{2(2n + l^n - k)}\right) + o((l^n - \lfloor \delta n^{1/2} \rfloor)^{-2})}{l^n(2n + l^n)^{-3/2} \exp\left(-\frac{(l^n)^2}{2(2n + l^n)}\right) + o((l^n)^{-2})} \mathbb{1}_{(k < 2n + l^n)}.$$

Since $(2n)^{-1/2}l^n \rightarrow \ell$, if we are given $\epsilon \in (0, 1)$, we can find $\zeta > 0$ and $n_0 \in \mathbb{N}$ such that for $n \geq n_0$, $1 \leq k \leq \zeta n$, and $0 < \delta_1 \leq \zeta$, we have $|f_{\delta_1}^n(k) - 1| \leq \frac{\epsilon}{4}$. Since $(2n)^{-1/2}C^n((2n)^{-1})$ converges in law in the uniform topology to an appropriate conditioned Brownian motion [7, Lemma 14], we can find $\delta_1 \in (0, \delta \wedge \zeta]$ and $n_* \geq n_0$ such that for $n \geq n_*$, it holds with probability at least $1 - \epsilon/4$ that $I(\delta_1 n^{1/2}) \leq \zeta n$. If $n \geq n_*$, then except on an event of probability at most $1 - q/4$, the Radon-Nikodym derivative of the law of $C^n|_{[0, I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}}$ with respect to the law of $C_\infty|_{[0, I_\infty(\delta_1 n^{1/2})]_{\mathbb{Z}}}$ lies in $[1 - \epsilon/4, 1 + \epsilon/4]$. Hence we can couple these restricted processes together so that with probability at least $1 - \epsilon/2$,

$$I^n(\delta_1 n^{1/2}) = I_\infty(\delta_1 n^{1/2}) \quad \text{and} \quad C^n|_{[0, I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}} = C_\infty|_{[0, I_\infty(\delta_1 n^{1/2})]_{\mathbb{Z}}}. \tag{4.12}$$

Since the conditional law of the shifted label function $L^{0,n}|_{[0, I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}}$ given $C^n|_{[0, I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}}$ coincides with the conditional law of $L_\infty^0|_{[0, I_\infty(\delta_1 n^{1/2})]_{\mathbb{Z}}}$ given $C_\infty|_{[0, I_\infty(\delta_1 n^{1/2})]_{\mathbb{Z}}}$, we also obtain a coupling of $(C^n, L^{0,n})$ with (C_∞, L_∞^0) such that with probability at least $1 - \epsilon/2$,

$$(C^n, L^{0,n})|_{[0, I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}} = (C_\infty, L_\infty^0)|_{[0, I_\infty(\delta_1 n^{1/2})]_{\mathbb{Z}}}. \tag{4.13}$$

Recall that L^n (resp. L_∞) is obtained from (C_∞, L_∞^0) and the bridge $b^{0,n}$ (resp. (C_∞, L_∞^0) and the walk b_∞^0) in the manner described in Section 3.1 (resp. Section 3.2). Recall also the processes b^n and b_∞ obtained from $b^{0,n}$ and b_∞^0 , respectively, by considering only times when the path makes a downward step. A similar absolute continuity argument to the one given above shows that after possibly increasing n_* , we can find $\delta_0 \in (0, \delta_1]$ and a coupling of $b^{0,n}$ with b_∞^0 such that with probability at least $1 - \epsilon/2$,

$$b^n(k) - b^n(\lfloor \delta_0 n^{1/2} \rfloor) = b_\infty(k) - b_\infty(\lfloor \delta_0 n^{1/2} \rfloor), \quad \forall k \in [\delta_0 n^{1/2}, \delta_1 n^{1/2}]_{\mathbb{Z}}. \tag{4.14}$$

The pair $(C^n, L^{0,n})$ (resp. (C_∞, L_∞^0)) is independent from $b^{0,n}$ (resp. b_∞^0), so for $n \geq n_*$, we can couple $(C^n, L^{0,n}, b^{0,n})$ with $(C_\infty, L_\infty^0, b_\infty^0)$ in such a way that (4.13) and (4.14) hold simultaneously with probability at least $1 - \epsilon$. Such a coupling satisfies the conditions in the statement of the lemma. \square

Proof of Proposition 4.5. The proof is essentially identical to that of Proposition 4.2, but we give the details for the sake of completeness. By Lemma 4.7, we can find $0 < \delta_0 < \delta_1 < \ell$ and $n_* \in \mathbb{N}$ such that for $n \geq n_*$, there exists a coupling of (C^n, L^n) with (C_∞, L_∞) such that with probability at least $1 - \epsilon/4$, $I^n(m) = I_\infty(m)$ for each $m \in [0, \delta_1 n^{1/2}]$, $C^n(i) = C_\infty(i)$ for each $i \in [0, I^n(\delta_1 n^{1/2})]$, and for each $i \in [I^n(\delta_0 n^{1/2}), I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}$,

$$L^n(i) - L^n(I^n(\delta_0 n^{1/2})) = L_\infty(i) - L_\infty(I_\infty(\delta_0 n^{1/2})). \tag{4.15}$$

Note that with b^n (resp. b_∞) the process from Section 3.1 (resp. Section 3.1), we have $L^n(I^n(m)) = b^n(m + 1)$ (resp. $L_\infty(I_\infty(m)) = b_\infty(m + 1)$), so (4.2) implies that also $b^n(m) - b^n(\lceil \delta_0 n^{1/2} \rceil) = b_\infty(m) - b_\infty(\lceil \delta_0 n^{1/2} \rceil)$ for each $m \in [\delta_0 n^{1/2}, \delta_1 n^{1/2}]$.

By choosing δ_0 sufficiently small and then δ_1 sufficiently close to δ_0 and possibly increasing n_* (c.f. the argument right before (4.10) in the proof of Proposition 4.5), we

can arrange that for $n \geq n_*$, it holds except on an event of probability at most $1 - \epsilon/4$ that

$$\max \left\{ \min_{m \in [0, \delta_1 n^{1/2}]_{\mathbb{Z}}} b^n(m), \min_{m \in [\delta_1 n^{1/2}, l^n]_{\mathbb{Z}}} b^n(m) \right\} < \min_{i \in [I^n(\delta_0 n^{1/2}), I^n(\delta_1 n^{1/2})]_{\mathbb{Z}}} L^n(i). \quad (4.16)$$

Let

$$m_* := \left\lfloor \frac{1}{2}(\delta_0 + \delta_1)n^{1/2} \right\rfloor.$$

Also define

$$r := - \left(\min_{m \in [\delta_0 n^{1/2}, m_* - 1]_{\mathbb{Z}}} (b_\infty(m) - b_\infty(\lceil \delta_0 n^{1/2} \rceil)) \vee \min_{m \in [m_* + 1, \delta_1 n^{1/2}]_{\mathbb{Z}}} (b_\infty(m) - b_\infty(\lceil \delta_0 n^{1/2} \rceil)) \right),$$

so that $r \geq 0$, and let

$$r' := b_\infty(m_*) - b_\infty(\lceil \delta_0 n^{1/2} \rceil) + r.$$

Since $m \mapsto b_\infty(m) - b_\infty(\lceil \delta_0 n^{1/2} \rceil)$ is obtained from a simple random walk by skipping its upward steps, we can find $\alpha > 0$ such that for large enough n , it holds with probability at least $1 - \epsilon/2$ that $\frac{1}{4}r' - 1 \geq \alpha n^{1/4}$.

Recall the contour functions $\mathbb{p}^n : [0, 2n + l^n]_{\mathbb{Z}} \rightarrow \mathcal{V}(Q^n)$ and $\mathbb{p}_\infty : \mathbb{Z} \rightarrow \mathcal{V}(Q_\infty)$. Let $\tilde{v}_0^n := \mathbb{p}^n(I^n(m_*))$ and $\tilde{v}_\infty := \mathbb{p}_\infty(I_\infty(m_*))$. If we let \tilde{e}_0^n (resp. \tilde{e}_∞) be the edge of ∂Q^n (resp. ∂Q_∞) immediately to the left of \tilde{v}_0^n (resp. \tilde{v}_∞), with one of the two possible orientations chosen with probability $1/2$ each. Then by re-rooting invariance,

$$(Q^n, \tilde{e}_0^n, v_*^n) \stackrel{d}{=} (Q^n, e_0^n, v_*^n) \quad \text{and} \quad (Q, \tilde{e}_\infty) \stackrel{d}{=} (Q, e_\infty).$$

Hence it suffices to show that whenever (4.15) and (4.16) occur, it holds that $B_{\frac{1}{4}(r'-3)}(\tilde{v}_0^n; Q^n)$ and $B_{\frac{1}{2}(r'-3)}(\tilde{v}_\infty; Q_\infty)$ are isomorphic via a graph isomorphism satisfying the conditions in the statement of the proposition.

Henceforth assume that (4.15) and (4.16) occur. Let $m_0 \in [\delta_0 n^{1/2}, m_* - 1]_{\mathbb{Z}}$ and $m_1 \in [i_* + 1, \delta_1 n^{1/2}]_{\mathbb{Z}}$ be chosen so that

$$b_\infty(m_0) - b_\infty(\lceil \delta_0 n^{1/2} \rceil) \leq -r \quad \text{and} \quad b_\infty(m_1) - b_\infty(\lceil \delta_0 n^{1/2} \rceil) \leq -r.$$

By Lemma 3.7 together with (4.15) and (4.16) (the latter equation is used to deal with the second term in the maximum in the estimate of Lemma 3.7), if $i \in \mathbb{N}_0$ with $i \geq I^n(m_1)$ then

$$\begin{aligned} \text{dist}(\mathbb{p}^n(i), \tilde{v}_0^n; Q^n) &\geq L^n(i) + b^n(m_*) - 2 \max \left\{ \inf_{j \in J^n(I^n(m_*), i)} L^n(j), \inf_{j \in (J')^n(I^n(m_*), i)} L^n(j) \right\} \\ &\geq L^n(i) + b^n(m_*) - L^n(i) + r - b^n(\lceil \delta_0 n^{1/2} \rceil) \geq r'. \end{aligned}$$

We have a similar estimate if $i \leq I^n(m_0)$. Again using Lemma 3.7, we obtain that if $i \in [0, 2n]_{\mathbb{Z}} \setminus [I^n(m_0), I^n(m_1)]_{\mathbb{Z}}$ then

$$\text{dist}(\mathbb{p}_\infty(i), \tilde{v}_\infty; Q_\infty) \geq r'.$$

Hence $\mathcal{V}(B_{r'}(\tilde{v}_\infty; Q_\infty)) \subset \mathbb{p}_\infty([I_\infty(m_0), I_\infty(m_1)]_{\mathbb{Z}})$ and $\mathcal{V}(B_{r'}(\tilde{v}_0^n; Q^n)) \subset \mathbb{p}^n([I^n(m_0), I^n(m_1)]_{\mathbb{Z}})$. By the local nature of the Schaeffer bijection and (4.15) it holds that $B_{\frac{1}{4}r'-1}(\tilde{v}_0^n; Q^n)$ and $B_{\frac{1}{4}r'-1}(\tilde{v}_\infty; Q_\infty)$ are isomorphic as graphs via an isomorphism satisfying the conditions in the proposition statement. The triangle inequality implies that any such isomorphism is an isometry when these balls are equipped with the metrics they inherit from Q^n and Q_∞ , respectively. \square

4.4 Proof of Theorem 1.11

We will deduce the scaling limit statement for the UIHPQ from Theorem 4.1 and the coupling results (Propositions 4.2 and 4.5).

For $R > 0$, let (H_R, d_R) be a Brownian disk with area R and boundary length $R^{1/2}$ and let μ_R and $\eta_R : [0, R] \rightarrow \partial H_R$ be its area measure and boundary path, respectively. Let

$$\mathfrak{H}_R := (H_R, d_R, \mu_R, \eta_R).$$

By Brownian scaling and the definition of the Brownian disk from Section 3.5, \mathfrak{H}_R has the same law as $(H_1, R^{1/4}d_1, R\mu_1, \eta_1(R^{-1/2}\cdot))$.

Let $(Q_R^n, e_R^n, v_{*,R}^n)$ be sampled uniformly from the set of boundary-rooted pointed quadrangulations with $\lfloor Rn \rfloor$ interior faces and perimeter $\lfloor 2^{3/2}(Rn)^{1/2} \rfloor$. View Q_R^n as a topological space in the manner of Remark 1.2. Let d_R^n be $(9/8)^{1/4}n^{-1/4}$ times the graph distance on Q_R^n and let μ_R^n be the measure on $\mathcal{V}(Q_R^n)$ which assigns mass to each vertex equal to $(4n)^{-1}$ times its degree. Let $\lambda_R^n : [0, \lfloor 2^{3/2}(Rn)^{1/2} \rfloor] \rightarrow \partial Q_R^n$ be the boundary path of Q_R^n started from e_R^n (extended to \mathbb{R} by linear interpolation) and let $\eta_R^n(t) := \lambda_R^n(2^{3/2}n^{1/2}t)$ for $t \in [0, R^{1/2}]$. Let

$$\mathfrak{Q}_R^n := (Q_R^n, d_R^n, \mu_R^n, \eta_R^n).$$

By Theorem 4.1 and the aforementioned scaling relation between \mathfrak{H}_R and \mathfrak{H}_1 , we find that $\mathfrak{Q}_R^n \rightarrow \mathfrak{H}_R$ in the GHPU topology for each $R > 0$.

Now fix $r > 0$ and $\epsilon \in (0, 1)$. We recall the r -truncation operator \mathfrak{B}_r from Definition 1.6. By Proposition 4.2 and the scale invariance of the law of the Brownian half-plane, we can find $R > r$ and a coupling of (H_∞, d_∞) with (H_R, d_R) such that with probability at least $1 - \epsilon$, $\mathfrak{B}_r \mathfrak{H}_\infty$ and $\mathfrak{B}_r \mathfrak{H}_R$ agree as elements of $\overline{\mathfrak{M}}^{\text{GHPU}}$.

By Proposition 4.5 (applied with $\lfloor Rn \rfloor$ in place of n), after possibly increasing R , we can find $n_* \in \mathbb{N}$ such that for $n \geq n_*$, we can couple (Q_∞, e_∞) with $(Q_R^n, e_R^n, v_{*,R}^n)$ in such a way that with probability at least $1 - \epsilon$, $\mathfrak{B}_r \mathfrak{Q}_\infty^n$ and $\mathfrak{B}_r \mathfrak{Q}_R^n$ agree as elements of $\overline{\mathfrak{M}}^{\text{GHPU}}$.

Since $\mathfrak{Q}_R^n \rightarrow \mathfrak{H}_R$ in the GHPU topology and a.s. r is a good radius for \mathfrak{H}_R (Definition 2.9), Lemma 2.11 implies that $\mathfrak{B}_r \mathfrak{Q}_R^n \rightarrow \mathfrak{B}_r \mathfrak{H}_R$ in law in the GHPU topology. Since $\epsilon \in (0, 1)$ is arbitrary, the existence of the above couplings implies that $\mathfrak{B}_r \mathfrak{Q}_\infty^n \rightarrow \mathfrak{B}_r \mathfrak{H}_\infty$ in law in the GHPU topology. Since $r > 0$ is arbitrary and by Lemma 2.11, it follows that $\mathfrak{Q}_\infty^n \rightarrow \mathfrak{H}_\infty$ in law in the local GHPU topology. \square

4.5 Proof of Theorem 1.12

In this section we will deduce Theorem 1.12 from Theorem 1.11. We define the UIHPQ_S (Q_S, e_S) and its associated elements of $\overline{\mathfrak{M}}_\infty^{\text{GHPU}}$, \mathfrak{Q}_S^n for $n \in \mathbb{N}$, as in the discussion just above Theorem 1.12.

We assume throughout this section that we have coupled (Q_S, e_S) with an instance of the UIHPQ (Q_∞, e_∞) in the manner described in Section 3.4, so that Q_S is obtained from Q_∞ by pruning the dangling quadrangulations from Q_∞ . Let $\{(q_v, e_v)\}_{v \in \mathcal{V}(\partial Q_S)}$ be these dangling quadrangulations (with their oriented boundary root edges). Then the pairs (q_v, e_v) except for the one dangling from the right endpoint of e_S are i.i.d. samples from the free Boltzmann distribution on quadrangulations with general boundary (Section 3.4) and

$$Q_\infty = Q_S \cup \bigcup_{v \in \mathcal{V}(\partial Q_S)} q_v.$$

Since each of the quadrangulations q_v for $v \in \partial Q_S$ is disconnected from Q_S by removing the single vertex v , it follows that no geodesic between vertices of Q_S enters $q_v \setminus \{v\}$. Hence for $r > 0$ and $v \in \mathcal{V}(Q_S)$,

$$B_r(v; Q_S) = B_r(v; Q_\infty) \cap Q_S. \tag{4.17}$$

In what follows we will bound the diameters, areas, and boundary lengths of the extra quadrangulations q_v , with the eventual goal of showing that Q_S and Q_∞ have the same scaling limit in the GHPU topology (up to multiplying boundary lengths by a constant factor).

Before proving these bounds, it will be convenient to have a general lemma for how many vertices along the boundary can be contained in a metric ball. Let $\{v_k\}_{k \in \mathbb{Z}}$ be the vertices of ∂Q_S , listed in the order in which they are hit by the boundary path λ_S and enumerated so that v_0 is the right endpoint of e_S (so that q_{v_0} is the one dangling quadrangulation which does not agree in law with the others; recall Section 3.4).

Let K_r^n be the largest $k \in \mathbb{N}$ for which either v_k or v_{-k} belongs to $B_{rn^{1/4}}(e_S; Q_S)$. By (4.17), K_r^n is also the largest $k \in \mathbb{N}$ for which either v_k or v_{-k} belongs to $B_{rn^{1/4}}(e_S; Q_\infty)$.

Lemma 4.8. *For each $r > 0$ and each $\epsilon \in (0, 1)$, there exists $C = C(r, \epsilon) > 0$ such that for each $n \in \mathbb{N}$,*

$$\mathbb{P}\left[K_r^n \leq Cn^{1/2}\right] \geq 1 - \epsilon.$$

Proof. Let $(C_\infty, L_\infty, b_\infty)$ be the Schaeffer encoding process of (Q_∞, e_∞) as in Section 3.2. Since b_∞ is obtained from a two-sided simple random walk by skipping the upward steps, we can find $C_0 = C_0(r, \epsilon)$ such that with probability at least $1 - \epsilon/4$,

$$\min_{k \in [0, C_0 n^{1/2}]_{\mathbb{Z}}} b_\infty(k) \leq -rn^{1/4} \quad \text{and} \quad \min_{k \in [-C_0 n^{1/2}, 0]_{\mathbb{Z}}} b_\infty(k) \leq -rn^{1/4}.$$

By Lemma 3.7, if this is the case then (with p_∞ the contour function as in Section 3.2)

$$\begin{aligned} \text{dist}(p_\infty(0), p_\infty(i); Q_\infty) &\geq L_\infty(p_\infty(0)) + L_\infty(p_\infty(i)) - \left(L_\infty(p_\infty(i)) + \min_{k \in [0, C_0 n^{1/2}]_{\mathbb{Z}}} b_\infty(k) \right) \\ &\geq rn^{1/4} \end{aligned} \tag{4.18}$$

whenever $i \geq I_\infty(\lfloor C_0 n^{1/2} \rfloor)$. We have a similar bound when $i \leq I_\infty(-\lfloor C_0 n^{1/2} \rfloor)$.

The number of vertices of ∂Q_∞ in $p_\infty([I_\infty(-\lfloor C_0 n^{1/2} \rfloor), I_\infty(\lfloor C_0 n^{1/2} \rfloor)]_{\mathbb{Z}})$ is at most the sum of the quantities $|b_\infty(k) - b_\infty(k-1)| + 1$ for $k \in [-C_0 n^{1/2}, C_0 n^{1/2}]_{\mathbb{Z}}$ (c.f. Remark 3.4). By the law of large numbers and (4.18) we can find $C_1 = C_1(r, \epsilon) > C_0$ such that with probability at least $1 - \epsilon/2$, there are at most $C_1 n^{1/2}$ vertices of ∂Q_∞ in $B_{rn^{1/4}}(e_\infty; Q_\infty)$. Since the dangling quadrangulation q_{v_0} is a.s. finite, applying the preceding bound with $r - A$ in place of r for A a deterministic, n -independent constant shows that we can find $C = C(r, \epsilon) > C_1$ such that with probability at least $1 - \epsilon$, there are at most $Cn^{1/2}$ vertices of ∂Q_∞ in $B_{rn^{1/4}}(e_S; Q_\infty)$.

Our coupling of Q_∞ and Q_S implies that in this case, there are at most $Cn^{1/2}$ vertices of ∂Q_S in $B_{rn^{1/4}}(e_S; Q_S)$. The statement of the lemma follows. \square

We next prove a bound for the diameter of dangling quadrangulations.

Lemma 4.9. *For each $r > 0$ and each $\delta > 0$, it holds with probability tending to 1 as $n \rightarrow \infty$ that*

$$\max_{v \in \mathcal{V}(\partial Q_S \cap B_{rn^{1/4}}(e_S; Q_S))} \text{diam}(q_v) \leq \delta n^{1/4}$$

where the diameter is taken with respect to the internal graph metric on q_v .

We will deduce Lemma 4.9 from Proposition 4.5 and an analogous bound for finite-volume quadrangulations. Fix $\ell > 0$ and a sequence of positive integers $\{l^n\}_{n \in \mathbb{N}}$ with $l^n \rightarrow \ell$. For $n \in \mathbb{N}$, let (Q^n, e_0^n, v_*^n) be sampled uniformly $\mathcal{Q}^\bullet(n, l^n)$. Let $\text{Core}(Q^n)$ be the quadrangulation obtained by removing from Q^n each vertex and each edge which can be disconnected from v_*^n by deleting a single vertex of ∂Q^n . Let \mathcal{C}^n be the set of connected components of the set of vertices and edges removed in this manner plus the vertices

which can be deleted to disconnect these vertices and edges from v_*^n . Then \mathcal{C}^n is a set of quadrangulations with general boundary which “dangle” from $\text{Core}(Q^n)$.

Lemma 4.10. *For each $\delta > 0$, it holds with probability tending to 1 as $n \rightarrow \infty$ that*

$$\max_{q \in \mathcal{C}^n} \text{diam}(q) \leq \delta n^{1/4}$$

where here the diameter is taken with respect to the internal graph metric on q .

Proof. This is essentially proven in [10, Section 5] but for the sake of clarity we explain how the precise statement of the lemma follows from existing results in the literature. For $n \in \mathbb{N}$ let d^n be the graph distance on Q^n , rescaled by $(9/8)^{1/4} n^{-1/4}$. Also let (H, d, z) be a Brownian disk with area 1 and boundary length ℓ together with a marked interior point, sampled uniformly from its area measure. We know from [10, Theorem 1] (c.f. Theorem 4.1) that $(Q^n, d^n, v_*^n) \rightarrow (H, d, z)$ in the pointed Gromov-Hausdorff topology. By the Skorokhod representation theorem, we can couple so that this convergence occurs a.s. By the analog of [23, Lemma A.1] for the pointed Gromov-Hausdorff topology (which follows from Proposition 1.5), we can find a compact metric space (W, D) and isometric embeddings $Q^n \rightarrow W$ and $H \rightarrow W$ such that the following is true. If we identify Q^n and H with their embeddings into W , then we a.s. have $Q^n \rightarrow H$ in the D -Hausdorff distance and $D(v_*^n, z) \rightarrow 0$.

Suppose now by way of contradiction that the statement of the lemma is false. Then we can find $\delta > 0$ such that with positive probability, there is an infinite sequence \mathcal{N} of positive integers such that for $n \in \mathcal{N}$, there is a $q^n \in \mathcal{C}^n$ with d^n -diameter $\geq \delta$. Let v^n be the vertex of ∂Q^n with the property that removing v^n from ∂Q^n disconnects q^n from v_*^n . By possibly replacing \mathcal{N} with a further subsequence, we can arrange that $v^n \rightarrow x \in H$ as $\mathcal{N} \ni n \rightarrow \infty$.

For $n \in \mathcal{N}$ and $\zeta \in (0, \delta)$, let $A_\zeta^n := q^n \setminus B_\zeta(v^n; d^n)$ and let $U_\zeta^n := Q^n \setminus (q^n \cup B_\zeta(v^n; d^n))$. Then for each $n \in \mathbb{N}$ we have $d^n(A_\zeta^n, U_\zeta^n) \geq \zeta$ and $A_\zeta^n \cup U_\zeta^n \cup B_\zeta(v^n; d^n) = \mathcal{V}(Q^n)$. For each rational $\zeta > 0$, we can find a subsequence \mathcal{N}_ζ of \mathcal{N} along which $A_\zeta^n \rightarrow A_\zeta \subset H$ and $U_\zeta^n \rightarrow U_\zeta \subset H$ in the D -Hausdorff distance. The sets A_ζ and U_ζ lie at distance at least ζ from each other (so are disjoint) and $A_\zeta \cup U_\zeta \cup B_\zeta(x; d) = H$. Since v_*^n lies at uniformly positive d^n -distance from ∂Q^n with probability tending to 1 as $n \rightarrow \infty$ and by our choice of the vertices v^n , it follows that it is a.s. the case that for small enough $\zeta > 0$, each of A_ζ and U_ζ is non-empty. Hence for each sufficiently small $\zeta > 0$, removing $B_\zeta(x; d)$ disconnects H into two non-empty components. This contradicts the fact that H has the topology of a disk [8, Theorem 2]. \square

Proof of Lemma 4.9. Let $\{v_k\}_{k \in \mathbb{Z}}$ and K_r^n be as in the discussion just before Lemma 4.8. Also fix $\epsilon \in (0, 1)$ and let $C = C(r, \epsilon)$ be chosen as in that lemma. By Theorem 1.11, we can find $\rho = \rho(r, \epsilon) > r$ such that with probability at least $1 - \epsilon$, each v_k for $k \in [-Cn^{1/2}, Cn^{1/2}]_{\mathbb{Z}}$ is contained in $B_{\rho n^{1/4}}(\mathfrak{e}_S, Q_S)$. Then with probability at least $1 - 2\epsilon$,

$$\mathcal{V}(\partial Q_S \cap B_{\rho n^{1/4}}(\mathfrak{e}_S, Q_S)) \subset \{v_k : k \in [-Cn^{1/2}, Cn^{1/2}]_{\mathbb{Z}}\} \subset \mathcal{V}(\partial Q_\infty \cap B_{\rho n^{1/4}}(\mathfrak{e}_\infty, Q_\infty)). \tag{4.19}$$

By Proposition 4.5, we can find $R = R(r, \epsilon, \delta) \in \mathbb{N}$ with $R > \rho + \delta$ such that for large enough $n \in \mathbb{N}$, we can couple $(Q^{Rn}, \mathfrak{e}_0^{Rn}, v_*^{Rn})$ with $(Q_\infty, \mathfrak{e}_\infty)$ in such a way that with probability at least $1 - \epsilon$, the graph metric balls $B_{(\rho+\delta)n^{1/4}}(\mathfrak{e}_\infty; \mathbb{Q}_\infty)$ and $B_{\rho+\delta}(\mathfrak{e}_0^{Rn}; R^{-1/4} d^{Rn})$ equipped with the restricted graph metrics are isometric via a graph isomorphism which preserves the intersection of these metric balls with ∂Q_∞ and ∂Q^{Rn} , respectively. If this is the case and (4.19) holds, then each q_v for $v \in \mathcal{V}(\partial Q_S \cap B_{\rho n^{1/4}}(\mathfrak{e}_S, Q_S))$ with internal diameter $\geq \delta n^{1/4}$ corresponds to a unique dangling quadrangulation in \mathcal{C}^{Rn} with internal diameter $\geq \delta n^{1/4}$. By Lemma 4.10, the probability that such a quadrangulation exists tends to 0 and $n \rightarrow \infty$. Since $\epsilon \in (0, 1)$ was arbitrary, we conclude. \square

Next we turn our attention to a bound for the areas of the dangling quadrangulations.

Lemma 4.11. *For each $r > 0$ and each $\epsilon > 0$, there exists $A = A(r, \epsilon) > 0$ such that for each $n \in \mathbb{N}$, it holds with probability at least $1 - \epsilon$ that*

$$\sum_{v \in \mathcal{V}(\partial Q_S) \cap B_{r n^{1/4}}(\mathfrak{e}_S; Q_S)} \mu_\infty^n(q_v) \leq A n^{-1/3},$$

where here we recall that μ_∞^n is the measure which assigns each vertex of Q_∞^n a mass equal to $(4n)^{-1}$ times its degree.

Proof. Let $\{v_k\}_{k \in \mathbb{Z}}$ and K_r^n be as in the discussion just before Lemma 4.8. By Lemma 4.8, for each $\epsilon > 0$ there exists $C = C(r, \epsilon) > 0$ such that with probability at least $1 - \epsilon$, we have $K_r^n \leq C n^{1/2}$, in which case

$$\sum_{v \in \mathcal{V}(\partial Q_S) \cap B_{r n^{1/4}}(\mathfrak{e}_S; Q_S)} \mu_\infty^n(q_v) \leq n^{-1} \sum_{k=-\lfloor C n^{1/2} \rfloor}^{\lfloor C n^{1/2} \rfloor} \#\mathcal{E}(q_{v_k})$$

with universal implicit constant.

By [15, Equation (24)], we have the tail estimate $\mathbb{P}[\#\mathcal{E}(q_v) > m] \sim c m^{-3/4}$ for $v \neq v_0$, for a universal constant $c > 0$. By the heavy-tailed central limit theorem, the random variables

$$n^{-2/3} \sum_{k=-\lfloor C n^{1/2} \rfloor}^{\lfloor C n^{1/2} \rfloor} \#\mathcal{E}(q_{v_k})$$

converge in law to a non-degenerate limiting distribution. The statement of the lemma follows. \square

It remains to prove a bound for the boundary length of the dangling quadrangulations (which will explain why we use a different scaling in the definitions of the re-scaled boundary paths η_∞^n and η_S^n). For $t \geq 0$, let $\sigma_S^n(t)$ be equal to $n^{-1/2}$ times the sum of the boundary lengths of the dangling quadrangulations q_v attached at vertices v which are hit by η_S^n between time 0 and time t plus $n^{-1/2}$ times the total number of such quadrangulations, so that for $t \geq 0$,

$$\eta_S^n(t) = \eta_\infty^n \left(2^{-3/2} \sigma_S^n(t) + O_n(n^{-1/2}) \right). \tag{4.20}$$

Also let $\tilde{\sigma}_S^n(-t)$ be equal to $n^{-1/2}$ times the sum of the boundary lengths of the dangling quadrangulations q_v attached at vertices v which are hit by η_S^n between time $-t$ and time 0 plus $n^{-1/2}$ times the total number of such quadrangulations so $\tilde{\sigma}_S^n$ satisfies an analog of (4.20) for negative times.

Lemma 4.12. *We have $\sigma_S^n \rightarrow (t \mapsto 2^{3/2}t)$ and $\tilde{\sigma}_S^n \rightarrow (t \mapsto 2^{3/2}t)$ in law with respect to the topology of uniform convergence on compact subsets of $[0, \infty)$.*

Proof. We will prove the statement for σ_S^n ; the statement for $\tilde{\sigma}_S^n$ is proven identically. Let $\{v_k\}_{k \in \mathbb{Z}}$ be as in the discussion just before Lemma 4.8. Then q_{v_0} is a.s. finite and by [15, Equation (23)], for $k \in \mathbb{N}$ we have $\mathbb{E}[\text{Perim}(q_{v_k}) + 1] = 3$, where here Perim denotes the perimeter. The random variables $\text{Perim}(q_{v_k}) + 1$ for $k \in \mathbb{N}$ are i.i.d. and $\text{Perim}(q_{v_0})$ is a.s. finite, so by the strong law of large numbers there a.s. exists a random $M \in \mathbb{N}$ such that for $m \geq M$,

$$\left| \frac{1}{m} \sum_{k=0}^m (\text{Perim}(q_{v_k}) + 1) - 3 \right| \leq \epsilon. \tag{4.21}$$

The law of M does not depend on n , so we can find a deterministic $m_0 \in \mathbb{N}$ such that with probability at least $1 - \epsilon$, (4.21) holds for each $m \geq m_0$. Let E be the event that this is the case. Since $\sum_{k=0}^{m_0} (\text{Perim}(q_{v_k}) + 1) < \infty$ a.s., we can find a deterministic $C > 0$, independent from n , such that with probability at least $1 - \epsilon$ this sum is at most C . Let F be the event that this is the case.

We have

$$\sigma_S(t) = n^{-1/2} \sum_{k=0}^{\lfloor \frac{2^{3/2}}{3} n^{1/2} t \rfloor} (\text{Perim}(q_{v_k}) + 1)$$

so if E occurs and $\lfloor \frac{2}{3} n^{1/2} t \rfloor \geq m_0$,

$$\left| \sigma_S(t) - 2^{3/2} t \right| \leq \epsilon t.$$

Hence if $T > 0$, then on $E \cap F$,

$$\sup_{t \in [0, T]} \left| \sigma_S(t) - 2^{3/2} t \right| \leq \epsilon T + O_n(n^{-1/2}).$$

Since $\epsilon > 0$ is arbitrary, the statement of the lemma follows. □

Proof of Theorem 1.12. Let $\{((Q_S^n, \mathfrak{e}_S^n), (Q_\infty^n, \mathfrak{e}_\infty^n))\}_{n \in \mathbb{N}}$ be a sequence of copies of the coupling $((Q_S, \mathfrak{e}_S), (Q_\infty, \mathfrak{e}_\infty))$ used in this section and let $\{q_v^n\}_{n \in \mathbb{N}}$ be the associated dangling quadrangulations. Define the elements $\mathfrak{Q}_\infty^n = (Q_\infty^n, d_\infty^n, \mu_\infty^n, \eta_\infty^n)$ and $\mathfrak{Q}_S^n = (Q_S^n, d_S^n, \mu_S^n, \eta_S^n)$ of $\mathbb{M}_\infty^{\text{GHPU}}$ as in Section 1.6 with respect to the n th pair in this sequence.

By Theorem 1.11 and the Skorokhod representation theorem, we can find a coupling of the sequence $\{((Q_S^n, \mathfrak{e}_S^n), (Q_\infty^n, \mathfrak{e}_\infty^n))\}_{n \in \mathbb{N}}$ and (H_∞, d_∞) such that a.s. $\mathfrak{Q}_\infty^n \rightarrow \mathfrak{H}_\infty$ in the local GHPU topology. By Proposition 1.9, we can a.s. find a random boundedly compact metric space (W, D) and isometric embeddings of (Q_∞^n, d_∞^n) for $n \in \mathbb{N}$ and (H_∞, d_∞) into (W, D) such that if we identify these spaces with their embeddings, then a.s. $\mathfrak{Q}_\infty^n \rightarrow \mathfrak{H}_\infty$ in the D -local HPU topology (Definition 1.8).

Now fix a deterministic $r > 0$ and recall (4.17). By Lemma 4.9, in any such coupling

$$B_r(\eta_S^n(0); d_S^n) \rightarrow B_r(\eta_\infty(0); d_\infty)$$

in probability with respect to the D -Hausdorff metric. By Lemma 4.11, in any such coupling

$$\mu_S^n|_{B_r(\eta_S^n(0); d_S^n)} \rightarrow \mu_S|_{B_r(\eta_\infty(0); d_\infty)}$$

in probability with respect to the D -Prokhorov metric. By (4.20) and Lemmas 4.9 and 4.12, in any such coupling

$$\mathfrak{B}_r \eta_S^n \rightarrow \mathfrak{B}_r \eta_\infty$$

in probability with respect to the D -uniform metric, where \mathfrak{B}_r is as in Definition 1.6. Therefore $\mathfrak{B}_r \mathfrak{Q}_S^n \rightarrow \mathfrak{B}_r \mathfrak{H}_\infty$ in law in the GHPU topology, so by Lemma 2.11 $\mathfrak{Q}_S^n \rightarrow \mathfrak{H}_\infty$ in law in the local GHPU topology. □

References

- [1] Céline Abraham. Rescaled bipartite planar maps converge to the Brownian map. *Ann. Inst. Henri Poincaré Probab. Stat.*, 52(2):575–595, 2016. MR-3498001
- [2] Romain Abraham, Jean-François Delmas, and Patrick Hoscheit. A note on the Gromov-Hausdorff-Prokhorov distance between (locally) compact metric measure spaces. *Electron. J. Probab.*, 18:no. 14, 21, 2013. MR-3035742
- [3] L. Addario Berry and M. Albenque. The scaling limit of random simple triangulations and random simple quadrangulations. *ArXiv e-prints*, June 2013. MR-3706731

- [4] E. Baur, G. Miermont, and G. Ray. Classification of scaling limits of uniform quadrangulations with a boundary. *ArXiv e-prints*, August 2016.
- [5] Itai Benjamini and Oded Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 13 pp. (electronic), 2001. MR-1873300
- [6] J. Bertoin and R. A. Doney. On conditioning a random walk to stay nonnegative. *Ann. Probab.*, 22(4):2152–2167, 1994. MR-1331218
- [7] Jérémie Bettinelli. Scaling limits for random quadrangulations of positive genus. *Electron. J. Probab.*, 15:no. 52, 1594–1644, 2010. MR-2735376
- [8] Jérémie Bettinelli. Scaling limit of random planar quadrangulations with a boundary. *Ann. Inst. Henri Poincaré Probab. Stat.*, 51(2):432–477, 2015. MR-3335010
- [9] Jérémie Bettinelli, Emmanuel Jacob, and Grégory Miermont. The scaling limit of uniform random plane maps, via the Ambjørn-Budd bijection. *Electron. J. Probab.*, 19:no. 74, 16, 2014. MR-3256874
- [10] Jérémie Bettinelli and Grégory Miermont. Compact Brownian surfaces I: Brownian disks. *Probab. Theory Related Fields*, 167(3-4):555–614, 2017. MR-3627425
- [11] G. Borot, J. Bouttier, and E. Guitter. A recursive approach to the $O(n)$ model on random maps via nested loops. *Journal of Physics A Mathematical General*, 45(4):045002, February 2012. MR-2874232
- [12] J. Bouttier, P. Di Francesco, and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.*, 11(1):Research Paper 69, 27, 2004. MR-2097335
- [13] J. Bouttier and E. Guitter. Distance statistics in quadrangulations with a boundary, or with a self-avoiding loop. *J. Phys. A*, 42(46):465208, 44, 2009. MR-2552016
- [14] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A course in metric geometry*, volume 33 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2001. MR-1835418
- [15] A. Caraceni and N. Curien. Geometry of the Uniform Infinite Half-Planar Quadrangulation. *ArXiv e-prints*, August 2015.
- [16] A. Caraceni and N. Curien. Self-Avoiding Walks on the UIPQ. *ArXiv e-prints*, September 2016.
- [17] Alessandra Caraceni. *The geometry of large outerplanar and half-planar maps*. PhD thesis, Scuola Normale Superiore, 2015.
- [18] Nicolas Curien and Jean-François Le Gall. The Brownian plane. *J. Theoret. Probab.*, 27(4):1249–1291, 2014. MR-3278940
- [19] Nicolas Curien and Grégory Miermont. Uniform infinite planar quadrangulations with a boundary. *Random Structures Algorithms*, 47(1):30–58, 2015. MR-3366810
- [20] B. Duplantier, J. Miller, and S. Sheffield. Liouville quantum gravity as a mating of trees. *ArXiv e-prints*, September 2014.
- [21] Bertrand Duplantier and Scott Sheffield. Liouville quantum gravity and KPZ. *Invent. Math.*, 185(2):333–393, 2011. MR-2819163
- [22] B. V. Gnedenko and A. N. Kolmogorov. *Limit distributions for sums of independent random variables*. Addison-Wesley Publishing Company, Inc., Cambridge, Mass., 1954. Translated and annotated by K. L. Chung. With an Appendix by J. L. Doob. MR-0062975
- [23] Andreas Greven, Peter Pfaffelhuber, and Anita Winter. Convergence in distribution of random metric measure spaces (Λ -coalescent measure trees). *Probab. Theory Related Fields*, 145(1-2):285–322, 2009. MR-2520129
- [24] Misha Gromov. *Metric structures for Riemannian and non-Riemannian spaces*, volume 152 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, 1999. Based on the 1981 French original [MR0682063 (85e:53051)], With appendices by M. Katz, P. Pansu and S. Semmes, Translated from the French by Sean Michael Bates. MR-1699320
- [25] E. Gwynne, N. Holden, and X. Sun. Joint scaling limit of a bipolar-oriented triangulation and its dual in the peanosphere sense. *ArXiv e-prints*, March 2016.
- [26] E. Gwynne, N. Holden, and X. Sun. Full scaling limit of site percolation on random triangulations toward SLE_6 on $\sqrt{8/3}$ -Liouville quantum gravity in the metric and peanosphere sense. In preparation, 2017.

- [27] E. Gwynne, A. Kassel, J. Miller, and D. B. Wilson. Active spanning trees with bending energy on planar maps and SLE-decorated Liouville quantum gravity for $\kappa \geq 8$. *ArXiv e-prints*, March 2016.
- [28] E. Gwynne, C. Mao, and X. Sun. Scaling limits for the critical Fortuin-Kasteleyn model on a random planar map I: cone times. *ArXiv e-prints*, February 2015.
- [29] E. Gwynne and J. Miller. Convergence of the self-avoiding walk on random quadrangulations to $\text{SLE}_{8/3}$ on $\sqrt{8/3}$ -Liouville quantum gravity. *ArXiv e-prints*, August 2016.
- [30] E. Gwynne and J. Miller. Metric gluing of Brownian and $\sqrt{8/3}$ -Liouville quantum gravity surfaces. *ArXiv e-prints*, August 2016.
- [31] E. Gwynne and J. Miller. Convergence of percolation on uniform quadrangulations with boundary to SLE_6 on $\sqrt{8/3}$ -Liouville quantum gravity. *ArXiv e-prints*, January 2017.
- [32] E. Gwynne and J. Miller. Convergence of the free Boltzmann quadrangulation with simple boundary to the Brownian disk. *ArXiv e-prints*, January 2017.
- [33] E. Gwynne and X. Sun. Scaling limits for the critical Fortuin-Kasteleyn model on a random planar map III: finite volume case. *ArXiv e-prints*, October 2015.
- [34] E. Gwynne and X. Sun. Scaling limits for the critical Fortuin-Kasteleyn model on a random planar map II: local estimates and empty reduced word exponent. *Electronic Journal of Probability*, 22:Paper No. 45, 1–56, 2017. MR-3661659
- [35] R. Kenyon, J. Miller, S. Sheffield, and D. B. Wilson. Bipolar orientations on planar maps and SLE_{12} . *ArXiv e-prints*, November 2015.
- [36] Jean-François Le Gall. The topological structure of scaling limits of large planar maps. *Invent. Math.*, 169(3):621–670, 2007. MR-2336042
- [37] Jean-François Le Gall. Uniqueness and universality of the Brownian map. *Ann. Probab.*, 41(4):2880–2960, 2013. MR-3112934
- [38] Jean-François Le Gall and Grégory Miermont. Scaling limits of random trees and planar maps. In *Probability and statistical physics in two and more dimensions*, volume 15 of *Clay Math. Proc.*, pages 155–211. Amer. Math. Soc., Providence, RI, 2012. MR-3025391
- [39] Jiří Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002. MR-1899299
- [40] Grégory Miermont. Tessellations of random maps of arbitrary genus. *Ann. Sci. Éc. Norm. Supér. (4)*, 42(5):725–781, 2009. MR-2571957
- [41] Grégory Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *Acta Math.*, 210(2):319–401, 2013. MR-3070569
- [42] J. Miller and S. Sheffield. An axiomatic characterization of the Brownian map. *ArXiv e-prints*, June 2015.
- [43] J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map I: The $\text{QLE}(8/3,0)$ metric. *ArXiv e-prints*, July 2015.
- [44] J. Miller and S. Sheffield. Liouville quantum gravity spheres as matings of finite-diameter trees. *ArXiv e-prints*, June 2015.
- [45] J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map II: geodesics and continuity of the embedding. *ArXiv e-prints*, May 2016.
- [46] J. Miller and S. Sheffield. Liouville quantum gravity and the Brownian map III: the conformal structure is determined. *ArXiv e-prints*, August 2016.
- [47] Jason Miller and Scott Sheffield. Imaginary geometry I: interacting SLEs. *Probab. Theory Related Fields*, 164(3-4):553–705, 2016. MR-3477777
- [48] Jason Miller and Scott Sheffield. Quantum Loewner evolution. *Duke Math. J.*, 165(17):3241–3378, 2016. MR-3572845
- [49] Oded Schramm. Scaling limits of loop-erased random walks and uniform spanning trees. *Israel J. Math.*, 118:221–288, 2000. MR-1776084
- [50] Oded Schramm and Scott Sheffield. A contour line of the continuum Gaussian free field. *Probab. Theory Related Fields*, 157(1-2):47–80, 2013. MR-3101840

Scaling limit of the UIHPQ in the GHPU topology

- [51] Scott Sheffield. Gaussian free fields for mathematicians. *Probab. Theory Related Fields*, 139(3-4):521–541, 2007. MR-2322706
- [52] Scott Sheffield. Quantum gravity and inventory accumulation. *Ann. Probab.*, 44(6):3804–3848, 2016. MR-3572324

Electronic Journal of Probability

Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards
- Free for authors, free for readers
- Quick publication (no backlog)
- Secure publication (LOCKSS¹)
- Easy interface (EJMS²)

Economical model of EJP-ECP

- Non profit, sponsored by IMS³, BS⁴, ProjectEuclid⁵
- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund⁶ (click here to donate!)
- Submit your best articles to EJP-ECP
- Choose EJP-ECP over for-profit journals

¹LOCKSS: Lots of Copies Keep Stuff Safe <http://www.lockss.org/>

²EJMS: Electronic Journal Management System <http://www.vtex.lt/en/ejms.html>

³IMS: Institute of Mathematical Statistics <http://www.imstat.org/>

⁴BS: Bernoulli Society <http://www.bernoulli-society.org/>

⁵Project Euclid: <https://projecteuclid.org/>

⁶IMS Open Access Fund: <http://www.imstat.org/publications/open.htm>