

MEMOIRS OF THE EUROPEAN MATHEMATICAL SOCIETY

—
Jérémie Bettinelli
Grégory Miermont

Compact Brownian Surfaces II
Orientable Surfaces

MEMS Vol. 27 / 2026

EM
S ■
PRESS

EM
S ■
PRESS

Memoirs of the European Mathematical Society

Edited by

Anton Alekseev (Université de Genève)

Hélène Esnault (Freie Universität Berlin)

Gerard van der Geer (Universiteit van Amsterdam)

Ari Laptev (Imperial College London)

Laure Saint-Raymond (Institut des Hautes Études Scientifiques)

Susanna Terracini (Università degli Studi di Torino)

The *Memoirs of the European Mathematical Society* publish outstanding research contributions in individual volumes, in all areas of mathematics and with a particular focus on works that are longer and more comprehensive than usual research articles.

The collection's editorial board consists of the editors-in-chief of the *Journal of the European Mathematical Society* and the *EMS Surveys in Mathematical Sciences*, along with editors of book series of the publishing house of the EMS as well as other distinguished mathematicians.

All submitted works go through a highly selective peer-review process.

Latest titles in this series:

- A. Carey, F. Gesztesy, G. Levitina, R. Nichols, F. Sukochev, D. Zanin, *The Limiting Absorption Principle for Massless Dirac Operators, Properties of Spectral Shift Functions, and an Application to the Witten Index of Non-Fredholm Operators*
- J. Kigami, *Conductive Homogeneity of Compact Metric Spaces and Construction of p -Energy*
- A. Buium, L. E. Miller, *Purely Arithmetic PDEs Over a p -Adic Field: δ -Characters and δ -Modular Forms*
- M. Duerinckx, A. Gloria, *On Einstein's Effective Viscosity Formula*
- R. Willett, G. Yu, *The Universal Coefficient Theorem for C^* -Algebras with Finite Complexity*
- B. Janssens, K.-H. Neeb, *Positive Energy Representations of Gauge Groups I. Localization*
- S. Dipierro, G. Giacomini, E. Valdinoci, *The Lévy Flight Foraging Hypothesis in Bounded Regions*
- A. Naor, *Extension, Separation and Isomorphic Reverse Isoperimetry*
- N. Lerner, *Integrating the Wigner Distribution on Subsets of the Phase Space, a Survey*
- B. Adcock, S. Brugiapaglia, N. Dexter, S. Moraga, *On Efficient Algorithms for Computing Near-Best Polynomial Approximations to High-Dimensional, Hilbert-Valued Functions from Limited Samples*
- J. J. Carmona, K. Fedorovskiy, *Carathéodory Sets in the Plane*
- D. Abramovich, Q. Chen, M. Gross, B. Siebert, *Punctured Logarithmic Maps*
- T. Oh, M. Okamoto, L. Tolomeo, *Stochastic Quantization of the Φ_3^3 -Model*
- H. Sasahira, M. Stoffregen, *Seiberg–Witten Floer Spectra for $b_1 > 0$*
- J. KAAD, D. Kyed, *The Quantum Metric Structure of Quantum $SU(2)$*
- B. Zavyalov, *Almost Coherent Modules and Almost Coherent Sheaves*
- J-P. Labesse, B. Lemaire, *La formule des traces tordue pour un corps global de caractéristique $p > 0$*
- S. Ma, *Vector-Valued Orthogonal Modular Forms*
- Y. Almog, B. Helffer, *On the Stability of Symmetric Flows in a Two-Dimensional Channel*
- A. Bors, D. Panario, Q. Wang, *Functional Graphs of Generalized Cyclotomic Mappings of Finite Fields*
- P. Schneider, O. Venjakob, *Reciprocity Laws for (φ_L, Γ_L) -Modules over Lubin–Tate Extensions*
- A.-L. Dalibard, F. Marbach, J. Rax, *Linear and Nonlinear Parabolic Forward-Backward Problems*
- L. Ertzbischoff, D. Han-Kwan, *On Well-Posedness for Thick Spray Equations*



Jérémie Bettinelli
Grégory Miermont

Compact Brownian Surfaces II

Orientable Surfaces



EM
S ■
PRESS

Authors

Jérémie Bettinelli
Laboratoire d'informatique (LIX)
CNRS and École Polytechnique
1 rue Honoré d'Estienne d'Orves
Bâtiment Alan Turing
91120 Palaiseau, France
Email: jeremie.bettinelli@normalesup.org

Grégory Miermont
Unité de Mathématiques Pures et Appliquées
École Normale Supérieure de Lyon and
Institut Universitaire de France
46 allée d'Italie
69364 Lyon, France
Email: gregory.miermont@ens-lyon.fr

Each volume of the *Memoirs of the European Mathematical Society* is available individually or as part of an annual subscription. It may be ordered from your bookseller, subscription agency, or directly from the publisher via subscriptions@ems.press.

ISSN 2747-9080, eISSN 2747-9099

ISBN 978-3-98547-102-7, eISBN 978-3-98547-602-2, DOI 10.4171/MEMS/27

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available on the Internet at <http://dnb.dnb.de>.

Published by EMS Press, an imprint of the

European Mathematical Society – EMS – Publishing House GmbH
Institut für Mathematik
Technische Universität Berlin
Straße des 17. Juni 136
10623 Berlin, Germany

Email: info@ems.press

<https://ems.press>

© 2026 European Mathematical Society

Typesetting: Vyacheslav Boyko and Galyna Popovych, Kyiv, Ukraine

Printed in Germany

♻️ Printed on acid free paper

Abstract

Fix an arbitrary compact orientable surface with a boundary and consider a uniform bipartite random quadrangulation of this surface with n faces and boundary component lengths of order \sqrt{n} or of lower order. Endow this quadrangulation with the usual graph metric renormalized by $n^{-1/4}$, mark it on each boundary component, and endow it with the counting measure on its vertex set renormalized by n^{-1} , as well as the counting measure on each boundary component renormalized by $n^{-1/2}$. We show that, as n goes to infinity, this random marked measured metric space converges in distribution for the Gromov–Hausdorff–Prokhorov topology, toward a random limiting marked measured metric space called a *Brownian surface*.

This extends known convergence results of uniform random planar quadrangulations with at most one boundary component toward the *Brownian sphere* and toward the *Brownian disk*, by considering the case of quadrangulations on general compact orientable surfaces. Our approach consists in cutting a Brownian surface into elementary pieces that are naturally related to the Brownian sphere and the Brownian disk and their noncompact analogs, the Brownian plane and the Brownian half-plane, and to prove convergence results for these elementary pieces, which are of independent interest.

Keywords: random maps, scaling limits, Brownian sphere, Brownian disk, Brownian surface, Gromov–Hausdorff–Prokhorov topology

Mathematics Subject Classification (2020): 60F17 (primary); 05C10 (secondary)

Acknowledgments. We thank G. Chapuy for stimulating discussions during the elaboration of this work, and in particular for his encouragements to deal with general compact surfaces with a boundary. Thanks to J. Bouttier for bringing our attention to the fact that the semigroup property of discrete slices might have an interesting continuum counterpart. Thanks are also due to X. Sun for discussions around the problem of defining Brownian surfaces in the context of Liouville quantum gravity and the question of the random modulus of these surfaces, and also for sharing the work [10]. We finally thank N. Holden and J. Miller for their comments on a first version of this memoir, as well as anonymous referees for their thorough reading.

Funding. We acknowledge financial support from the projects ANR-20-CE48-0018 *3DMaps*, ANR-21-CE48-0007 *IsOMa*, ANR-23-CE48-0018 *CartesEtPlus*.

Contents

1	Introduction	1
1.1	Context	1
1.2	Generalities and terminology on maps	4
1.3	The Gromov–Hausdorff–Prokhorov topology	5
1.4	The main convergence result	7
1.5	Scaling limits of Boltzmann quadrangulations	12
1.6	Perspectives	15
1.7	Organization of the memoir	16
2	Variants of the Cori–Vauquelin–Schaeffer bijection	17
2.1	Basic construction	17
2.2	The generalized Chapuy–Marcus–Schaeffer bijection	19
2.3	Composite slices	24
2.4	Quadrilaterals with geodesic sides	27
2.5	Scaling limits of elementary pieces	29
3	Marking and gluing along geodesics	31
3.1	Useful facts on the Gromov–Hausdorff–Prokhorov topology and markings	31
3.2	Geodesics in metric spaces	34
3.3	Gluing along geodesics	36
3.4	Proof of Theorem 1.1	41
3.5	Gluing quadrangulations from elementary pieces	42
3.6	Scaling limit of the collection of elementary pieces	48
3.7	Gluing pieces together	51
3.8	Topology and Hausdorff dimension	52
4	Convergence of composite slices	55
4.1	Metric spaces coded by real functions	55
4.2	Random continuum composite slices	61
4.3	The Brownian half-plane, and its embedded slices	63
4.4	The uniform infinite half-planar quadrangulation	66
4.5	Scaling limit of conditioned slices	72
5	Convergence of quadrilaterals with geodesic sides	77
5.1	Quadrilaterals coded by two functions	77

5.2	Random continuum quadrilaterals	81
5.3	The Brownian plane, and its embedded quadrilaterals	82
5.4	The uniform infinite planar quadrangulation	87
5.5	Discrete quadrilaterals in the uniform infinite planar quadrangulation	89
5.6	Scaling limit of conditioned quadrilaterals	94
6	Construction from a continuous unicellular map	99
A	Technical lemmas on the Brownian plane	105
A.1	Equivalence of definitions of the Brownian plane	105
A.2	Convergence of the uniform infinite planar quadrangulation to the Brownian plane	106
B	Scaling limit of size parameters in labeled maps	111
B.1	Preliminaries	111
B.2	Asymptotics of the scheme	113
B.3	Asymptotics of the size parameters	119
B.4	Boltzmann quadrangulations	121
	References	123

Chapter 1

Introduction

1.1 Context

Random maps, seen as discrete models of random two-dimensional geometries, have generated a sustained interest in the last couple of decades. An important instance of this line of research are the results by Le Gall [67] and the second author [80], showing that a uniform random quadrangulation of the sphere with n faces, seen as a random finite metric space by endowing its vertex set with the usual graph metric renormalized by $n^{-1/4}$, converges in distribution toward the so-called *Brownian sphere*, or *Brownian map*. The aim of the present work is to generalize this result to the case of general compact orientable surfaces. Let us start with some elements of context.

Random surfaces as scaling limits of random maps. While the idea that continuum random geometries should be obtained as scaling limits of random maps originates from the physics literature on two-dimensional quantum gravity [44, 59, 88], this question was first approached in the mathematical literature in the pioneering work of Chassaing and Schaeffer [37], who studied the model of uniformly chosen random quadrangulations of the sphere, and found in particular that the proper scaling factor in this case was $n^{-1/4}$. Marckert and Mokkadem [76] then constructed a candidate limiting space today called the *Brownian sphere*, and showed the convergence toward it in another topology than the Gromov–Hausdorff topology. Le Gall [65] later showed that the sequence of rescaled metric spaces associated with uniform random quadrangulations of the sphere was relatively compact. Finally, Le Gall [67] and the second author [80] showed by two independent approaches that the previous sequence converges toward the Brownian sphere.

It is known that the Brownian sphere arises as a universal scaling limit for many models of planar maps *that are uniformly chosen in a certain class, given their face degrees*, and provided that face degrees are typically all of the same order of magnitude; see [1, 3, 4, 16, 24, 41, 67, 77]. See also [72] for models of maps that fall out of this universality class.

The scaling limits of quadrangulations on surfaces that are more general than the sphere were considered by the first author in [21, 22], who showed similar results to the above, but only up to extraction of appropriate subsequences, leaving a gap that amounts to uniquely characterize the limit. This gap was filled in the particular case of the disk topology in our previous work [25]. In particular, we showed that a uniform quadrangulation of genus 0 with one boundary component having n internal

faces and perimeter $2l_n$ weakly converges, once scaled by the factor $n^{-1/4}$ and when $l_n \sim L\sqrt{2n}$, toward a random metric space called the *Brownian disk of perimeter L* . Two alternate constructions of Brownian disks were proposed by Le Gall [68, 70], allowing in particular to show that Brownian disks arise as connected components of the complement of metric balls in the Brownian sphere, conditionally given their areas and boundary lengths. See also [18, 74, 83].

Besides the case of the sphere and the disk, only a few results have been obtained for maps on compact surfaces. Namely, it has been shown that uniform quadrangulations of a given compact surface with a boundary exhibit scaling limits [19, 20, 22], all of the same topology as the considered surface, and geodesics to a uniformly chosen points were studied [22]. More recently, it was shown that uniformly distributed essentially simple toroidal triangulations (that is, triangulations of the torus without contractible loops or double edges forming cycles that are homotopic to 0) also exhibit scaling limits [15], which are believed to be the same as for random quadrangulations. See also [10] for a scaling limit result of Boltzmann random maps with annular topology.

There has also been a growing interest in noncompact versions of these models, especially as they bridge some Brownian surfaces with so-called *uniform infinite random maps*, which are maps with infinitely many faces that first arose in a work by Angel and Schramm [12], as *local limits* of random finite maps. Three main models of noncompact Brownian surfaces have been identified: the *Brownian plane* [40], the *Brownian half-plane* [13, 52], and the *infinite-volume Brownian disk* [13], which can be thought of as noncompact versions of the Brownian sphere and Brownian disks, either with unbounded or bounded boundary. See [75] for a framework unifying those objects. The first two of these models will play an important role in the current work.

This whole line of research crucially depends on strong combinatorial techniques, and in particular on bijective approaches [6, 28, 90] that allow to give very detailed quantitative information on the geodesic paths in random maps and their scaling limits. The present work is no exception. See, for instance, [11, 66, 71, 79, 81] for results related to the structure of geodesics in the Brownian sphere, [69] for a recent survey, and [39] for another approach called *peeling*. We note, however, that, so far, these methods are restricted to models of maps chosen uniformly, conditionally given their face degrees, as alluded to above.

Random surfaces via Liouville quantum gravity metrics. A line of research parallel to the above consists in building the limiting spaces directly as continuum random metrics in planar domains or Riemann surfaces. This approach also finds its roots in the physics theory of Liouville quantum gravity [88]. In the case of Brownian surfaces, this has first been implemented by Miller and Sheffield in a series of works [82–85], where they use a growth model called *quantum Loewner evolution (QLE)* to define a random metric on the plane, whose metric balls are described by QLE,

and whose law as an abstract metric space is equal to that of the Brownian plane. Local variants of the construction allow to define the Brownian sphere in this way. The Miller–Sheffield metric is in fact a special element of a one-parameter family of *Liouville quantum gravity* (LQG) metrics, that have been defined as scaling limits of first-passage percolation models in mollified exponentiated Gaussian free fields landscapes [46, 55]. See [47] for an overview of LQG metrics.

These constructions operate entirely in the continuum, and naturally ask whether canonical embeddings of random maps in the sphere are compatible with the convergence toward the Brownian sphere, in the sense that the metrics induced by the embedding converge to the random metric of Miller–Sheffield. Such a result was recently obtained by Holden and Sun [58] (which is the last piece of a vast research project, described in details in this reference), who showed the joint convergence of the metric and the area measure generated by a uniform plane triangulation embedded via the Cardy–Smirnov embedding in an equilateral triangle. We refer to the overview article [51].

The existence of a canonical conformal structure for Brownian surfaces was also approached in a more direct way by Gwynne, Miller and Sheffield in [56, 57]. Their method, which has been implemented so far for the plane, half-plane, sphere and disk topologies, consists in taking limits of discrete embeddings obtained directly from the continuum limit by considering Poisson–Voronoi tessellations with a finer and finer mesh, and showing that the random walk on the discrete approximation converges to Brownian motion in the plane. In passing, this allows one to define Brownian motion on the Brownian surfaces under consideration.

Random surfaces and conformal field theories. While the definition of LQG metrics applies to any field that “locally looks like” the Gaussian free field, the exact law of the latter is of crucial importance to obtain the exact law of random surfaces that arise as scaling limits of maps, and this law can be obtained from Liouville conformal field theory [88]. Here, rather than dealing with random metrics, one is rather interested in the computation of partition functions defined from the field, and it has been shown recently in a rich body of work – see [45, 50, 61] and references therein – that this theory has a probabilistic interpretation in terms of Gaussian multiplicative chaoses, which are random measures defined in terms of the Gaussian free field. This approach has unveiled fundamental integrability properties for planar Gaussian multiplicative chaoses, which can be used to provide exact distributions for various quantities related to the LQG metrics, hence to the scaling limits of random maps. For instance, in [10], the authors compute the law of the conformal modulus of a Brownian annulus, which is a member of the family of Brownian surfaces described in the present work.

The interplay between these approaches provides a wealth of methods to prove various properties of random surfaces [91], and the geometric properties of the

Brownian surfaces, as well as the other LQG metrics, are the object of intensive current research.

1.2 Generalities and terminology on maps

Surface with a boundary. Recall that a *surface with a boundary* is a nonempty Hausdorff topological space in which every point has an open neighborhood homeomorphic to some open subset of $\mathbb{R} \times \mathbb{R}_{\geq 0}$. Its *boundary* is the set of points having a neighborhood homeomorphic to a neighborhood of $(0, 0)$ in $\mathbb{R} \times \mathbb{R}_{\geq 0}$. When it is nonempty, this set forms a one-dimensional topological manifold. In this work, we will only consider *orientable* compact connected surfaces with a (possibly empty) boundary. By the classification theorem, these are characterized up to homeomorphisms by two nonnegative integers, the genus g and the number b of connected components of the boundary. We denote by $\Sigma_b^{[g]}$ the compact orientable surface of genus g with b boundary components, which is unique up to homeomorphisms. It can be obtained from the connected sum of g tori, or from the sphere in the case $g = 0$, by removing b disjoint open disks whose boundaries are pairwise disjoint circles.

Map. A *map* is a proper cellular embedding of a finite graph, possibly with multiple edges and loops, into a compact connected orientable surface *without* boundary. Here, the word *proper* means that edges can intersect only at vertices, and *cellular* means that the connected components of the complement of the edges, which are called the *faces* of the map, are homeomorphic to two-dimensional open disks. Maps will always be considered up to orientation-preserving homeomorphisms of the surface into which they are embedded. The *genus* of a map is defined as the genus of the surface into which it is embedded; we speak of *plane* maps when the genus is 0. We call *half-edge* an oriented edge in a map. With every half-edge, we may associate in a one-to-one way a *corner*, which is the angular sector lying to its left at the origin of the half-edge. Note that this makes sense because the surfaces we are considering are orientable. We say that a corner, or the corresponding half-edge, is *incident* to a face f if it lies into f . We also say that the face is incident to the corner or the half-edge in this case. The number of half-edges (or equivalently, of corners) incident to a face is called its *degree*.

A map is *rooted* if it comes with a distinguished corner – or, equivalently, a half-edge – called the *root*. Rooting is a very useful notion as it allows to kill the symmetries of a map. In fact, when dealing with nonrooted maps, we will systematically count them by weighting each map \mathbf{m} by a factor $1/\text{Aut}(\mathbf{m})$, where $\text{Aut}(\mathbf{m})$ denotes the number of automorphisms of \mathbf{m} . The latter is also equal to $2|E(\mathbf{m})|/R(\mathbf{m})$, where $E(\mathbf{m})$ is the edge set of \mathbf{m} , and $R(\mathbf{m})$ is the number of distinct rooted maps that can be obtained from the nonrooted map \mathbf{m} . Therefore, with this convention, the weighted

number of nonrooted maps in a given family of maps with a given number e of edges is simply the cardinality of the set of rooted maps from this same family, divided by $2e$.

Map with holes. We will consider maps with pairwise distinct distinguished elements, generically denoted by h_1, h_2, \dots, h_k , that can be either *faces* or *vertices*. These distinguished elements are called the *holes* of the map, a given hole being called either an *external face* or an *external vertex*, depending on its nature. The nondistinguished faces and vertices are called the *internal faces* and *internal vertices*. The degree of a hole, also called its *perimeter*, is defined as 0 for an external vertex or as the degree of the face for an external face. Beware that the boundaries of the external faces are in general neither simple curves, nor pairwise disjoint. As a result, the object obtained by removing them from the surface in which the map is embedded is not necessarily a surface. Note that, however, removing from every external face an open disk whose closure is included in the (open) face results in a surface with a boundary.

Bipartite map. Finally, we say that a map is *bipartite* if its vertex set can be partitioned into two subsets such that no edge links two vertices of the same subset.

Tuples. The many tuples considered in this work will conventionally be denoted by a boldface font letter (possibly with a subscript) and their coordinates with the same letter in a normal font, with the index written as a superscript, as in $\mathbf{x} = (x^1, \dots, x^r)$, for instance. When \mathbf{x} is a tuple of real nonnegative numbers, we set $\|\mathbf{x}\| = \sum_{i=1}^r x^i$. We denote by $\mathbf{x}y$ the concatenation of \mathbf{x} with y . Finally, when concatenating with a 1-tuple, we often identify it with its unique coordinate, writing, for instance, $\mathbf{x}0 = (x^1, \dots, x^r, 0)$.

1.3 The Gromov–Hausdorff–Prokhorov topology

In this memoir, a *metric measure space* is a triple $(\mathcal{X}, d_{\mathcal{X}}, \mu_{\mathcal{X}})$, where $(\mathcal{X}, d_{\mathcal{X}})$ is a nonempty *compact* metric space and $\mu_{\mathcal{X}}$ is a finite Borel measure on \mathcal{X} . We say that two metric measure spaces $(\mathcal{X}, d_{\mathcal{X}}, \mu_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mu_{\mathcal{Y}})$ are *isometry-equivalent* if there exists an isometry ϕ from $(\mathcal{X}, d_{\mathcal{X}})$ onto $(\mathcal{Y}, d_{\mathcal{Y}})$ such that $\mu_{\mathcal{Y}} = \phi_*\mu_{\mathcal{X}}$. This defines an equivalence relation on the class of all metric measure spaces. If $(\mathcal{X}, d_{\mathcal{X}}, \mu_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mu_{\mathcal{Y}})$ are two metric measure spaces, the *Gromov–Hausdorff–Prokhorov metric* (*GHP metric* for short) is defined by

$$\begin{aligned} d_{\text{GHP}}((\mathcal{X}, d_{\mathcal{X}}, \mu_{\mathcal{X}}), (\mathcal{Y}, d_{\mathcal{Y}}, \mu_{\mathcal{Y}})) \\ = \inf_{\substack{\phi: \mathcal{X} \rightarrow \mathcal{Z} \\ \psi: \mathcal{Y} \rightarrow \mathcal{Z}}} \{d_{\mathcal{Z}}^{\text{H}}(\phi(\mathcal{X}), \psi(\mathcal{Y})) \vee d_{\mathcal{Z}}^{\text{P}}(\phi_*\mu_{\mathcal{X}}, \psi_*\mu_{\mathcal{Y}})\}, \end{aligned} \quad (1.1)$$

where the infimum is taken over all choices of compact metric spaces $(\mathcal{Z}, d_{\mathcal{Z}})$, and all isometric maps ϕ, ψ from \mathcal{X}, \mathcal{Y} to \mathcal{Z} , where $d_{\mathcal{Z}}^{\text{H}}$ is the *Hausdorff metric* on compact subsets of \mathcal{Z} , and $d_{\mathcal{Z}}^{\text{P}}$ is the *Prokhorov metric* on finite positive measures on \mathcal{Z} , defined as follows. First, for any $\varepsilon > 0$ and any closed subset $A \subseteq \mathcal{Z}$, we denote by

$$A^\varepsilon = \{z \in \mathcal{Z} : \inf_{y \in A} d_{\mathcal{Z}}(z, y) < \varepsilon\}$$

its ε -enlargement. Then, for any compact subsets $A, B \subseteq \mathcal{Z}$,

$$d_{\mathcal{Z}}^{\text{H}}(A, B) = \inf\{\varepsilon > 0 : A \subseteq B^\varepsilon \text{ and } B \subseteq A^\varepsilon\},$$

and, for any finite Borel measures μ, ν on \mathcal{Z} ,

$$d_{\mathcal{Z}}^{\text{P}}(\mu, \nu) = \inf\{\varepsilon > 0 : \text{for all closed } A \subseteq \mathcal{Z}, \mu(A) \leq \nu(A^\varepsilon) + \varepsilon \\ \text{and } \nu(A) \leq \mu(A^\varepsilon) + \varepsilon\}.$$

Equation (1.1) defines a metric on the set \mathbb{M} of isometry-equivalence classes of metric measure spaces, making it a complete and separable metric space. The references [93, Chapter 27] as well as [2, 68] discuss relevant aspects of the GHP topology, with some variations, as the exact definition of the metric may differ from place to place.

More generally, for $\ell, m \geq 0$, we will consider ℓ -marked, m -measured metric spaces of the form $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \boldsymbol{\mu}_{\mathcal{X}})$, where

- $(\mathcal{X}, d_{\mathcal{X}})$ is a nonempty compact metric space,
- \mathbf{A} is an ℓ -tuple, called *marking*, of nonempty compact subsets of \mathcal{X} , called *marks*,
- $\boldsymbol{\mu}_{\mathcal{X}}$ is an m -tuple of finite Borel measures on \mathcal{X} .

We often consider marks that are singletons; in this case, we identify the singleton with the point it contains. We define the ℓ -marked, m -measured *Gromov–Hausdorff–Prokhorov metric* (still *GHP metric* for short) on such spaces by

$$d_{\text{GHP}}^{(\ell, m)}((\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \boldsymbol{\mu}_{\mathcal{X}}), (\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \boldsymbol{\mu}_{\mathcal{Y}})) \\ = \inf_{\substack{\phi: \mathcal{X} \rightarrow \mathcal{Z} \\ \psi: \mathcal{Y} \rightarrow \mathcal{Z}}} \left\{ d_{\mathcal{Z}}^{\text{H}}(\phi(\mathcal{X}), \psi(\mathcal{Y})) \vee \max_{1 \leq i \leq \ell} d_{\mathcal{Z}}^{\text{H}}(\phi(A^i), \psi(B^i)) \right. \\ \left. \vee \max_{1 \leq j \leq m} d_{\mathcal{Z}}^{\text{P}}(\phi_* \mu_{\mathcal{X}}^j, \psi_* \mu_{\mathcal{Y}}^j) \right\},$$

where the infimum is taken over the same family as in (1.1). Again, this defines a complete and separable metric on the set $\mathbb{M}^{(\ell, m)}$ of isometry-equivalence classes of ℓ -marked, m -measured metric spaces, where $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \boldsymbol{\mu}_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \boldsymbol{\mu}_{\mathcal{Y}})$ are isometry-equivalent if there exists an isometry ϕ from \mathcal{X} onto \mathcal{Y} such that $\phi(A^i) = B^i$ for $1 \leq i \leq \ell$ and $\phi_* \mu_{\mathcal{X}}^j = \mu_{\mathcal{Y}}^j$ for $1 \leq j \leq m$. Note that we have $d_{\text{GHP}}^{(0, 1)} = d_{\text{GHP}}$. Finally,

the space $(\mathbb{M}^{(\ell)}, d_{\text{GH}}^{(\ell)}) = (\mathbb{M}^{(\ell,0)}, d_{\text{GHP}}^{(\ell,0)})$ of ℓ -marked compact metric spaces without measures is the so-called ℓ -marked Gromov–Hausdorff metric (GH metric for short).

As a first example, we will sometimes use in the present work the *point space* $\{\varrho\}$ consisting of a single point, seen as the element $(\{\varrho\}, (\varrho, \dots, \varrho), (0, \dots, 0)) \in \mathbb{M}^{(\ell,m)}$ for any values of ℓ and m .

In what follows, we will often simply use the terminology “marked” or “measured” instead of “ ℓ -marked” or “ m -measured” if the numbers ℓ or m are clear from the context. Furthermore, when $m \geq 2$, we will often single out the first measure by writing it as a separate coordinate, writing $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}, \nu_{\mathcal{X}})$, for instance. The reason is that this first measure will often be an *area* measure whereas the other will be *boundary* measures, and these have different natural scales, as we will see shortly.

1.4 The main convergence result

Brownian surfaces. For $k \geq 0$, a *quadrangulation with k holes* is a *bipartite* map having k holes h_1, \dots, h_k and whose internal faces are all of degree 4. For¹ $n \in \mathbb{Z}_{\geq 0}$ and $\mathbf{l} = (l^1, \dots, l^k) \in (\mathbb{Z}_{\geq 0})^k$ (with the convention that $(\mathbb{Z}_{\geq 0})^0 = \{\emptyset\}$), we define the set $\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}$ of all genus g *rooted* quadrangulations with k holes having n internal faces, and whose holes h_1, \dots, h_k are of respective degrees $2l^1, \dots, 2l^k$; see Figure 1.1 for an example.

Likewise, we denote by $\mathbf{Q}_{n,\mathbf{l}}^{[g]}$ the set of *nonrooted* quadrangulations of genus g with n internal faces and half-perimeters given by \mathbf{l} . Since maps are counted with an inverse factor given by the number of automorphisms, the weighted cardinality of this set is

$$\sum_{\mathbf{q} \in \mathbf{Q}_{n,\mathbf{l}}^{[g]}} \frac{1}{\text{Aut}(\mathbf{q})} = \frac{|\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}|}{4n + 2\|\mathbf{l}\|}, \quad (1.2)$$

where $|\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}|$ is the cardinality of $\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}$, and $4n + 2\|\mathbf{l}\|$ is the number of oriented edges, hence of potential roots, in any element of $\mathbf{Q}_{n,\mathbf{l}}^{[g]}$.

It will be useful to notice for further reference that the quadrangulations with k holes in $\vec{\mathbf{Q}}_{n,\mathbf{l}}^{[g]}$ or in $\mathbf{Q}_{n,\mathbf{l}}^{[g]}$ all have the same number of internal vertices, namely

$$n + \|\mathbf{l}\| + 2 - 2g - k. \quad (1.3)$$

Indeed, let us consider such a map, and denote by v, e, f , its number of vertices, edges, faces. The number of external faces is thus $f - n$ so that the desired number

¹We will write $\mathbb{Z}_{\geq 0} = \{0, 1, 2, \dots\}$ the set of nonnegative integers, as well as $\mathbb{N} = \{1, 2, 3, \dots\}$ the set of positive integers.

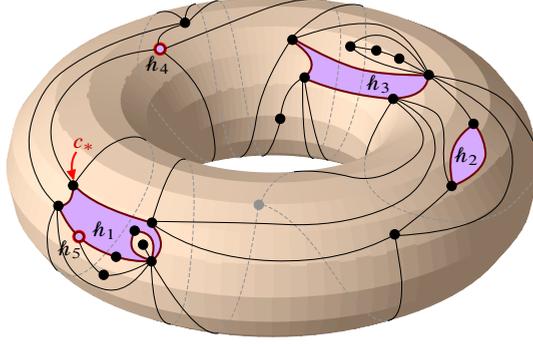


Figure 1.1. A quadrangulation from $\overline{\mathcal{Q}}_{19,(4,1,2,0,0)}^{[1]}$. The root is the corner c_* . Here, h_1 , h_2 , and h_3 are external faces, while h_4 and h_5 are external vertices.

is $v - k + f - n$. Furthermore, counting the corners yields $2e = 4n + 2\|I\|$, and the result follows from the Euler characteristic formula $v - e + f = 2 - 2g$.

If \mathbf{q} is a quadrangulation with k holes, we can view it as a k -marked, $(k + 1)$ -measured metric space, in the following way. We let $V(\mathbf{q})$ be the vertex set of \mathbf{q} , and $d_{\mathbf{q}}$ the graph metric on this set. We let

$$\partial\mathbf{q} = (V(h_1), \dots, V(h_k)),$$

where for $1 \leq i \leq k$, $V(h_i)$ is either $\{h_i\}$ if h_i is an external vertex, or the set of vertices incident to h_i if it is an external face. We let $\mu_{\mathbf{q}}$ and $\nu_{\partial\mathbf{q}}$ be the measures on $V(\mathbf{q})$ and the elements of $\partial\mathbf{q}$ defined by

$$\mu_{\mathbf{q}} = \sum_{v \in V(\mathbf{q})} \delta_v, \quad \nu_{\partial\mathbf{q}}^i = \sum_{v \in V(h_i)} m_v \delta_v,$$

where m_v , the *multiplicity* of v , is the number of corners of the face h_i that are incident to v (by convention, we set $m_{h_i} = 1$ for an external vertex h_i). These are respectively called the *area measure* and *boundary measures*. While we believe that our results also hold when $\nu_{\partial\mathbf{q}}^i$ is replaced by the counting measure on $V(h_i)$ (without multiplicities), it turns out that the above definition makes matter simpler. We associate with the quadrangulation \mathbf{q} the space

$$(V(\mathbf{q}), d_{\mathbf{q}}, \partial\mathbf{q}, \mu_{\mathbf{q}}, \nu_{\partial\mathbf{q}}) \in \mathbb{M}^{(k, k+1)}.$$

Our main result exhibits a family

$$\mathbb{S}_L^{[g]}, \quad g \geq 0, \quad L \in \bigsqcup_{k \geq 0} [0, \infty)^k,$$

of random marked measured metric spaces, where $\mathbb{S}_{\mathbf{L}}^{[g]}$ will be called the *Brownian surface of genus g with boundary perimeter vector \mathbf{L} and unit area*. The latter family describes the scaling limits of uniform random elements of $\vec{\mathbf{Q}}_{n, \mathbf{l}_n}^{[g]}$, in the following sense. Define the scaling operator Ω_n by

$$\Omega_n(\mathbf{q}) = \left(V(\mathbf{q}), \left(\frac{9}{8n} \right)^{1/4} d_{\mathbf{q}}, \partial \mathbf{q}, \frac{1}{n} \mu_{\mathbf{q}}, \frac{1}{\sqrt{8n}} \nu_{\partial \mathbf{q}} \right). \quad (1.4)$$

The scaling constants $(8/9)^{1/4}$ and $\sqrt{8}$ are here to make the upcoming description of $\mathbb{S}_{\mathbf{L}}^{[g]}$ simpler in Sections 3.4 to Chapter 6. Our main result is the following.

Theorem 1.1. *Fix $g, k \geq 0$. Let $\mathbf{L} = (L^1, \dots, L^k)$ be a k -tuple of nonnegative real numbers and, for $n \geq 1$, let $\mathbf{l}_n = (l_n^1, \dots, l_n^k) \in (\mathbb{Z}_{\geq 0})^k$ be such that $l_n^i / \sqrt{2n} \rightarrow L^i$ as $n \rightarrow \infty$ for $1 \leq i \leq k$. Let Q_n be a random variable that is uniformly distributed over $\vec{\mathbf{Q}}_{n, \mathbf{l}_n}^{[g]}$. Then*

$$\Omega_n(Q_n) \xrightarrow[n \rightarrow \infty]{(d)} \mathbb{S}_{\mathbf{L}}^{[g]},$$

where the convergence holds in distribution in the space $(\mathbb{M}^{(k, k+1)}, d_{\text{GHP}}^{(k, k+1)})$.

By our discussion on nonrooted maps, note that the same statement holds if Q_n is rather distributed over the set $\mathbf{Q}_{n, \mathbf{l}_n}^{[g]}$ of nonrooted maps, with a probability proportional to the inverse of the number of automorphisms. Note however that this automorphism number is equal to 1 for the vast majority of maps [89], so we expect that our results also hold for genuine uniform random nonrooted maps.

If $\mathbb{S}_{\mathbf{L}}^{[g]} = (\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}, \nu_{\mathcal{X}})$, we will call $\mu_{\mathcal{X}}$ the *area measure*, and $\nu_{\mathcal{X}}$ the *boundary measures*. Note that $\mu_{\mathcal{X}}$ is a probability measure, since (1.3) implies that $|V(Q_n)| \sim n$ as $n \rightarrow \infty$, while $\nu_{\mathcal{X}}^i$ has total mass L^i for $1 \leq i \leq k$, so $\nu_{\mathcal{X}}^i$ is the trivial zero measure if $L^i = 0$.

Note that, for $(g, k) = (0, 0)$, the above result amounts to the aforementioned convergence of plane quadrangulations to the Brownian sphere [67, 80], while for $(g, k) = (0, 1)$ with $L^1 > 0$, it corresponds to the convergence of quadrangulations with a boundary to the Brownian disk [25]. Note however that the statement of the present memoir is slightly stronger, since it is formulated in terms of the marked GHP topology rather than the weaker GH topology. In the case $(g, k) = (0, 0)$ of the Brownian sphere, it amounts to the GHP topology since there are no marks and only one measure; this stronger form appears, for instance, in [5, Theorem 1.2] and [68, Theorem 7].

Topology and Hausdorff dimension. Let us also list some basic properties of the limiting metric spaces, which justify the terminology of *Brownian surfaces*. We say that a metric space is *locally of Hausdorff dimension d* if any nontrivial ball has Hausdorff dimension d .

Proposition 1.2. *Let $L = (L^1, \dots, L^k)$ be fixed and let b denote the number of positive coordinates of L . Almost surely, the random metric space $\mathcal{S}_L^{[g]}$ is homeomorphic to $\Sigma_b^{[g]}$, is locally of Hausdorff dimension 4, and, if $b > 0$, each of the b connected components of its boundary, considered as a metric space by restriction of the metric on $\mathcal{S}_L^{[g]}$, is locally of Hausdorff dimension 2.*

This statement is an immediate corollary of a result from [22], showing that, in the case $b = k$, any subsequential limit in distribution of $n^{-1/4}Q_n$ satisfies the stated properties. The case $b < k$ is easily obtained from there by the observation concerning null perimeters at the end of this section. However, our method of proof of Theorem 1.1 will also provide an alternative and rather transparent proof of Proposition 1.2, once an analogous statement has been established for the noncompact analogs of the cases of the sphere and disk, namely, the Brownian plane and the Brownian half-plane [13, 40, 52] (see Section 3.8). We also mention that the case $(g, k, b) = (0, 1, 0)$ was obtained in [21] (see also [29]).

A comment on notation. Throughout this memoir, we will often work in fixed topology and consistently use the following pieces of notation, as in the above statement:

- g for the genus of the surface,
- k for the size of the boundary perimeter vector, that is the number of holes in the discrete maps,
- b , as in *boundary*, for the number of nonzero coordinates in the boundary perimeter vector,
- p , as in *puncture*, for the number of null coordinates in the boundary perimeter vector.

Beware that the latter two numbers do not always correspond to the numbers of external faces and external vertices in the discrete maps, since we only require that $L_n/\sqrt{2n} \rightarrow L$. However, for n sufficiently large, the b holes corresponding to the b nonzero coordinates in the boundary perimeter vector are external faces; each of the p remaining holes can be either a vertex or a face but, in the latter case, it should be thought of as a “small face” in the sense that its perimeter is of order $o(\sqrt{n})$, and we will see that this implies a diameter of order $o(n^{1/4})$.

Method of proof. We prove Theorem 1.1 by some surgical methods, and from the known cases $g = 0$ and $k \in \{0, 1\}$. Heuristically, we will cut Q_n along well-chosen geodesics into a *finite number* of elementary pieces of planar topology, to which we can apply a variant of the cases $(g, k) \in \{(0, 0), (0, 1)\}$ of Theorem 1.1. The idea of cutting quadrangulations along geodesics into so-called slices appears in Bouttier and Guitter [29, 30]. The use of these slices and the study of their scaling limits play an important role in Le Gall’s proof [67] of the uniqueness of the Brownian sphere (they are called *maps with a piecewise geodesic boundary* in this reference)

and are crucial to our study [25] in the case of the disk. More specifically, in the latter reference, we view Brownian disks as a continuum version of the slice decomposition.

The proof of Theorem 1.1 relies on similar but yet different ideas, and will require the introduction of other types of surgeries on objects that we call (*composite*) *slices* and *quadrilaterals (with geodesic sides)*. The core of the proof of Theorem 1.1 consists in showing scaling limit results for these elementary pieces, as stated in Theorems 2.6 and 2.8. We believe that these results are of independent interest, as elementary pieces and their scaling limits might serve as building blocks in other models of random surfaces. In order to prove this result, it turns out that it is simpler to view the discrete and continuum elementary pieces as embedded into noncompact version of the Brownian sphere and disk, namely the Brownian plane and half-plane defined in [13, 40, 52]. We stress that the description of the Brownian half-plane in terms of gluing of composite slices considered in Section 4.3 below is related to the *slice decomposition of metric bands* property used by Miller and Qian [81] for studying geodesic stars in the Brownian sphere.

Theorem 1.1 generalizes the case of the sphere at two different levels, one given by the positive genus and one given by the addition of a boundary. Although these two levels of generalization rely to some extent on similar ideas, the difficulties that they generate are of quite different nature. The case of the disk, which was the focus of [25], relied on relatively well-understood objects, but required gluing an infinite number of such objects, which in principle could create problems in the limit. On the other hand, the surgery involved in the general case consists in gluing a bounded number of objects, but the objects themselves will turn out to be of a more complicated nature.

Null perimeter coordinates. We end this section by the following observation relating Brownian surfaces in case of null perimeter coordinates. The operations of adding or removing a mark used in the following proposition are given by Lemmas 3.1 and 3.4 in Chapter 3.

Proposition 1.3. *Let $\mathbf{L} = (L^1, \dots, L^k) \in [0, \infty)^k$, and $\mathbf{L}0 = (L^1, \dots, L^k, 0)$. Then the space $\mathbb{S}_{\mathbf{L}0}^{[g]}$ has same distribution as the space $\mathbb{S}_{\mathbf{L}}^{[g]}$, where, denoting by μ the area measure of the latter space, a random μ -distributed point has been added to the set of marks of $\mathbb{S}_{\mathbf{L}}^{[g]}$ in $(k + 1)$ -th position (and the zero measure has been added as a trivial $(k + 1)$ -th boundary measure).*

Consequently, if $L^i = 0$ for some given $i \in \{1, 2, \dots, k\}$, and if $\hat{\mathbf{L}}$ denotes the vector \mathbf{L} with i -th coordinate removed, then $\mathbb{S}_{\hat{\mathbf{L}}}^{[g]}$ has same distribution as the space $\mathbb{S}_{\mathbf{L}}^{[g]}$ with its i -th mark and (trivial) i -th boundary measure removed.

Proof. Let us fix $\mathbf{l}_n = (l_n^1, \dots, l_n^k) \in (\mathbb{Z}_{\geq 0})^k$ such that $l_n^j \sim \sqrt{2n}L^j$ for $1 \leq j \leq k$, and let Q_n be uniformly distributed over $\vec{\mathbb{Q}}_{n, \mathbf{l}_n}^{[g]}$. Setting $\mathbf{l}_n 0 = (l_n^1, \dots, l_n^k, 0)$,

a uniformly distributed random variable Q'_n in $\vec{\mathbf{Q}}_{n, \mathbf{l}_n 0}^{[g]}$ may be obtained by choosing an extra distinguished external vertex h_{k+1} uniformly at random among the internal vertices of Q_n , that is, according to the measure μ_{Q_n} conditioned on the set of internal vertices. Since the number of distinguished vertices in Q_n is at most k , while the total number of vertices is asymptotically equivalent to n , the GHP limit of the quadrangulation Q'_n rescaled as in Theorem 1.1 is the same as if we had chosen h_{k+1} uniformly at random among the set of all vertices of Q_n . By Theorem 1.1 applied to Q_n and Lemma 3.1 below, we obtain the result. The second part of the statement is obtained by permuting or removing marks and measures appropriately, as discussed in Lemma 3.4 below. ■

As an example, the Brownian sphere $\mathbb{S}_\emptyset^{[0]}$ can be seen as $\mathbb{S}_{(0,0)}^{[0]}$ by forgetting its two marks. Anticipating on the construction of the Brownian surfaces in Section 3.4, this provides a nontrivially equivalent construction of the Brownian sphere as the gluing of one quadrilateral with geodesic sides, rather than the one from [67, 80].

1.5 Scaling limits of Boltzmann quadrangulations

We may also consider scaling limits for models of quadrangulations with holes in which the area and perimeters are not fixed, but rather weighted by Boltzmann factors. We introduce the following sets of nonrooted maps:

$$\mathbf{Q}_l^{[g]} = \bigsqcup_{n \geq 0} \mathbf{Q}_{n, l}^{[g]} \quad \text{for } g \geq 0, \quad l \in \bigsqcup_{k \geq 0} (\mathbb{Z}_{\geq 0})^k,$$

and

$$\mathbf{Q}^{[g]}(b, p) = \bigsqcup_{l \in \mathbb{N}^b} \mathbf{Q}_{l \mathbf{0}^p}^{[g]} \quad \text{for } g, b, p \geq 0,$$

where $\mathbf{l0}^p$ stands for the sequence \mathbf{l} to which we append p terms equal to 0.

We then let \mathcal{W} be the σ -finite measure on the set of nonrooted quadrangulations with an arbitrary number of holes and arbitrary genus, given by

$$\mathcal{W}(\mathbf{q}) = \frac{1}{\text{Aut}(\mathbf{q})} 12^{-|\mathbf{q}|} 8^{-\|\partial \mathbf{q}\|},$$

where $|\mathbf{q}|$ is the number of internal faces of \mathbf{q} , and $\|\partial \mathbf{q}\|$ is the sum of the perimeters of its holes. The reason for the choice of the weights $1/12$ and $1/8$ for the internal faces and perimeters comes from the following enumeration result, which will be proved in an extended form in Proposition B.1, in Appendix B.

Proposition 1.4. Fix $b \geq 0$ and $\mathbf{L} \in (0, \infty)^b$. Let $(\mathbf{l}_n, n \geq 0)$ be a sequence of integers such that $l_n^i \sim \sqrt{2n}L^i$ as $n \rightarrow \infty$ for $1 \leq i \leq b$. Then there exists a continuous function t_g of \mathbf{L} such that

$$\mathcal{W}(\mathbf{Q}_{n, \mathbf{l}_n}^{[g]}) \underset{n \rightarrow \infty}{\sim} t_g(\mathbf{L}) n^{(5g-7)/2+3b/4}.$$

The function $t_g(\mathbf{L})$ is related to the so-called *double scaling limit* of maps, as described in [48, Chapter 5], and its Laplace transform can be computed by solving Eynard and Orantin’s topological recursion. The method presented in Appendix B is based on the bijections presented in Chapter 2.

For any $g \geq 0$, $p \geq 0$, $\mathbf{L} \in [0, \infty)^p$ and $A > 0$, if $(\mathcal{X}, d, \mathbf{A}, \mu)$ is a random variable with same law as $\mathbb{S}_{\mathbf{L}/\sqrt{A}}^{[g]}$, we define the Brownian surface of genus g with boundary perimeter vector \mathbf{L} and area A as a random variable $\mathbb{S}_{A, \mathbf{L}}^{[g]}$ with same law as $(\mathcal{X}, A^{1/4}d, \mathbf{A}, A\mu)$. If $\mathbf{L} \in \bigsqcup_{b \geq 0} (0, \infty)^b$ and $p \geq 0$, we let $\mathbf{L}\mathbf{0}^p \in [0, \infty)^{b+p}$ be the sequence \mathbf{L} to which we append p terms equal to 0.

For integers $g, b, p \geq 0$, and for $\mathbf{L} \in (0, \infty)^b$, setting $k = b + p$, we define a σ -finite measure on $\mathbb{M}^{(k, k+1)}$ by the formula

$$\mathbb{S}_{\mathbf{L}, p}^{[g]}(\cdot) = \int_{(0, \infty)} dA A^{(5g-7)/2+3b/4+p} t_g\left(\frac{\mathbf{L}}{\sqrt{A}}\right) \mathbb{P}(\mathbb{S}_{A, \mathbf{L}\mathbf{0}^p}^{[g]} \in \cdot).$$

The measure $\mathbb{S}_{\mathbf{L}, p}^{[g]}$ is a σ -finite measure that “randomizes” the area measure of the Brownian surface of genus g with b boundary components of lengths given by \mathbf{L} , as well as p marked vertices, in the sense that its conditional law given having total area A is that of $\mathbb{S}_{A, \mathbf{L}\mathbf{0}^p}^{[g]}$.

Recall that the scaling operator Ω_n is defined by (1.4); here, we use it for any $n \in (0, \infty)$.

Theorem 1.5. Let $g, b, p \in \mathbb{Z}_{\geq 0}$, let $k = b + p$, let $\mathbf{L} \in (0, \infty)^b$, let $K > 0$, and let $F: \mathbb{M}^{(k, k+1)} \rightarrow \mathbb{R}$ be a continuous and bounded function that is supported on the set of spaces $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}, \nu_{\mathcal{X}})$ such that $\mu_{\mathcal{X}}(\mathcal{X}) \in [1/K, K]$. Let $(\mathbf{l}_a, a > 0)$ be a family where $\mathbf{l}_a \in \mathbb{N}^b$ is such that $l_a^i \sim \sqrt{2/a}L^i$ for $1 \leq i \leq b$. Then, it holds that

$$a^{5(g-1)/2+3b/4+p} \mathcal{W}(F(\Omega_{a^{-1}}(Q)) \mathbf{1}_{\mathbf{Q}_{\mathbf{l}_a \mathbf{0}^p}^{[g]}}) \xrightarrow{a \downarrow 0} \mathbb{S}_{\mathbf{L}, p}^{[g]}(F).$$

Note that our main result, Theorem 1.1, can be seen as a “local limit” version of Theorem 1.5, in the sense that it gives the conditional statement of this last result given $\mathbf{Q}_{a^{-1}, \mathbf{l}_a \mathbf{0}^p}^{[g]}$, taking $a = 1/n$. There is also a version of this theorem, where the perimeters given by \mathbf{L} are left free as well. For $g, b, p \geq 0$, we define the σ -finite measure

$$\mathbb{S}_{b, p}^{[g]}(\cdot) = \int_{(0, \infty)^b} d\mathbf{L} \mathbb{S}_{\mathbf{L}, p}^{[g]}(\cdot).$$

Corollary 1.6. *Let $g, b, p \in \mathbb{Z}_{\geq 0}$, $k = b + p$, $K > 0$, and let $F: \mathbb{M}^{(k, k+1)} \rightarrow \mathbb{R}$ be a continuous and bounded function that is supported on the set of spaces $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}, \nu_{\mathcal{X}})$ such that $\mu_{\mathcal{X}}(\mathcal{X})$ and $\nu_{\mathcal{X}}^i(A^i)$, $1 \leq i \leq b$, all lie in $[1/K, K]$. Then it holds that*

$$2^{b/2} a^{5(g-1)/2+5b/4+p} \mathcal{W}(F(\Omega_{a^{-1}}(Q))) \mathbf{1}_{\mathcal{Q}^{[g]}(b,p)} \xrightarrow{a \downarrow 0} \mathcal{S}_{b,p}^{[g]}(F).$$

Interestingly, the measure $\mathcal{S}_{L,p}^{[g]}$ is finite in the particular cases $g = 0, b = 1$ and $p \in \{0, 1\}$, or $g = 0, b = 2$ and $p = 0$; it can be checked that it is infinite in all other cases. By computing the functions $t_0(L)$ in the case $b \in \{1, 2\}$, we obtain three probability distributions by normalizing the measures $\mathcal{S}_{(L),0}^{[0]}$, $\mathcal{S}_{(L),1}^{[0]}$, $\mathcal{S}_{(L,L'),0}^{[0]}$. Those are the law of the *free Brownian disk* of perimeter $L \in (0, \infty)$:

$$\text{FBD}_L = \int_0^\infty dA \frac{L^3}{\sqrt{2\pi A^5}} \exp\left(-\frac{L^2}{2A}\right) \mathbb{P}(\mathcal{S}_{(L),A}^{[0]} \in \cdot),$$

the law of the *free pointed Brownian disk* of perimeter $L \in (0, \infty)$:

$$\text{FBD}_L^\bullet = \int_0^\infty dA \frac{L}{\sqrt{2\pi A^3}} \exp\left(-\frac{L^2}{2A}\right) \mathbb{P}(\mathcal{S}_{(L),0,A}^{[0]} \in \cdot),$$

and the law of the *free Brownian annulus* of boundary perimeters L, L' :

$$\text{FBA}_{L,L'} = \int_0^\infty dA \frac{(L + L')}{\sqrt{2\pi A^3}} \exp\left(-\frac{(L + L')^2}{2A}\right) \mathbb{P}(\mathcal{S}_{(L,L'),A}^{[0]} \in \cdot).$$

Note in particular that

$$\text{FBD}_L^\bullet = \lim_{\varepsilon \downarrow 0} \text{FBA}_{L,\varepsilon}.$$

These laws, as well as the associated σ -finite measures $\mathcal{S}_{1,0}^{[0]}$, $\mathcal{S}_{1,1}^{[0]}$, $\mathcal{S}_{2,0}^{[0]}$, play an important role in [10].

In the case $b = 0$, the two previous statements are in fact the same, since $\mathcal{S}_{\emptyset,p}^{[g]} = \mathcal{S}_{0,p}^{[g]}$. This measure describes the scaling limit of quadrangulations with no boundary, p marked vertices, and free area measure. In this case, the quantity $t_g(\emptyset)$ is equal to the classical universal constant t_g arising in map enumeration; see [17, 62]. Explicitly, the numbers $\tau_g = 2^{5g-2} \Gamma((5g-1)/2) t_g$ satisfy $\tau_0 = -1$ and the recursion

$$\tau_{g+1} = \frac{(5g+1)(5g-1)}{3} \tau_g + \frac{1}{2} \sum_{h=1}^g \tau_h \tau_{g+1-h}, \quad g \geq 0.$$

In this case, we thus have the following formula:

$$\mathcal{S}_{0,p}^{[g]}(\cdot) = t_g \int_{(0,\infty)} dA A^{(5g-7)/2+p} \mathbb{P}(\mathcal{S}_{A,0^p}^{[g]} \in \cdot).$$

1.6 Perspectives

A natural question, which we plan to investigate in future works, is to derive the analog of Theorem 1.1 for bipartite quadrangulations on *nonorientable compact surfaces*, using the bijective techniques developed in [23, 34]. The first step of showing the existence of subsequential limits for nonorientable quadrangulations without boundary has been taken in [34]. Addressing this question would complete the catalog of compact Brownian surfaces.

As mentioned in the first section of this introduction, an important aspect is that of *universality* of the spaces $\mathbb{S}_L^{[g]}$. In fact, we expect these spaces to be the scaling limits of many other models of random maps on surfaces. In the case of the Brownian sphere $\mathbb{S}_\emptyset^{[0]}$, this was indeed verified for several models; see the references mentioned above. In the case of Brownian disks, we showed in [25] that the spaces $\mathbb{S}_{(L,1)}^{[0]}$ appear as scaling limits of many conditioned Boltzmann models. This approach to universality should generalize to our context, at the price of some specific technicalities. We will not address this question here, but will comment more on this in Chapter 6.

It would be most interesting to complete the bridge between Brownian surfaces and LQG metrics and CFT (conformal field theory). As was pointed to us by Jason Miller, in order to define a canonical conformal structure and Brownian motion on Brownian surfaces, it would be natural to investigate whether the construction of general Brownian surfaces given in the present memoir, by gluing elementary pieces of disk topologies along geodesic boundaries, can be made compatible with the approach of [56, 57] mentioned in the introduction. Knowing that such a structure exists, one can try to delve even further into its integrability properties. The works [45, 50] state precise conjectures linking Liouville CFT with scaling limits of the area measure of random maps (without boundary) after suitable uniformization. In a nutshell, the LQG metrics are local objects that can be defined globally by using charts and atlases on general Riemann surfaces. However, fixing a surface amounts to fixing the conformal modulus of the LQG metric, while Brownian surfaces have a random modulus. Hence, the computation of the law of this modulus is an important question, which has been solved by [10] in the case of the annular topology. It seems that the case of general compact surfaces should be approachable as well given the recent developments on conformal bootstrap in Liouville CFT [49, 94].

We also mention that random surfaces with boundaries of the type studied in this memoir are related to the study of self-avoiding paths in random geometries. See [53, 54] for more on this in the case of the gluing of two Brownian half-planes or disks. It would be interesting to explicitly describe the scaling limits of self-avoiding paths and loops on maps of fixed topologies as gluings of Brownian surfaces along boundaries. As Nina Holden pointed to us, this would involve presumably difficult computations of the partition functions for self-avoiding loops in fixed classes of the

fundamental group of the surface, although this problem simplifies in the case of the self-avoiding loop on a Brownian sphere [9].

1.7 Organization of the memoir

In Chapter 2, we present the extension of the famous Cori–Vauquelin–Schaeffer bijection allowing to encode a quadrangulation with a simpler tree-like structure carrying integer labels on its vertices. We also present a variant of the bijection, which leads to the definition of the elementary pieces into which we decompose a quadrangulation. We finally state the relevant scaling limit results for these elementary pieces. In Chapter 3, we present the surgical operation we need in order to reconstruct a metric space from its elementary pieces, namely gluing along geodesic segments. The proofs of Theorem 1.1 and Proposition 1.2 are given in Section 3.4. In Chapters 4 and 5, we present the metric spaces forming the continuum elementary pieces into consideration and explain how they are natural building blocks of the Brownian plane and half-plane, which are the noncompact analogs of the Brownian sphere and disk, and we tweak known convergence results to these noncompact Brownian surfaces to prove that the continuum elementary pieces are the scaling limits of the discrete elementary pieces. Finally, we give in Chapter 6 an alternate description of Brownian surfaces that does not involve gluing operations and that is closer to the usual definition of the Brownian sphere and disks.

Chapter 2

Variants of the Cori–Vauquelin–Schaeffer bijection

As is customary when studying on scaling limits of maps, this work strongly relies on powerful encodings of discrete maps by tree-like objects. We now present variants of the famous Cori–Vauquelin–Schaeffer (CVS) bijection [38,90] between plane quadrangulations and so-called *well-labeled trees*, and its generalizations by Chapuy–Marcus–Schaeffer [35] for higher genera and by Bouttier–Di Francesco–Guitter [28] for plane maps with faces of arbitrary degrees. We only give the constructions from the encoding objects to the considered maps and refer the reader to the aforementioned works for converse constructions and proofs.

2.1 Basic construction

Let \mathbf{m} be a map, rooted or not, and f be a face of \mathbf{m} . Starting from a choice of a corner c_0 in f , we index the subsequent corners of f in counterclockwise order as $(c_i, i \in \mathbb{Z})$ (forming a periodic sequence). Let $\lambda: V(\mathbf{m}) \rightarrow \mathbb{Z}$ be a labeling of the vertices of \mathbf{m} by integers. We extend the definition of λ to the corners of the map by setting $\lambda(c) = \lambda(v)$ if v is the vertex incident to the corner c . In what follows, we will either consider that λ is defined up to addition of a constant, or that the value of λ at some corner is fixed, for instance, that $\lambda(c_0) = 0$.

We say that (\mathbf{m}, λ) is *well labeled* inside f if $\lambda(c_{i+1}) \geq \lambda(c_i) - 1$ for every $i \geq 0$. In particular, if (\mathbf{m}, λ) is well labeled inside f and e is a half-edge of \mathbf{m} such that both e and its reverse \bar{e} are incident to f , then $|\lambda(e^+) - \lambda(e^-)| \leq 1$, where e^- , e^+ denote the origin and end of e . Note that this will be the case for every edge when \mathbf{m} is a map with a single face.

Let (\mathbf{m}, λ) be well labeled inside f . With the above notation, let us define $s(i) = \inf\{j > i : \lambda(c_j) = \lambda(c_i) - 1\} \in \mathbb{Z} \cup \{\infty\}$ and the *successor* of c_i as $s(c_i) = c_{s(i)}$, where c_∞ is by convention the unique corner incident to a vertex v_* that is added in the interior of f , and which naturally carries the label $\lambda(v_*) = \min\{\lambda(c_i), i \in \mathbb{Z}\} - 1$. Clearly, $s(c_i)$ is then well defined for all corners (distinct from c_∞), and only depends on the corner c_i and not on the particular choice of the index i . The *CVS construction inside the face f* consists in

- linking by an arc every corner c incident to f to its successor $s(c)$, in such a way that arcs do not cross, which is always possible due to the well-labeling condition,
- deleting all the edges of \mathbf{m} .

This construction results in an embedded graph¹ denoted by $\text{CVS}(\mathbf{m}, \lambda; f)$, whose vertex set consists of v_* and the vertices of \mathbf{m} incident to f , and whose edges are the arcs between the corners of f and their successors. By construction, the edges of $\text{CVS}(\mathbf{m}, \lambda; f)$ are in bijection with the corners of \mathbf{m} incident to f . If \mathbf{m} is rooted inside f , say at the corner c_i , then $\text{CVS}(\mathbf{m}, \lambda; f)$ naturally inherits a root at the corner preceding the arc linking c_i to $s(c_i)$. Note that the well-labeling condition, as well as the output $\text{CVS}(\mathbf{m}, \lambda; f)$, are invariant under addition of a constant to λ , as they should.

We will also need an *interval variant* of this construction, where we fix a sequence, also referred to as an *interval*, of subsequent corners $I = \{c_0, c_1, \dots, c_r\}$ of a face f of \mathbf{m} , and only ask that (\mathbf{m}, λ) is *well labeled* on I in the sense that $\lambda(c_{i+1}) \geq \lambda(c_i) - 1$ for $0 \leq i \leq r - 1$. In this case, we set

$$\lambda_* = \min\{\lambda(c_i), 0 \leq i \leq r\} - 1 \quad \text{and} \quad \ell = \lambda(c_r) - \lambda_*.$$

Instead of a single extra corner c_∞ , we introduce inside f a sequence of distinct consecutive corners $c_{r+1}, c_{r+2}, \dots, c_{r+\ell}$, incident to new vertices $v_{r+1}, v_{r+2}, \dots, v_{r+\ell}$ with labels $\lambda(c_r) - 1, \lambda(c_r) - 2, \dots, \lambda_*$. The successor mapping s is then defined for all corners except $c_{r+\ell}$. We let $\text{CVS}(\mathbf{m}, \lambda; I)$ be the resulting (nonrooted) embedded graph whose edges are the arcs. In this embedded graph, the following are of particular interest:

- (1) the *apex* $v_{r+\ell}$, which will usually be denoted with a subscript $*$;
- (2) the *maximal geodesic*, which is the chain of arcs linking $c_0, s(c_0), s(s(c_0)), \dots, c_{r+\ell}$, and which will always be denoted with the letter γ and depicted in red (darker color) in the figures;
- (3) the *shuttle*, which is the chain of arcs linking $c_r, c_{r+1}, \dots, c_{r+\ell}$, and which will always be denoted with the letter ξ and depicted in green (lighter color) in the figures.

Note that the two latter are paths from the first and last corners of I to the apex.

The construction generalizes to several intervals I, J, \dots that pairwise share at most one extremity. In the case of a shared extremity, say $I = \{c_0, c_1, \dots, c_r\}$ and $J = \{c'_0, c'_1, \dots, c'_r\}$ with $c_r = c'_0$, one first duplicates the common corner before applying the construction, in the sense that the copy c_r is used in the shuttle of I and c'_0 is used in the maximal geodesic of J ; see Figure 2.1. In this construction, each interval yields a distinct apex, maximal geodesic, and shuttle, and the construction results in an embedded graph denoted by $\text{CVS}(\mathbf{m}, \lambda; I, J, \dots)$. Plainly, the ordering of the intervals does not affect the construction.

¹In general, this embedded graph is not a map of the surface into consideration. In all the constructions we will use in this work, it will, however, always turn out to be a map.

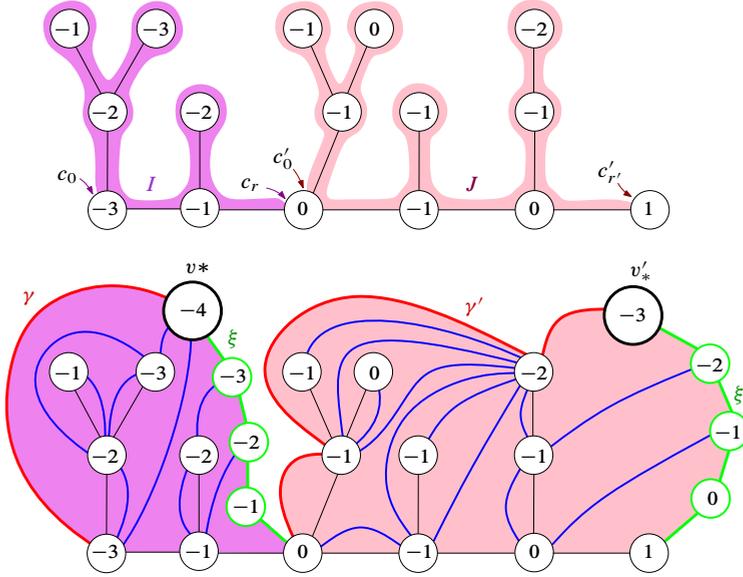


Figure 2.1. Performing the interval variant of the Cori–Vauquelin–Schaeffer bijection with two intervals sharing an extremity. The interval I consists of the corners in the purple (darker) area, starting with c_0 and ending with c_r , while J consists of the corners in the red (lighter) area, starting with c'_0 and ending with $c'_{r'}$. The interval I yields the apex v_* , maximal geodesic γ , and shuttle ξ , while J yields v'_* , γ' , and ξ' , respectively. As will be the case in all the figures, the maximal geodesics are in red (darker colored boundary) and the shuttles in green (lighter colored boundary).

We make the important observation that any chain $c, s(c), \dots, s^i(c)$ of consecutive successors induces a geodesic chain for the graph metric in the resulting embedded graph $\text{CVS}(\mathbf{m}, \lambda; I, J, \dots)$, that is, a path of minimal length between its extremities. This is simply because, by construction, any arc of the resulting embedded graph links two vertices u and v such that $|\lambda(u) - \lambda(v)| = 1$, and because λ decreases by 1 at every step on a chain of consecutive successors. In particular, the maximal geodesics and shuttles of $\text{CVS}(\mathbf{m}, \lambda; I, J, \dots)$ are geodesic chains.

2.2 The generalized Chapuy–Marcus–Schaeffer bijection

Encoding quadrangulations. As a first example, let us perform this construction on a particular class of maps. For $n \in \mathbb{Z}_{\geq 0}$ and $\mathbf{l} = (l^1, \dots, l^k) \in (\mathbb{Z}_{\geq 0})^k$, we let $\vec{\mathbf{M}}_{n, \mathbf{l}}^{[g]}$ be the set of labeled rooted maps (\mathbf{m}, λ) satisfying the following properties:

- \mathbf{m} is a map of genus g with $n + \sum_{i=1}^k l^i$ edges, one internal face f_* and k holes h_1, \dots, h_k , rooted at a corner of its internal face f_* ,

- for all i , the hole h_i is of degree l^i ; if it is an external face, then it has a simple boundary,²
- for any $i \neq j$, if h_i and h_j are faces, then they are not incident to any common edge,
- (\mathbf{m}, λ) is well labeled inside f_* .

We similarly define the set $\mathbf{M}_{n,I}^{[g]}$ of labeled nonrooted maps. Setting

$$I0 = (l^1, \dots, l^k, 0),$$

the CVS construction applied to the internal face f_* provides a bijection between $\mathbf{M}_{n,I}^{[g]}$ and $\mathbf{Q}_{n,I0}^{[g]}$, through which the k first holes correspond, while the extra hole h_{k+1} of the quadrangulation is the extra vertex v_* of the construction. In case of rooted maps, it yields a one-to-two correspondence³ between $\vec{\mathbf{M}}_{n,I}^{[g]}$ and $\vec{\mathbf{Q}}_{n,I0}^{[g]}$.

Decomposition into elementary pieces. Let us now perform the construction on the same set of maps $\mathbf{M}_{n,I}^{[g]}$ but with well-chosen intervals. We will decompose a map of $\mathbf{M}_{n,I}^{[g]}$ into a collection of labeled forests indexed by an underlying structure called the *scheme*. For the remainder of this section, we exclude the cases $(g, k) \in \{(0, 0), (0, 1)\}$ leading to encoding objects not entering the upcoming framework. We fix $(\mathbf{m}, \lambda) \in \mathbf{M}_{n,I}^{[g]}$.

Let $\tilde{\mathbf{m}}$ be the nonrooted map obtained from \mathbf{m} by iteratively removing all its vertices of degree 1 that are not holes. The resulting map $\tilde{\mathbf{m}}$ may be seen as a submap of \mathbf{m} : the map \mathbf{m} is obtained from $\tilde{\mathbf{m}}$ by appending rooted labeled trees at its corners. We call *nodes* of \mathbf{m} the following vertices:

- the external vertices of \mathbf{m} ,
- the vertices of \mathbf{m} having degree 3 or more in $\tilde{\mathbf{m}}$.

These nodes are linked in $\tilde{\mathbf{m}}$ by maximal chains of edges not containing any nodes other than their extremities. Replacing every such chain with a single edge yields a nonrooted map \mathbf{s} , called the *scheme* of \mathbf{m} . It has one internal face, still denoted by f_* , and k holes, still denoted by h_1, \dots, h_k ; see Figure 2.2.

We denote by $\vec{E}(\mathbf{s})$ the set of half-edges incident to the internal face of \mathbf{s} ; this set is partitioned into the set $\vec{I}(\mathbf{s})$ of half-edges whose reverses belong to $\vec{E}(\mathbf{s})$ as well, and the set $\vec{B}(\mathbf{s})$ of half-edges whose reverses do not belong to $\vec{E}(\mathbf{s})$. (We used the letter I for *internal* and B for *boundary*.) The set $\vec{B}(\mathbf{s})$ is further partitioned as

$$\vec{B}(\mathbf{s}) = \bigsqcup_{1 \leq r \leq k} \vec{B}_r(\mathbf{s}),$$

²A face has a *simple boundary* if it is incident to as many vertices as its degree.

³The factor 2 comes from the fact that the corners of f_* correspond to the edges of the resulting map, each edge corresponding to 2 half-edges. We refer the interested reader to [22, Section 3.1] for a presentation of the reverse mapping.

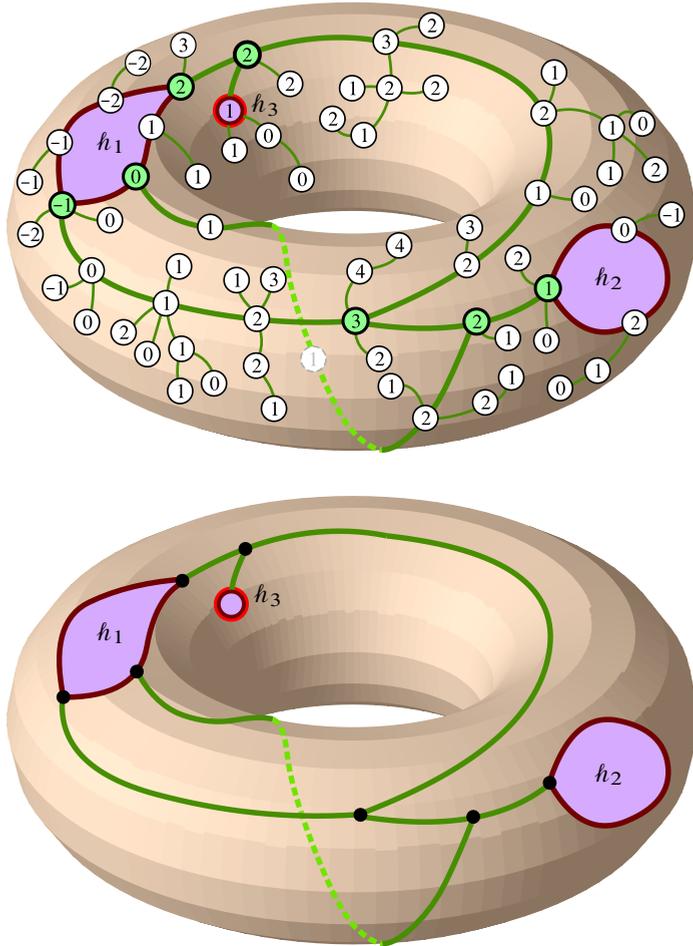


Figure 2.2. *Top:* A labeled map from $\mathbf{M}_{63, (6,3,0)}^{[1]}$. The outlined vertices are its nodes and the thicker edges correspond to the map $\tilde{\mathbf{m}}$. *Bottom:* The corresponding scheme.

where $\vec{B}_r(\mathbf{s})$ is either empty if h_r is a vertex, or the set of half-edges of $\vec{E}(\mathbf{s})$ whose reverses are incident to h_r if it is a face. We consider $e \in \vec{E}(\mathbf{s})$. It corresponds to a chain e_1, \dots, e_j of half-edges in \mathbf{m} . Let us denote by c_e and c'_e the corners of $\tilde{\mathbf{m}}$ preceding e_1 and succeeding e_j in the contour order. In \mathbf{m} , there are several corners that make up c_e and c'_e . The corner interval I_e is the interval of corners of \mathbf{m} from the first corner corresponding to c_e to the first corner corresponding to c'_e . Observe that, in \mathbf{m} , the tree grafted at c_e is thus covered by I_e , whereas the tree grafted at c'_e is not.

By construction, $\bigcup_{e \in \vec{E}(\mathbf{s})} I_e$ is equal to the set of corners of f_* and each extremity of these intervals is shared by exactly two such intervals. More precisely, the intervals I_e , $e \in \vec{E}(\mathbf{s})$, with their last corner removed give a partition of the corners

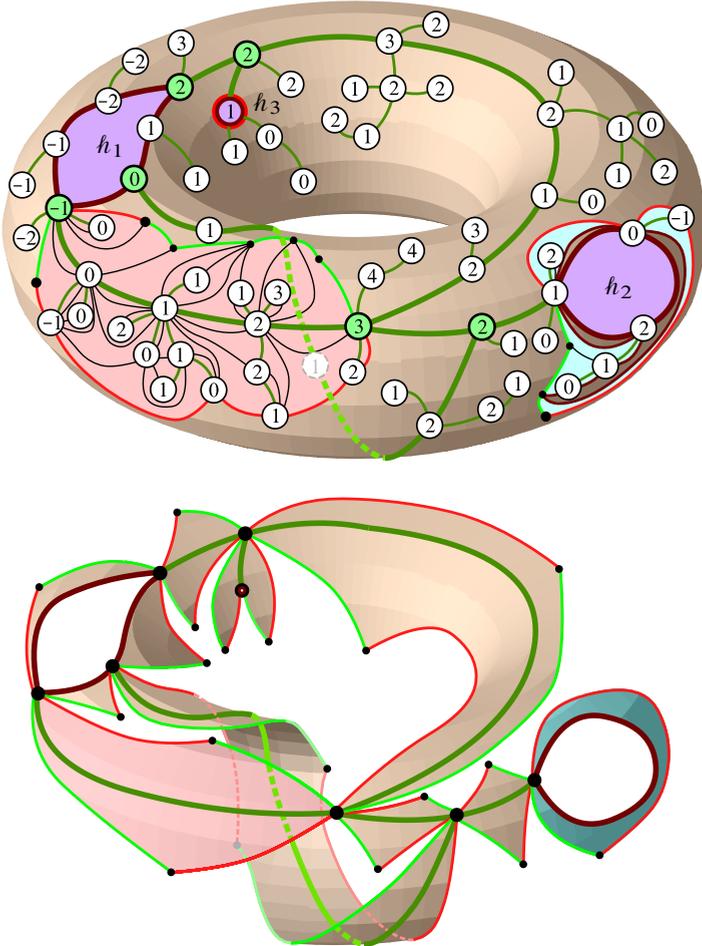


Figure 2.3. Performing the interval bijection on the labeled map from Figure 2.2. *Top.* Two elementary pieces are represented: one quadrilateral with geodesic sides in red (light color, bottom left), and one composite slice in blue (even lighter color, around h_2). *Bottom.* The interval bijection yields a decomposition into 4 composite slices and 7 quadrilaterals with geodesic sides. Here, only the maximal geodesics and shuttles are depicted. We let the edges of the original scheme figure on this output map, but these are neither edges nor chains of edges of this output map (remember that the edges of the original map are never edges of the output map).

of f_* . Applying the interval CVS construction $\text{CVS}(\mathbf{m}, \lambda; \{I_e, e \in \vec{E}(\mathbf{s})\})$ gives a natural decomposition of the quadrangulation $(\mathbf{q}, v_*) = \text{CVS}(\mathbf{m}, \lambda; f_*)$ into submaps, whose study, starting in the next section, is the key to this work; see Figure 2.3. These submaps are called the *elementary pieces* of (\mathbf{q}, v_*) and are of two types: the ones corresponding to half-edges of $\vec{B}(\mathbf{s})$ are called (*composite*) *slices* and the ones corresponding to half-edges of $\vec{I}(\mathbf{s})$ are called *quadrilaterals (with geodesic sides)*. They

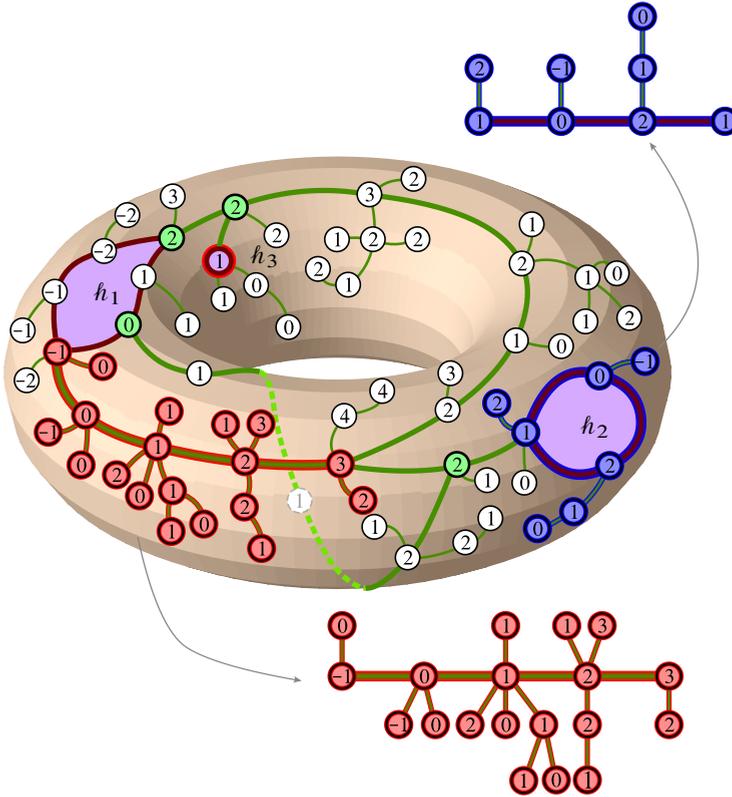


Figure 2.4. The parts of the labeled map from Figure 2.2 encoding the elementary pieces. The two parts (in red and blue) corresponding to the (red) quadrilateral with geodesic sides and the (blue) composite slice depicted in Figure 2.3 are extracted.

are not rooted and come with distinguished vertices on their boundaries that will be discussed later on.

The elementary piece corresponding to the half-edge $e \in \vec{E}(\mathbf{s})$ is encoded by the part of the labeled map (\mathbf{m}, λ) corresponding to

- either the interval I_e if $e \in \vec{B}(\mathbf{s})$,
- or the union $I_e \cup I_{\bar{e}}$ if $e \in \vec{I}(\mathbf{s})$, where \bar{e} denotes the reverse of e .

These encoding parts are depicted in Figure 2.4. Note that, when $e \in \vec{I}(\mathbf{s})$, the elementary pieces corresponding to e and to its reverse \bar{e} are the same map; only the distinguished vertices on the boundary will differ (more precisely be given in a different order). We refer the reader to [22, Section 3.4.1] for more on this decomposition, keeping in mind that, in the latter reference, the maps are rooted and the root is encoded in the scheme, which essentially amounts in seeing the root of the map as an extra external vertex.

Finiteness of the number of schemes. We will elaborate more on elementary pieces in the next two sections and end this one with a simple combinatorial lemma. We say that a map with holes is a *scheme* if it has one internal face, all its external faces have a simple boundary and do not share a common incident edge, and all its internal vertices have degree 3 or more.

Lemma 2.1. *For fixed values of $(g, k) \notin \{(0, 0), (0, 1)\}$, there are finitely many genus g schemes with k holes and these have at most $3(2g + k - 1)$ edges and $2(2g + k - 1)$ vertices.*

Proof. As there is a finite number of maps with a given number of edges, the bound on the number of edges yields the finiteness of the considered set.

Let v, e, f be the number of vertices, edges and faces of a given scheme as in the statement, and let b be its number of external faces (so that $p = k - b$ are external vertices). By construction, we have $f = b + 1$ and the vertices are all of degree at least 3, except possibly up to p of them, which have degree at least 1. The sum of the degrees of the vertices being twice the number of edges, we obtain $2e \geq 3(v - p) + p$, and we see by the Euler characteristic formula $v - e + f = 2 - 2g$ that the considered scheme has at most $6g + 3k - p - 3 \leq 3(2g + k - 1)$ edges, and at most $2(2g + k - 1)$ vertices, using that $k = p + b$. ■

2.3 Composite slices

We call *plane forest* a collection $\mathbf{f} = (\mathbf{t}^0, \dots, \mathbf{t}^{l-1}, \rho^l)$, for some $l \geq 1$, of rooted plane trees (the last one being reduced to the vertex-tree), which we view systematically as a map by taking an embedding of every \mathbf{t}^i in the upper half-plane $\mathbb{R} \times \mathbb{R}_{\geq 0}$, with root ρ^i at the point $(i, 0)$, and in which ρ^i is linked to ρ^{i-1} by the line segment between $(i, 0)$ and $(i - 1, 0)$ for $1 \leq i \leq l$. The union of these line segments is called the *floor* of the forest. The resulting embedded graph, which we still denote by \mathbf{f} (see, e.g., the left of Figure 2.5), is a nonrooted plane map coming with the two distinguished vertices $\rho = \rho^0$ and $\bar{\rho} = \rho^l$; it is in fact a plane tree, but we insist on calling it a forest.

We let a be the total number of edges of $\mathbf{t}^0, \dots, \mathbf{t}^{l-1}$, and $I = \{c_0, c_1, \dots, c_{2a+l}\}$ be the interval of corners of \mathbf{f} that are incident to the upper half-plane (hence excluding the corners that are “below” the floor), starting from the root corner of \mathbf{t}^0 and ending with the only corner incident to ρ^l , arranged in the usual contour order. We now equip \mathbf{f} with an integer-valued labeling function $\lambda: V(\mathbf{f}) \rightarrow \mathbb{Z}$, again defined up to addition of a constant, that we require to satisfy the well-labeling condition in the interval I . In this case, it means that

- $\lambda(u) - \lambda(v) \in \{-1, 0, 1\}$ when u and v are neighboring vertices of the same tree,
- for $1 \leq i \leq l$, we have $\lambda(\rho^i) \geq \lambda(\rho^{i-1}) - 1$.

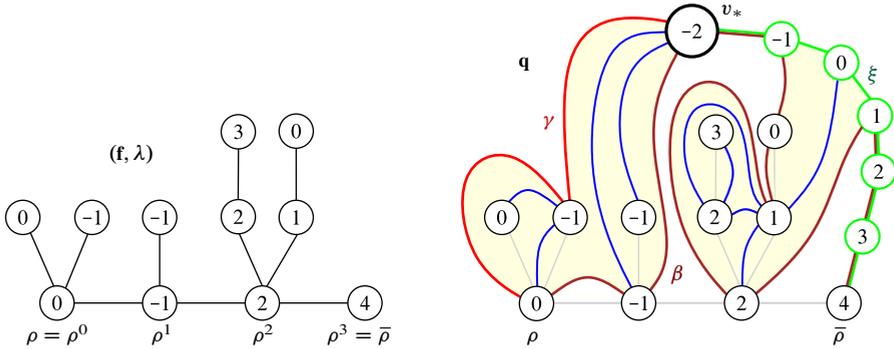


Figure 2.5. The interval Cori–Vauquelin–Schaeffer bijection giving composite slices. On this example, the forest has $a = 7$ edges and $l = 3$ trees (remember that the last vertex-tree does not count as a “real” tree). The boundary of \mathbf{q} has three parts: the maximal geodesic γ (in red), the shuttle ξ (in green) and the base β (in burgundy). Its tilt is 4.

The map $\mathbf{sl} = \text{CVS}(\mathbf{f}, \lambda; I)$ is then a nonrooted plane quadrangulation with one hole having a internal faces. Setting $\lambda_* = \min\{\lambda(v) : v \in V(\mathbf{f})\} - 1$, we see that the boundary of \mathbf{sl} has length $2(\lambda(\rho^l) - \lambda_* + I)$. It contains three distinguished vertices: ρ , the apex v_* (the extra vertex with label λ_*), and $\bar{\rho}$, as well as three distinguished paths:

- (1) the maximal geodesic γ , which has length $\lambda(\rho^0) - \lambda_*$;
- (2) the shuttle ξ , which has length $\lambda(\rho^l) - \lambda_*$;
- (3) the remaining boundary segment, called the *base* and denoted by β , consisting in the arcs connecting the root vertices of the trees. More precisely, if c_j denotes the last corner of the tree \mathbf{t}^j , then the part of the boundary of \mathbf{sl} between ρ^j and ρ^{j+1} consists in the arc linking c_j to $s(c_j)$ and the successive arcs linking $c_{j+1}, s(c_{j+1}), s(s(c_{j+1})), \dots, s(c_j)$. As a result, this base has length $\lambda(\rho^l) - \lambda(\rho^0) + 2l$. Moreover, any vertex of the base is at distance at most $\max_{1 \leq i \leq l} |\lambda(\rho^i) - \lambda(\rho^{i-1})| + 1$ from some element of the set $\{\rho^0, \dots, \rho^l\}$.

Note that, as is the case in Figure 2.5, the base may overlap with the other distinguished paths. Furthermore, as noted at the end of Section 2.1, the maximal geodesic and the shuttle are geodesic chains. On the contrary, the base is not a geodesic in general.

Definition 2.2. A map obtained by this construction will be called a *discrete composite slice*, or simply *slice* for short: its *area* is the integer a , its *width* is the integer l and its *tilt* is defined as the integer

$$\delta = \lambda(\bar{\rho}) - \lambda(\rho).$$

The terminology of composite slices, width and tilt are borrowed from [27]; however, the reader should mind that our exact definitions differ slightly from those in that reference.⁴ Note also that, in the present work, we use the simplified terminology of *slice* in order to designate a composite slice. Beware that these have not to be confused with similar objects existing in the literature, in particular in our previous work [25], called *elementary slices* or also slices for short; they actually correspond to composite slices of width 0, objects that we do not consider here.

We record the following useful counting result.

Proposition 2.3. *The number of slices with area a , width l and tilt δ is equal to*

$$3^a \frac{l}{2a+l} \binom{2a+l}{a} \binom{2l+\delta-1}{l-1},$$

which can also be recast as

$$12^a 8^l 2^\delta Q_l(2a+l) P_l(\delta),$$

where $Q_\ell(u)$ is the probability that a simple random walk hits $-\ell$ for the first time at time u , and $P_\ell(j) = \mathbb{P}(G_1 + \dots + G_\ell = j)$, where G_1, G_2, \dots are independent random variables with shifted Geometric(1/2) law, i.e., such that $\mathbb{P}(G_1 = j) = 2^{-j-2}$ for $j \geq -1$.

Proof. The term $\frac{l}{2a+l} \binom{2a+l}{a}$ is the number of forests with l trees and a nonfloor edges, the term $\binom{2l+\delta-1}{l-1}$ counts the number of possible ways to well label the roots, and the term 3^a counts the number of ways to well label the other vertices, since it amounts to choosing a label difference in $\{-1, 0, 1\}$ along each edge.

The probabilistic form is a simple exercise using the encoding of forests and geometric walks by simple walks, yielding

$$\frac{l}{2a+l} \binom{2a+l}{a} = 2^{2a+l} Q_l(2a+l) \quad \text{and} \quad \binom{2l+\delta-1}{l-1} = 2^{2l+\delta} P_l(\delta).$$

See Section 4.4 and [21, Lemma 6]. ■

An important feature of the construction is that the labels on $V(\mathbf{sl})$ inherited from those on $V(\mathbf{f})$ are exactly the relative distances to v_* in \mathbf{sl} :

$$d_{\mathbf{sl}}(v, v_*) = \lambda(v) - \lambda_*, \quad v \in V(\mathbf{sl}),$$

and that the following bound holds:

$$d_{\mathbf{sl}}(c_i, c_j) \leq \lambda(c_i) + \lambda(c_j) - 2 \min_{i \leq r \leq j} \lambda(c_r) + 2, \quad i \leq j. \quad (2.1)$$

⁴In particular, in [27], the width is the length of the base, equal to $2l + \delta$ in our notation, and the tilt is the opposite $-\delta$ of what we call the tilt in this memoir.

If \mathbf{sl} is a slice, using a slightly different convention from that of Section 1.4, we view it as the marked measured metric space in $\mathbb{M}^{(5,2)}$ given by

$$(V(\mathbf{sl}), d_{\mathbf{sl}}, \partial\mathbf{sl}, \mu_{\mathbf{sl}}, \nu_{\beta}) \quad \text{with } \partial\mathbf{sl} = (\beta, \rho, \gamma, \bar{\rho}, \xi), \quad (2.2)$$

where each boundary part is identified with the vertices it contains, where $\mu_{\mathbf{sl}}$ is the counting measure on the vertices of \mathbf{sl} that *do not belong to the shuttle*, and where ν_{β} is the counting measure (with multiplicities) on $\beta \setminus \{\bar{\rho}\}$. The measures $\mu_{\mathbf{sl}}$ and ν_{β} are respectively called the *area measure* and the *base measure* of the slice. It might be surprising at this point to include ρ and $\bar{\rho}$ in the marking as these can be found from the other three marks; they are here to enter the framework of geodesic marks introduced in Section 3.2. The idea is that the data of (ρ, γ) suffice to recover the maximal geodesic as an *oriented* path, whereas the data of γ (as a set of vertices) do not give the orientation of the path.

2.4 Quadrilaterals with geodesic sides

Consider a *double forest*, that is, a pair $(\mathbf{f}, \bar{\mathbf{f}})$ of plane forests with the same number of trees. Let $h \geq 1$ denote this common number of trees and recall that this means that \mathbf{f} and $\bar{\mathbf{f}}$ have h trees plus an additional vertex-tree. Similarly to the previous section, we represent it by letting

- the floors be both sent to the chain linking the points $(i, 0) \in \mathbb{R}^2$, where $0 \leq i \leq h$,
- the trees of \mathbf{f} be contained in the upper half-plane $\mathbb{R} \times \mathbb{R}_{\geq 0}$, the i -th tree attached to $(i - 1, 0)$ for $1 \leq i \leq h$,
- and the trees of $\bar{\mathbf{f}}$ be contained in the lower half-plane, the i -th tree attached to $(h - i + 1, 0)$ for $1 \leq i \leq h$.

We obtain a nonrooted plane map, which we denote by $\mathbf{f} \cup \bar{\mathbf{f}}$, coming with the two distinguished vertices $\rho = (0, 0)$ and $\bar{\rho} = (h, 0)$. Here also, it is in fact a plane tree having two distinguished vertices.

We let $I = \{c_0, c_1, \dots, c_{2a+h}\}$ be the interval of corners of $\mathbf{f} \cup \bar{\mathbf{f}}$, in facial order, that are incident to the upper half-plane, and $\bar{I} = \{\bar{c}_0, \bar{c}_1, \dots, \bar{c}_{2\bar{a}+h}\}$ the interval of those incident to the lower half-plane, where a (resp. \bar{a}) is the number of edges in the trees of the upper (resp. lower) half-plane. As mentioned during Section 2.1, we use the slightly unusual convention that $c_{2a+h} \neq \bar{c}_0$ (and similarly $\bar{c}_{2\bar{a}+h} \neq c_0$): this means that the first corner incident to ρ is “split” in two corners, one in the upper half-plane and one in the lower half-plane.

Finally, assume that, in its unique face, the map $\mathbf{f} \cup \bar{\mathbf{f}}$ is well labeled by an integer function $\lambda: V(\mathbf{f} \cup \bar{\mathbf{f}}) \rightarrow \mathbb{Z}$ defined up to addition of a constant: this simply means that $\lambda(u) - \lambda(v) \in \{-1, 0, 1\}$ whenever u and v are neighboring vertices. See Figure 2.6 for an example. Note that, equivalently, a well-labeled double forest $((\mathbf{f}, \bar{\mathbf{f}}), \lambda)$ can be

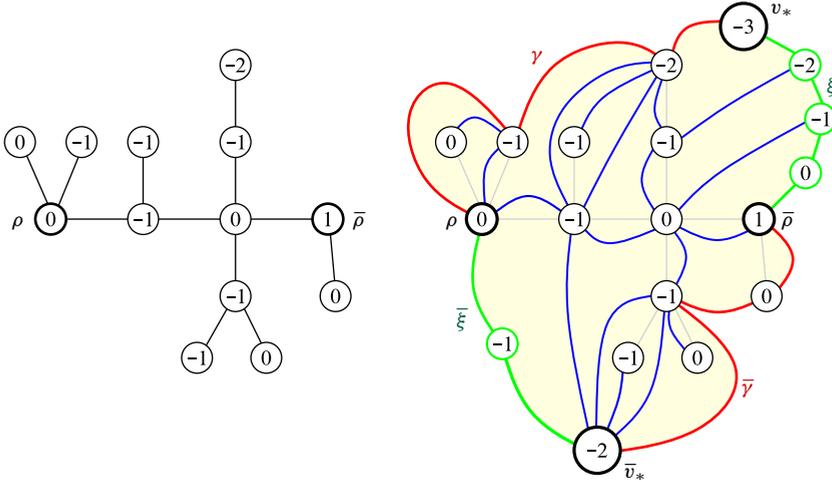


Figure 2.6. The interval Cori–Vauquelin–Schaeffer bijection giving quadrilaterals with geodesic sides. The quadrilateral with geodesic sides has half-areas 5 and 4, width 3, and tilt 1.

seen as a well-labeled *vertebrate*, that is, a well-labeled tree with two distinct distinguished vertices $\rho, \bar{\rho}$, where the interval I corresponds to the consecutive corners in the contour order from ρ to $\bar{\rho}$, which contains all the corners incident to ρ and stops at the first corner incident to $\bar{\rho}$, and \bar{I} is defined similarly with the roles of $\rho, \bar{\rho}$ exchanged.

The map $\mathbf{qd} = \text{CVS}(\mathbf{f} \cup \bar{\mathbf{f}}, \lambda; I, \bar{I})$ is then a nonrooted plane quadrangulation with one hole having $a + \bar{a} + h$ internal faces. Its boundary contains the four distinguished vertices ρ , the apex v_* associated with $I, \bar{\rho}$, and the apex \bar{v}_* associated with \bar{I} , as well as the maximal geodesics $\gamma, \bar{\gamma}$ and shuttles $\xi, \bar{\xi}$, with obvious notation.

Definition 2.4. The *quadrilateral with geodesic sides*, or simply *quadrilateral* for short, associated with $((\mathbf{f}, \bar{\mathbf{f}}), \lambda)$ is by definition $\mathbf{qd} = \text{CVS}(\mathbf{f} \cup \bar{\mathbf{f}}, \lambda; I, \bar{I})$. Its *width*, *half-areas*, and *tilt* are respectively the numbers h, a and $\bar{a}, \lambda(\bar{\rho}) - \lambda(\rho)$.

Observe that the parameters of a quadrilateral can be recovered from the map \mathbf{qd} and the distinguished vertices $\rho, \bar{\rho}$. Furthermore, the quadrilateral associated with $((\bar{\mathbf{f}}, \mathbf{f}), \lambda)$ is obtained from the one associated with $((\mathbf{f}, \bar{\mathbf{f}}), \lambda)$ simply by switching the distinguished elements ρ with $\bar{\rho}$, γ with $\bar{\gamma}$, and ξ with $\bar{\xi}$. It has the same width, its half-areas are switched and its tilt is reversed.

Proposition 2.5. *The number of quadrilaterals with half-areas a, \bar{a} , width h and tilt δ is equal to*

$$12^{a+\bar{a}+h} Q_h(2a+h) Q_h(2\bar{a}+h) M_h(\delta),$$

where Q_ℓ has been defined in Proposition 2.3, and $M_\ell(j) = \mathbb{P}(U_1 + \dots + U_\ell = j)$, where U_1, U_2, \dots are independent uniform random variables in $\{-1, 0, 1\}$.

Proof. As in the proof of Proposition 2.3, the number of forests with h floor edges and α nonfloor edges is $2^{2\alpha+h} Q_h(2\alpha+h)$, so the number of double forests with proper parameters is $4^{a+\bar{a}+h} Q_h(2a+h) Q_h(2\bar{a}+h)$. Then, the number of possible labelings of the floor vertices is the number of walks with h steps in $\{-1, 0, 1\}$ going from 0 to δ , which equals $3^h M_h(\delta)$. The final term $3^{a+\bar{a}}$ counts the possible labelings of the nonroot vertices in the double forest. ■

If \mathbf{qd} is a quadrilateral, we will view it as a marked measured metric space in $\mathbb{M}^{(6,1)}$ given by

$$(V(\mathbf{qd}), d_{\mathbf{qd}}, \partial\mathbf{qd}, \mu_{\mathbf{qd}}) \quad \text{with } \partial\mathbf{qd} = (\rho, \gamma, \xi, \bar{\rho}, \bar{\gamma}, \bar{\xi}), \quad (2.3)$$

where each boundary part is identified with the vertices it contains, and where $\mu_{\mathbf{qd}}$ is the counting measure on the vertices of \mathbf{qd} that *do not belong to the shuttles*. We call this measure $\mu_{\mathbf{qd}}$ the *area measure* of the quadrilateral.

2.5 Scaling limits of elementary pieces

In this section, we state two important results that will be crucial in the proof of Theorem 1.1. These show that, under appropriate hypotheses, random discrete slices and quadrilaterals converge in distribution in the GHP topology toward “continuum analogs” of these objects.

We first fix three sequences $(a_n) \in (\mathbb{Z}_{\geq 0})^{\mathbb{N}}$, $(l_n) \in \mathbb{N}^{\mathbb{N}}$ and $(\delta_n) \in \mathbb{Z}^{\mathbb{N}}$ such that

$$\frac{a_n}{n} \xrightarrow{n \rightarrow \infty} A > 0, \quad \frac{l_n}{\sqrt{2n}} \xrightarrow{n \rightarrow \infty} L > 0 \quad \text{and} \quad \left(\frac{9}{8n}\right)^{1/4} \delta_n \xrightarrow{n \rightarrow \infty} \Delta \in \mathbb{R}. \quad (2.4)$$

Recall that a slice is seen as an element of $\mathbb{M}^{(5,2)}$ given by (2.2) and that Ω_n is the scaling operator defined in (1.4).

Theorem 2.6. *Let Sl_n be uniformly distributed among composite slices with area a_n , width l_n and tilt δ_n . Then we have the convergence*

$$\Omega_n(\text{Sl}_n) \xrightarrow[n \rightarrow \infty]{(d)} \text{Sl}_{A,L,\Delta}$$

in distribution in the space $(\mathbb{M}^{(5,2)}, d_{\text{GHP}}^{(5,2)})$. The limit is called a (continuum composite) slice with area A , width L and tilt Δ .

This theorem will be proved in Chapter 4, where a detailed characterization of the limiting object will be given. For the time being, this theorem should be taken as a definition of the spaces $\text{Sl}_{A,L,\Delta}$.

The following statement deals with the case of vanishing areas and widths, and will be useful in Section 3.7 below.

Corollary 2.7. *Let the sequences $(a_n) \in (\mathbb{Z}_{\geq 0})^{\mathbb{N}}$ and $(l_n) \in (\mathbb{Z}_{\geq 0})^{\mathbb{N}}$ satisfy $l_n = o(\sqrt{n})$ and $a_n + l_n = \Theta((l_n)^2)$. Let Sl_n be the vertex map whenever $l_n = 0$, or be uniformly distributed among slices with area a_n , width l_n and tilt 0 otherwise. Then we have the convergence toward the point space*

$$\Omega_n(\text{Sl}_n) \xrightarrow[n \rightarrow \infty]{(d)} \{\varrho\}$$

in distribution in the space $(\mathbb{M}^{(5,2)}, d_{\text{GHP}}^{(5,2)})$.

Let us sketch in a few lines why this is indeed a consequence of Theorem 2.6. By the assumption that $a_n + l_n = \Theta((l_n)^2)$, the sequence $((a_n + l_n)/(l_n)^2)$, restricted to the values of n for which $l_n \neq 0$, is bounded away from 0 and ∞ . This compact way of writing this property covers in fact the two following situations. If (l_n) is a bounded integer sequence, it simply means that (a_n) is a bounded integer sequence, in which case the statement becomes trivial. If (l_n) is unbounded, then it means that $(a_n/(l_n)^2)$ is bounded away from 0 and ∞ . In this case, Theorem 2.6 easily implies that the diameters of $\Omega_{a_n}(\text{Sl}_n)$ form a tight family of random variables. Since $a_n = o(n)$, the conclusion follows. Note that, up to extracting subsequences, we may always assume that we are in one of the two situations discussed above.

We will derive Theorem 2.6 from the known convergence of the uniform infinite half-planar quadrangulation toward the Brownian half-plane. The former naturally contains a family of slices and the latter contains a continuous “flow” of continuum slices. These consist in free versions of the objects considered here so that we will need to finish with a conditioning argument.

We now turn to quadrilaterals, which are seen as elements of $\mathbb{M}^{(6,1)}$ given by formula (2.3). We consider four sequences $(a_n), (\bar{a}_n) \in (\mathbb{Z}_{\geq 0})^{\mathbb{N}}$, $(h_n) \in \mathbb{N}^{\mathbb{N}}$ and $(\delta_n) \in \mathbb{Z}^{\mathbb{N}}$ such that, as $n \rightarrow \infty$,

$$\frac{a_n}{n} \rightarrow A > 0, \quad \frac{\bar{a}_n}{n} \rightarrow \bar{A} > 0, \quad \frac{h_n}{\sqrt{2n}} \rightarrow H > 0, \quad \left(\frac{9}{8n}\right)^{1/4} \delta_n \rightarrow \Delta \in \mathbb{R}. \quad (2.5)$$

Theorem 2.8. *Let Qd_n be a random variable uniformly distributed among quadrilaterals with half-areas a_n and \bar{a}_n , width h_n and tilt δ_n . Then we have the convergence*

$$\Omega_n(\text{Qd}_n) \xrightarrow[n \rightarrow \infty]{(d)} \text{Qd}_{A, \bar{A}, H, \Delta}$$

in distribution in the space $(\mathbb{M}^{(6,1)}, d_{\text{GHP}}^{(6,1)})$. The limit is called a continuum quadrilateral with half-areas A and \bar{A} , width H and tilt Δ .

As for slices, the proof of this result is postponed to Chapter 5, where a detailed characterization of the limiting object will be given. The idea of the proof will be similar to that of Theorem 2.6, using the uniform infinite planar quadrangulation and Brownian plane as reference spaces instead of the half-planar versions mentioned above.

Chapter 3

Marking and gluing along geodesics

In our previous work [25], we proved Theorem 1.1 in the case of disks (for the GH topology) by writing Q_n and $S_{(L^1)}^{[0]}$ as gluings of appropriate subspaces along geodesic segments, namely so-called *slices* in the discrete setting and their scaling limits in the continuum. The fact that the number of gluings needed was infinite caused some difficulties (which we mainly overcame by noticing that any geodesic between two typical points may be broken down to a finite number of pieces lying in different such subspaces). In contrast, in this work, we will only need to consider gluings of a *finite number* of subspaces along geodesic segments. As this operation is well behaved in a more general setting, we present it in this chapter. But first, we collect a number of useful lemmas on the GHP topology.

We will use the following notation. If $\mu_{\mathcal{X}}$ is a finite positive measure on a set \mathcal{X} , we let $\bar{\mu}_{\mathcal{X}} = \mu_{\mathcal{X}} / \mu_{\mathcal{X}}(\mathcal{X})$ be the normalized probability measure. If $\mu_{\mathcal{X}} = 0$, we use the convention $\bar{\mu}_{\mathcal{X}} = 0$. If $\boldsymbol{\mu} = (\mu^1, \dots, \mu^m)$ is a finite family of nonnegative measures, we let $\bar{\boldsymbol{\mu}} = (\bar{\mu}^1, \dots, \bar{\mu}^m)$.

3.1 Useful facts on the Gromov–Hausdorff–Prokhorov topology and markings

Recall the definitions of $(\mathbb{M}^{(\ell, m)}, d_{\text{GHP}}^{(\ell, m)})$ and $(\mathbb{M}^{(\ell)}, d_{\text{GH}}^{(\ell)})$ from Section 1.3. If the space $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \boldsymbol{\mu}_{\mathcal{X}})$ is an element of $\mathbb{M}^{(\ell, m)}$ and $\mathbf{r} \in (\mathbb{Z}_{\geq 0})^m$ is such that

$$r^j = 0 \quad \text{whenever } \mu_{\mathcal{X}}^j = 0,$$

we may consider the variable

$$(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}(x_1^1, \dots, x_{r^1}^1) \cdots (x_1^m, \dots, x_{r^m}^m))$$

taking values in $\mathbb{M}^{(\ell + \|\mathbf{r}\|)}$, where, for each $j \in \{1, \dots, m\}$, the points $x_1^j, \dots, x_{r^j}^j$ are i.i.d. sampled random variables with law $\bar{\mu}_{\mathcal{X}}^j$ (if the latter measure is 0, then this still makes sense since $r^j = 0$); we denote by $\text{Mark}_{\mathbf{r}}((\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \boldsymbol{\mu}_{\mathcal{X}}), \cdot)$ the law of this random marked metric space. Some care is actually needed here since we are considering isometry classes of metric measure spaces. See [79] for an accurate definition of this notion, which is immediately generalized to our setting where we incorporate the extra marks given by \mathbf{A} , and several measures. The following lemma states that one can formulate the GHP convergence entirely in terms of the GH convergence of randomly marked spaces.

Lemma 3.1. *Let $(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n, \mu_{\mathcal{X}_n})$, $n \geq 1$, and $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}})$ be elements of $\mathbb{M}^{(\ell, m)}$. The following statements are equivalent:*

- (i) *The space $(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n, \mu_{\mathcal{X}_n})$ converges to the space $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}})$ in $(\mathbb{M}^{(\ell, m)}, d_{\text{GHP}}^{(\ell, m)})$.*
- (ii) *One has $\mu_{\mathcal{X}_n}(\mathcal{X}_n) \rightarrow \mu_{\mathcal{X}}(\mathcal{X})$ coordinatewise as $n \rightarrow \infty$ and, for every $\mathbf{r} \in (\mathbb{Z}_{\geq 0})^m$ such that $r^j = 0$ whenever $\mu_{\mathcal{X}}^j = 0$, it holds that*

$$\text{Mark}_{\mathbf{r}}((\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n, \mu_{\mathcal{X}_n}), \cdot) \xrightarrow{n \rightarrow \infty} \text{Mark}_{\mathbf{r}}((\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}), \cdot)$$

for weak convergence of probability measures on $(\mathbb{M}^{(\ell + \|\mathbf{r}\|)}, d_{\text{GH}}^{(\ell + \|\mathbf{r}\|)})$.

Proof. The implication (i) \Rightarrow (ii) is an easy generalization of known results. See [79, Proposition 10] for the case where the measures are probability measures, and [68, Section 2.2] for a generalized context with finite measures; our extended context of marked measured metric spaces adds no difficulty. To show the converse implication, we argue as follows. By taking the trivial case $\mathbf{r} = \mathbf{0}^m$ of (ii), we obtain that $\{(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n), n \geq 1\}$ is relatively compact in $(\mathbb{M}^{(\ell)}, d_{\text{GH}}^{(\ell)})$. Since the sequences $(\mu_{\mathcal{X}_n}^j(\mathcal{X}_n), n \geq 1)$ are bounded, this implies that $\{(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n, \mu_{\mathcal{X}_n}), n \geq 1\}$ is relatively compact in $(\mathbb{M}^{(\ell, m)}, d_{\text{GHP}}^{(\ell, m)})$. So let $(\mathcal{X}', d_{\mathcal{X}'}, \mathbf{A}', \mu_{\mathcal{X}'})$ be a limit in $\mathbb{M}^{(\ell, m)}$ along some subsequence of $(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n, \mu_{\mathcal{X}_n})$. By using the implication (i) \Rightarrow (ii), we obtain that, for every \mathbf{r} such that $r^j = 0$ whenever $\mu_{\mathcal{X}}^j = 0$,

$$\text{Mark}_{\mathbf{r}}((\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}), \cdot) = \text{Mark}_{\mathbf{r}}((\mathcal{X}', d_{\mathcal{X}'}, \mathbf{A}', \mu_{\mathcal{X}'}), \cdot).$$

Now, let m' be the number of nonzero elements of $\mu_{\mathcal{X}}$, fix $r > 0$ and set $r^j = r \mathbf{1}_{\{\mu_{\mathcal{X}}^j \neq 0\}}$. Let $(\mathcal{X}, d_{\mathcal{X}}, (A^1, \dots, A^\ell, x_1^1, \dots, x_{r-1}^1, \dots, x_1^m, \dots, x_{r-1}^m))$ be the $(\ell + rm')$ -marked metric space with law $\text{Mark}_{\mathbf{r}}((\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}), \cdot)$, and set $\theta_r = (\theta_r^j, 1 \leq j \leq m)$, where $\theta_r^j = r^{-1} \sum_{i=1}^r \delta_{x_i^j}$ if $\mu_{\mathcal{X}}^j \neq 0$ and $\theta_r^j = 0$ if $\mu_{\mathcal{X}}^j = 0$. It is a consequence of the law of large numbers that $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \theta_r)$ converges almost surely in $\mathbb{M}^{(\ell, m)}$, as $r \rightarrow \infty$, to $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \bar{\mu}_{\mathcal{X}})$; see, for instance, [68, Lemma 5]. Applying this same result to $(\mathcal{X}', d_{\mathcal{X}'}, \mathbf{A}', \mu_{\mathcal{X}'})$ allows us to show that $(\mathcal{X}', d_{\mathcal{X}'}, \mathbf{A}', \bar{\mu}'_{\mathcal{X}'})$ is isometry-equivalent to $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \bar{\mu}_{\mathcal{X}})$. Since $\mu_{\mathcal{X}}(\mathcal{X}) = \mu_{\mathcal{X}'}(\mathcal{X}')$ is the limit of $\mu_{\mathcal{X}_n}(\mathcal{X}_n)$, we conclude. \blacksquare

We also recall that, often, the most useful way to estimate GH distances is via the notion of *distortion of a correspondence*. A *correspondence* between two sets \mathcal{X} and \mathcal{Y} is a subset $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{Y}$ whose coordinate projections are \mathcal{X} and \mathcal{Y} . We will often write $x \mathcal{R} y$ instead of $(x, y) \in \mathcal{R}$. If \mathcal{X} and \mathcal{Y} are endowed with the metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$, the *distortion* of the correspondence \mathcal{R} is the number

$$\text{dis}(\mathcal{R}) = \sup\{|d_{\mathcal{X}}(x, x') - d_{\mathcal{Y}}(y, y')| : x \mathcal{R} y, x' \mathcal{R} y'\}.$$

If $\mathbf{A} = (A^1, \dots, A^\ell)$ and $\mathbf{B} = (B^1, \dots, B^\ell)$ are markings of \mathcal{X} and of \mathcal{Y} , we say that the correspondence \mathcal{R} between \mathcal{X} and \mathcal{Y} is *compatible* with the markings if for every $1 \leq i \leq \ell$, $\mathcal{R} \cap (A^i \times B^i)$ is a correspondence between A^i and B^i .

Lemma 3.2 ([79, Section 6.4]). *It holds that*

$$d_{\text{GH}}^{(\ell)}((\mathcal{X}, \mathbf{A}, d_{\mathcal{X}}), (\mathcal{Y}, \mathbf{B}, d_{\mathcal{Y}})) = \frac{1}{2} \inf_{\mathcal{R}} \text{dis}(\mathcal{R}),$$

where the infimum is taken over correspondences compatible with the markings.

Correspondences are also useful for estimating GHP distances when used together with the notion of couplings, which are measures on the product of the two spaces to be compared. The following is a direct adaptation of [68, Lemma 4], which treats the case of $\mathbb{M}^{(0,1)}$.

Lemma 3.3. *Let $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \mu_{\mathcal{Y}})$ be elements of $\mathbb{M}^{(\ell,m)}$ for some $\ell, m \geq 0$. Let $\varepsilon > 0$, and let \mathcal{R} be a correspondence between \mathcal{X} and \mathcal{Y} compatible with the markings and of distortion bounded above by ε . For $1 \leq j \leq m$, let ν^j be a finite measure on the product $\mathcal{X} \times \mathcal{Y}$ such that $\nu^j(\mathcal{R}^c) < \varepsilon$ and, letting $p_{\mathcal{X}}, p_{\mathcal{Y}}$ be the coordinate projections onto \mathcal{X} and \mathcal{Y} ,*

$$d_{\mathcal{X}}^{\text{P}}(\mu_{\mathcal{X}}^j, (p_{\mathcal{X}})_* \nu^j) \vee d_{\mathcal{Y}}^{\text{P}}(\mu_{\mathcal{Y}}^j, (p_{\mathcal{Y}})_* \nu^j) < \varepsilon.$$

Then $d_{\text{GHP}}^{(\ell,m)}((\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}), (\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \mu_{\mathcal{Y}})) \leq 3\varepsilon$.

Finally, we state an elementary lemma whose proof is straightforward and omitted.

Lemma 3.4. *The mappings*

$$\begin{aligned} (\mathcal{X}, d_{\mathcal{X}}, (A^1, \dots, A^\ell), \mu_{\mathcal{X}}) &\mapsto (\mathcal{X}, d_{\mathcal{X}}, (A^1 \cup A^2, A^3, \dots, A^\ell), \mu_{\mathcal{X}}), \\ (\mathcal{X}, d_{\mathcal{X}}, (A^1, \dots, A^\ell), \mu_{\mathcal{X}}) &\mapsto (\mathcal{X}, d_{\mathcal{X}}, (A^1, \dots, A^{\ell-1}), \mu_{\mathcal{X}}) \end{aligned}$$

are 1-Lipschitz from $(\mathbb{M}^{(\ell,m)}, d_{\text{GHP}}^{(\ell,m)})$ to $(\mathbb{M}^{(\ell-1,m)}, d_{\text{GHP}}^{(\ell-1,m)})$; the mappings

$$\begin{aligned} (\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, (\mu_{\mathcal{X}}^1, \dots, \mu_{\mathcal{X}}^m)) &\mapsto (\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, (\mu_{\mathcal{X}}^1 + \mu_{\mathcal{X}}^2, \mu_{\mathcal{X}}^3, \dots, \mu_{\mathcal{X}}^m)), \\ (\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, (\mu_{\mathcal{X}}^1, \dots, \mu_{\mathcal{X}}^m)) &\mapsto (\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, (\mu_{\mathcal{X}}^1, \dots, \mu_{\mathcal{X}}^{m-1})) \end{aligned}$$

are 2-Lipschitz and 1-Lipschitz from $(\mathbb{M}^{(\ell,m)}, d_{\text{GHP}}^{(\ell,m)})$ to $(\mathbb{M}^{(\ell,m-1)}, d_{\text{GHP}}^{(\ell,m-1)})$; and, for every permutation σ of $\{1, 2, \dots, \ell\}$ and τ of $\{1, 2, \dots, m\}$,

$$\begin{aligned} (\mathcal{X}, d_{\mathcal{X}}, (A^1, \dots, A^\ell), (\mu_{\mathcal{X}}^1, \dots, \mu_{\mathcal{X}}^m)) \\ \mapsto (\mathcal{X}, d_{\mathcal{X}}, (A^{\sigma(1)}, \dots, A^{\sigma(\ell)}), (\mu_{\mathcal{X}}^{\tau(1)}, \dots, \mu_{\mathcal{X}}^{\tau(m)})) \end{aligned}$$

is an isometry from $(\mathbb{M}^{(\ell,m)}, d_{\text{GHP}}^{(\ell,m)})$ onto itself.

3.2 Geodesics in metric spaces

We now discuss the important notion of geodesics in metric spaces, as well as its relations with GHP limits.

In a metric space $(\mathcal{X}, d_{\mathcal{X}})$, compact or not, a *geodesic* is a mapping $\chi: [0, \ell] \rightarrow \mathcal{X}$ defined on some compact interval¹ $[0, \ell]$ and that is isometric, i.e., satisfies

$$d_{\mathcal{X}}(\chi(s), \chi(t)) = |t - s|, \quad 0 \leq s, t \leq \ell. \quad (3.1)$$

The points $\chi(0)$, $\chi(\ell)$ are called the *extremities* of χ , and $\ell = d_{\mathcal{X}}(\chi(0), \chi(\ell))$ is the *length* of the geodesic, denoted by $\text{length}_{d_{\mathcal{X}}}(\chi)$, or simply $\text{length}(\chi)$ when there is little risk of ambiguity. The space $(\mathcal{X}, d_{\mathcal{X}})$ is called a *geodesic space* if, for every pair of points $x, y \in \mathcal{X}$, there exists a geodesic with extremities x, y .

The range $\chi([0, \text{length}(\chi)])$ of a geodesic path is called a *geodesic segment*. An *oriented geodesic segment* is a pair $(\chi(0), \chi([0, \text{length}(\chi)]))$ made of a geodesic segment and a distinguished extremity, called its *origin*. Note that an oriented geodesic segment uniquely determines the geodesic χ , since $\chi(t)$ is the unique point at distance t away from the origin. For this reason, we will systematically identify geodesics with oriented geodesic segments and use the same piece of notation for both of them.

In a marked measured metric space $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}})$, some pairs (A^i, A^j) of marks might be oriented geodesic segments; such pairs are called *geodesic marks*.

Geodesic marks and GHP limits. The following proposition states that geodesic marks nicely pass to the limit in the GHP topology.

Proposition 3.5. *Let $(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n, \mu_{\mathcal{X}_n})$, $n \geq 1$, be a sequence of marked measured compact metric spaces that converges to some limit $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}})$ in the GHP topology. Suppose that i, j are fixed and that, for every n , the pair of marks $(A_n^i, A_n^j) = \gamma_n$ is a geodesic mark. Then the pair of marks $(A^i, A^j) = \gamma$ of \mathbf{A} is also a geodesic mark. Moreover, it holds that*

$$\text{length}(\gamma) = \lim_{n \rightarrow \infty} \text{length}(\gamma_n).$$

Proof. The wanted property deals only with the marks and not with the measures, so it suffices to establish the proposition in the space $\mathbb{M}^{(\ell)}$ of marked, nonmeasured spaces. Without loss of generality (by Lemma 3.4), we may and will assume that $i = 1$ and $j = 2$.

By Lemma 3.2, we may find a sequence of correspondences \mathcal{R}_n between \mathcal{X}_n and \mathcal{X} that is compatible with the markings \mathbf{A}_n and \mathbf{A} , and whose distortion

¹We allow $\ell = 0$ in this definition.

$\varepsilon_n := \text{dis}(\mathcal{R}_n)$ goes to zero. From now on, we will never need to refer to marks other than the first two.

Let $y, z \in A^1$. Since A_n^1 contains a single point, which we denote by $x_n = \gamma_n(0)$, we have $x_n \mathcal{R}_n y$ and $x_n \mathcal{R}_n z$, so that $d_{\mathcal{X}}(y, z) \leq \varepsilon_n$ for every $n \geq 1$, entailing $y = z$. So A^1 is a singleton, which we denote by $A^1 = \{x\}$.

Next, let $a \in A^2$ and $a_n \in A_n^2$ be such that $a_n \mathcal{R}_n a$. Then

$$|d_{\mathcal{X}}(x, a) - d_{\mathcal{X}_n}(x_n, a_n)| \leq \varepsilon_n,$$

which implies that $d_{\mathcal{X}_n}(x_n, a_n) \rightarrow d_{\mathcal{X}}(x, a)$ as $n \rightarrow \infty$, and in particular, $d_{\mathcal{X}}(x, a) \leq \liminf_{n \rightarrow \infty} \text{length}(\gamma_n)$, and therefore

$$\max_{a \in A^2} d_{\mathcal{X}}(x, a) \leq \liminf_{n \rightarrow \infty} \text{length}(\gamma_n).$$

In the other direction, let $t \leq \limsup_{n \rightarrow \infty} \text{length}(\gamma_n)$. We claim that there exists at least a point $c_t \in A^2$ such that $d_{\mathcal{X}}(x, c_t) = t$. This will entail that $\text{length}(\gamma_n)$ converges to $l = \max_{a \in A^2} d_{\mathcal{X}}(x, a)$. To see the claim, observe that, at least along a suitable extraction, there exists a sequence $t_n \rightarrow t$ such that $\text{length}(\gamma_n) \geq t_n$. Along this extraction, let $\gamma_n(t_n)$ be the unique point of γ_n such that $t_n = d_{\mathcal{X}_n}(x_n, \gamma_n(t_n))$, and let g_n be an element of A^2 such that $\gamma_n(t_n) \mathcal{R}_n g_n$. Then $|d_{\mathcal{X}}(x, g_n) - t_n| \leq \varepsilon_n$, so that, possibly by further extracting, (g_n) converges to a limit $c_t \in A^2$. It then holds that $d_{\mathcal{X}}(x, c_t) = t$, as claimed.

Now fix $s, t \in [0, l]$ with $s \leq t$, and let $a, b \in A^2$ be such that $d_{\mathcal{X}}(x, a) = s$ and $d_{\mathcal{X}}(x, b) = t$. By the triangle inequality, we have $d_{\mathcal{X}}(a, b) \geq t - s$, and, on the other hand, if $a_n \mathcal{R}_n a$ and $b_n \mathcal{R}_n b$ with $a_n, b_n \in A_n^2$, then

$$\begin{aligned} d_{\mathcal{X}}(a, b) &\leq d_{\mathcal{X}_n}(a_n, b_n) + \varepsilon_n = |d_{\mathcal{X}_n}(x_n, b_n) - d_{\mathcal{X}_n}(x_n, a_n)| + \varepsilon_n \\ &\leq |d_{\mathcal{X}}(x, b) - d_{\mathcal{X}}(x, a)| + 3\varepsilon_n \\ &= t - s + 3\varepsilon_n, \end{aligned}$$

where in the second line, we have used the fact that a_n, b_n lie on a geodesic having x_n as one of its extremities. Letting $n \rightarrow \infty$, this shows that $d_{\mathcal{X}}(a, b) = t - s$, and in particular, taking $t = s$ shows that the point c_t of the preceding paragraph is the unique point of A^2 at distance t from x . We conclude that γ is an oriented geodesic segment with length l and origin x . ■

Maps as compact geodesic metric spaces. So far, we have been seeing maps as finite metric spaces. We may also interpret a map \mathbf{m} as a compact geodesic metric space, by viewing each edge as isometric to a real segment of length 1 (this is called the metric graph [32] associated with \mathbf{m}). Note that the restriction of the metric to the subset corresponding to the vertex set of \mathbf{m} is the graph metric, so that the two metric spaces corresponding to \mathbf{m} are at d_{GH} -distance less than $1/2$. In the scaling limit, this bears no effects.

With this point of view on maps, note that, in the notation of Sections 2.3 and 2.4,

- (ρ, γ) and $(\bar{\rho}, \xi)$ are geodesic marks of **sl**,
- (ρ, γ) , $(\bar{\rho}, \xi)$, $(\bar{\rho}, \bar{\gamma})$, and $(\rho, \bar{\xi})$ are geodesic marks of **qd**.

3.3 Gluing along geodesics

Quotient pseudometrics. Let (\mathcal{X}, d) be a pseudometric space, that is, a set equipped with a symmetric function $d: \mathcal{X}^2 \rightarrow \mathbb{R}_{\geq 0} \sqcup \{\infty\}$ that vanishes on the diagonal and satisfies the triangle inequality. Then $\{d = 0\}$ is an equivalence relation on \mathcal{X} , and the quotient set $\mathcal{X}/\{d = 0\}$ equipped with the function induced by d (still denoted by d for simplicity), is a true metric space, meaning that d is also separated.

Let R be an equivalence relation on \mathcal{X} . Let d/R be the largest pseudometric on \mathcal{X} such that $d/R \leq d$ and that satisfies $d/R(x, y) = 0$ as soon as $x R y$. By [32, Theorem 3.1.27], it is given by the formula

$$d/R(x, y) = \inf \left\{ \sum_{i=1}^m d(x_i, y_i) : m \geq 1, x_1, \dots, x_m, y_1, \dots, y_m \in \mathcal{X}, x_1 = x, \right. \\ \left. y_m = y, y_i R x_{i+1} \text{ for } i \in \{1, \dots, m-1\} \right\}. \quad (3.2)$$

In this setting, the set $\{d/R = 0\}$ is another equivalence relation on \mathcal{X} that contains R , possibly strictly. We let $(\mathcal{X}, d)/R = (\mathcal{X}/\{d/R = 0\}, d/R)$ and call it the gluing of (\mathcal{X}, d) along R .

A simple observation is that if R_1, R_2 are two equivalence relations on \mathcal{X} , then we have the equality of pseudometrics on \mathcal{X}

$$(d/R_1)/R_2 = (d/R_2)/R_1 = d/R, \quad (3.3)$$

where R is the coarsest equivalence relation containing $R_1 \cup R_2$. This expression is indeed a direct consequence of (3.2) and the fact that $x R y$ if and only if there exist some integer m and points $x_0 = x, x_1, \dots, x_m = y$ such that $(x_{i-1}, x_i) \in R_1 \cup R_2$ for every $i \in \{1, 2, \dots, m\}$.

Gluing two spaces along geodesics. Let $(\mathcal{X}, d_{\mathcal{X}}), (\mathcal{Y}, d_{\mathcal{Y}})$ be two pseudometric spaces and γ, ξ be two geodesics in \mathcal{X} and \mathcal{Y} , respectively, where the definition of a geodesic given by (3.1) is naturally extended to pseudometric spaces. The pseudometric of the disjoint union $\mathcal{X} \sqcup \mathcal{Y}$ is defined by

$$d_{\mathcal{X} \sqcup \mathcal{Y}}(x, y) = \begin{cases} d_{\mathcal{X}}(x, y) & \text{if } x, y \in \mathcal{X}, \\ d_{\mathcal{Y}}(x, y) & \text{if } x, y \in \mathcal{Y}, \\ \infty & \text{otherwise.} \end{cases}$$

We define the metric gluing of \mathcal{X} and \mathcal{Y} along γ and ξ by letting $l = \text{length}(\gamma) \wedge \text{length}(\xi)$ and by setting

$$G(\mathcal{X}, \mathcal{Y}; \gamma, \xi) = (\mathcal{X} \sqcup \mathcal{Y}, d_{\mathcal{X} \sqcup \mathcal{Y}}) / R, \quad (3.4)$$

where R is the coarsest equivalence relation satisfying $\gamma(t) R \xi(t)$ for every $t \in [0, l]$.

In this particular case, the fact that γ and ξ are geodesics greatly simplifies (3.2). Indeed, $y_i R x_{i+1}$ and $y_{i+1} R x_{i+2}$ imply that $d_{\mathcal{X} \sqcup \mathcal{Y}}(y_i, x_{i+2}) = d_{\mathcal{X} \sqcup \mathcal{Y}}(x_{i+1}, y_{i+1})$. In other words, using twice the relation R does not create shortcuts. As a result, the pseudometric of the gluing is the function d_G whose restrictions to $\mathcal{X} \times \mathcal{X}$ and $\mathcal{Y} \times \mathcal{Y}$ are $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$, respectively, and

$$d_G(x, y) = d_G(y, x) = \inf_{t \in [0, l]} \{d_{\mathcal{X}}(x, \gamma(t)) + d_{\mathcal{Y}}(\xi(t), y)\} \quad \text{if } x \in \mathcal{X}, y \in \mathcal{Y}. \quad (3.5)$$

Remark 3.6. In fact, equation (3.5) holds in the more general setting of gluing along isometric subspaces [31, Chapter I.5], the underlying isometry in our context being $\gamma(t) \mapsto \xi(t)$. We will not need this level of generality here.

If $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}}) \in \mathbb{M}^{(\ell, m)}$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \mu_{\mathcal{Y}}) \in \mathbb{M}^{(\ell', m')}$ are marked measured metric spaces, we may view $G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$ as an element of $\mathbb{M}^{(\ell + \ell', m + m')}$ by assigning marks and measures

$$(\mathbf{p}(A^1), \dots, \mathbf{p}(A^\ell), \mathbf{p}(B^1), \dots, \mathbf{p}(B^{\ell'}))$$

and

$$(\mathbf{p} * \mu_{\mathcal{X}}^1, \dots, \mathbf{p} * \mu_{\mathcal{X}}^m, \mathbf{p} * \mu_{\mathcal{Y}}^1, \dots, \mathbf{p} * \mu_{\mathcal{Y}}^{m'}),$$

where $\mathbf{p}: \mathcal{X} \sqcup \mathcal{Y} \rightarrow G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$ is the canonical projection. With a slight abuse of notation, we will keep denoting these by \mathbf{AB} and $\mu_{\mathcal{X} \sqcup \mathcal{Y}}$. Observe that γ and ξ may themselves be part of the marking, in which case they induce the same marks $\mathbf{p}(\gamma) = \mathbf{p}(\xi)$ in the glued space. Observe also that geodesic marks in \mathbf{A} or in \mathbf{B} remain geodesic marks in \mathbf{AB} , due to the fact that, by definition, $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ are isometrically embedded in $G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$.

Finally, observe that the gluing of the point space as an element of $\mathbb{M}^{(\ell, m)}$ with $(\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \mu_{\mathcal{Y}})$ along ξ only has the effect of prepending ℓ times $\xi(0)$ to \mathbf{B} and m times the zero measure to $\mu_{\mathcal{Y}}$.

Gluing two geodesics in the same space. A similar gluing procedure² can be defined for two geodesics γ, ξ in the same pseudometric space $(\mathcal{X}, d_{\mathcal{X}})$. We again set $l = \text{length}(\gamma) \wedge \text{length}(\xi)$ and then define

$$G(\mathcal{X}; \gamma, \xi) = (\mathcal{X}, d_{\mathcal{X}}) / R,$$

²In fact, since we are allowing points at infinite distance, the gluing $G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$ could be seen as a particular case of gluing of a single space along two geodesics, but we refrain to do so as we will mostly be interested in gluing true metric spaces.

where R is the coarsest equivalence relation satisfying $\gamma(t) R \xi(t)$ for every $t \in [0, l]$. The quotient pseudometric $d_G(x, y)$ may be condensed into

$$\begin{aligned} d_{\mathcal{X}}(x, y) \wedge \inf_{t \in [0, l]} \{d_{\mathcal{X}}(x, \gamma(t)) + d_{\mathcal{X}}(\xi(t), y)\} \\ \wedge \inf_{t \in [0, l]} \{d_{\mathcal{X}}(x, \xi(t)) + d_{\mathcal{X}}(\gamma(t), y)\}. \end{aligned} \quad (3.6)$$

Similarly to the above, the space $G(\mathcal{X}; \gamma, \xi)$ naturally inherits the marking \mathbf{A} and measures $\mu_{\mathcal{X}}$ that \mathcal{X} may be endowed with, simply by pushing those forward by the canonical projection $\mathcal{X} \rightarrow G(\mathcal{X}; \gamma, \xi)$; by a slight abuse of notation, we keep the piece of notation \mathbf{A} , $\mu_{\mathcal{X}}$ for these inherited objects. Note however that it is *not true in general* that geodesic marks in $(\mathcal{X}, d_{\mathcal{X}})$ remain geodesic marks in $G(\mathcal{X}; \gamma, \xi)$. Let us state a useful comparison result between $d_{\mathcal{X}}$ and d_G .

Lemma 3.7. *Let $(\mathcal{X}, d_{\mathcal{X}})$ be a pseudometric space with two distinguished geodesics γ, ξ . Denote by d_G the pseudometric on $G(\mathcal{X}; \gamma, \xi)$ as in (3.6).*

(i) *For every $x, y \in \mathcal{X}$,*

$$d_G(x, y) \leq d_{\mathcal{X}}(x, y) \leq d_G(x, y) + R(\gamma, \xi),$$

where $R(\gamma, \xi)$ is the Hausdorff distance in the space $(\mathcal{X}, d_{\mathcal{X}})$ between the initial segments of γ, ξ that are glued together, i.e., of length l .

(ii) *For every $\varepsilon > 0$ and $x, y \in \mathcal{X}$, if $d_{\mathcal{X}}(x, y) < \varepsilon$ and $d_{\mathcal{X}}(x, \gamma) \wedge d_{\mathcal{X}}(y, \gamma) > \varepsilon$ (or if $d_{\mathcal{X}}(x, \xi) \wedge d_{\mathcal{X}}(y, \xi) > \varepsilon$), it holds that*

$$d_G(x, y) = d_{\mathcal{X}}(x, y).$$

Proof. Let us first prove (i). The first inequality is a direct consequence of (3.6). To prove the other bound, simply observe that for every $t \in [0, l]$ we have

$$\begin{aligned} d_{\mathcal{X}}(x, y) &\leq d_{\mathcal{X}}(x, \gamma(t)) + d_{\mathcal{X}}(\gamma(t), \xi(t)) + d_{\mathcal{X}}(\xi(t), y) \\ &\leq d_{\mathcal{X}}(x, \gamma(t)) + d_{\mathcal{X}}(\xi(t), y) + R(\gamma, \xi). \end{aligned}$$

Taking the infimum over t , and then applying the same reasoning with the roles of γ, ξ interchanged, we obtain the result by (3.6).

The proof of (ii) is even more straightforward. Under our assumptions, it holds that both $d_{\mathcal{X}}(x, \gamma(t)) + d_{\mathcal{X}}(y, \xi(t))$ and $d_{\mathcal{X}}(x, \xi(t)) + d_{\mathcal{X}}(y, \gamma(t))$ are greater than $\varepsilon > d_{\mathcal{X}}(x, y)$ for every choice of t , so that $d_G(x, y)$ must be equal to $d_{\mathcal{X}}(x, y)$. ■

Gluing and GHP limits. Henceforth, we mostly focus on compact geodesic spaces. The gluing of one or two geodesic spaces along geodesics is again a geodesic space by general results presented in [32]. The gluing operation also preserves the compactness of the spaces that are glued together. Furthermore, if a compact metric space is the

Gromov–Hausdorff limit of a sequence of compact geodesic spaces, then it is also a compact geodesic space [32, Theorem 7.5.1].

The next result shows that the gluing operations behave well with respect to the GHP metric. For simplicity, we state it with the first marks of the markings but it obviously holds up to index permutations, using Lemma 3.4, for instance.

Proposition 3.8. *Let $(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n, \mu_{\mathcal{X}_n})$ and $(\mathcal{Y}_n, d_{\mathcal{Y}_n}, \mathbf{B}_n, \mu_{\mathcal{Y}_n})$ be geodesic marked measured metric spaces that converge in the marked GHP topology to $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}})$, $(\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B}, \mu_{\mathcal{Y}})$. Assume that the first pairs of marks $(A_n^1, A_n^2) = \gamma_n$ and $(B_n^1, B_n^2) = \xi_n$ of \mathcal{X}_n and of \mathcal{Y}_n are geodesic marks for every $n \geq 0$. Then the first two marks of \mathbf{A} and of \mathbf{B} are geodesic marks γ , ξ , and*

$$G(\mathcal{X}_n, \mathcal{Y}_n; \gamma_n, \xi_n) \xrightarrow{n \rightarrow \infty} G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$$

in the GHP topology.

Similarly, if we now assume that the first four marks are such that $(A_n^1, A_n^2) = \gamma_n$ and $(A_n^3, A_n^4) = \xi_n$ are geodesic marks, then the same holds for the first marks γ , ξ of \mathbf{A} , and

$$G(\mathcal{X}_n; \gamma_n, \xi_n) \xrightarrow{n \rightarrow \infty} G(\mathcal{X}; \gamma, \xi)$$

in the GHP topology.

In order to prove this proposition, we are first going to state and prove a useful lemma that allows one to deal only with the situations where the geodesics along which the spaces of interest are glued have the same lengths. While Lemma 3.4 showed that the operation of merging two marks is continuous on $(\mathbb{M}^{(\ell, m)}, d_{\text{GHP}}^{(\ell, m)})$, this lemma states that in the case of geodesic marks, the natural splitting operation is continuous. If γ is a geodesic mark and $r \in [0, \text{length}(\gamma)]$, the *splitting* of γ at level r is the two geodesic marks $(\gamma(0), \{\gamma(t) : 0 \leq t \leq r\})$, $(\gamma(r), \{\gamma(t), r \leq t \leq \text{length}(\gamma)\})$.

Lemma 3.9. *Let $(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n, \mu_{\mathcal{X}_n})$, $n \geq 0$, be a sequence of geodesic marked measured metric spaces converging in the GHP topology toward a geodesic marked measured metric space $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A}, \mu_{\mathcal{X}})$, and assume that, say, the first pairs of marks, are geodesic marks γ_n and γ . Denote by $l_n = \text{length}(\gamma_n)$ and $l = \text{length}(\gamma)$ their lengths. Let $r_n \in (0, l_n)$ be real numbers such that $r_n \rightarrow r \in (0, l)$. Then the convergence $\mathcal{X}_n \rightarrow \mathcal{X}$ still holds in the GHP topology after replacing the marks γ_n and γ in \mathbf{A}_n and \mathbf{A} , with their splittings γ'_n, γ''_n and γ', γ'' at levels r_n and r , respectively.*

Proof. Since the desired property does not involve the measures, it suffices to establish it in the space of marked, nonmeasured spaces. Let \mathcal{R} be a correspondence between \mathcal{X}_n and \mathcal{X} compatible with the markings. We fix $\varepsilon > \text{dis}(\mathcal{R})$ and consider the enlarged correspondence

$$\mathcal{R}^\varepsilon = \{(x, y) \in \mathcal{X}_n \times \mathcal{X} : \exists(x', y') \in \mathcal{R}, d_{\mathcal{X}_n}(x, x') \vee d_{\mathcal{X}}(y, y') < \varepsilon\}.$$

By the triangle inequality, the distortion of \mathcal{R}^ε is at most $\text{dis}(\mathcal{R}) + 4\varepsilon$. Moreover, we claim that for $s \in [0, l_n]$ and $t \in [0, l]$ such that $|t - s| < \varepsilon - \text{dis}(\mathcal{R})$, it holds that $\gamma_n(s) \mathcal{R}^\varepsilon \gamma(t)$. Indeed, since \mathcal{R} is compatible with the markings, for every s, t as above, there exists u such that $\gamma_n(s) \mathcal{R} \gamma(u)$, so

$$|s - u| = |d_{\mathcal{X}_n}(\gamma_n(s), \gamma_n(0)) - d_{\mathcal{X}}(\gamma(u), \gamma(0))| \leq \text{dis} \mathcal{R},$$

and therefore $|d_{\mathcal{X}}(\gamma(t), \gamma(u))| = |t - u| \leq |t - s| + |s - u| < \varepsilon$, as wanted. From this, we conclude that \mathcal{R}^ε is compatible with the geodesic marks γ'_n, γ' and γ''_n, γ'' , as soon as $\varepsilon > |r_n - r| \vee |l_n - l|$. Choosing a sequence of correspondences \mathcal{R}_n with vanishing distortion and taking $\varepsilon_n = (\text{dis}(\mathcal{R}_n) \vee |r_n - r| \vee |l_n - l|) + 1/n$ yields the result. \blacksquare

Proof of Proposition 3.8. With the help of Lemma 3.9, we may and will assume that γ_n, ξ_n have same length for every n . The fact that the limiting marks are geodesics marks with same length comes from Proposition 3.5. We will first establish the result for marked, nonmeasured spaces, from which we will deduce the full result thanks to Lemma 3.1.

Let \mathcal{R} and \mathcal{R}' be two correspondences between \mathcal{X}_n and \mathcal{X} and between \mathcal{Y}_n and \mathcal{Y} , compatible with the considered markings. Identifying \mathcal{X}_n and \mathcal{Y}_n (resp. \mathcal{X} and \mathcal{Y}) with their canonical embeddings into $G(\mathcal{X}_n, \mathcal{Y}_n; \gamma_n, \xi_n)$ (resp. into $G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$), we consider $\mathcal{R}'' = \mathcal{R} \cup \mathcal{R}'$ as a correspondence between $G(\mathcal{X}_n, \mathcal{Y}_n; \gamma_n, \xi_n)$ and $G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$, obviously compatible with the other marks. In order to bound its distortion, let us take $x_n \mathcal{R} x$ and $y_n \mathcal{R}' y$.

We let $l_n = \text{length}(\gamma_n) = \text{length}(\xi_n)$ and $l = \text{length}(\gamma) = \text{length}(\xi)$ be the lengths of the considered geodesics and we denote by d_n and d the metrics in the previous gluings. From (3.5) and by compactness, there exists $t \in [0, l]$ such that

$$d(x, y) = d_{\mathcal{X}}(x, \gamma(t)) + d_{\mathcal{Y}}(\xi(t), y).$$

Then there exist $t_n, t'_n \in [0, l_n]$ such that $\gamma_n(t_n) \mathcal{R} \gamma(t)$ and $\xi_n(t'_n) \mathcal{R}' \xi(t)$. As a result,

$$\begin{aligned} d_n(x_n, y_n) &\leq d_{\mathcal{X}_n}(x_n, \gamma_n(t_n)) + d_n(\gamma_n(t_n), \xi_n(t'_n)) + d_{\mathcal{Y}_n}(\xi_n(t'_n), y_n) \\ &\leq d_{\mathcal{X}}(x, \gamma(t)) + \text{dis}(\mathcal{R}) + d_n(\gamma_n(t_n), \xi_n(t'_n)) + d_{\mathcal{Y}}(\xi(t), y) + \text{dis}(\mathcal{R}') \\ &\leq d(x, y) + \text{dis}(\mathcal{R}) + \text{dis}(\mathcal{R}') + d_n(\gamma_n(t_n), \xi_n(t'_n)). \end{aligned}$$

Using the facts that $\gamma_n, \xi_n, \gamma, \xi$ are geodesics and $\gamma_n(0) \mathcal{R} \gamma(0)$, $\xi_n(0) \mathcal{R}' \xi(0)$, we easily obtain

$$\begin{aligned} d_n(\gamma_n(t_n), \xi_n(t'_n)) &= |d_{\mathcal{X}_n}(\gamma_n(0), \gamma_n(t_n)) - d_{\mathcal{Y}_n}(\xi_n(0), \xi_n(t'_n))| \\ &\leq |d_{\mathcal{X}}(\gamma(0), \gamma(t)) - d_{\mathcal{Y}}(\xi(0), \xi(t))| + \text{dis}(\mathcal{R}) + \text{dis}(\mathcal{R}') \\ &= \text{dis}(\mathcal{R}) + \text{dis}(\mathcal{R}'). \end{aligned}$$

Using a symmetric argument, we obtain

$$|d_n(x_n, y_n) - d(x, y)| \leq 2(\text{dis}(\mathcal{R}) + \text{dis}(\mathcal{R}')).$$

Adding to this the simpler cases where the pairs of points we compare belong both to \mathcal{R} or both to \mathcal{R}' , we obtain

$$\text{dis}(\mathcal{R}'') \leq 2(\text{dis}(\mathcal{R}) + \text{dis}(\mathcal{R}'))$$

and the first statement easily follows for the GH topology (without the measures).

Let us show that the result still holds when considering the measures. We assume for simplicity that the terms of $\mu_{\mathcal{X}}$, $\mu_{\mathcal{Y}}$ are all nonzero, since the case of a vanishing measure, say $\mu_{\mathcal{X}}^j$, is equivalent to the fact that $\mu_{\mathcal{X}_n}^j(\mathcal{X}_n) \rightarrow 0$. Denote by m, m' the numbers of coordinates of $\mu_{\mathcal{X}}$, $\mu_{\mathcal{Y}}$ and sample $r(m + m')$ independent points $\mathbf{x} = (x_i^j, 1 \leq i \leq r, 1 \leq j \leq m)$ and $\mathbf{y} = (y_i^j, 1 \leq i \leq r, 1 \leq j \leq m')$, where x_i^j has law $\bar{\mu}_{\mathcal{X}}^j$ and y_i^j has law $\bar{\mu}_{\mathcal{Y}}^j$. We identify these points with their images in the glued space by the canonical projection $\mathcal{X} \sqcup \mathcal{Y} \rightarrow G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$. We assume that $\mu_{\mathcal{X}_n}(\mathcal{X}_n) > 0$ and $\mu_{\mathcal{Y}_n}(\mathcal{Y}_n) > 0$, which hold for n sufficiently large since, as $n \rightarrow \infty$, $\mu_{\mathcal{X}_n}(\mathcal{X}_n) \rightarrow \mu_{\mathcal{X}}(\mathcal{X}) > 0$ and $\mu_{\mathcal{Y}_n}(\mathcal{Y}_n) \rightarrow \mu_{\mathcal{Y}}(\mathcal{Y}) > 0$. We proceed similarly to sample $r(m + m')$ random points $\mathbf{x}_n = (x_{n,i}^j)$, $\mathbf{y}_n = (y_{n,i}^j)$ in $G(\mathcal{X}_n, \mathcal{Y}_n; \gamma_n, \xi_n)$ with laws $\bar{\mu}_{\mathcal{X}_n}^j$ and $\bar{\mu}_{\mathcal{Y}_n}^j$ as appropriate. Lemma 3.1 guarantees that the marked spaces $(\mathcal{X}_n, d_{\mathcal{X}_n}, \mathbf{A}_n \mathbf{x}_n)$ and $(\mathcal{Y}_n, d_{\mathcal{Y}_n}, \mathbf{B}_n \mathbf{y}_n)$ converge to $(\mathcal{X}, d_{\mathcal{X}}, \mathbf{A} \mathbf{x})$ and $(\mathcal{Y}, d_{\mathcal{Y}}, \mathbf{B} \mathbf{y})$ in distribution in the GH topology. Applying the result of Proposition 3.8 proved above in the case without measures, we obtain the convergence in distribution of the glued space $G(\mathcal{X}_n, \mathcal{Y}_n; \gamma_n, \xi_n)$ with markings $\mathbf{A}_n \mathbf{B}_n \mathbf{x}_n \mathbf{y}_n$ to $G(\mathcal{X}, \mathcal{Y}; \gamma, \xi)$ with the marking $\mathbf{A} \mathbf{B} \mathbf{x} \mathbf{y}$. Since $\mathbf{x}_n, \mathbf{y}_n$ and \mathbf{x}, \mathbf{y} are also independent samples from the renormalized measures $\bar{\mu}_{\mathcal{X}_n}, \bar{\mu}_{\mathcal{Y}_n}$ and $\bar{\mu}_{\mathcal{X}}, \bar{\mu}_{\mathcal{Y}}$ viewed as measures on the glued spaces, an application of the converse implication of Lemma 3.1 implies the result. ■

The second part of the statement dealing with metric spaces that are glued along two marked geodesics is shown in a similar fashion. We leave the details to the reader. ■

3.4 Proof of Theorem 1.1

Let $g, k \in \mathbb{Z}_{\geq 0}$ be fixed; as in Section 2.2, we exclude the cases $(g, k) \in \{(0, 0), (0, 1)\}$ of the sphere, the pointed sphere, or the disk. Recall that the case $(g, k) = (0, 0)$ of the sphere is already known. The case $(g, b, k) = (0, 1, 1)$ of the disk is partially known but has not been treated in the complete setting of Theorem 1.1. In fact, it may actually enter the following framework: in this case, the decomposition of Section 2.2 yields only one well-labeled forest (and thus one unique slice) indexed

by a degenerate scheme with one external face having one unique vertex with a self-loop edge. This creates a small ambiguity coming from the choice of the first tree in the forest, which can be overcome by randomization when considering random maps. The case $(g, b, k) = (0, 0, 1)$ of the pointed sphere would yield an even more degenerate decomposition with a one-vertex map as a scheme and a unique composite slice of width 0, object that is not introduced in this work.

Instead of considering these extensions and objects, we rather obtain these cases by using Proposition 1.3 at the end of Section 1.4. More precisely, provided Theorem 1.1 holds for $(g, k) = (0, 2)$, we infer the case $(g, k) = (0, 1)$ as follows. Let $l_n^1 \in \mathbb{Z}_{\geq 0}$ be such that $l_n^1/\sqrt{2n} \rightarrow L^1$ as $n \rightarrow \infty$ and Q_n be uniform in $\vec{\mathcal{Q}}_{n, (l_n^1)}^{[0]}$. Then choose a vertex uniformly at random among the internal vertices of Q_n and denote by Q_n^* the map obtained from Q_n by declaring the chosen vertex as a second hole. Since Q_n^* is clearly uniform in $\vec{\mathcal{Q}}_{n, (l_n^1, 0)}^{[0]}$, the case $(g, k) = (0, 2)$ of Theorem 1.1 implies the convergence of the corresponding metric measure space toward $\mathcal{S}_{(L^1, 0)}^{[0]}$, and finally the convergence of the metric measure space corresponding to Q_n toward the space $\mathcal{S}_{(L^1)}^{[0]}$, defined as $\mathcal{S}_{(L^1, 0)}^{[0]}$ with its second mark and second boundary measure forgotten.

3.5 Gluing quadrangulations from elementary pieces

We start by interpreting the observations of Section 2.2 in the light of the previous section for a deterministic map. Let $n \in \mathbb{Z}_{\geq 0}$ and $l \in \mathbb{N}^k$ be fixed, let $\mathbf{q} \in \mathcal{Q}_{n, l}^{[g]}$ and $v_* \in V(\mathbf{q})$. The CVS construction being one-to-one, there is a unique labeled map $(\mathbf{m}, \lambda) \in \mathcal{M}_{n, l}^{[g]}$ that corresponds to (\mathbf{q}, v_*) . We denote by \mathbf{s} the scheme of \mathbf{m} and by $(\text{EP}^e, e \in \vec{E}(\mathbf{s}))$ the collection of elementary pieces of (\mathbf{q}, v_*) . We emphasize that the decomposition strongly depends on the distinguished point v_* and not only on \mathbf{q} .

If $e \in \vec{B}(\mathbf{s})$, we let γ^e, ξ^e and β^e be the maximal geodesic, shuttle, and base of the slice EP^e , and we let μ^e, ν^e be the associated area and base measures defined at (2.2). If $e \in \vec{I}(\mathbf{s})$, we let $\gamma^e, \xi^e, \bar{\gamma}^e, \bar{\xi}^e$ be the maximal geodesics and shuttles of EP^e , where the first two correspond to I_e and the latter two to $I_{\bar{e}}$, in the notation of Section 2.2; we also let μ^e be the associated area measure defined at (2.3). Note that $\text{EP}^{\bar{e}}$ yields the same map as EP^e with the same measure; only the marks are ordered differently, namely $\bar{\gamma}^e, \bar{\xi}^e, \gamma^e, \xi^e$.

The construction of \mathbf{q} from (\mathbf{m}, λ) consists in connecting every corner of \mathbf{m} to its successor, and the paths following consecutive successors are geodesic paths all aiming toward v_* . On the other hand, the construction of the elementary pieces from (\mathbf{m}, λ) consists in performing the interval CVS bijection on every interval I_e , in the notation of Section 2.2. The only difference between these constructions lies on the shuttles of these elementary pieces: if c is a corner in some interval I_e whose

successor in the interval bijection belongs to the shuttle ξ^e , then in order to obtain \mathbf{q} we should rather connect c to its successor $s(c)$ in the contour order around f_* . This successor will belong to some interval $I_{e'}$ arriving later in contour order around f_* – note that this interval can be I_e itself. Note also that this successor $s(c)$ belongs to the maximal geodesic $\gamma^{e'}$. Moreover, interpreting \mathbf{q} and its elementary pieces as compact geodesic marked metric spaces, it is straightforward to see that the identifications correspond to metric gluings along geodesics.

Iterative gluing procedure. In order to reconstruct \mathbf{q} from EP^e , $e \in \vec{E}(\mathbf{s})$, rather than connecting the shuttle vertices to their actual successors all at once, we will proceed progressively by first connecting only those whose successors belong to the maximal geodesic of the elementary piece that arrives immediately after in contour order around f_* .

We now formalize this idea. Let κ denote the cardinality of $\vec{E}(\mathbf{s})$. We arrange the half-edges e_1, \dots, e_κ incident to the internal face of \mathbf{s} according to the contour order, starting at an arbitrarily chosen half-edge. While following the contour of the internal face f_* of \mathbf{m} , we successively visit the elementary pieces EP^{e_i} , $1 \leq i \leq \kappa$, which are themselves viewed as marked measured geodesic metric spaces. The reconstruction of \mathbf{q} will be done recursively in κ steps, resulting in a sequence of marked $(k+1)$ -measured metric spaces $\mathbf{q}_0, \dots, \mathbf{q}_\kappa$. At the i -th step, \mathbf{q}_{i+1} will be obtained from \mathbf{q}_i by gluing EP^{e_i} along (part of) its marked maximal geodesic γ^{e_i} . At the same time, we will do some operations on the markings and measures, namely re-orderings, unions of marks and sums of measures, which are all continuous by Lemma 3.4.

We need to keep track of the boundary marks and the geodesics yet to be glued as marks. More precisely, the marking of \mathbf{q}_i is $(\gamma_i^0, \xi_i^0, \gamma_i^1, \xi_i^1, \dots, \gamma_i^{u_i}, \xi_i^{u_i}, \beta_i^1, \dots, \beta_i^k)$, where

- γ_i^j, ξ_i^j , $0 \leq j \leq u_i$, are geodesic marks,
- $\beta_i^1, \dots, \beta_i^k$ are called the *boundary marks*. By convention, certain of these marks may be empty, in which case they are simply discarded from the marks.

The mark ξ_i^0 is the mark along which the subsequent gluing producing \mathbf{q}_{i+1} will occur, and u_i represents the number of quadrilaterals that have been involved only once in the gluing procedures up to the i -th step. Each of these quadrilateral yields two marks γ_i^j, ξ_i^j for some $j \in \{1, \dots, u_i\}$, corresponding to the unvisited half of the quadrilateral, which will have to be glued at a further step. Finally, each \mathbf{q}_i will come with measures μ_i, ν_i , where μ_i is called the *area measure* and $\nu_i = (\nu_i^1, \dots, \nu_i^k)$ is the k -tuple of *boundary measures*.

We initiate the construction by letting $u_0 = 0$, $\mathbf{q}_0 \in \mathbb{M}^{(2, k+1)}$ be the point space with the two marks γ_0^0, ξ_0^0 being the unique point, and measures $\mu_0 = 0$, $\nu_0 = \mathbf{0}^k$. We also let all the boundary marks be empty.

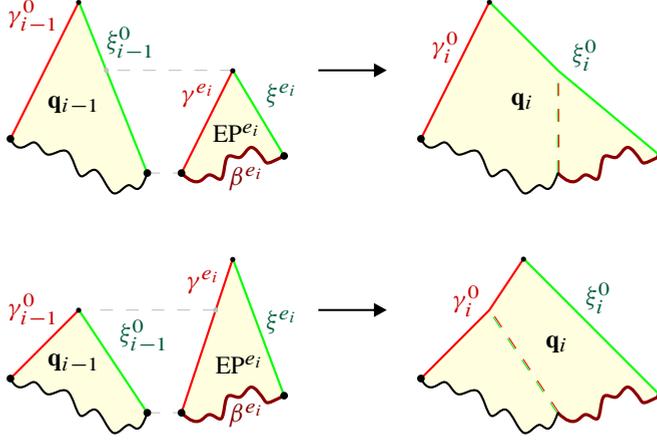


Figure 3.1. The gluing procedure in the case where EP^{e_i} is a slice. In this picture and the following ones, the black wiggly curve depicts all the marks different from $\gamma_{i-1}^0, \xi_{i-1}^0$. The reader should bear in mind that, in general, \mathbf{q}_{i-1} has no reason to present a planar topology as in these pictures. The boundary ξ_{i-1}^0 of \mathbf{q}_{i-1} is glued to the maximal geodesic γ^{e_i} , and the base of EP^{e_i} is added to the r -th boundary mark of \mathbf{q}_{i-1} whenever \bar{e}_i is incident to h_r . The first two geodesic marks are updated according to (3.7), which leads to the two alternative situations described in this figure, depending on which of ξ_{i-1}^0 and γ^{e_i} is the longest: the unglued part of these geodesics becomes part of ξ_i^0 or of γ_i^0 .

Next, provided that \mathbf{q}_{i-1} has been constructed for some $i \in \{1, \dots, \kappa\}$, we define \mathbf{q}_i by considering the following cases, depicted in Figures 3.1–3.3.

- If $e_i \in \bar{B}_r(\mathbf{s})$ for some $r \in \{1, \dots, k\}$, meaning in particular that EP^{e_1} is a slice, we set

$$\mathbf{q}_i = G(\mathbf{q}_{i-1}, EP^{e_i}; \xi_{i-1}^0, \gamma^{e_i}),$$

and mark it as follows. We update the boundary marks by setting

$$\beta_i^r = \beta_{i-1}^r \cup \beta^{e_i}, \quad \beta_i^{r'} = \beta_{i-1}^{r'} \text{ for } r' \in \{1, \dots, k\} \setminus \{r\}.$$

We update the geodesic marks by letting³

$$\gamma_i^0 = \gamma_{i-1}^0 \cup (\gamma^{e_i} \setminus \xi_{i-1}^0), \quad \xi_i^0 = \xi^{e_i} \cup (\xi_{i-1}^0 \setminus \gamma^{e_i}), \quad (3.7)$$

and, setting $u_i = u_{i-1}$, we let $\gamma_i^j = \gamma_{i-1}^j$ and $\xi_i^j = \xi_{i-1}^j$ for $1 \leq j \leq u_i$. Finally, we update the measures by

$$\mu_i = \mu_{i-1} + \mu^{e_i}, \quad \nu_i^r = \nu_{i-1}^r + \nu^{e_i}, \quad \nu_i^{r'} = \nu_{i-1}^{r'} \text{ for } r' \in \{1, \dots, k\} \setminus \{r\}.$$

This case is illustrated in Figure 3.1.

³In (3.7), we use the convention set in Section 3.3 for marks in a glued space: in particular, after gluing, one of the marks $\gamma^{e_i}, \xi_{i-1}^0$ is contained in the other, depending on which is longest.

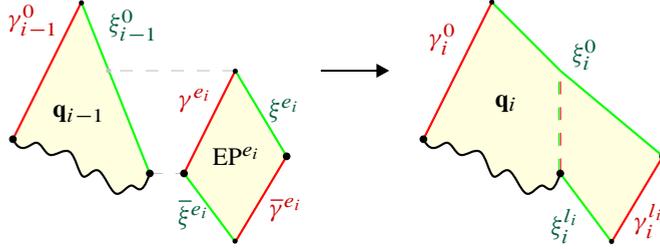


Figure 3.2. The gluing procedure in the case where EP^{e_i} is a quadrilateral that was not involved previously in the construction. The boundary ξ_{i-1}^0 of \mathbf{q}_{i-1} is glued to the maximal geodesic γ^{e_i} . The first two geodesic marks are again updated according to (3.7), leading to two possible situations depending on which of ξ_{i-1}^0 and γ^{e_i} is the longest. Only one of these situations is represented on this figure. In this case, the two geodesic boundary marks $\bar{\gamma}^{e_i}$ and $\bar{\xi}^{e_i}$ are added to the marking; they will be involved in a later construction step.

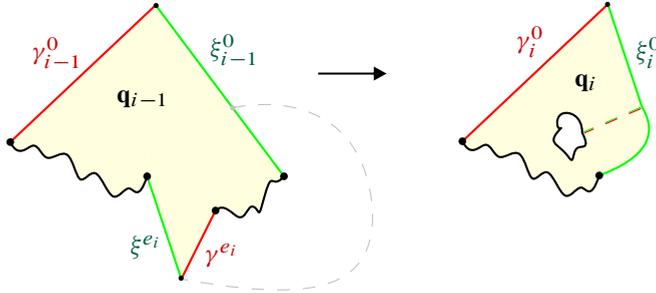


Figure 3.3. The gluing procedure in the case where EP^{e_i} is a quadrilateral, one side of which was already involved in a previous construction step. The boundary ξ_{i-1}^0 of \mathbf{q}_{i-1} is glued to the maximal geodesic γ^{e_i} , which had been introduced as a geodesic mark in this previous construction step. The first two geodesic marks are again updated according to (3.7), and the geodesic marks γ^{e_i} , ξ^{e_i} are removed from the remaining marks.

- If $e_i \in \bar{I}(\mathbf{s})$, meaning in particular that EP^{e_1} is a quadrilateral, we keep the boundary marks unchanged by setting $\beta_i^r = \beta_{i-1}^r$ for $1 \leq r \leq k$, we set $\mathbf{v}_i = \mathbf{v}_{i-1}$, and consider the following two possible situations.
 - If $e_i \notin \{\bar{e}_j, 1 \leq j < i\}$, that is, if the unoriented edge corresponding to e_i is visited for the *first time*, we let again

$$\mathbf{q}_i = G(\mathbf{q}_{i-1}, EP^{e_i}; \xi_{i-1}^0, \gamma^{e_i}),$$

and update its geodesic marks as follows. We update the first two geodesic marks by (3.7). We set $u_i = u_{i-1} + 1$ and let $\gamma_i^j = \gamma_{i-1}^j$ and $\xi_i^j = \xi_{i-1}^j$ for $1 \leq j \leq u_i - 1$. Finally, we set $\gamma_i^{u_i} = \bar{\gamma}^{e_i}$, $\xi_i^{u_i} = \bar{\xi}^{e_i}$, and $\mu_i = \mu_{i-1} + \mu^{e_i}$. This case is illustrated in Figure 3.2.

- If $e_i \in \{\bar{e}_j, 1 \leq j < i\}$, say $e_i = \bar{e}_\ell$, that is, if the unoriented edge corresponding to e_i is visited for the *second time*, then $\gamma^{e_i} = \bar{\gamma}^{e_\ell}$ is a mark of \mathbf{q}_{i-1} : it is the mark $\gamma^{e_i} = \gamma_\ell^{u_\ell}$ of \mathbf{q}_ℓ and stays a mark of the subsequent spaces $\mathbf{q}_{\ell+1}, \dots, \mathbf{q}_{i-1}$. Similarly, $\xi^{e_i} = \bar{\xi}^{e_\ell}$ is a mark of \mathbf{q}_{i-1} . We let

$$\mathbf{q}_i = G(\mathbf{q}_{i-1}; \xi_{i-1}^0, \gamma^{e_i}),$$

we update the first two geodesic marks by (3.7), and, setting $u_i = u_{i-1} - 1$, we let $(\gamma_i^j, \xi_i^j, 1 \leq j \leq u_i)$ be the sequence $(\gamma_{i-1}^j, \xi_{i-1}^j, 1 \leq j \leq u_{i-1})$ from which the terms γ^{e_i} and ξ^{e_i} have been removed. Finally, we set $\mu_i = \mu_{i-1}$. This case is illustrated in Figure 3.3.

It is important to notice that, in \mathbf{q}_i , all the marks $\gamma_i^j, \xi_i^j, 0 \leq j \leq u_i$, are geodesic marks. Indeed, each of these paths always take the form of a chain of consecutive successors, which therefore must be a geodesic; more precisely, these are the maximal geodesics and shuttles of the interval CVS bijection on the intervals $I_{e_1} \cup \dots \cup I_{e_i}$ and $I_{\bar{e}_j}$ for each $j \leq i$ such that $e_j \in \bar{I}(\mathbf{s}) \setminus \{\bar{e}_1, \dots, \bar{e}_i\}$.

At the end of this inductive procedure, we have connected all shuttle corners of some interval I_{e_i} , to their actual successors in \mathbf{m} whenever these lie on some $I_{e_{i'}}$, with $1 \leq i < i' \leq \kappa$. It remains to connect the shuttle corners in some I_{e_i} whose actual successor in \mathbf{m} lies in some $I_{e_{i'}}$, with $1 \leq i' \leq i \leq \kappa$. But one can observe that $u_\kappa = 0$, so that \mathbf{q}_κ carries exactly two geodesic marks $\gamma_\kappa^0, \xi_\kappa^0$. The shuttle corners yet to be connected are exactly those of ξ_κ^0 , and should be matched to the successive corners of γ_κ^0 . Therefore, as marked metric spaces, we have $\mathbf{q} = G(\mathbf{q}_\kappa; \gamma_\kappa^0, \xi_\kappa^0)$, with marks $\beta_\kappa^r, 1 \leq r \leq k$, which are precisely the connected components of the boundary $\partial \mathbf{q}$, ordered as they should.

It is also possible to view all these gluing operations at once, as shown in Figure 3.4.

Measures. We claim that the previous equality $\mathbf{q} = G(\mathbf{q}_\kappa; \gamma_\kappa^0, \xi_\kappa^0)$ only holds as k -marked, $(k + 1)$ -measured metric spaces up to a difference in the supports consisting of a bounded number of vertices. When considering rescaled measures through the operator Ω_n , this small difference will be of no importance in our limiting arguments.

First of all, the boundary measures of the external faces match. This is because the boundary of a given external face of \mathbf{q} is made of the bases, say $\beta^{e_{i_1}}, \dots, \beta^{e_{i_j}}$, of several slices that satisfy $\beta^{e_{i_\ell}} \cap \beta^{e_{i_{\ell+1}}} = \{\bar{\rho}^{e_{i_\ell}}\}$ for $1 \leq \ell \leq j$ with $i_{j+1} = i_1$ and where we denoted by $\bar{\rho}^e$ the final point of the base of EP^e . Since the base measure of the slice EP^e is the counting measure on $\beta^e \setminus \{\bar{\rho}^e\}$, the resulting measure in the gluing is the counting measure on the boundary of the considered external face of \mathbf{q} , as desired.

For the boundary measures of the external vertices, the measure in \mathbf{q} is the counting measure on the singleton consisting of the external vertex, while the corresponding boundary measure in the gluing is the zero measure.

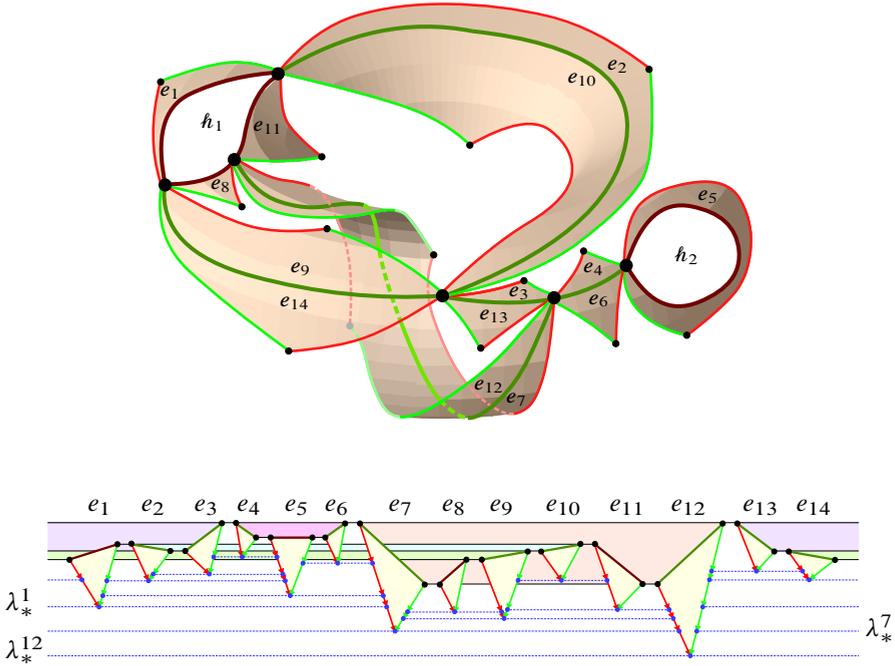


Figure 3.4. Reconstructing q by gluing its elementary pieces along geodesics. On this example, we have $\kappa = 14$. Although we used the same scheme as in Figure 2.2 without h_3 (for lower complexity), beware that the labels here do not match those from the top of Figure 2.2. On the top, the half-edges of $\vec{E}(s)$ are arranged according to the contour order. With each one of them corresponds a “triangle,” which is either a slice (here, with e_1, e_5, e_8, e_{11}) or “half” a quadrilateral, depending on whether the half-edge belongs to $\vec{B}(s)$ or $\vec{I}(s)$. The five matchings of the corresponding “halves” of quadrilaterals, that is, the matchings of half-edges in $\vec{I}(s)$, are represented with light colors. The vertices of these triangles are represented at a height corresponding to their label, where $\lambda_*^i = \min_{I_{e_i}} \lambda - 1$. The geodesic marks of the pieces may be involved in multiple gluings. For instance, the shuttle of the leftmost triangle is involved in three gluings; it is split in three parts, corresponding to the triangles that can be “seen” to its right (those corresponding to e_2, e_5 , and e_7).

For the area measure, by convention, we decided that in the elementary pieces, the measures were taken on vertices *outside of the shuttles*. In doing so, after each gluing operation where a piece of a shuttle is glued to a piece of a maximal geodesic, the corresponding vertices are counted only once, as they should, *except possibly for the vertices $\rho, \bar{\rho}$ of the quadrilaterals*, since they lie both on a maximal geodesic and on a shuttle, and are therefore not part of the counting measures by convention. Therefore, the final gluing is naturally equipped with an area measure that is the counting measure on all but at most $2(2g + k - 1)$ vertices, which is an upper bound on the number of vertices of the scheme by Lemma 2.1.

Conclusion. We finally observe that the number κ of gluing operations necessary to obtain \mathbf{q} is uniformly bounded in n . Indeed, the number of edges of \mathbf{s} is, by Lemma 2.1, smaller than $3(2g + k - 1)$.

As a result, Theorem 1.1 will directly follow from subsequent applications of Proposition 3.8 once we will have shown that, after a proper scaling, the elementary pieces of a uniform quadrangulation jointly converge in distribution in the GHP topology.

3.6 Scaling limit of the collection of elementary pieces

We now fix $L = (L^1, \dots, L^k) \in [0, \infty)^k$. We let b and p be the numbers of indices i such that $L^i > 0$ and $L^i = 0$ respectively. In order to ease notation, we assume that $L^1, \dots, L^b > 0$ while $L^{b+1}, \dots, L^k = 0$.

Limiting measure for size parameters. We denote by $\vec{\mathbf{S}}^*$ the set of rooted genus g schemes with k holes, h_1, \dots, h_b being faces, h_{b+1}, \dots, h_k being vertices of degree 1, and whose internal vertices are all of degree exactly 3. These are called *dominant* schemes. Let $\mathbf{s} \in \vec{\mathbf{S}}^*$ be fixed and denote its root vertex by v_0 . We let $\mathcal{T}_{\mathbf{s}}$ be the set of tuples

$$((\mathbf{a}^e)_{e \in \vec{E}(\mathbf{s})}, (\mathbf{h}^e)_{e \in \vec{I}(\mathbf{s})}, (\mathbf{l}^e)_{e \in \vec{B}(\mathbf{s})}, (\lambda^v)_{v \in V(\mathbf{s})})$$

in $(\mathbb{R}_{\geq 0})^{\vec{E}(\mathbf{s})} \times (\mathbb{R}_{\geq 0})^{\vec{I}(\mathbf{s})} \times \mathbb{R}^{\vec{B}(\mathbf{s})} \times \mathbb{R}^{V(\mathbf{s})}$ such that

- $\sum_{e \in \vec{E}(\mathbf{s})} \mathbf{a}^e = 1$,
- $\mathbf{h}^{\bar{e}} = \mathbf{h}^e$ for all $e \in \vec{I}(\mathbf{s})$,
- $\sum_{e \in \vec{B}_i(\mathbf{s})} \mathbf{l}^e = L^i$ for $1 \leq i \leq b$,
- $\lambda^{v_0} = 0$.

There is a natural Lebesgue measure $\mathcal{L}_{\mathbf{s}}$ on $\mathcal{T}_{\mathbf{s}}$ defined as follows. First, if J is a finite set, and $L > 0$ a positive real number, we let Δ_J^L be the Lebesgue measure on the simplex $\{(x^j, j \in J) \in (\mathbb{R}_{\geq 0})^J : \sum_{j \in J} x^j = L\}$. The latter measure can be defined as the image of the measure $\bigotimes_{j \in J'} dx^j \mathbf{1}_{\{\sum_{j \in J'} x^j < L\}}$, where J' is obtained from J by removing one arbitrary element j' , by the mapping $(x^j, j \in J') \mapsto (x^j, j \in J)$, where $x^{j'} = 1 - \sum_{j \in J'} x^j$.

Next, let $I(\mathbf{s})$ be an orientation of $\vec{I}(\mathbf{s})$, that is, a set containing exactly one element from $\{e, \bar{e}\}$ for every $e \in \vec{I}(\mathbf{s})$. We let $\mathcal{L}_{I(\mathbf{s})}^+$ be the measure $\bigotimes_{e \in I(\mathbf{s})} dh^e \mathbf{1}_{\{h^e \geq 0\}}$. Similarly, we let $\mathcal{L}_{V(\mathbf{s})}$ be the measure $\bigotimes_{v \in V'(\mathbf{s})} d\lambda^v$, where $V'(\mathbf{s}) = V(\mathbf{s}) \setminus \{v_0\}$.

Finally, the measure $\mathcal{L}_{\mathbf{s}}$ is the image measure of

$$\Delta_{\vec{E}(\mathbf{s})}^1 \otimes \mathcal{L}_{I(\mathbf{s})}^+ \otimes \bigotimes_{i=1}^b \Delta_{\vec{B}_i(\mathbf{s})}^{L^i} \otimes \mathcal{L}_{V(\mathbf{s})}$$

by the mapping that associates with

$$((a^e)_{e \in \vec{E}(s)}, (h^e)_{e \in I(s)}, ((l^e)_{e \in \vec{B}_i(s)}, 1 \leq i \leq b), (\lambda^v)_{v \in V'(s)})$$

the unique compatible element of \mathcal{T}_s , that is, such that $h^e = h^{\bar{e}}$ for every $e \in \vec{I}(s)$, and such that $\lambda^{v_0} = 0$.

We let $\text{Param}_{\mathcal{L}}$ be the probability measure on $\bigcup_{s \in \vec{\mathcal{S}}^*} \{s\} \times \mathcal{T}_s$ whose density with respect to the measure $\sum_{s \in \vec{\mathcal{S}}^*} \delta_s \otimes \mathcal{L}_s$ is

$$\frac{1}{\mathcal{Z}_{\mathcal{L}}} \prod_{e \in \vec{I}(s)} q_{h^e}(a^e) \prod_{e \in \vec{B}(s)} q_{l^e}(a^e) \prod_{e \in I(s)} p_{h^e}(\delta \lambda^e) \prod_{e \in \vec{B}(s)} p_{3l^e}(\delta \lambda^e), \quad (3.8)$$

where $p_t(x) = \exp(-x^2/(2t))/\sqrt{2\pi t}$ is the Gaussian density, $q_x(t) = (x/t)p_t(x) \times \mathbf{1}_{\{t>0\}}$ is the (stable 1/2) density for the hitting time of level $-x$ by standard Brownian motion, $\delta \lambda^e = \lambda^{e^+} - \lambda^{e^-}$ for $e \in \vec{E}(s)$, and $\mathcal{Z}_{\mathcal{L}}$ is a normalizing constant, equal to the integral of the remaining display. Beware that the third product is over $I(s)$, not $\vec{I}(s)$.

Scaling limits for size parameters. Next, let $(l_n) = (l_n^1, \dots, l_n^k)$ and Q_n be as in the statement of Theorem 1.1. We let v_n^* be uniformly distributed over the set of internal vertices of Q_n , whose cardinality given by (1.3) only depends on the parameters. Consequently, (Q_n, v_n^*) is uniformly distributed over the set of quadrangulations from $\vec{\mathcal{Q}}_{n, l_n^0}^{[g]}$, seeing v_n^* as a $(k+1)$ -th hole. The rooted labeled map (M_n, λ_n) corresponding via the CVS correspondence is thus uniformly distributed over $\vec{\mathcal{M}}_{n, l_n}^{[g]}$. We denote by S_n the scheme of the nonrooted map corresponding to M_n , and we root S_n uniformly at random among its half-edges, incident to internal or external faces, but such that the corresponding edge does not belong to the boundary of h_{b+1}, \dots, h_k . Note that we could have rooted S_n from the root of M_n by asking that the root of M_n belongs to the forest indexed by the root of S_n but this would have introduced an undesirable bias. Here instead, from the unrooted map corresponding to M_n , the map M_n is rooted at a uniform corner incident to its internal face. Furthermore, the boundaries of the holes h_{b+1}, \dots, h_k are excluded from the possible rootings of S_n since they should be thought of as having null length in the limit.

For $i \in \{b+1, \dots, k\}$, the hole h_i of M_n is called a *vanishing face* if it is a face, that is, if $l_n^i > 0$. The corresponding hole h_i of the scheme S_n is called a *tadpole* if it is made of a single self-loop edge incident to a single vertex of degree 3. We let S_n° be the map S_n in which every tadpole corresponding to a vanishing face h_i has been shrunk into a single vertex, still denoted by h_i , in the sense that the corresponding self-loop has been removed. Note that the root of S_n is never removed in this operation, so that S_n° is always rooted.

Forgetting the root of Q_n , we let $(\text{EP}_n^e, e \in \vec{E}(S_n))$ be the collection of elementary pieces of (Q_n, v_n^*) . For the half-edges $e \in \vec{B}(S_n)$, we let A_n^e, L_n^e be the area and width of the slice EP_n^e . For $e \in \vec{I}(S_n)$, we let A_n^e, H_n^e be the first half-area and width

of the quadrilateral EP_n^e ; note that $A_n^{\bar{e}}$ is the second half-area of EP_n^e . Recall that the vertices of S_n are in one-to-one correspondence with the nodes of M_n : for every $v \in V(S_n)$, we denote by Λ_n^v the label of the corresponding node, where we choose for the labeling function λ_n the representative giving label 0 to the root vertex of S_n .

Proposition 3.10. *As $n \rightarrow \infty$, with probability tending to one, every vanishing face of M_n induces a tadpole in S_n , and, on this likely event, for every $e \in \bigsqcup_{i=b+1}^k \bar{B}_i(S_n)$, it holds that $A_n^e + L_n^e = \Theta((L_n^e)^2)$ in probability.*

Moreover, the following convergence in distribution holds:

$$\begin{aligned} & \left(S_n^\circ, \left(\frac{A_n^e}{n} \right)_{e \in \bar{E}(S_n^\circ)}, \left(\frac{H_n^e}{\sqrt{2n}} \right)_{e \in \bar{I}(S_n^\circ)}, \left(\frac{L_n^e}{\sqrt{2n}} \right)_{e \in \bar{B}(S_n^\circ)}, \left(\left(\frac{9}{8n} \right)^{1/4} \Lambda_n^v \right)_{v \in V(S_n^\circ)} \right) \\ & \xrightarrow[n \rightarrow \infty]{(d)} (S, (A^e)_{e \in \bar{E}(S)}, (H^e)_{e \in \bar{I}(S)}, (L^e)_{e \in \bar{B}(S)}, (\Lambda^v)_{v \in V(S)}), \end{aligned} \quad (3.9)$$

where the limiting random variable has the law $\text{Param}_{\mathbf{L}}$ described in the previous paragraph.

This proposition is a generalization of [22, Proposition 15]. Given its technical nature, we postpone its proof to Appendix B.

Scaling limits of the elementary pieces. As (M_n, λ_n) is uniformly distributed over the set $\vec{\mathbf{M}}_{n, I_n}^{[g]}$, conditionally given (3.9), the random variables EP_n^e , $e \in \vec{E}(S_n)$, are only dependent through the relations linking EP_n^e with $\text{EP}_n^{\bar{e}}$ for $e \in \vec{I}(S_n)$. Moreover,

- if $e \in \vec{B}(S_n)$, then EP_n^e is uniformly distributed among slices with area A_n^e , width L_n^e and tilt $\Lambda_n^{e^+} - \Lambda_n^{e^-}$,
- if $e \in \vec{I}(S_n)$, then EP_n^e is uniformly distributed among quadrilaterals with half-areas A_n^e and $A_n^{\bar{e}}$, width H_n^e and tilt $\Lambda_n^{e^+} - \Lambda_n^{e^-}$.

Applying the Skorokhod representation theorem, we may and will assume that the convergence of (3.9) holds almost surely. Since, by Lemma 2.1, there are finitely many possible schemes, this furthermore implies that $S_n^\circ = S$ for n sufficiently large. Together with Theorems 2.6 and 2.8, the above observations entail that the collection of rescaled random metric spaces $(\Omega_n(\text{EP}_n^e), e \in \vec{E}(S_n^\circ))$, converge in distribution in the GHP topology toward a family $(\text{EP}^e, e \in \vec{E}(S))$ of continuum elementary pieces with the following law conditionally given the right-hand side of (3.9):

- if $e \in \vec{B}(S)$, then EP^e is a continuum slice with area A^e , width H^e and tilt $\Lambda^{e^+} - \Lambda^{e^-}$,
- if $e \in \vec{I}(S)$, then EP^e is a continuum quadrilateral with half-areas A^e and $A^{\bar{e}}$, width H^e and tilt $\Lambda^{e^+} - \Lambda^{e^-}$.

We write $\gamma(\text{EP}^e)$, $\xi(\text{EP}^e)$, $\mu(\text{EP}^e)$, and either $\beta(\text{EP}^e)$, $\nu(\text{EP}^e)$ or $\bar{\gamma}(\text{EP}^e)$, $\bar{\xi}(\text{EP}^e)$ the marks and measures of EP^e , with an obvious choice of notation. As continuum

elementary pieces have only been defined as limits in the GHP topology of discrete elementary pieces so far, these marks and measures are the limits of the corresponding marks and measures of the discrete pieces.

We treat the elementary pieces corresponding to the vanishing faces thanks to Corollary 2.7. We define, for every hole \hat{h}_i of S_n with $b + 1 \leq i \leq k$, the elementary piece $\text{EP}_n^{\hat{h}_i}$ as EP_n^e if $\vec{B}_i(S_n) = \{e\}$ or the vertex map otherwise, marked five times at its unique vertex and endowed twice with the zero measure (thinking of it as an “empty slice”). By Proposition 3.10 and Corollary 2.7, with probability tending to 1, each one of the rescaled random metric spaces $\Omega_n(\text{EP}_n^{\hat{h}_i})$, $b + 1 \leq i \leq k$, converges in distribution in the GHP topology toward the point space (note that the tilt of $\text{EP}_n^{\hat{h}_i}$ is always equal to 0).

3.7 Gluing pieces together

We can now complete the proof of Theorem 1.1 by applying Proposition 3.8 at every step of the inductive construction of Section 3.5.

We work on the event of asymptotic full probability of Proposition 3.10 and assume that n is large enough so that $S_n^\circ = S$. We denote by $\kappa = |\vec{E}(S)| + p$ and let e_1, \dots, e_κ be the sequence made of the half-edges of $\vec{E}(S)$, as well as the external vertices of S , listed in contour order. Since S is dominant, these external vertices are $\hat{h}_{b+1}, \dots, \hat{h}_k$. Applying the construction of Section 3.5 to the random quadrangulation Q_n , up to adding the gluing of the point space for each external vertex (not changing the markings and measures), we obtain a sequence $Q_{n,1}, \dots, Q_{n,\kappa}$ of marked measured metric spaces.

The limiting space $S_L^{[g]}$ is obtained from $(S, (\text{EP}^e, e \in \vec{E}(S)))$ by recursively defining a sequence of marked measured metric spaces S_0, \dots, S_κ , in the following way. For $0 \leq i \leq \kappa$, the marked measured metric space S_i will carry geodesic marks γ_i^j, ξ_i^j , $0 \leq j \leq u_i$, boundary marks $\beta_i^1, \dots, \beta_i^k$, an area measure μ_i , and boundary measures ν_i^1, \dots, ν_i^k .

We initiate the construction by letting $u_0 = 0$, $S_0 \in \mathbb{M}^{(2,k+1)}$ be the point space with the two marks γ_0^0, ξ_0^0 being the unique point, and measures $\mu_0 = 0$, $\nu_0 = \mathbf{0}^k$. We also let all the boundary marks be empty.

Next, given S_{i-1} for some $i \in \{1, \dots, \kappa\}$, consider the following cases:

- If $e_i \in \vec{B}_r(S)$ for some $r \in \{1, \dots, k\}$, set

$$\begin{aligned} S_i &= G(S_{i-1}, \text{EP}^{e_i}; \xi_{i-1}^0, \gamma(\text{EP}^{e_i})), \\ \beta_i^r &= \beta_{i-1}^r \cup \beta(\text{EP}^{e_i}), & \beta_i^{r'} &= \beta_{i-1}^{r'} \text{ for } r' \in \{1, \dots, k\} \setminus \{r\}, \\ \gamma_i^0 &= \gamma_{i-1}^0 \cup (\gamma(\text{EP}^{e_i}) \setminus \xi_{i-1}^0), & \xi_i^0 &= \xi(\text{EP}^{e_i}) \cup (\xi_{i-1}^0 \setminus \gamma(\text{EP}^{e_i})), \end{aligned} \quad (3.10)$$

and, setting $u_i = u_{i-1}$, let $\gamma_i^j = \gamma_{i-1}^j$ and $\xi_i^j = \xi_{i-1}^j$ for $1 \leq j \leq u_i$. Finally, we let $\mu_i = \mu_{i-1} + \mu(\text{EP}^{e_i})$, $\nu_i^r = \nu_{i-1}^r + \nu(\text{EP}^{e_i})$ and $\nu_i^{r'} = \nu_{i-1}^{r'}$ for $r' \neq r$.

- If $e_i \in \vec{I}(S)$, set $\beta_i^r = \beta_{i-1}^r$, $\nu_i^r = \nu_{i-1}^r$ for $1 \leq r \leq k$, and consider the following two possible situations:

- If $e_i \notin \{\bar{e}_j, 1 \leq j < i\}$, let

$$S_i = G(S_{i-1}, \text{EP}^{e_i}; \xi_{i-1}^0, \gamma(\text{EP}^{e_i})),$$

update the first two geodesic marks by (3.10), and, setting $u_i = u_{i-1} + 1$, let $\gamma_i^j = \gamma_{i-1}^j$ and $\xi_i^j = \xi_{i-1}^j$ for $1 \leq j \leq u_i - 1$, and $\gamma_i^{u_i} = \bar{\gamma}(\text{EP}^{e_i})$, $\xi_i^{u_i} = \bar{\xi}(\text{EP}^{e_i})$. Finally, set $\mu_i = \mu_{i-1} + \mu(\text{EP}^{e_i})$.

- If $e_i \in \{\bar{e}_j, 1 \leq j < i\}$, let

$$S_i = G(S_{i-1}; \xi_{i-1}^0, \gamma(\text{EP}^{e_i})),$$

update the first two geodesic marks by (3.10), and, setting $u_i = u_{i-1} - 1$, let $(\gamma_i^j, \xi_i^j, 1 \leq j \leq u_i)$ be the sequence $(\gamma_{i-1}^j, \xi_{i-1}^j, 1 \leq j \leq u_{i-1})$ from which the terms $\gamma(\text{EP}^{e_i})$ and $\xi(\text{EP}^{e_i})$ have been removed. Finally, set $\mu_i = \mu_{i-1}$.

- If e_i is an external vertex of S , set $S_i = S_{i-1}$.

Finally, we let $S_L^{[g]} = G(S_\kappa; \xi_\kappa^0, \gamma_\kappa^0)$, seen as an element of $\mathbb{M}^{(k, k+1)}$, equipped with the marking $(\beta_\kappa^1, \dots, \beta_\kappa^k)$ and measures μ_κ, ν_κ . An application of Proposition 3.8 and of Lemma 3.4 at every step of the construction shows that, for every $i \in \{1, \dots, \kappa\}$, the rescaled marked measured metric space $\Omega_n(Q_{n,i})$ converges to S_i in the marked GHP topology, and finally $\Omega_n(Q_n)$ converges to $S_L^{[g]}$ by a final application of Proposition 3.8 and the observation regarding the measure supports in the final gluing mentioned at the end of Section 3.5. This completes the proof of Theorem 1.1.

3.8 Topology and Hausdorff dimension

In this section, we derive from Proposition 1.2 in the case $(g, k) \in \{(0, 0), (0, 1)\}$ an alternate proof of Proposition 1.2 in the other cases. In the spherical case, Proposition 1.2 was obtained by Le Gall and Paulin [73] thanks to a theorem of Moore by seeing the Brownian sphere as a rather wild quotient of the sphere by some equivalence relation. The same result was later obtained in [78] through the theory of regularity of sequences developed by Begle and studied by Whyburn. The latter approach was generalized in [20–22] in order to obtain the general cases.

Our approach of decomposition into elementary pieces gives a rather direct and transparent proof of Proposition 1.2 in the case $(g, k) \notin \{(0, 0), (0, 1)\}$ provided the following lemma, which will be obtained in Chapters 4 and 5, and which amounts to Proposition 1.2 for the noncompact analogs of the cases $(g, k) \in \{(0, 0), (0, 1)\}$.

Lemma 3.11. *Almost surely, a slice or a quadrilateral is homeomorphic to a disk and is locally of Hausdorff dimension 4. Its boundary as a topological manifold consists in the union of its marks, the intersection of any two marks being empty or a singleton. Furthermore, in the case of a slice, the base is locally of Hausdorff dimension 2.*

Now, a quadrangulation from $\mathbf{Q}_{n, \mathbf{I}_n}^{[g]}$ is “not far” from being homeomorphic to $\Sigma_{b_n}^{[g]}$, where b_n is the number of external faces, in the sense that we can “fill in” the internal faces with small topological disks and “fill in” the external faces by thin topological annuli without altering the metric and in such a way that the resulting object is homeomorphic to $\Sigma_{b_n}^{[g]}$ and at bounded GH distance from the quadrangulation (see [22, Section 4.3.3] for more details about this procedure). The decomposition of the quadrangulation into elementary pieces gives a decomposition of this surface into pieces that are no other than the elementary pieces of the quadrangulation with faces filled in and a thin rectangle added on the bases of the slices; see Figure 3.5.

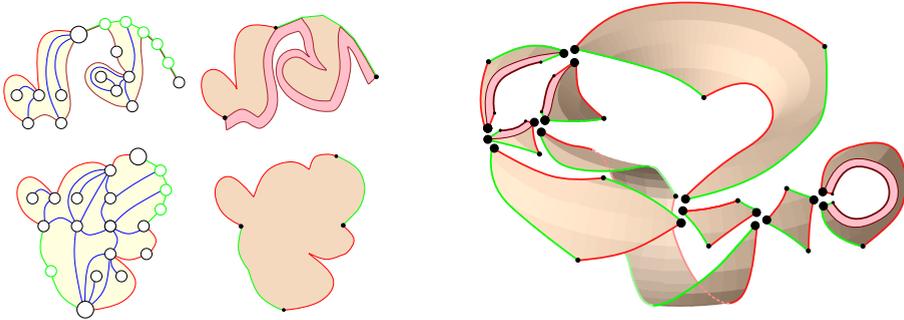


Figure 3.5. *Left.* Topological disk corresponding to an elementary piece of a quadrangulation. *Right.* Decomposition of the surface associated with a quadrangulation into surfaces homeomorphic to disks. Here also, we used the same scheme as in Figure 2.2 without h_3 .

With the notation of the previous section, Q_n yields a surface homeomorphic to $\Sigma_{b_n}^{[g]}$ and its elementary pieces EP_n^e , $e \in \vec{E}(S_n)$, yield surfaces homeomorphic to disks.

- If $e \in \vec{I}(S_n)$, then the boundary of the surface associated with EP_n^e consists in the two maximal geodesics and the two shuttles of EP_n^e .
- If $e \in \vec{B}(S_n)$, then the boundary of the surface associated with EP_n^e consists in the maximal geodesic, the shuttle, as well as three sides of the added thin rectangle.

Then gluing back these disks along their boundaries in the same way as they were cut gives back the surface $\Sigma_{b_n}^{[g]}$. Forgetting the slice corresponding to a tadpole topologically amounts to fill in the corresponding vanishing face. Assuming that n is sufficiently large, all the vanishing faces correspond to tadpoles in the scheme. So the gluing of the elementary pieces without the slices corresponding to tadpoles yields $\Sigma_b^{[g]}$.

In the limit, we glue topological disk exactly in the same way (using markings that are topologically equivalent), so we obtain the same surface. The result about the topology follows.

The statement about the Hausdorff dimension is even more straightforward as we glue along geodesics a finite number of objects that are locally of dimension 4 and the boundary is the union of the bases of the slices, which are all locally of Hausdorff dimension 2.

Chapter 4

Convergence of composite slices

The goal of this chapter is to prove Theorem 2.6 on the convergence of slices to their limiting slices. To this end, we are first going to derive a “free” version of this result by finding slices with a free area and tilt within the uniform infinite half-planar quadrangulation. The latter is known to converge to the Brownian half-plane, which itself contains a “flow” of continuum slices with free areas and tilts; these are shown to be the scaling limits of the discrete slices. We conclude by a conditioning argument to pass from free to fixed area and tilt. First, let us start with deterministic considerations.

4.1 Metric spaces coded by real functions

Here we borrow some material from [13, Section 2.1], with however several slight differences, in order to describe in a unified fashion the various random metric spaces we will use. Let \mathcal{C} (resp. $\mathcal{C}^{(2)}$) be the set of continuous functions of one variable (resp. of two variables) defined on some nonempty closed interval:

$$\mathcal{C} = \bigsqcup_{\substack{I \text{ closed interval} \\ I \neq \emptyset}} \mathcal{C}(I, \mathbb{R})$$

and

$$\mathcal{C}^{(2)} = \bigsqcup_{\substack{I \text{ closed interval} \\ I \neq \emptyset}} \mathcal{C}(I^2, \mathbb{R}).$$

For a function $f \in \mathcal{C}(I, \mathbb{R})$, we denote by $I(f) = I$ its interval of definition and by $\bar{\tau}(f) = \inf I$ and $\tau(f) = \sup I$ its extremities. The set \mathcal{C} is naturally equipped with the topology of uniform convergence over compact subsets of \mathbb{R} ; more precisely, the topology induced by the following metric:

$$\begin{aligned} \text{dist}_{\mathcal{C}}(f, g) &= |\arctan(\bar{\tau}(f)) - \arctan(\bar{\tau}(g))| + |\arctan(\tau(f)) - \arctan(\tau(g))| \\ &\quad + \sum_{n \geq 1} \frac{1}{2^n} \sup_{t \in [-n, n]} |f(\bar{\tau}(f) \vee t \wedge \tau(f)) - g(\bar{\tau}(g) \vee t \wedge \tau(g))|. \end{aligned}$$

We also equip $\mathcal{C}^{(2)}$ with a straightforward adaptation $\text{dist}_{\mathcal{C}^{(2)}}$.

4.1.1 \mathbb{R} -trees coded by functions

\mathbb{R} -trees. For $f \in \mathcal{C}$ and $s, t \in I = I(f)$ with $s \leq t$, set

$$\underline{f}(s, t) = \inf_{[s, t]} f \quad (4.1)$$

and, for $s, t \in I$, set

$$d_f(s, t) = f(s) + f(t) - 2\underline{f}(s \wedge t, s \vee t). \quad (4.2)$$

This formula defines a pseudometric on I , which is continuous as a function from I^2 to $\mathbb{R}_{\geq 0}$, since $d_f(s, t) \leq 2\omega(f; [s \wedge t, s \vee t])$, where $\omega(f; J) = \sup_J f - \inf_J f$. We let $\mathcal{T}_f = (I/\{d_f = 0\}, d_f)$ be the associated quotient space, and $\mathbf{p}_f: I \rightarrow \mathcal{T}_f$ be the canonical projection, which is continuous since d_f is. The space \mathcal{T}_f is a so-called \mathbb{R} -tree, that is, satisfies the following:

- For every two points $a, b \in \mathcal{T}_f$, there exists a geodesic from a to b , that is, an isometric mapping $\chi_{a,b}: [0, d_f(a, b)] \rightarrow \mathcal{T}_f$ with $\chi_{a,b}(0) = a$ and $\chi_{a,b}(d_f(a, b)) = b$.
- The image of the path $\chi_{a,b}$, which we denote by $\llbracket a, b \rrbracket_f$, is the image of any injective path from a to b .

If I is compact, we let $a_*(f) = \mathbf{p}_f(t_*)$, where t_* is any point at which f attains its overall minimum. In this case, for $t \in I$ and $a = \mathbf{p}_f(t)$, the geodesic segment $\llbracket a, a_*(f) \rrbracket_f$ is given by

$$\llbracket a, a_*(f) \rrbracket_f = \mathbf{p}_f(\{s \in [t \wedge t_*, t \vee t_*] : f(s) \leq f(u), \forall u \in [s \wedge t, s \vee t]\}).$$

In the case where I is unbounded, we will systematically make the extra assumption that

$$\begin{cases} \text{when } \bar{\tau}(f) = -\infty, & \inf_{t \leq 0} f(t) = -\infty \text{ or } \lim_{t \rightarrow -\infty} f(t) = \infty, \\ \text{when } \tau(f) = \infty, & \inf_{t \geq 0} f(t) = -\infty \text{ or } \lim_{t \rightarrow \infty} f(t) = \infty. \end{cases} \quad (4.3)$$

In particular, it holds that

$$\forall s \in I, \quad \lim_{|t| \rightarrow \infty, t \in I} d_f(s, t) = \infty, \quad (4.4)$$

which implies that \mathcal{T}_f is locally compact, as the reader may easily check.

Gluing two \mathbb{R} -trees. Next, given two functions $f, g \in \mathcal{C}$ with common interval of definition

$$I = I(f) = I(g)$$

both satisfying (4.3), we define another pseudometric on I as the quotient pseudometric (defined by (3.2))

$$D_{f,g}(s, t) = d_g/\{d_f = 0\}, \quad (4.5)$$

and equip the quotient set $M_{f,g} = I/\{D_{f,g} = 0\}$ with the metric $D_{f,g}$. Note that $D_{f,g}: I^2 \rightarrow \mathbb{R}_{\geq 0}$ is continuous since

$$\begin{aligned} |D_{f,g}(s, t) - D_{f,g}(s', t')| &\leq D_{f,g}(s, s') + D_{f,g}(t, t') \\ &\leq 2\omega(g; [s \wedge s', s \vee s']) + 2\omega(g; [t \wedge t', t \vee t']). \end{aligned}$$

For this reason, the canonical projection $\mathbf{p}_{f,g}: I \rightarrow M_{f,g}$ is continuous. We may view $(M_{f,g}, D_{f,g})$ as gluing the \mathbb{R} -tree \mathcal{T}_g along the equivalence relation defining the \mathbb{R} -tree \mathcal{T}_f . In fact, since either of $d_f(s, t) = 0$ or $d_g(s, t) = 0$ implies $D_{f,g} = 0$, the canonical projection $\mathbf{p}_{f,g}$ factorizes as

$$\mathbf{p}_{f,g} = \pi_f \circ \mathbf{p}_f = \pi_g \circ \mathbf{p}_g,$$

where $\pi_f: \mathcal{T}_f \rightarrow M_{f,g}$ and $\pi_g: \mathcal{T}_g \rightarrow M_{f,g}$ are two surjective maps. Note that these functions are continuous: if $a_n = \mathbf{p}_f(t_n)$ converges to some point a , then, up to taking extractions (and using (4.4) if I is unbounded), we may assume that t_n converges to some limit t , and then $\mathbf{p}_f(t) = a$ by continuity of \mathbf{p}_f , while $\pi_f(a_n) = \mathbf{p}_{f,g}(t_n)$ converges to $\mathbf{p}_{f,g}(t) = \pi_f(a)$. As a consequence, every geodesic segment $\llbracket a, b \rrbracket_f$ in \mathcal{T}_f , and every geodesic $\llbracket c, d \rrbracket_g$ in \mathcal{T}_g is “immersed” into $M_{f,g}$ via the mappings π_f, π_g .

4.1.2 Composite slices coded by two functions

Slice trajectory. We say that (f, g) is a *slice trajectory* if $f, g \in \mathcal{C}$ have common interval of definition I ,

$$\forall s, t \in I, \quad d_f(s, t) = 0 \Rightarrow g(s) = g(t), \quad (4.6)$$

if $\inf I = -\infty$, then

$$\lim_{t \rightarrow -\infty} f(t) = +\infty \quad \text{and} \quad \inf_{t \leq 0} g(t) = -\infty,$$

and, if $\sup I = \infty$, then

$$\inf_{t \geq 0} f(t) = -\infty \quad \text{and} \quad \inf_{t \geq 0} g(t) = -\infty.$$

In particular, f and g both satisfy (4.3) in the case where I is noncompact, and the quantity $f(\inf I) \in \mathbb{R} \cup \{+\infty\}$ is always well defined.

In the remainder of this section, we fix a slice trajectory (f, g) , and call the metric space

$$\text{Sl}_{f,g} = (M_{f,g}, D_{f,g})$$

the *slice coded by* (f, g) . For the moment, we focus on deterministic considerations; the functions f, g will be randomized in the following section.

Marks and measures. The slice $\text{Sl}_{f,g}$ naturally comes with the following distinguished elements.

Geodesics sides. For every $t \in I$, we set

$$\begin{aligned}\Gamma_t(r) &= \inf\{s \geq t : g(s) = g(t) - r\} \quad \text{for } r \in \mathbb{R}_{\geq 0} \text{ such that } \inf_{\substack{s \geq t \\ s \in I}} g(s) \leq g(t) - r, \\ \Xi_t(r) &= \sup\{s \leq t : g(s) = g(t) - r\} \quad \text{for } r \in \mathbb{R}_{\geq 0} \text{ such that } \inf_{\substack{s \leq t \\ s \in I}} g(s) \leq g(t) - r.\end{aligned}$$

In particular, we have $d_g(\Gamma_t(r), \Xi_t(r)) = 0$ for every $t \in I$ and every r satisfying both inequalities above.

We extend the definition given in Section 3.2 of geodesics to paths $\chi: [0, \infty) \rightarrow \mathcal{X}$ that satisfy (3.1) for every $s, t \in \mathbb{R}_{\geq 0}$. In this case, the point $\chi(0)$ is called the *origin* of χ , its length is set to $\text{length}(\chi) = \infty$ by convention, and the range of χ is called a *geodesic ray*. The geodesic ray uniquely determines the geodesic χ by the same argument as for finite length, since the origin of a geodesic ray is the unique point a such that, for any $s > 0$, the number of points in the ray at distance s from a is one.

We observe that Γ_t and Ξ_t are geodesics (possibly of infinite length) from t for the pseudometrics d_g and $D_{f,g}$, in the sense that, for every r, r' such that $\Gamma_t(r)$ and $\Gamma_t(r')$ are defined,

$$d_g(\Gamma_t(r), \Gamma_t(r')) = D_{f,g}(\Gamma_t(r), \Gamma_t(r')) = |r' - r|, \quad (4.7)$$

and the same holds with Ξ_t in place of Γ_t . This fact is immediate for d_g by definition. In fact, when I is compact, one checks that the images of $\mathbf{p}_g \circ \Gamma_t$ and $\mathbf{p}_g \circ \Xi_t$ are the geodesic segments $\llbracket \mathbf{p}_g(t), a_*(g|_{[t, \sup I]}) \rrbracket_g$ and $\llbracket \mathbf{p}_g(t), a_*(g|_{[\inf I, t]}) \rrbracket_g$ in \mathcal{T}_g .

For $D_{f,g}$, this fact follows from the bound

$$|g(s) - g(s')| \leq D_{f,g}(s, s') \leq d_g(s, s'), \quad s, s' \in I,$$

where the first inequality is an easy consequence of the fact that (f, g) is a slice trajectory.

Therefore, for $t \in I$, the paths defined by

$$\begin{aligned}\gamma_t(r) &= \mathbf{p}_{f,g}(\Gamma_t(r)), \quad 0 \leq r \leq g(t) - \underline{g}(t, \sup I), \quad r \in \mathbb{R}, \\ \xi_t(r) &= \mathbf{p}_{f,g}(\Xi_t(r)), \quad 0 \leq r \leq g(t) - \underline{g}(\inf I, t), \quad r \in \mathbb{R},\end{aligned}$$

are two geodesics (possibly of infinite length) from $\mathbf{p}_{f,g}(t)$, sharing a common initial part. As mentioned in Section 3.2, we will often identify these paths with the pairs formed by their origins and image sets, the latter being $\pi_g(\llbracket \mathbf{p}_g(t), a_*(g|_{[t, \sup I]}) \rrbracket_g)$ and $\pi_g(\llbracket \mathbf{p}_g(t), a_*(g|_{[\inf I, t]}) \rrbracket_g)$ when I is compact.

The slice $M_{f,g}$ comes with zero, one, or two geodesic sides. If $\inf I > -\infty$, then the geodesic $\gamma = \gamma_{\inf I}$ is called the *maximal geodesic* of $M_{f,g}$, and, if $\sup I < \infty$, the

geodesic $\xi = \xi_{\sup I}$ is called the *shuttle* of $M_{f,g}$. If $\inf I = -\infty$ (resp. $\sup I = \infty$), we let $\gamma_{-\infty}$ (resp. ξ_{∞}) be the empty set. If I is a bounded interval, then the paths $\gamma_{\inf I}$ and $\xi_{\sup I}$ have a common endpoint at the *apex*

$$x_* = \mathbf{p}_{f,g}(s_*) = \pi_g(a_*(g)),$$

where s_* denotes any point s in I such that $g(s) = \inf_I g$.

Base. For $x \in \mathbb{R}$, we define

$$T_x = \inf\{t \in I : f(t) = -x\} \in \mathbb{R} \cup \{\infty\},$$

the hitting time of level $-x$ by the function f , with the convention that $\inf \emptyset = \infty$. Note that, for $x \in \mathbb{R}$, $T_x \neq -\infty$ because of the fact that (f, g) is admissible. By convention, we also set $T_\infty = -T_{-\infty} = \infty$. The *base* of $\text{Sl}_{f,g}$ is the set

$$\beta = \mathbf{p}_{f,g}(\{T_x : -f(\inf I) \leq x \leq -\inf_I f\} \cap \mathbb{R}).$$

Note that the set inside brackets projects via \mathbf{p}_f to a geodesic in \mathcal{T}_f . When I is compact, the base is the path $\pi_f(\llbracket \mathbf{p}_f(T_{-f(\inf I)}), \mathbf{p}_f(T_{-\inf_I f}) \rrbracket)$, and in general, it is the increasing union of the paths

$$\pi_f(\llbracket \mathbf{p}_f(T_x), \mathbf{p}_f(T_y) \rrbracket_f), \quad -f(\inf I) \leq x < y \leq -\inf_I f, \quad x, y \in \mathbb{R}.$$

Measures. Finally, denoting by Leb_J the Lebesgue measure on the interval J , the slice $\text{Sl}_{f,g}$ is endowed with the following measures:

- the *area measure* $\mu = (\mathbf{p}_{f,g})_* \text{Leb}_I$,
- the *base measure* ν , defined as the pushforward of $\text{Leb}_{[-f(\inf I), -\inf_I f] \cap \mathbb{R}}$ by the mapping $x \mapsto \mathbf{p}_{f,g}(T_x)$.

Gluing slices. In what follows, we will make a slight abuse of notation and identify intervals of the form $[a, \infty]$, $[-\infty, a]$ for $a \in \mathbb{R}$ and $[-\infty, \infty]$ with the intervals $[a, \infty)$, $(-\infty, a]$ and \mathbb{R} , respectively. For L, L' in the extended line $\mathbb{R} \cup \{\pm\infty\}$ such that $-f(\inf I) \leq L \leq L' \leq -\inf_I f$, we define the restrictions $f^{(L,L')}$ and $g^{(L,L')}$ of f and g to the interval $[T_L, T_{L'}] \cap I$, yielding also a slice trajectory. We may therefore define the slice coded by $(f^{(L,L')}, g^{(L,L')})$ and denote it by

$$\text{Sl}^{(L,L')} = (M^{(L,L')}, D^{(L,L')}) = (M_{f^{(L,L')}, g^{(L,L')}}), D_{f^{(L,L')}, g^{(L,L')}}).$$

We let $\mathbf{p}^{(L,L')}: [T_L, T_{L'}] \rightarrow M^{(L,L')}$ be the canonical projection, $\gamma^{(L,L')}$, $\xi^{(L,L')}$, $\beta^{(L,L')}$ be the maximal geodesic, shuttle, and base, and $\mu^{(L,L')}$, $\nu^{(L,L')}$ be the area and base measures of $\text{Sl}^{(L,L')}$.

This family of metric spaces is compatible with the gluing operation in the following sense, illustrated in Figure 4.1.

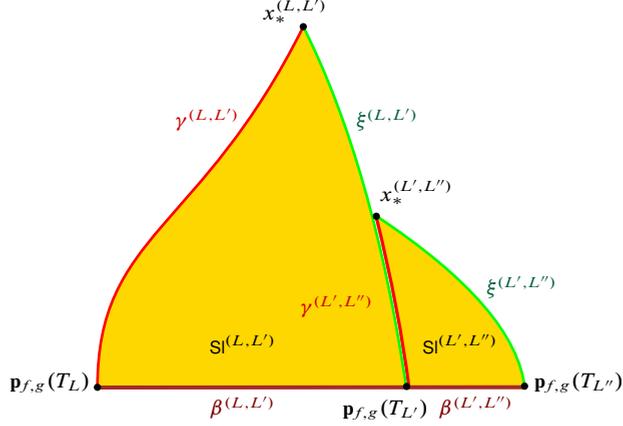


Figure 4.1. Gluing slices encoded by a slice trajectory: the gluing of $\text{SI}^{(L,L')}$ with $\text{SI}^{(L',L'')}$ results in $\text{SI}^{(L,L'')}$. Here, $T_L > -\infty$ and $T_{L''} < \infty$. We denoted by $x_*^{(L',L'')}$ the apex of $\text{SI}^{(L',L'')}$ and $x_*^{(L,L')}$ the apex of $\text{SI}^{(L,L')}$, which, on this example, is also the apex of $\text{SI}^{(L,L'')}$. Consequently, the shuttle $\xi^{(L,L'')}$ is obtained by the union of $\xi^{(L',L'')}$ and the part of $\xi^{(L,L')}$ that is not glued to $\gamma^{(L',L'')}$, whereas the maximal geodesic $\gamma^{(L,L'')} = \gamma^{(L,L')}$, as stated at the end of Proposition 4.1. The bases and measures simply add up. The fact that the slices depicted here are topological disks does not hold true in general; it will, however, be the case for the random processes we will consider in the upcoming sections.

Proposition 4.1. *Let $-f(\inf I) \leq L < L' < L'' \leq -\inf_I f$ be in the extended real line. Then*

$$\text{SI}^{(L,L'')} = G(\text{SI}^{(L,L')}, \text{SI}^{(L',L'')}; \xi^{(L,L')}, \gamma^{(L',L'')}).$$

Moreover, the marks and measures satisfy

$$\begin{aligned} \gamma^{(L,L'')} &= \gamma^{(L,L')} \cup (\gamma^{(L',L'')} \setminus \xi^{(L,L')}), \\ \xi^{(L,L'')} &= \xi^{(L',L'')} \cup (\xi^{(L,L')} \setminus \gamma^{(L',L'')}), \\ \beta^{(L,L'')} &= \beta^{(L,L')} \cup \beta^{(L',L'')}, \\ \mu^{(L,L'')} &= \mu^{(L,L')} + \mu^{(L',L'')}, \\ \nu^{(L,L'')} &= \nu^{(L,L')} + \nu^{(L',L'')}, \end{aligned}$$

with the convention that, in the right-hand side, sets and measures are identified with their images and pushforwards by the canonical projections in $\text{SI}^{(L,L'')}$.

Proof. In the disjoint union $[T_L, T_{L'}] \sqcup [T_{L'}, T_{L''}]$, in order to avoid ambiguities due to the fact that the point $T_{L'}$ belongs to both intervals (thus should be duplicated), we use a superscript 0 for points in the first interval and a superscript 1 for points in the second interval. We observe that $d_{g^{(L,L'')}}$ can be seen as a quotient pseudometric d/R_1 , where d is the disjoint union pseudometric on $[T_L, T_{L'}] \sqcup [T_{L'}, T_{L''}]$

given by $d(s, t) = d_g(s, t)$ if s, t belong to the same of the two intervals above and $d(s, t) = \infty$ otherwise, and R_1 is the coarsest equivalence relation containing

$$\{(\Xi_{T_{L'}}(r)^0, \Gamma_{T_{L'}}(r)^1), 0 \leq r \leq g(T_{L'}) - \underline{g}(T_L, T_{L'}) \vee \underline{g}(T_{L'}, T_{L''})\}.$$

Note also that, as $T_{L'}$ is a hitting time, the equivalence relation $\{d_{f^{(L, L'')}} = 0\}$ factorizes over these two intervals, in the sense that if $d_{f^{(L, L'')}}(s, t) = 0$ with $s \neq t$, then s, t must belong to the same interval $[T_L, T_{L'}]$ or $[T_{L'}, T_{L''}]$. So if R_2 is the equivalence relation on the above disjoint union given by $(s^i, t^j) \in R_2$ if and only if $d_f(s, t) = 0$ and $i = j \in \{0, 1\}$, using (3.3), we have

$$D^{(L, L'')} = (d/R_1)/R_2 = (d/R_2)/R_1 = (D^{(L, L')} \sqcup D^{(L', L'')})/R_1,$$

which is precisely the quotient metric of $G(\text{Sl}^{(L, L')}, \text{Sl}^{(L', L'')}; \xi^{(L, L')}, \gamma^{(L', L'')})$.

Checking the claimed formulas for the marks and measures of $\text{Sl}^{(L, L'')}$ is straightforward. ■

We finish this paragraph with a very strong identity, saying that the distances in a slice $\text{Sl}^{(L, L')}$ encoded by a restriction of the slice trajectory (f, g) are in fact the restrictions of the distances in the “whole” slice $\text{Sl}_{f, g}$.

Corollary 4.2. *Let (f, g) be a slice trajectory on the interval I , and $-f(\inf I) \leq L \leq L' \leq -\inf_I f$. Then $D^{(L, L')}$ is the restriction of the function $D_{f, g}$ to $[T_L, T_{L'}]$.*

Proof. This is a direct consequence of the preceding proposition, which entails that $D_{f, g}$ is the pseudometric obtained by gluing $\text{Sl}^{(L, L')}$ with $\text{Sl}^{(L', \sup I)}$ along $\xi^{(L, L')}$ and $\gamma^{(L', \sup I)}$, and then by gluing the resulting space $\text{Sl}^{(L, \sup I)}$ with $\text{Sl}^{(\inf I, L)}$ along $\gamma^{(L, \sup I)}$ and $\xi^{(\inf I, L)}$. Since at each stage, the spaces that are glued together are isometrically embedded in the resulting gluing, we obtain that $\text{Sl}^{(L, L')}$ is isometrically embedded in $\text{Sl}^{(\inf I, \sup I)} = \text{Sl}_{f, g}$. ■

4.2 Random continuum composite slices

We now randomize the functions f, g considered in the preceding section in various ways to construct random spaces of interest. For a fixed continuous function $f \in \mathcal{C}$ with $0 \in I(f)$, the *snake*¹ driven by f is a random centered Gaussian process $(Z_t^f, t \in I(f))$ with $Z_0^f = 0$ and with covariance function specified by

$$\mathbb{E}[(Z_t^f - Z_s^f)^2] = d_f(s, t), \quad s, t \in I(f). \quad (4.8)$$

¹Literally, this is rather called the “head of the snake driven by f ”; see [63].

As soon as f is Hölder continuous, which will always be the case in this memoir, this process admits a continuous modification; we systematically consider this continuous modification of Z^f . If now Y is a (almost surely Hölder continuous) random function, then the random snake driven by Y is defined as the Gaussian process Z^Y conditionally given Y .

By (4.8), it holds that $Z_s^f = Z_t^f$ whenever $d_f(s, t) = 0$, so that, provided f satisfies the required limit conditions if $I(f)$ is noncompact, the pair (f, Z^f) is a slice trajectory. In what follows, we will let $(X, W): (f, g) \mapsto (f, g)$ be the canonical process on \mathcal{C}^2 .

Below and throughout this work, we use, for any process Y defined on an interval I , the piece of notation $\underline{Y}_t = \inf_{s \leq t, s \in I} Y_s$.

Let us proceed to the definition of continuum slices, which arise in Theorem 2.6. Fix $A, L \in (0, \infty)$ and $\Delta \in \mathbb{R}$. We let $\mathbf{Slice}_{A,L,\Delta}$ be the probability distribution under which

- the process X is a first-passage bridge² of standard Brownian motion from 0 to $-L$ with duration A ,
- conditionally given X , the process W has the same law as $(Z_t + \zeta_{-\underline{X}_t}, 0 \leq t \leq A)$, where Z is the random snake driven by $X - \underline{X}$, and $\zeta/\sqrt{3}$ is a standard Brownian bridge of duration L and terminal value $\Delta/\sqrt{3}$, independent of X and Z .

To be more precise, the process $\zeta = (\zeta_t, 0 \leq t \leq L)$ is Gaussian, with $\mathbb{E}[\zeta_t] = t\Delta/L$ for $t \in [0, L]$ and

$$\text{Cov}(\zeta_s, \zeta_t) = 3 \frac{s(L-t)}{L}, \quad 0 \leq s \leq t \leq L.$$

With this definition, it is simple to see that $\mathbf{Slice}_{A,L,\Delta}$ is indeed carried by slice trajectories defined on the interval $I = [0, A]$.

Definition 4.3. The (composite) slice with area A , width L and tilt Δ , generically denoted by $\text{Sl}_{A,L,\Delta}$, is the 5-marked³ 2-measured metric space $\text{Sl}_{X,W}$ under the law $\mathbf{Slice}_{A,L,\Delta}$, endowed with the marking

$$\partial \text{Sl}_{A,L,\Delta} = (\beta, \gamma_0, \xi_A)$$

comprising its base and two geodesic marks, namely its maximal geodesic and its shuttle, as well as its area and base measures μ, ν .

The piece of notation $\partial \text{Sl}_{A,L,\Delta}$ for the marking comes from the fact that the union of the three marks gives the topological boundary of $\text{Sl}_{A,L,\Delta}$, as stated in Lemma 3.11.

²A first-passage bridge of Brownian motion can be defined from Brownian motion stopped when first hitting $-L$ by absolute continuity; see [25].

³Recall from Section 3.2 that each geodesic mark counts for 2 marks, the first one being its marked extremity.

4.3 The Brownian half-plane, and its embedded slices

There is a natural relation between slices and the Brownian half-plane [13,52], which we now introduce. Let $(B_t, t \geq 0)$, $(B'_t, t \geq 0)$ be two independent standard Brownian motions, and let $(\Pi_t = B'_t - 2 \inf_{\{0 \leq s \leq t\}} B'_s, t \geq 0)$ be the so-called *Pitman transform* of B' , which is a three-dimensional Bessel process. Recall the piece of notation $\underline{X}_t = \inf_{s \leq t} X_s$. We let **Half** be the probability distribution on \mathcal{C}^2 under which

- the process X has same distribution as $(B_t \mathbf{1}_{\{t \geq 0\}} + \Pi_{-t} \mathbf{1}_{\{t < 0\}}, t \in \mathbb{R})$, and
- conditionally given X , the process W has same distribution as $(Z_t + \zeta_{-\underline{X}_t}, t \in \mathbb{R})$, where Z is the random snake driven by $X - \underline{X}$, and $\zeta/\sqrt{3}$ is a two-sided standard Brownian motion,⁴ independent of X and Z .

The measure **Half** is carried by slice trajectories defined on the interval \mathbb{R} . We note that we can make this definition more symmetric using standard excursion theory, in a way similar to the encoding triples of [75]. For this, we let $T_{L-} = \lim_{L' \uparrow L} T_{L'}$ and denote by

$$\begin{aligned} X^{(L)} &= (L + X_{T_{L-}+t}, 0 \leq t \leq T_L - T_{L-}), \\ W^{(L)} &= (W_{T_{L-}+t}, 0 \leq t \leq T_L - T_{L-}) \end{aligned}$$

the excursion of X above its past infimum at level $-L$ and the corresponding piece of W , respectively. Note first that the process $\zeta_L = W_{T_L}$, $L \in \mathbb{R}$, is under **Half** a standard two-sided Brownian motion multiplied by $\sqrt{3}$. Then, conditionally given ζ , the point measure on $\mathbb{R} \times \mathcal{C} \times \mathcal{C}$ given by

$$\mathcal{M}(dL dX dW) = \sum_{L \in \mathbb{R}: T_L \neq T_{L-}} \delta_{(L, X^{(L)}, W^{(L)})}$$

is a Poisson measure with intensity $2dL \mathbb{N}_{\zeta_L}(d(X, W))$, where \mathbb{N}_x is the σ -finite “law” of the lifetime process and head of the Brownian snake (started at x) driven by the Itô measure of positive excursions of Brownian motion. The process (X, W) is then a measurable function of ζ and \mathcal{M} by Itô’s reconstruction theory of Brownian motion from its excursions.

Definition 4.4. The *Brownian half-plane*, which is generically denoted by BHP, is the 1-marked 2-measured metric space $\text{Sl}_{X,W}$ considered under **Half**, endowed with the one mark $\partial\text{BHP} = \beta$, its area measure μ and its base measure ν .

There is only one mark here, the base; there is no maximal geodesic nor shuttle since the interval of definition is \mathbb{R} . The name comes from the fact that BHP is

⁴This means that $(\zeta_x/\sqrt{3}, x \geq 0)$ and $(\zeta_{-x}/\sqrt{3}, x \geq 0)$ are independent (one-dimensional) standard Brownian motions issued from 0.

homeomorphic to the half-plane $\mathbb{R} \times \mathbb{R}_{\geq 0}$, its boundary as a topological manifold being equal to the base; see [13, Corollary 3.8].

In the light of Proposition 4.1, the Brownian half-plane can be seen to have a natural Markov property. First, let $\theta_t: f \mapsto f(t + \cdot) - f(t)$ be the translation operator on \mathcal{C} . We claim that **Half** is invariant under θ_{T_L} , since its action simply consists in translating by L the time in process ζ , and the first coordinate of \mathcal{M} , which leaves their laws invariant. For similar reasons, for every $L \in \mathbb{R}$, the processes $(X^{(0,L)}, W^{(0,L)})$, $(X^{(-\infty,0)}, W^{(-\infty,0)})$ and $(\theta_{T_L} X^{(L,+\infty)}, \theta_{T_L} W^{(L,+\infty)})$ are independent under **Half**, since they are respectively functionals of the independent random elements

$$\begin{aligned} (\zeta_x, 0 \leq x \leq L), & \quad \mathcal{M}((0, L] \times \mathcal{C} \times \mathcal{C}), \\ (\zeta_x, x \leq 0), & \quad \mathcal{M}((-\infty, 0] \times \mathcal{C} \times \mathcal{C}), \\ (\zeta_{L+x} - \zeta_L, x \geq 0), & \quad \mathcal{M}((L, \infty) \times \mathcal{C} \times \mathcal{C}). \end{aligned}$$

Free slices. Note that, under **Half**, the process $X^{(0,L)}$ is simply a standard Brownian motion killed at its first hitting time of $-L$, while the process $(W_{T_x}^{(0,L)}/\sqrt{3}, 0 \leq x \leq L)$ is a standard Brownian motion killed at time L . For this reason, the law of $(X^{(0,L)}, W^{(0,L)})$ under **Half** is the mixture

$$\mathbf{FSlice}_L = \int_0^\infty q_L(A) dA \int_{\mathbb{R}} p_{3L}(\Delta) d\Delta \mathbf{Slice}_{A,L,\Delta}, \quad (4.9)$$

where p_t, q_x are defined after (3.8). In what follows, a random metric space with same law as $\mathbf{Sl}^{(0,L)}$ under \mathbf{FSlice}_L will be referred to as a *free (composite) slice* of width L . This, together with Proposition 4.1, yields the following result.

Proposition 4.5. *Fix $L < L' < L''$ in the extended line. Then, under **Half**, it holds that $\mathbf{Sl}^{(L,L'')} = G(\mathbf{Sl}^{(L,L')}, \mathbf{Sl}^{(L',L'')}; \xi^{(L,L')}, \gamma^{(L',L'')})$, where the spaces $\mathbf{Sl}^{(L,L')}$, $\mathbf{Sl}^{(L',L'')}$ are independent. Moreover, if L and L' are finite, then $\mathbf{Sl}^{(L,L')}$ is a free slice of width $L' - L$.*

Recall that this result is illustrated in Figure 4.1, which can be completed by extending the brown segment into a line, letting the half-plane above be BHP, the line being its base $\beta = \beta^{(-\infty, \infty)}$. This also suggests that $\mathbf{Sl}^{(L,L')}$ is the bounded connected component of the complement of $\gamma^{(L,L')} \cup \xi^{(L,L')}$ in BHP. More precisely, the following holds.

Proposition 4.6. *For every $L < L'$ in \mathbb{R} , almost surely under **Half**, the geodesics $\gamma^{(L,L')}$ and $\xi^{(L,L')}$ meet only at the apex $x_*^{(L,L')}$, and meet the base β only at their respective origins $\mathbf{p}_{X,W}(T_L)$ and $\mathbf{p}_{X,W}(T_{L'})$. Moreover, $\mathbf{Sl}^{(L,L')}$ is the closure of the bounded connected component of the complement of the union of these two paths in BHP. It is therefore homeomorphic to the closed unit disk, with boundary given by the union of the three sets $\beta^{(L,L')}$, $\gamma^{(L,L')}$ and $\xi^{(L,L')}$, which meet only at $\mathbf{p}_{X,W}(T_L)$, $\mathbf{p}_{X,W}(T_{L'})$ and $x_*^{(L,L')}$.*

Proof. This proposition is proved in the same way as [13, Lemma 6.15]. Let us recall briefly the ideas. For any point $t \in \mathbb{R}$, we let $\Sigma_t(r) = \inf\{s \geq t : X_s = X_t - r\}$ for $0 \leq r \leq X_t - \underline{X}_t$, so that the range of $\mathbf{p}_X \circ \Sigma_t$ is the geodesic path $\llbracket \mathbf{p}_X(t), \mathbf{p}_X(T_{-\underline{X}_t}) \rrbracket_X$ in \mathcal{T}_X . Its image by π_X defines a path σ_t starting at $\mathbf{p}_{X,W}(t)$ and ending at the point $\mathbf{p}_{X,W}(T_{-\underline{X}_t})$ of the base. Moreover, almost surely, any path σ_t , $t \in \mathbb{R}$, do not intersect a geodesic γ_s , $s \in \mathbb{R}$, except possibly at the starting point of either $\mathbf{p}_{X,W}(s)$ or $\mathbf{p}_{X,W}(t)$. This implies that any point $\mathbf{p}_{X,W}(t)$ of $\text{Sl}^{(L,L')}$ that is not in the union $\gamma^{(L,L')} \cup \xi^{(L,L')}$ can be linked to the bounded segment $\beta^{(L,L')}$ of the base of BHP by the path σ_t without intersecting $\gamma^{(L,L')} \cup \xi^{(L,L')}$ except perhaps at its endpoint. This latest possibility can be discarded by noting that, with probability 1, we have $T_{-\underline{X}_t} \notin \{T_L, T_{L'}\}$. Similarly, a point $\mathbf{p}_{X,W}(t)$ of BHP outside of $\text{Sl}^{(L,L')}$ is linked to the unbounded set $\beta \setminus \beta^{(L,L')}$ of the base of BHP by the path σ_t , which does not intersect $\gamma^{(L,L')} \cup \xi^{(L,L')}$. This means that $\text{Sl}^{(L,L')}$ is the closure of the bounded connected component of BHP minus $\gamma^{(L,L')} \cup \xi^{(L,L')}$. ■

The above discussion shows that the Brownian half-plane contains a natural “flow” of free slices. We can also link directly the slices of Section 4.2 with the Brownian half-plane via an absolute continuity argument. Recall the definitions of p_t and q_x after (4.9).

Lemma 4.7. *Let us fix $0 < K < L$, as well as $A > 0$ and $\Delta \in \mathbb{R}$. For every non-negative function G that is measurable with respect to the σ -algebra generated by $(X^{(0,K)}, W^{(0,K)})$, we have*

$$\mathbf{Slice}_{A,L,\Delta}[G] = \mathbf{Half}[\varphi_{A,L,\Delta}(T_K, K, W_{T_K}) \cdot G],$$

where

$$\varphi_{A,L,\Delta}(A', L', \Delta') = \frac{q_{L-L'}(A - A')}{q_L(A)} \frac{p_{3(L-L')}(\Delta - \Delta')}{p_{3L}(\Delta)}. \quad (4.10)$$

Proof. This comes from similar statements for Brownian bridges and first-passage bridges; see, for instance, [19, (18) and (19)]. For bounded measurable functions f, g on \mathcal{C} , for $0 < A' < A$ and $0 < K < L$,

$$\begin{aligned} & \mathbf{Slice}_{A,L,\Delta}[f(X|_{[0,A']}) \cdot g(\zeta|_{[0,K]})] \\ &= \mathbf{Half} \left[f(X|_{[0,A']}) \frac{q_{L-X_{A'}}(A - A')}{q_L(A)} \mathbf{1}_{\{\underline{X}_{A'} > -L\}} \cdot g(\zeta|_{[0,K]}) \frac{p_{3(L-K)}(\Delta - \zeta_K)}{p_{3L}(\Delta)} \right]. \end{aligned} \quad (4.11)$$

Here, the factor 3 in the index of the Gaussian density function comes from the fact that $\zeta/\sqrt{3}$ is a bridge of standard Brownian motion. We replace A' by T_K by a standard argument, writing

$$f(X^{(0,K)}) = \lim_{n \rightarrow \infty} \sum_{i \geq 0} \mathbf{1}_{\{(i-1)2^{-n} < T_K \leq i2^{-n}\}} f(X|_{[0,i2^{-n}]})$$

using dominated convergence and applying the above equality (4.11) to $A' = i2^{-n}$, noting that $\mathbf{1}_{\{(i-1)2^{-n} < T_K \leq i2^{-n}\}}$, $f(X|_{[0, i2^{-n}]})$ is a function of $X|_{[0, i2^{-n}]}$.

The result follows by noting that $W^{(0,K)}$ is built in the same way from $X^{(0,K)}$ and $\zeta|_{[0,K]}$ under **Slice** $_{A,L,\Delta}$ as from $X^{(0,K)}$ and $\zeta|_{[0,K]}$ under **Half**. ■

We may now prove the statement about the topology and Hausdorff dimension of a slice.

Proof of Lemma 3.11 for slices. First, almost surely, the Brownian half-plane is homeomorphic to the half-plane [13, Corollary 3.8], is locally of Hausdorff dimension 4 and its boundary is locally of Hausdorff dimension 2. The latter facts are obtained from similar statements for Brownian disks [21] thanks to [13, Theorem 3.7] allowing us to couple arbitrary balls of BHP centered at the root $\mathbf{p}_{X,W}(0)$ with balls of large enough Brownian disks, centered at a point on the boundary.

Hence, under the probability distribution **Half**, for any $L < L'$, by Corollary 4.2, the metric space $\text{Sl}^{(L,L')}$ is almost surely locally of Hausdorff dimension 4 and its base $\beta^{(L,L')}$ is locally of Hausdorff dimension 2. Furthermore, it is homeomorphic to the disk by Proposition 4.6 and its boundary is the union of its three marks $\beta^{(L,L')}$, $\gamma^{(L,L')}$ and $\xi^{(L,L')}$, whose pairwise intersections are identified singletons.

Now, arguing under **Slice** $_{A,L,\Delta}$, we use the fact from Proposition 4.1 that $\text{Sl}^{(0,L)} = G(\text{Sl}^{(0,L/2)}, \text{Sl}^{(L/2,L)}, \xi^{(0,L/2)}, \gamma^{(L/2,L)})$. Lemma 4.7 entails that, almost surely, under this probability distribution, the law of $\text{Sl}^{(0,L/2)}$ is absolutely continuous with respect to that of the same random variable under **Half**, and so is homeomorphic to a disk. Now, we observe that, under **Slice** $_{A,L,\Delta}$, the process $\theta_{T_{L/2}}(X^{(L/2,L)}, W^{(L/2,L)})$ has same distribution as $(X^{(0,L/2)}, W^{(0,L/2)})$, which we leave as an exercise to the reader. Therefore, under this law, $\text{Sl}^{(L/2,L)}$ has same distribution as $\text{Sl}^{(0,L/2)}$ and both are homeomorphic to a disk. We conclude that the same is true for $\text{Sl}^{(0,L)}$ since it is obtained by gluing two topological disks along two segments of their boundaries. The identification of the marks given in Proposition 4.1 easily yields the desired property on the marks of $\text{Sl}^{(0,L)}$. The facts on the local Hausdorff dimension are obtained similarly. ■

4.4 The uniform infinite half-planar quadrangulation

We now define a slight variant of the classical uniform infinite half-planar quadrangulation (UIHPQ) [13, 14, 33, 43], the half-planar version of the uniform infinite random planar quadrangulation, in the following way. Let $F_\infty = (\mathbf{T}^k, k \in \mathbb{Z})$ be a two-sided sequence of independent Bienaymé–Galton–Watson trees with a geometric offspring distribution of parameter $1/2$. Conditionally on F_∞ , we let λ_∞^0 be a uniformly chosen well-labeling condition function, meaning that every tree \mathbf{T}^k is assigned a well-labeling condition function giving label 0 to its root vertex, independently, uniformly

at random. Lastly, and independently of F_∞ and λ_∞^0 , we let $(b_k, k \in \mathbb{Z})$ be a doubly-infinite walk with shifted geometric steps, meaning that $b_0 = 0$ almost surely, and that $b_k - b_{k-1}, k \in \mathbb{Z}$, are independent and identically distributed random variables with $\mathbb{P}(b_1 = r) = 2^{-r-2}$ for every $r \in \{-1, 0, 1, 2, \dots\}$. For a vertex $v \in \mathbf{T}^k$, we let $\lambda_\infty(v) = b_k + \lambda_\infty^0(v)$, and call $(F_\infty, \lambda_\infty)$ the *infinite random well-labeled forest*. We then embed F_∞ in the plane in such a way that all trees are contained in the upper half-plane, and the root ρ^k of \mathbf{T}^k is located at the point $(k, 0) \in \mathbb{R}^2$. We also link consecutive roots ρ^k, ρ^{k+1} by a line segment. We then let $(c_i, i \in \mathbb{Z})$ be the sequence of corners of the upper half-plane part of the resulting map, in contour order from left to right, with origin the first corner c_0 incident to ρ^0 . The UIHPQ is then the infinite map Q_∞ obtained by applying the CVS construction to $(F_\infty, \lambda_\infty)$, that is, by linking every corner to its successor as defined in Section 2.1, and removing all edges of the forest afterward. The root of Q_∞ is defined as the corner preceding the arc from c_0 to its successor. Note that, in this case, there is no need to add an extra vertex with a corner c_∞ .

Remark 4.8. The difference between this definition of the UIHPQ and the one appearing in the mentioned references is a slight rooting bias. Indeed, the simplest way to obtain the usual definition is to consider a two-sided simple random walk $(z_i, i \in \mathbb{Z})$ and construct the sequence $(b_k, k \in \mathbb{Z})$ from it as follows. Let $S^\downarrow = \{i \in \mathbb{Z} : z_{i+1} - z_i = -1\}$ be the set of descending steps of $(z_i, i \in \mathbb{Z})$ and $i_0 = \sup(S^\downarrow \cap \mathbb{Z}_{\leq 0})$ the index of the descending step preceding 0. Then we define the sequence $(b_k, k \in \mathbb{Z})$ by reindexing $(z_i - z_{i_0}, i \in S^\downarrow)$ with \mathbb{Z} in such a way that i_0 corresponds to the index 0. The UIHPQ is then constructed as above with this bridge but rooted at the corner preceding the arc linking $s^{-i_0}(c_0)$ to its successor $s^{-i_0+1}(c_0)$ instead of the convention we presented. Apart from this slight root shift, the resulting law of $(b_k, k \in \mathbb{Z})$ is not exactly that of a doubly-infinite bridge with shifted geometric steps. The first step gets a size-biased distribution $\mathbb{P}(b_1 = r) = (r + 2)2^{-r-3}$, $r \geq -1$, whereas all other steps get the desired shifted geometric distribution. See the discussion in [13, Section 4.5.2] for more information.

The construction we use here has the advantage of making the law of the slices invariant by translation.

Convergence toward the Brownian half-plane. Denoting by v_i the vertex of F_∞ incident to c_i and by $\Upsilon(i) \in \mathbb{Z}$ the index of the tree to which v_i belongs, we define the *contour* and *label processes* on \mathbb{R} by

$$C(i) = d_{\mathbf{T}^{\Upsilon(i)}}(v_i, \rho^{\Upsilon(i)}) - \Upsilon(i) \quad \text{and} \quad \Lambda(i) = \lambda_\infty(v_i), \quad i \in \mathbb{Z},$$

and by linear interpolation between integer values; see Figure 4.2. As is well known, the part of the contour process corresponding to \mathbf{T}^k (counting the edge linking ρ^k to ρ^{k+1}) has the same distribution as a simple random walk started at $-k$ and killed

We then define the renormalized versions of C , Λ , and D_∞ : for every $s, t \in \mathbb{R}$, we set

$$C_{(n)}(s) = \frac{C(2ns)}{\sqrt{2n}}, \quad \Lambda_{(n)}(s) = \frac{\Lambda(2ns)}{(8n/9)^{1/4}}, \quad D_{(n)}(s, t) = \frac{D_\infty(2ns, 2nt)}{(8n/9)^{1/4}}. \quad (4.13)$$

The next result can be seen as a reformulation of [52, Theorem 1.11] or [13, Theorem 3.6], proving the convergence of the UIHPQ to the Brownian half-plane defined in Section 4.3.

Proposition 4.9. *On $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)}$, it holds that*

$$(C_{(n)}, \Lambda_{(n)}, D_{(n)}) \xrightarrow[n \rightarrow \infty]{(d)} (X, W, D_{X,W}), \quad (4.14)$$

where the limiting triple is understood under **Half**.

This statement does not appear in this exact form in the aforementioned references, which do not explicitly focus on the processes $C_{(n)}$, $\Lambda_{(n)}$, X , W . In [13, Remark 6.16], it was however mentioned how to extend the results therein in order to take into account these processes, so we will follow the line of reasoning sketched in that work.

Proof. The proof proceeds via established convergence results for random quadrangulations with one external face to Brownian disks. Fix some number $K > 0$. We will sample a quadrangulation with one external face, whose areas and perimeters are so large that, in a neighborhood of 0 of amplitude K , this rescaled large quadrangulation and its limit, a Brownian disk of large area and perimeter, are indistinguishable from the rescaled UIHPQ and the Brownian half-plane, in a sense to be made precise. In the following, we will use for all the objects related to the quadrangulation with one external face or the limiting Brownian disk a similar notation as for those related to the UIHPQ or the Brownian half-plane, only with a superscript prime symbol \prime .

Fix $L > 0$, which should be thought of as being large. For $n \geq 1$, we sample the aforementioned quadrangulation Q'_n with one external face as follows. First, consider a uniform random element (M'_n, λ'_n) of $\vec{\mathbf{M}}_{a_n, (l_n)}^{[0]}$, where

$$a_n = \lfloor nL \rfloor \quad \text{and} \quad l_n = \lfloor L\sqrt{2n} \rfloor.$$

We can view this as a labeled forest (F'_n, λ'_n) with l_n trees arranged in a circle, and rooted at $\rho^0, \dots, \rho^{l_n-1}$, where ρ^0 is the root of the tree containing the root corner of f_* . We let C'_n, Λ'_n be the contour and label process of this forest, defined as above, starting from the tree rooted at ρ^0 . We let $Q'_n = \text{CVS}(M'_n, \lambda'_n; f_*)$ be the rooted quadrangulation encoded by (M'_n, λ'_n) , and we let $D'_n(i, j) = d_{Q'_n}(v'_i, v'_j)$ for $0 \leq i, j \leq 2a_n + l_n$, where v'_i is the i -th visited vertex of F'_n in contour order, viewed as a vertex of Q'_n . As usual, we extend D'_n into a continuous function on

$[0, 2a_n + l_n]^2$. Finally, we extend the definition of these processes to the interval $[-2a_n - l_n, 2a_n + l_n]$ by the simple translation formulas

$$\begin{aligned} C'_n(t) &= C'_n(t + 2a_n + l_n) + l_n, & \Lambda'_n(t) &= \Lambda'_n(t + 2a_n + l_n), \\ t &\in [-2a_n - l_n, 0], \end{aligned} \quad (4.15)$$

$$\begin{aligned} D'_n(s, t) &= D'_n(s + (2a_n + l_n)\mathbf{1}_{\{s < 0\}}, t + (2a_n + l_n)\mathbf{1}_{\{t < 0\}}) \\ s, t &\in [-2a_n - l_n, 2a_n + l_n]. \end{aligned} \quad (4.16)$$

The idea behind this extension is that we are going to consider these processes in neighborhoods of 0, so that we are really interested in the behavior of these processes when the argument is close from 0 or from $2a_n + l_n$.

Define their rescaled versions: for $s, t \in [-2a_n - l_n, 2a_n + l_n]$,

$$C'_{(n)}(s) = \frac{C'_n(2ns)}{\sqrt{2n}}, \quad \Lambda'_{(n)}(s) = \frac{\Lambda'_n(2ns)}{(8n/9)^{1/4}}, \quad D'_{(n)}(s, t) = \frac{D'_n(2ns, 2nt)}{(8n/9)^{1/4}}. \quad (4.17)$$

Then by [25, (26) and Theorem 20], one has the joint convergence

$$(C'_{(n)}, \Lambda'_{(n)}, D'_{(n)}) \xrightarrow[n \rightarrow \infty]{(d)} (X', W', D') \quad (4.18)$$

in distribution in $\mathcal{C}([0, L]) \times \mathcal{C}([0, L]) \times \mathcal{C}([0, L]^2)$, where (X', W', D') is an explicit limiting process, which is the encoding process of the Brownian disk of area L and width \sqrt{L} . In particular, the process D' is a measurable function of the pair (X', W') . Due to the formulas in (4.15) and (4.16), this implies the convergence of these processes on $\mathcal{C}([-L, L]) \times \mathcal{C}([-L, L]) \times \mathcal{C}([-L, L]^2)$, where (X', W', D') are extended to functions on $[-L, L]$ or $[-L, L]^2$ in a similar way as above. Note that we choose to omit the dependence of (X', W', D') on L for lighter notation, but we will need later to choose L appropriately.

Now recall that $K > 0$ is a fixed number. The first crucial observation is that we may choose L large enough, so that with high probability, the laws of the restrictions of (X', W', D') and $(X, W, D_{X,W})$ to the interval $[-K, K]$ are very close. More precisely, given $\varepsilon \in (0, 1)$, fix $r > 0$ and $A > 0$ such that

$$\mathbb{P}\left(\max_{-K \leq t \leq K} D_{X,W}(0, t) > r\right) < \frac{\varepsilon}{3}, \quad \mathbb{P}(T_{-A} < -K < K < T_A) \geq 1 - \frac{\varepsilon}{3}.$$

Then [13, Proposition 6.6] and its proof (Lemmas 6.7 and 6.8) show that there exists $L_0 > 0$ such that, for $L > L_0$, the two processes (X, W) and (X', W') can be coupled in such a way that on some event \mathcal{F} of probability $\mathbb{P}(\mathcal{F}) \geq 1 - \varepsilon/3$, we have

$$X_t = X'_t, \quad W_t = W'_t, \quad D_{X,W}(s, t) = D'(s, t), \quad (4.19)$$

for every $s, t \in [T_{-A}, T_A]$ such that $\max(D_{X,W}(0, t), D_{X,W}(0, s)) \leq r$. Given our choice of r, A , we see that with probability at least $1 - \varepsilon$, (4.19) holds for every $s, t \in [-K, K]$.

Our second important observation is that, still with K and ε fixed, and possibly up to choosing L even larger than the above, albeit in a way that does not depend on n , the laws of $(C_{(n)}, \Lambda_{(n)}, D_{(n)})$ and $(C'_{(n)}, \Lambda'_{(n)}, D'_{(n)})$ in restriction to the interval $[-K, K]$ are also close, in the sense that they can be coupled in such a way that these restrictions coincide with probability at least $1 - \varepsilon$. This follows from the proof of [13, Theorem 3.6], a minor difference being that this proposition establishes that the balls of radius $(8n/9)^{1/4}r$ centered at the root in Q_∞ and Q'_n are isometric, rather than giving a statement on $D_{(n)}$ and $D'_{(n)}$. Therefore, in order to show that the latter coincide on $[-K, K]$, one again has to choose in the first place a radius $r > 0$ so that uniformly over n , with probability at least $1 - \varepsilon/3$, the vertices v'_i for integers i lying in $[-2Kn, 2Kn]$ (where we naturally let $v'_i = v'_{i+2a_n+l_n}$ for $i \leq 0$), all belong to this ball. The existence of such an r is guaranteed by the convergence (4.18) and the continuity of D' . Finally, we see that both sides of (4.18) can be coupled in such a way that with probability at least $1 - \varepsilon$, they coincide with both sides of (4.14). Since ε was arbitrary, we conclude that (4.14) holds in restriction to $[-K, K]$. Since K was arbitrary, this concludes the proof. ■

Seeing a slice as part of the UIHPQ. We consider a fixed $L > 0$ and a sequence $(l_n) \in \mathbb{N}^{\mathbb{N}}$ such that

$$\frac{l_n}{\sqrt{2n}} \xrightarrow{n \rightarrow \infty} L$$

and, for each n , we let (F_n, λ_n) be the random well-labeled forest obtained by keeping only the labeled trees $\mathbf{T}^0, \dots, \mathbf{T}^{l_n-1}$ of the infinite random well-labeled forest $(F_\infty, \lambda_\infty)$, as well as the root ρ^{l_n} of the tree \mathbf{T}^{l_n} . In particular, the forest F_n has l_n independent Bienaymé–Galton–Watson trees with Geometric(1/2) offspring distribution, and the labels of the root vertices of the trees (including ρ^{l_n}) follow a random walk of length l_n whose step distribution is a shifted Geometric(1/2) given by $\mathbb{P}(\cdot = r) = 2^{-r-2}$ for $r \geq -1$.

Recall that $(c_i, i \in \mathbb{Z})$ denotes the sequence of corners of the infinite random well-labeled forest $(F_\infty, \lambda_\infty)$ and that v_i is the vertex of F_∞ incident to c_i . According to the construction of Section 2.3, (F_n, λ_n) encodes a slice Q_n , which is part of the UIHPQ Q_∞ constructed from the whole infinite forest $(F_\infty, \lambda_\infty)$. More precisely, the maximal geodesic (resp. shuttle) can be read inside the UIHPQ as the chain of arcs linking c_0 (resp. c_{τ_n}) to its subsequent successors⁵ and the edges of the slice are given by the arcs from c_i to $s(c_i)$ for $0 \leq i < \tau_n$. As a consequence, the vertex v_i can be seen both as a vertex of Q_∞ and as a vertex of Q_n for $0 \leq i \leq \tau_n$.

Furthermore, we can check that Q_n is in fact isometrically embedded in Q_∞ in the sense that, whenever $0 \leq i, j \leq \tau_n$, it holds that $d_{Q_n}(v_i, v_j) = D_\infty(i, j)$.

⁵Recall Section 2.1.

Indeed, similarly to Proposition 4.1, Q_∞ can be obtained as the gluing of the infinite quadrangulation corresponding to the trees \mathbf{T}^k , $k < 0$, of $(F_\infty, \lambda_\infty)$, with Q_n and then with the infinite quadrangulation corresponding to the trees \mathbf{T}^k , $k \geq l_n$, of $(F_\infty, \lambda_\infty)$ along the proper shuttles and maximal geodesics. Alternatively, one may also argue that there are no shortcuts outside Q_n : for $0 \leq i, j \leq \tau_{l_n}$, any path linking v_i to v_j in Q_∞ may be shortened to a path that stays within Q_n since the maximal geodesic and shuttle are geodesics and since the path $c_0 \rightarrow s(c_0) \rightarrow s^2(c_0) \rightarrow \dots$ is a geodesic ray that disconnects Q_∞ .

The contour function, label function and pseudometric corresponding to Q_n are thus obtained by restricting to $[0, \tau_{l_n}]$ the analog functions corresponding to Q_∞ . After rescaling, their joint limit is a direct consequence of Proposition 4.9.

Corollary 4.10. *On $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)}$, it holds that*

$$(C_{(n)}|_{[0, \tau_{l_n}/2n]}, \Lambda_{(n)}|_{[0, \tau_{l_n}/2n]}, D_{(n)}|_{[0, \tau_{l_n}/2n]^2}) \xrightarrow[n \rightarrow \infty]{(d)} (X^{(0,L)}, W^{(0,L)}, D^{(0,L)}),$$

where we used the notation of Section 4.1, that is,

- (a) the pair $(X^{(0,L)}, W^{(0,L)})$ is the restriction to the interval $[0, T_L]$ of (X, W) distributed under **Half**,
- (b) $D^{(0,L)} = D_{X^{(0,L)}, W^{(0,L)}}$ is the random pseudometric on \mathbb{R} defined by (4.5).

Proof. By the Skorokhod representation theorem, we may and will assume that the convergence (4.14) holds almost surely. Classically, the almost sure path properties of X at time T_L , namely, the fact that X immediately visits the interval $(-L - \varepsilon, -L)$ after time T_L , yield that $\tau_{l_n}/2n$ almost surely converges to T_L . Corollary 4.2 then yields the result. ■

4.5 Scaling limit of conditioned slices

We now derive Theorem 2.6 from the results of the previous section by standard conditioning arguments.

Convergence of the encoding processes. First, without loss of generality, we may assume that the contour and label processes (C, Λ) of the infinite random well-labeled forest defined in Section 4.4 are the canonical processes, considered under the probability distribution \mathbb{P}_∞ on the canonical space. Next, for $a, l \in \mathbb{N}$ and $\delta \in \mathbb{Z}$, we denote by $\mathbb{P}_{a,l,\delta}$ the distribution of $(C|_{[0, 2a+l]}, \Lambda|_{[0, 2a+l]})$, where (C, Λ) is distributed under $\mathbb{P}_\infty[\cdot \mid \tau_l = 2a + l, \Lambda(\tau_l) = \delta]$. The corresponding forest encoded by this random process is thus composed of l Bienaymé–Galton–Watson trees with Geometric(1/2) offspring distribution and uniform admissible labels, conditioned on the fact that the total number of edges in the trees is a and the label of the root of the last vertex-tree is δ . Similarly to the slice it encodes, we will say that the forest has *tilt* δ .

For every measurable nonnegative functional G , it thus holds that

$$\mathbb{E}_{a,l,\delta}[G] = \mathbb{E}_\infty[G((C(k), \Lambda(k)), 0 \leq k \leq \tau_l) \mid \tau_l = 2a + l, \Lambda(\tau_l) = \delta].$$

Let $(\mathcal{F}_k, k \geq 0)$ be the natural filtration associated with the canonical process (C, Λ) . Note that $((C(k), \Lambda(k)), 0 \leq k \leq \tau_l)$ is the pair of contour and label processes of the first l trees in the forest, and that \mathcal{F}_{τ_l} is the σ -algebra generated by these first l trees (with their labels and that of the root ρ^l). Recall from Proposition 2.3 the definitions of Q_ℓ, P_ℓ .

Lemma 4.11. *Fix $0 < k < l$, as well as $a \in \mathbb{N}$ and $\delta \in \mathbb{Z}$. For every nonnegative functional G that is \mathcal{F}_{τ_k} -measurable, we have*

$$\mathbb{E}_{a,l,\delta}[G] = \mathbb{E}_\infty[\Phi_{a,l,\delta}(\tau_k, k, \Lambda(\tau_k)) \cdot G],$$

where

$$\Phi_{a,l,\delta}(t, l', j) = \frac{Q_{l-l'}(2a + l - t) P_{l-l'}(\delta - j)}{Q_l(2a + l) P_l(\delta)}.$$

Proof. It suffices to prove the result when G is the indicator of the contour and label processes of a given well-labeled forest with l' trees, $(t - l')/2$ edges, and tilt j . In this case, $\mathbb{E}_{a,l,\delta}[G]$ is equal to the number of ways in which one can complete this labeled forest into a well-labeled forest with l trees, a edges and tilt δ , which is the number of forests with $l - l'$ trees, $a + (l' - t)/2$ edges and tilt $\delta - j$, divided by the number of well-labeled forests with l trees, a edges and tilt δ . We conclude by Proposition 2.3. \blacksquare

In addition to the already fixed sequence (l_n) , we consider two sequences (a_n) , (δ_n) satisfying (2.4). We will need the following direct consequence of the local limit theorem [26, Theorem 8.4.1]. Recall the definition of $\varphi_{A,L,\Delta}$ given in (4.10).

Lemma 4.12. *If the integer-valued sequence (l'_n) satisfies $l'_n/\sqrt{2n} \rightarrow L' \in (0, L)$, it holds that*

$$\sup_{0 \leq t \leq a_n, j \in \mathbb{Z}} \left| \Phi_{a_n, l_n, \delta_n}(t, l'_n, j) - \varphi_{A,L,\Delta}\left(\frac{t}{n}, L', \left(\frac{9}{8n}\right)^{1/4} j\right) \right| \xrightarrow{n \rightarrow \infty} 0.$$

We start with the following conditioned version of Corollary 4.10.

Proposition 4.13. *On $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)}$, the triple $(C(n), \Lambda(n), D(n)|_{[0, \tau_n/2n]^2})$ considered under $\mathbb{P}_{a_n, l_n, \delta_n}$ converges in distribution to $(X, W, D_{X,W})$, considered under $\text{Slice}_{A,L,\Delta}$.*

Proof. The joint convergence of the first two coordinates is standard; see, for example, [19, Corollary 16]. Let us fix $\varepsilon \in (0, L)$, define $l_n^\varepsilon = l_n - \lfloor \varepsilon \sqrt{2n} \rfloor$, so that $l_n^\varepsilon/\sqrt{2n} \rightarrow L - \varepsilon$, and set $D_{(n)}^\varepsilon = D(n)|_{[0, \tau_n^\varepsilon/2n]^2}$ and $D_{(n)}^0 = D(n)|_{[0, \tau_n/2n]^2}$.

By the usual bound (2.1), for every $i, j \in [0, \tau_{l_n}]$,

$$\begin{aligned} |D_\infty(i \wedge \tau_{l_n}^\varepsilon, j \wedge \tau_{l_n}^\varepsilon) - D_\infty(i, j)| &\leq D_\infty(i, i \wedge \tau_{l_n}^\varepsilon) + D_\infty(j, j \wedge \tau_{l_n}^\varepsilon) \\ &\leq 4(\omega(\Lambda_n; \tau_{l_n} - \tau_{l_n}^\varepsilon) + 1), \end{aligned}$$

where $\omega(f; \cdot)$ denotes the modulus of continuity of f . This implies that

$$\text{dist}_{\mathcal{E}^{(2)}}(D_{(n)}^\varepsilon, D_{(n)}^0) \leq \frac{\tau_{l_n} - \tau_{l_n}^\varepsilon}{2n} + 4\omega\left(\Lambda_{(n)}, \frac{\tau_{l_n} - \tau_{l_n}^\varepsilon}{2n}\right) + \mathcal{O}(n^{-1/4}).$$

From the joint convergence of the first two coordinates, we have, for every $\eta > 0$,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathbb{P}_{a_n, l_n, \delta_n}(\text{dist}_{\mathcal{E}^{(2)}}(D_{(n)}^\varepsilon, D_{(n)}^0) \geq \eta) \\ \leq \mathbf{Slice}_{A, L, \Delta}(A - T_{L-\varepsilon} + 4\omega(W; A - T_{L-\varepsilon}) \geq \eta), \end{aligned}$$

which tends to 0 as $\varepsilon \rightarrow 0$ since $T_{L-\varepsilon} \rightarrow T_L = A$ almost surely under $\mathbf{Slice}_{A, L, \Delta}$. We next show that $D_{(n)}^\varepsilon$ under $\mathbb{P}_{a_n, l_n, \delta_n}$ converges in distribution to $D^{(0, L-\varepsilon)}$ under $\mathbf{Slice}_{A, L, \Delta}$, and use the principle of accompanying laws [92, Theorem 9.1.13] to conclude that, jointly with the convergence of $(C_{(n)}, \Lambda_{(n)})$ to (X, W) , the process $D_{(n)}^0$ converges to the distributional limit of $D^{(0, L-\varepsilon)}$ as $\varepsilon \rightarrow 0$, which is none other than $D^{(0, L)}$, due to Corollary 4.2.

To prove the claimed convergence of $D_{(n)}^\varepsilon$ to $D^{(0, L-\varepsilon)}$, we denote by $C_{(n)}^\varepsilon$ and $\Lambda_{(n)}^\varepsilon$ the restrictions of $C_{(n)}$ and $\Lambda_{(n)}$ to $[0, \tau_{l_n}^\varepsilon/2n]$ and let F be a nonnegative bounded continuous function. Using Lemma 4.11, then Corollary 4.10 (for the choice of $L - \varepsilon$ instead of L) and Lemma 4.12 gives

$$\begin{aligned} \mathbb{E}_{a_n, l_n, \delta_n}[F(C_{(n)}^\varepsilon, \Lambda_{(n)}^\varepsilon, D_{(n)}^\varepsilon)] &= \mathbb{E}_\infty[\Phi_{a_n, l_n, \delta_n}(\tau_{l_n}^\varepsilon, l_n^\varepsilon, \Lambda(\tau_{l_n}^\varepsilon))F(C_{(n)}^\varepsilon, \Lambda_{(n)}^\varepsilon, D_{(n)}^\varepsilon)] \\ &\xrightarrow{n \rightarrow \infty} \mathbf{Half}[\varphi_{A, L, \Delta}(T_{L-\varepsilon}, L - \varepsilon, W_{T_{L-\varepsilon}})F(X^{(0, L-\varepsilon)}, W^{(0, L-\varepsilon)}, D^{(0, L-\varepsilon)})], \end{aligned}$$

the latter being equal to $\mathbf{Slice}_{A, L, \Delta}[F(X^{(0, L-\varepsilon)}, W^{(0, L-\varepsilon)}, D^{(0, L-\varepsilon)})]$ according to Lemma 4.7. \blacksquare

GHP convergence. We infer from Proposition 4.13 the GHP convergence of Theorem 2.6 by a standard method. First, by Skorokhod's representation theorem, we may assume that we are working on a probability space on which the convergence of Proposition 4.13 is almost sure. We let $\text{Sl}_{A, L, \Delta}$ be the continuum slice coded by the limiting process, and Sl_n be the slice encoded by the forest whose rescaled contour and label processes make up the pair $(C_{(n)}, \Lambda_{(n)})$. As mentioned before Corollary 4.10, Sl_n is isometrically embedded in \mathcal{Q}_∞ , so that the process $D_{(n)}|_{[0, \tau_{l_n}/2n]^2}$ under $\mathbb{P}_{a_n, l_n, \delta_n}$ projects into the metric of $\Omega_n(\text{Sl}_n)$.

Then, from this almost sure convergence, we easily deduce that the distortion of the correspondence \mathcal{R}_n given by

$$\mathcal{R}_n = \{(v_{\lfloor (2a_n + l_n)s \rfloor}, \mathbf{p}_{X, W}(As)) : s \in [0, 1]\}$$

between $\Omega_n(\text{Sl}_n)$ minus its shuttle and $\text{Sl}_{A,L,\Delta}$ tends to 0 as $n \rightarrow \infty$. Forgetting the marks and measures, this gives the desired convergence in the 0-marked Gromov–Hausdorff topology.

In order to include the marking and measures, we use the technique of enlargement of correspondences already used in the proof of Lemma 3.9. Namely, we fix $\varepsilon > 0$ and let $\mathcal{R}_n^\varepsilon$ be the set of points of the form (v, x) in $\text{Sl}_n \times \text{Sl}_{A,L,\Delta}$ such that there exists $(w, y) \in \mathcal{R}_n$ satisfying $d_{\text{Sl}_n}(v, w) < (8n/9)^{1/4}\varepsilon$ and $D(x, y) < \varepsilon$. As before, the distortion of $\mathcal{R}_n^\varepsilon$ is at most $\text{dis}(\mathcal{R}_n) + 4\varepsilon$. Let us start with the marks.

Marks. For a function $f \in \mathcal{C}$ defined over the interval I , we say that $s \in I$ is a *left-minimum* of f if $f(t) \geq f(s)$ for every $t \leq s$ in I , and we call it *strict* if $f(t) > f(s)$ for $t < s$ in I . Note that the points of the form v_i and $\mathbf{p}_{X,W}(s)$, where i and s are left-minimums of Λ_n and W respectively belong to the maximal geodesics of Sl_n and $\text{Sl}_{A,L,\Delta}$, and that all points in these sets are in fact of this form, where we can even require the stronger property that i and s are strict left-minimums.

By the uniform convergence of $\Lambda_{(n)}$ toward W , for every $\eta > 0$, the following holds provided $n \geq n_0$ for some n_0 : every strict left-minimum of $\Lambda_{(n)}$ is at distance at most $\eta/2$ from some (not necessarily strict) left-minimum of W , and vice-versa, exchanging the roles of $\Lambda_{(n)}$ and W . Up to increasing n_0 , we furthermore assume that $|(2a_n + l_n)/2n - A| < \eta/2$ as soon as $n \geq n_0$. Choosing η small enough so that $|D_{(n)}(s, t) - D_{(n)}(s', t')| \leq \varepsilon$ for every $n, |s - s'| \leq \eta, |t - t'| \leq \eta$, we deduce that the extended correspondence $\mathcal{R}_n^\varepsilon$ is compatible with the maximal geodesics for $n \geq n_0$.

The argument is similar for the shuttles. This time, we note that elements of the shuttle of $\text{Sl}_{A,L,\Delta}$ are of the form $\mathbf{p}_{X,W}(s)$, where s is a right-minimum of the function W (with an obvious definition), while elements of the shuttle of Sl_n are at distance 1 from points of the form v_i , where i is a right-minimum of the function Λ_n .

The mark corresponding to the base is also treated similarly. Recall from Section 2.3 that vertices of the base are at distance at most $B_n = \max_{1 \leq i \leq l_n} |\Lambda_n(\rho^i) - \Lambda_n(\rho^{i-1})| + 1$ from some element of the floor $\{\rho^0, \dots, \rho^{l_n}\}$ of the forest coding the slice. The process of labels $(\Lambda_n(\rho^i), 0 \leq i \leq l_n)$ forms a random walk with shifted Geometric(1/2) increments conditioned to be equal to δ_n at time l_n , so, under our assumptions, it converges, after rescaling by $\sqrt{2n}$ in time and $(8n/9)^{1/4}$ in space, to a continuous process (which is easily checked to be the Brownian bridge $\zeta = (W_{T_x}, 0 \leq x \leq L)$), so that $B_n = o(n^{1/4})$ almost surely. Therefore, the base of Sl_n is at Hausdorff distance $o(n^{1/4})$ from the floor $\{\rho^i, 0 \leq i \leq l_n\}$. In turn, these vertices are exactly those of the form v_i , where i is a left-minimum of the contour process C_n . Moreover, by definition, the base of $\text{Sl}_{A,L,\Delta}$ consists of the points $\mathbf{p}_{X,W}(s)$, where s is a left-minimum of the process X . Therefore, the same argument as for the maximal geodesic – replacing the processes Λ_n and W by C_n and X – shows that, almost surely, for every n large enough, the correspondence $\mathcal{R}_n^\varepsilon$ is also compatible with the bases of Sl_n and of $\text{Sl}_{A,L,\Delta}$.

Measures. Finally, let us deal with the convergence of the measures, starting with the area measure. To this end, note that, for t in $[0, 2a_n + l_n]$, the contour process C_n at time t has either a left derivative equal to $+1$ or to -1 . Letting $i_t^n = \lceil t \rceil$ in the former case and $i_t^n = \lfloor t \rfloor$ in the latter case, the image of $\text{Leb}_{[0, (2a_n + l_n)/2n]}$ by $t \mapsto v_{i_t^n}^{2nt}$ is the counting measure on the set of all nonfloor vertices of the encoding forest, divided by n . Since the number of floor vertices is $\mathcal{O}(\sqrt{n})$, the counting measure on all vertices of Sl_n (except on the shuttle) divided by n is at vanishing Prokhorov distance from the counting measure on nonfloor vertices of the forest, divided by n . Let ω_n be the image of the Lebesgue measure on $[0, A \wedge ((2a_n + l_n)/2n)]$ by the mapping $t \mapsto (v_{i_t^n}^{2nt}, \mathbf{p}_{X,W}(t))$. Then ω_n is carried by the correspondence $\mathcal{R}_n^\varepsilon$ for every n large enough, and its image measures on Sl_n and $\text{Sl}_{A,L,\Delta}$ by the coordinate projections are at vanishing Prokhorov distances from μ_{Sl_n}/n and $\mu^{(0,L)}$, respectively.

For the base measure, we let ω'_n be the image of the Lebesgue measure on $[0, L \wedge l_n/\sqrt{2n}]$ by the mapping $t \mapsto (v_{\tau_{\lfloor \sqrt{2nt} \rfloor}}, \mathbf{p}_{X,W}(T_t))$. Then ω'_n is carried by $\mathcal{R}_n^\varepsilon$, by the above discussion on the mark corresponding to the base. Moreover, the coordinate projections of ω'_n are at vanishing Prokhorov distance, respectively, from the counting measure on $\{\rho^0, \dots, \rho^{l_n}\}$ divided by $\sqrt{2n}$, and $\nu^{(0,L)}$. We now observe that, in turn, the counting measure on $\{\rho^0, \dots, \rho^{l_n}\}$ divided by $\sqrt{2n}$, is at vanishing Prokhorov distance from the renormalized counting measure (with multiplicities) $\nu_{\beta_n}/\sqrt{8n}$ of the base. To justify this, observe from Section 2.3 and the definition of the interval CVS bijection that the sequence $\Lambda_n(w_0), \dots, \Lambda_n(w_{2l_n + \delta_n})$ of labels of the vertices $w_0, \dots, w_{2l_n + \delta_n}$ of the base, taken in contour order, forms a simple random walk starting with a -1 step, and conditioned on hitting δ_n at time $2l_n + \delta_n$. Moreover, if we write the set $\{j \in \{0, \dots, 2l_n + \delta_n - 1\} : \Lambda_n(w_{j+1}) - \Lambda_n(w_j) = -1\}$ of down steps of this walk as $\{j_0, j_1, \dots, j_{l_n-1}\}$ with $0 = j_0 < j_1 < j_2 < \dots < j_{l_n-1}$, then the i -th root ρ^i is equal to w_{j_i} for $0 \leq i < l_n$. Now consider a uniform random variable U in $[0, 1)$. Then $w_{j_{\lfloor l_n U \rfloor}}$ is a uniformly chosen forest root, while $w_{\lfloor (2l_n + \delta_n) U \rfloor}$ is a vertex of the base chosen with probability proportional to its multiplicity (and excluding ρ^{l_n} in both cases). Moreover, a standard large deviation estimate entails that $\max_{0 \leq k < l_n} |j_k - 2k| = \mathcal{O}(\log n)$ in probability. In turn, this easily implies that $d_{\text{Sl}_n}(w_{j_{\lfloor l_n U \rfloor}}, w_{\lfloor (2l_n + \delta_n) U \rfloor}) = \mathcal{O}(\log n)$ in probability, showing that the uniform measure on the $l_n \sim L\sqrt{2n}$ elements of $\{\rho^0, \dots, \rho^{l_n-1}\}$ is at vanishing Prokhorov distance from the law of the vertex incident to a corner uniformly chosen among the $2l_n + \delta_n \sim L\sqrt{8n}$ corners incident to the base.

Conclusion. By Lemma 3.3, we finally obtain that

$$\limsup_{n \rightarrow \infty} d_{\text{GHP}}^{(5,2)}(\Omega_n(\text{Sl}_n), \text{Sl}_{A,L,\Delta}) \leq \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, this concludes the proof of Theorem 2.6.

Chapter 5

Convergence of quadrilaterals with geodesic sides

The general method to prove Theorem 2.8 is the same as for slices. We start by seeing a discrete quadrilateral as part of a discrete map that is known to converge to a Brownian surface, which in this case is the Brownian plane rather than the Brownian half-plane. However, the lack of an analog of Corollary 4.2, namely that quadrilaterals are only *locally* isometrically embedded in the Brownian plane, makes matters considerably more delicate. For this reason, we adapt the strategy we used in [25, Section 4] when treating the case of noncomposite slices. Beware that, in this chapter, part of the notation we will be using is slightly conflicting with that of Chapter 4: in particular, the random times T_x will be re-defined.

5.1 Quadrilaterals coded by two functions

In contrast with slices, which were coded by a pair of functions defined on a common interval I , a quadrilateral will be coded by a pair of functions defined on a common union of two intervals $I_- \cup I_+$, each interval accounting for one “half” of the quadrilateral. This leads to similar but slightly more intricate definitions. We start with the most convenient setting, asking that $\sup I_- = \inf I_+ = 0$.

Recall the notation of Section 4.1.1. We adapt Section 4.1.2 to quadrilaterals as follows. We now say that a pair $(f, g) \in \mathcal{C}^2$ of functions with common closed interval of definition I is a *quadrilateral trajectory* if they satisfy (4.6) and the following:

- the interval I contains 0 in its interior and is either bounded or equal to the whole real line \mathbb{R} , and, letting $I_+ = I \cap \mathbb{R}_{\geq 0}$ and $I_- = I \cap \mathbb{R}_{\leq 0}$,
- we have $\inf_{I_-} f = \inf_{I_+} f$, and
- if $I = \mathbb{R}$, then $\inf_{t \geq 0} f(t) = \inf_{t \leq 0} f(t) = \inf_{t \geq 0} g(t) = \inf_{t \leq 0} g(t) = -\infty$.

We may observe that $(f|_{I_+}, g|_{I_+})$ is a slice trajectory, a fact that will not be used here. For a quadrilateral trajectory (f, g) , we set

$$\hat{d}_g(s, t) = \begin{cases} d_g(s, t) & \text{for } s, t \in I_+ \text{ or } s, t \in I_-, \\ \infty & \text{for } st < 0 \end{cases}$$

and

$$\hat{D}_{f,g} = \hat{d}_g / \{d_f = 0\}. \quad (5.1)$$

Note that \hat{d}_g is the disjoint union pseudometric of the two \mathbb{R} -tree pseudometrics $d_g|_{I_+}$ and $d_g|_{I_-}$. Let $(\hat{M}_{f,g}, \hat{D}_{f,g})$ be the quotient space $I / \{\hat{D}_{f,g} = 0\}$ equipped with

the metric induced by $\widehat{D}_{f,g}$, still denoted by the same symbol. We call the metric space

$$\text{Qd}_{f,g} = (\widehat{M}_{f,g}, \widehat{D}_{f,g})$$

the *quadrilateral* coded by (f, g) .

We extend the above constructions to unions of two closed intervals $I = I_- \cup I_+$, where $I_- \subseteq \mathbb{R}_{\leq 0}$ and $I_+ \subseteq \mathbb{R}_{\geq 0}$, as follows. First, a pair of functions (f, g) defined on I is a *quadrilateral trajectory* if the pair (f', g') defined by

$$f'(t) = \begin{cases} f(t + \inf I_+) & \text{for } t \in I_+ - \inf I_+, \\ f(t + \sup I_-) & \text{for } t \in I_- - \sup I_-, \end{cases}$$

and similarly for g' , is a quadrilateral trajectory as defined above. Note that the continuity hypothesis on f' implies in particular that $f(\sup I_-) = f(\inf I_+)$, and similarly for g . We then define the quadrilateral coded by (f, g) using the exact same definitions as above. Note that the mapping $t \mapsto (t - \inf I_+) \mathbf{1}_{t \in I_+} + (t - \sup I_-) \mathbf{1}_{t \in I_-}$ induces an isometry from $(\widehat{M}_{f,g}, \widehat{D}_{f,g})$ onto $(\widehat{M}_{f',g'}, \widehat{D}_{f',g'})$.

From now on, we work in this extended framework and consider a fixed quadrilateral trajectory (f, g) .

Geodesic sides and area measure. For every $t \in I \setminus \{0\}$, we let $I_t = I_-$ if $t < 0$ or $I_t = I_+$ if $t > 0$, and set

$$\Gamma_t(r) = \inf\{s \geq t : g(s) = g(t) - r\} \quad \text{for } r \in \mathbb{R}_{\geq 0} \text{ such that } \inf_{\substack{s \geq t \\ s \in I_t}} g(s) \leq g(t) - r;$$

$$\Xi_t(r) = \sup\{s \leq t : g(s) = g(t) - r\} \quad \text{for } r \in \mathbb{R}_{\geq 0} \text{ such that } \inf_{\substack{s \leq t \\ s \in I_t}} g(s) \leq g(t) - r.$$

If $0 \in I$, we also define Γ_0 and Ξ_0 by the same definition, using $I_0 = I_+$ in the definition of Γ_0 , while using $I_0 = I_-$ in the definition of Ξ_0 . Observe that, in contrast with the definition for slices, the infimum of g is now taken on a subset of I_t . In particular, this implies that the ranges of Γ_t, Ξ_t are included in I_t . From the same discussion as the one around (4.7), we see that Γ_t, Ξ_t are geodesics for the pseudo-metrics \widehat{d}_g and $\widehat{D}_{f,g}$. In the case where $\sup I_+ = \infty$, then, for every $t \in I_+$, the range of the path Γ_t is a geodesic ray, and, in the case where $\inf I_- = -\infty$, then the same goes for Ξ_t for every $t \in I_-$. This allows us to define geodesic paths in $\text{Qd}_{f,g}$ by the formulas

$$\gamma_t(r) = \widehat{\mathbf{p}}_{f,g}(\Gamma_t(r)), \quad 0 \leq r \leq g(t) - \underline{g}(t, \sup I_t), \quad r \in \mathbb{R},$$

$$\xi_t(r) = \widehat{\mathbf{p}}_{f,g}(\Xi_t(r)), \quad 0 \leq r \leq g(t) - \underline{g}(\inf I_t, t), \quad r \in \mathbb{R},$$

where $\widehat{\mathbf{p}}_{f,g}: I \rightarrow \widehat{M}_{f,g}$ is the canonical projection and, as above, if $0 \in I$, $I_0 = I_+$ in the definition of γ_0 and $I_0 = I_-$ in that of ξ_0 . Note that the geodesics γ_t, ξ_t share a common initial part.

The quadrilateral $\text{Qd}_{f,g}$ comes with four or two geodesic sides, defined as follows. If I is bounded, the particular geodesics $\gamma = \gamma_{\inf I_+}$ and $\bar{\gamma} = \gamma_{\inf I_-}$ are called the *maximal geodesics* of $\text{Qd}_{f,g}$, while $\underline{\xi} = \xi_{\sup I_+}$ and $\bar{\xi} = \xi_{\sup I_-}$ are called the *shuttles* of $\text{Qd}_{f,g}$. In this case, γ, ξ (resp. $\bar{\gamma}, \bar{\xi}$) have a common endpoint $x_* = \hat{\mathbf{p}}_{f,g}(s_*)$ (resp. $\bar{x}_* = \hat{\mathbf{p}}_{f,g}(\bar{s}_*)$), where $s_* \in I_+$ is such that $g(s) = \inf_{I_+} g$ (resp. $\bar{s}_* \in I_-$ is such that $g(\bar{s}_*) = \inf_{I_-} g$). The points x_*, \bar{x}_* are called the *apexes* of $\text{Qd}_{f,g}$. If I is unbounded, then $\text{Qd}_{f,g}$ has one *maximal geodesic* $\gamma = \gamma_{\inf I_+}$ and one *shuttle* $\bar{\xi} = \xi_{\sup I_-}$; we set $\underline{\xi}_\infty = \gamma_{-\infty} = \emptyset$.

Finally, the *area measure* is defined as $\mu = (\hat{\mathbf{p}}_{f,g})_* \text{Leb}_I$.

Gluing quadrilaterals. For $x \in \mathbb{R}$, we let

$$\begin{aligned} T_x &= \inf\{t \in I_+ : f(t) = -x\} \in \mathbb{R}_{\geq 0} \cup \{+\infty\}, \\ \bar{T}_x &= \sup\{t \in I_- : f(t) = -x\} \in \mathbb{R}_{\leq 0} \cup \{-\infty\}, \end{aligned}$$

as well as $T_\infty = -\bar{T}_\infty = \infty$. Note that, here again, there is a slight difference with the definition of Chapter 4 since, now, $\mathbb{R}_{\leq 0}$ and $\mathbb{R}_{\geq 0}$ play different roles. Recall that $\inf_{I_-} f = \inf_{I_+} f = \inf_I f$, and let $H, H' \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ be such that

$$0 \leq H \leq H' \leq -\inf_I f.$$

We may define the restrictions $f^{(H,H')}, g^{(H,H')}$ of f and g to the union of intervals $I^{(H,H')} = [\bar{T}_{H'}, \bar{T}_H] \cup [T_H, T_{H'}]$, which is a subset of I . The pair $(f^{(H,H')}, g^{(H,H')})$ is another quadrilateral trajectory. The associated quadrilateral is defined as

$$\text{Qd}^{(H,H')} = (\hat{M}^{(H,H')}, \hat{D}^{(H,H')}) = \text{Qd}_{f^{(H,H')}, g^{(H,H')}}.$$

We let $\hat{\mathbf{p}}^{(H,H')} : I^{(H,H')} \rightarrow \hat{M}^{(H,H')}$ be the canonical projection, $\mu^{(H,H')}$ be the area measure of $\text{Qd}^{(H,H')}$, and, whenever they exist, $\gamma^{(H,H')}, \bar{\gamma}^{(H,H')}$ be the maximal geodesics, $\xi^{(H,H')}, \bar{\xi}^{(H,H')}$ be the shuttles.

We refer to Figure 5.1 for an illustration of the following proposition in the upcoming context of random quadrilaterals in the Brownian plane.

Proposition 5.1. *Let $0 \leq H < H' < H'' \leq -\inf_I f$ be in the extended positive real line. Then*

$$\text{Qd}^{(H,H'')} = G(G(\text{Qd}^{(H,H')}, \text{Qd}^{(H',H'')}; \xi^{(H,H')}, \gamma^{(H',H'')}); \bar{\gamma}^{(H,H')}, \bar{\xi}^{(H',H'')}), \quad (5.2)$$

and it holds that

$$\begin{aligned} \gamma^{(H,H'')} &= \gamma^{(H,H')} \cup (\gamma^{(H',H'')} \setminus \xi^{(H,H')}), \\ \xi^{(H,H'')} &= \xi^{(H',H'')} \cup (\xi^{(H,H')} \setminus \gamma^{(H',H'')}), \\ \bar{\gamma}^{(H,H'')} &= \bar{\gamma}^{(H',H'')} \cup (\bar{\gamma}^{(H,H')} \setminus \bar{\xi}^{(H',H'')}), \\ \bar{\xi}^{(H,H'')} &= \bar{\xi}^{(H,H')} \cup (\bar{\xi}^{(H',H'')} \setminus \bar{\gamma}^{(H,H')}). \end{aligned}$$

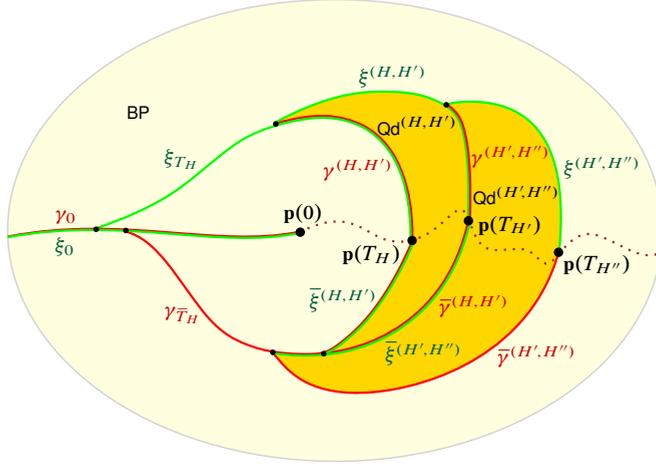


Figure 5.1. Seeing free quadrilaterals in the Brownian plane. The union of the dark yellow regions forms $\text{Qd}^{(H, H')}$. The dotted brown line is $\{\mathbf{p}(T_h) : h \geq 0\}$. Note that BP is obtained by gluing $\text{Qd}^{(0, \infty)}$ along the geodesics $\gamma^{(0, \infty)}$ and $\xi^{(0, \infty)}$, resulting in the geodesic $\gamma_0 = \xi_0$.

Observe that, after the first gluing operation is performed, the marks $\bar{\gamma}$ and $\bar{\xi}$ remain geodesic, as observed in Section 3.3. Note also that the order of the gluings in (5.2) is not important, due to (3.3).

Proof. The proof is similar to that of Proposition 4.1, and we only sketch the argument. Again, we view $I^{(H, H')}$ as a disjoint union $I^{(H, H')} \sqcup I^{(H', H')}$ (denoting elements of these sets with superscripts 0, 1, respectively), where the extremities $T_{H'}^0, T_{H'}^1$ and $\bar{T}_{H'}^0, \bar{T}_{H'}^1$ are identified. We then observe that the pseudometric \hat{d}_g can be viewed as a quotient d/R_1 , where d is the disjoint union metric on $I^{(H, H')} \sqcup I^{(H', H')}$ whose restriction to each interval composing this set equals the restriction of d_g to that interval, and R_1 is the coarsest equivalence relation containing

$$\begin{aligned} & \{(\Xi_{T_{H'}^0}(r)^0, \Gamma_{T_{H'}^1}(r)^1), 0 \leq r \leq g(T_{H'}) - \underline{g}(T_H, T_{H'}) \vee \underline{g}(T_{H'}, T_{H''})\}, \\ & \{(\Gamma_{\bar{T}_{H'}^0}(r)^0, \Xi_{\bar{T}_{H'}^1}(r)^1), 0 \leq r \leq g(\bar{T}_{H'}) - \underline{g}(\bar{T}_H, \bar{T}_{H'}) \vee \underline{g}(\bar{T}_{H'}, \bar{T}_{H''})\}. \end{aligned}$$

Moreover, the equivalence relation $\{d_f = 0\}$ factorizes in the sense that if $d_f(s, t) = 0$ with $s, t \in I^{(H, H')}$, then it must hold that s, t belong either both to $I^{(H, H')}$ or both to $I^{(H', H')}$. Therefore, setting R_2 as the equivalence relation on $I^{(H, H')} \sqcup I^{(H', H')}$ defined by $s^i R_2 t^j$ if and only if $d_f(s, t) = 0$ and $i = j \in \{0, 1\}$, we obtain

$$\hat{D}^{(H, H')} = (d/R_1)/R_2 = (d/R_2)/R_1.$$

We see that d/R_2 is the pseudometric of the disjoint union of $(I^{(H, H')}, \hat{D}^{(H, H')})$ and $(I^{(H', H')}, \hat{D}^{(H', H')})$, while R_1 can be seen as the coarsest equivalence relation obtained by first gluing $\xi^{(H, H')}$ with $\gamma^{(H', H')}$, and then $\bar{\gamma}^{(H, H')}$ with $\bar{\xi}^{(H', H')}$. ■

Due to the fact that the second gluing operation in Proposition 5.1 involves two geodesics belonging to the same space, there is no direct analog of Corollary 4.2. However, we have the following alternative, which is an immediate consequence of Lemma 3.7 and a crude estimate of the length of the path $\bar{\xi}^{(H',H'')}$.

Proposition 5.2. *Under the same assumptions as in Proposition 5.1, it holds that for every $s, t \in I^{(H,H')}$,*

$$\widehat{D}^{(H,H'')}(s, t) \leq \widehat{D}^{(H,H')}(s, t) \leq \widehat{D}^{(H,H'')}(s, t) + \omega(g; I^{(H',H'')}).$$

Finally, we observe that the metric space $(M_{f,g}, D_{f,g})$ obtained by metric gluing of the pseudometric d_g along the relation $\{d_f = 0\}$, rather than using \widehat{d}_g as in the definition of $\text{Qd}_{f,g}$, is related to the latter by a final gluing operation. The proof is analog to that of Proposition 5.1, noting that d_g is the gluing of \widehat{d}_g along the coarsest equivalence relation containing $\{(\Gamma_0(r), \Xi_0(r)), r \geq 0\}$.

Lemma 5.3. *One has $(M_{f,g}, D_{f,g}) = G(\text{Qd}_{f,g}; \gamma, \bar{\xi})$.*

5.2 Random continuum quadrilaterals

Let us now describe the limiting continuum quadrilaterals that appear in Theorem 2.8, by suitably randomizing the quadrilateral trajectory (f, g) . We let (X, W) be the canonical process defined on quadrilateral trajectories. We introduce, for any process Y defined on an interval containing 0, the piece of notation $\underline{Y}_t = \underline{Y}(0 \wedge t, 0 \vee t)$.

Let us fix $A, \bar{A}, H \in (0, \infty)$ and $\Delta \in \mathbb{R}$. We let $\mathbf{Quad}_{A, \bar{A}, H, \Delta}$ be the probability distribution under which

- $(X_t, 0 \leq t \leq A)$ and $(X_{-t}, 0 \leq t \leq \bar{A})$ are independent first-passage bridges of standard Brownian motion from 0 to $-H$, with durations A and \bar{A} ,
- conditionally given X , the process W has same law as $(Z_t + \zeta_{-\underline{X}_t}, -\bar{A} \leq t \leq A)$, where Z is the random snake driven by $X - \underline{X}$, and ζ is a standard Brownian bridge of duration H and terminal value Δ , independent of X and Z .

In this way, the probability distribution $\mathbf{Quad}_{A, \bar{A}, H, \Delta}$ is carried by quadrilateral trajectories on the interval $[-\bar{A}, A]$. We remark that, in fact, we can view W more directly as the random snake driven by X , conditioned on the event $\{W_A = \Delta\}$, a fact that we leave to the interested reader.

Definition 5.4. The *quadrilateral with half-areas A, \bar{A} , width H and tilt Δ* , generically denote by $\text{Qd}_{A, \bar{A}, H, \Delta}$, is the 6-marked 1-measured metric space $\text{Qd}_{X, W}$ under the law $\mathbf{Quad}_{A, \bar{A}, H, \Delta}$, endowed with its area measure μ , as well as the marking $\partial \text{Qd}_{A, \bar{A}, H, \Delta} = (\gamma, \xi, \bar{\gamma}, \bar{\xi})$, where $\gamma, \bar{\gamma}$ are geodesic marks as usual, while $\xi, \bar{\xi}$ are seen as (nonoriented) geodesic segments, that is, given without their origins.

As for slices, the piece of notation $\partial\text{Qd}_{A,\bar{A},H,\Delta}$ comes from Lemma 3.11, which we will prove at the end of the upcoming section. The boundary of the topological disk $\text{Qd}_{A,\bar{A},H,\Delta}$ is the union of $\gamma, \xi, \bar{\gamma}, \bar{\xi}$, which intersect only at the points $\gamma(0) = \bar{\xi}(0) = \hat{\mathbf{p}}_{f,g}(0), x_*, \bar{\gamma}(0) = \xi(0) = \hat{\mathbf{p}}_{f,g}(A) = \hat{\mathbf{p}}_{f,g}(-\bar{A}),$ and $\bar{x}_*.$

5.3 The Brownian plane, and its embedded quadrilaterals

Similarly to the fact that (free) slices can be found in the Brownian half-plane, one can obtain quadrilaterals from the Brownian plane, as we now explain. We let **Plane** be the probability distribution on \mathcal{C}^2 under which

- the process X is a two-sided standard Brownian motion,¹ and
- the process W is the random snake driven by $X.$

The measure **Plane** is carried by quadrilateral trajectories defined over $\mathbb{R}.$

Definition 5.5. The *Brownian plane*, generically denoted by BP, is the metric space $(M_{X,W}, D_{X,W})$ defined by (4.5), considered under **Plane**. Letting $\mathbf{p}: \mathbb{R} \rightarrow \text{BP}$ be the canonical projection, it is endowed with the area measure $\mu = \mathbf{p}_*\text{Leb}_{\mathbb{R}}.$

In this definition, beware that the metric is indeed defined by (4.5) rather than (5.1), which would produce the metric space $\text{Qd}_{X,W} = \text{Qd}^{(0,\infty)} = (\hat{M}_{X,W}, \hat{D}_{X,W}).$ Observe that, by Lemma 5.3,

$$\text{BP} = G(\text{Qd}^{(0,\infty)}; \gamma^{(0,\infty)}, \bar{\xi}^{(0,\infty)}); \tag{5.3}$$

see Figure 5.1 below for an illustration. Alternatively, the space $\text{Qd}^{(0,\infty)}$ can be seen as cutting the Brownian plane along the geodesic ray $\gamma_0 = \xi_0;$ we do not go into further details as we will not explicitly need this property.

Note also that, despite the similarity between this definition and that of the Brownian half-plane, there is no marking now because, as its name suggests, the Brownian plane is homeomorphic to \mathbb{R}^2 and so has an empty boundary as a topological surface.

One should finally mind that this definition is different from the original one given in [40], which will be recalled in Appendix A; in a nutshell, one goes from a definition to the other by changing X into the process obtained by taking its Pitman transform both on $\mathbb{R}_{\geq 0}$ and on $\mathbb{R}_{\leq 0}.$

Free quadrilaterals. Similarly to the discussion of Section 4.3, the Brownian plane satisfies a Markov property which can be interpreted as a “flow” of continuum quadrilaterals. Fixing $0 \leq H \leq H' \leq \infty,$ and denoting by

$$\vartheta_H: t \in [\bar{T}_{H'} - \bar{T}_H, T_{H'} - T_H] \mapsto (t + \bar{T}_H)\mathbf{1}_{t < 0} + (t + T_H)\mathbf{1}_{t \geq 0},$$

¹This means that $(X_t, t \geq 0)$ and $(X_{-t}, t \geq 0)$ are independent standard Brownian motions.

we see that the process $(X^{(H,H')} \circ \vartheta_H + H, W^{(H,H')} \circ \vartheta_H - W_{T_H})$ is independent of $(X^{(0,H)}, W^{(0,H)})$, $(X^{(H',+\infty)}, W^{(H',+\infty)})$, and has same distribution as the pair $(X^{(H'-H)}, W^{(H'-H)})$. This can be proved by excursion theory of (X, W) separately in positive and negative times; we omit the details, which are similar to those presented in Section 4.3.

Under **Plane**, the process $(X^{(0,H)})$ is a two-sided Brownian motion killed at its first hitting times T_H, \bar{T}_H of $-H$ in positive and negative times, respectively, while the process $(W_{T_x}^{(0,H)}, 0 \leq x \leq H)$ is a standard Brownian motion killed at time H . This implies that the law of $(X^{(0,H)}, W^{(0,H)})$ under **Plane** equals

$$\mathbf{FQuad}_H = \int_{(0,\infty)^2} q_H(A)q_H(\bar{A}) \, dA \, d\bar{A} \int_{\mathbb{R}} p_H(\Delta) \, d\Delta \, \mathbf{Quad}_{A,\bar{A},H,\Delta},$$

where the densities p_t, q_x are defined after (4.9). A random metric space with same law as $\mathbf{Qd}^{(0,H)}$ under \mathbf{FQuad}_H will be referred to as a *free (continuum) quadrilateral* of width H . From these considerations and Proposition 5.1, we obtain the following result.

Proposition 5.6. *Let $0 \leq H < H' < H'' \leq \infty$. Then, under **Plane**, it holds that*

$$\mathbf{Qd}^{(H,H'')} = G(G(\mathbf{Qd}^{(H,H')}, \mathbf{Qd}^{(H',H'')}; \xi^{(H,H')}, \gamma^{(H',H'')}); \bar{\gamma}^{(H,H')}, \bar{\xi}^{(H',H'')}),$$

where the glued spaces $\mathbf{Qd}^{(H,H')}$ and $\mathbf{Qd}^{(H',H'')}$ are independent. Moreover, the space $\mathbf{Qd}^{(H,H')}$ is a free continuum quadrilateral of width $H' - H$.

We refer to Figure 5.1 for an illustration, which suggests, as is proved in the following proposition, that quadrilaterals are topological disks bounded by their geodesic sides. In contrast with our treatment of slices, a difficulty arises from the fact that the quadrilaterals $\mathbf{Qd}^{(H,H')}$ are not isometrically embedded in BP, and, in general, not even locally isometrically embedded (think of a point of BP lying on the geodesic γ_0).

Proposition 5.7. *For every $H \in (0, \infty)$, almost surely under \mathbf{FQuad}_H , the quadrilateral $\mathbf{Qd}^{(0,H)}$ is a topological disk with boundary given by the geodesics $\gamma^{(0,H)}$, $\xi^{(0,H)}$, $\bar{\gamma}^{(0,H)}$ and $\bar{\xi}^{(0,H)}$, which pairwise meet only at the points $\gamma(0) = \bar{\xi}(0)$, $\xi(0) = \bar{\gamma}(0)$, and the apexes $x_*^{(0,H)}$ and $\bar{x}_*^{(0,H)}$.*

In order to prove this proposition and for later use, it will be important to characterize the set $\{D_{X,W} = 0\}$.

Lemma 5.8. *The following holds almost surely under **Plane**. For every $s, t \in \mathbb{R}$ such that $s \neq t$, it holds that*

$$D_{X,W}(s, t) = 0$$

if and only if either $d_X(s, t) = 0$ or $d_W(s, t) = 0$, these two cases being mutually exclusive.

Proof. According to [40, Proposition 11], it holds that $D_{X,W}(s, t) = 0$ implies that $df_X(s, t) = 0$ or $d_W(s, t) = 0$. The fact that these two properties are mutually exclusive is a consequence of the fact from [65, Lemma 2.2] that almost surely, if s is a point such that $X_u \geq X_s$ for every $u \in [s, s + \varepsilon]$ for some $\varepsilon > 0$, then it must hold that $\inf_{u \in [s, s + \delta]} W_u < W_s$ for every $\delta \in (0, \varepsilon)$. In fact, [65, Lemma 2.2] is proved when the process X is distributed as a standard Brownian excursion, and W as a random snake Z driven by this excursion. However, being a local property of the processes at hand, it extends easily to our setting by an absolute continuity argument. Details are left to the reader. \blacksquare

To the terminology of Section 4.1.1, we add

$$\llbracket a, b \rrbracket_f = \llbracket a, b \rrbracket_f \setminus \{b\} \quad \text{for } a, b \in \mathcal{T}_f.$$

The important consequence of this lemma for us is the following. Almost surely, if $a, b \in \mathcal{T}_X$ and $c, d \in \mathcal{T}_W$, then the paths $\pi_X(\llbracket a, b \rrbracket_X)$ and $\pi_W(\llbracket c, d \rrbracket_W)$ are simple paths. Furthermore, $\pi_X(\llbracket a, b \rrbracket_X)$ may intersect $\pi_W(\llbracket c, d \rrbracket_W)$ only if $\pi_X(a) = \pi_W(c)$, in which case these paths intersect at this point only. In particular, if we denote the geodesic ray $\mathbf{p}_X(\{s \geq t : X_s = \underline{X}(t, s)\})$ of \mathcal{T}_X by $\llbracket \mathbf{p}_X(t), \infty \rrbracket_X$, then $\pi_X(\llbracket \mathbf{p}_X(t), \infty \rrbracket_X)$ is a simple path in BP. For instance, in Figure 5.1, we represented the simple path $\pi_X(\llbracket \mathbf{p}_X(0), \infty \rrbracket_X)$ with a dotted brown line.

Proof of Proposition 5.7. Let us depart slightly from the setting of the statement and fix for now two numbers $0 \leq H < H' < \infty$.

Claim. *We assume that the geodesics $\gamma^{(H, H')}$ and $\bar{\xi}^{(H, H')}$ do not intersect γ_0 in BP. Then the following holds:*

- (i) *The geodesics $\gamma^{(H, H')}$, $\xi^{(H, H')}$, $\bar{\gamma}^{(H, H')}$, $\bar{\xi}^{(H, H')}$ intersect only at the points $x_*^{(H, H')}$, $\bar{\gamma}^{(H, H')}(0) = \xi^{(H, H')}(0)$, $\bar{x}_*^{(H, H')}$ and $\gamma^{(H, H')}(0) = \bar{\xi}^{(H, H')}(0)$ in this cyclic order, and their union forms a Jordan curve C .*
- (ii) *The set $\mathbf{p}(I^{(H, H')}) \subseteq \text{BP}$ is the closure of the bounded connected component of $\text{BP} \setminus C$.*

Indeed, note that $\mathbf{p}_X(T_H)$ and $\mathbf{p}_X(T_{H'})$ are two distinct points of $\llbracket \mathbf{p}_X(0), \infty \rrbracket_X$, so that their images by π_X are distinct in BP. Then the paths $\gamma^{(H, H')}$ and $\xi^{(H, H')}$ are the images by π_W of the two geodesic paths $\llbracket \mathbf{p}_W(T_H), a_*(W^{(H, H')}) \rrbracket_W$ and $\llbracket \mathbf{p}_W(T_{H'}), a_*(W^{(H, H')}) \rrbracket_W$ in \mathcal{T}_W , which by definition meet only at $a_*(W^{(H, H')})$, and their union is the geodesic $\llbracket \mathbf{p}_W(T_H), \mathbf{p}_W(T_{H'}) \rrbracket_W$ in \mathcal{T}_W , which is thus projected via π_W to a simple path in BP. Therefore, $\gamma^{(H, H')}$ and $\xi^{(H, H')}$ meet only at $x_*^{(H, H')} = \pi_W(a_*(W^{(H, H')}))$. The same reasoning shows that $\bar{\gamma}^{(H, H')}$ and $\bar{\xi}^{(H, H')}$ intersect only at $\bar{x}_*^{(H, H')}$, and gives that the points $x_*^{(H, H')}$ and $\bar{x}_*^{(H, H')}$ are distinct points (because they are distinct points in \mathcal{T}_W lying inside two geodesics).

Next, if the path $\gamma^{(H,H')}$ does not intersect $\gamma_0 = \pi_W(\llbracket \mathbf{p}_W(0), \infty \rrbracket_W)$, then necessarily the path $\llbracket \mathbf{p}_W(T_H), a_*(W^{(H,H')}) \rrbracket_W$ must be disjoint from $\llbracket \mathbf{p}_W(0), \infty \rrbracket_W$, which means that

$$\underline{W}(T_H, T_{H'}) > \underline{W}(0, T_H) \quad \text{and} \quad \underline{W}(\bar{T}_{H'}, \bar{T}_H) > \underline{W}(\bar{T}_H, 0).$$

This implies that $\llbracket \mathbf{p}_W(T_{H'}), a_*(W^{(H,H')}) \rrbracket_W$ is also disjoint from $\llbracket \mathbf{p}_W(0), \infty \rrbracket_W$, and by projecting by π_W , that $\xi^{(H,H')}$ is disjoint from γ_0 . A similar argument applies to $\bar{\gamma}^{(H,H')}$ and $\bar{\xi}^{(H,H')}$. Therefore, under the conditions of the claim, the paths $\llbracket \mathbf{p}_W(T_H), a_*(W^{(H,H')}) \rrbracket_W$ and $\llbracket \mathbf{p}_W(\bar{T}_H), \bar{a}_*(W^{(H,H')}) \rrbracket_W$ are disjoint paths in \mathcal{T}_W , and their projections $\gamma^{(H,H')}$ and $\bar{\xi}^{(H,H')}$ via π_W intersect, if at all, only at their extremities. It is indeed the case that $\mathbf{p}(T_H) = \mathbf{p}(\bar{T}_H)$, while, as we already saw, $x_*^{(H,H')} \neq \bar{x}_*^{(H,H')}$. This proves (i).

The argument for (ii) is similar to that in the proof of Proposition 4.6, where the role of the base is now played by the infinite path $\pi_X(\llbracket \mathbf{p}_X(0), \infty \rrbracket_X) = \{\mathbf{p}(T_h) : h \geq 0\}$. For any $t \in \mathbb{R}$, we let

$$\Sigma_t(r) = \inf\{s \geq t : X_s = X_t - r\} \quad \text{for } 0 \leq r \leq X_t - \underline{X}_t,$$

where we recall that $\underline{X}_t = \underline{X}(0 \wedge t, 0 \vee t)$. The range of $\mathbf{p}_X \circ \Sigma_t$ is the geodesic path $\llbracket \mathbf{p}_X(t), \mathbf{p}_X(T_{-\underline{X}_t}) \rrbracket_X$ in \mathcal{T}_X and its image by π_X defines a path $\sigma_t = \mathbf{p} \circ \Sigma_t$ from $\mathbf{p}(t)$ to $\mathbf{p}(T_{-\underline{X}_t})$. Moreover, by Lemma 5.8, the paths σ_t , $t \in \mathbb{R}$, do not intersect any of the geodesics γ_s , $s \in \mathbb{R}$, except possibly at their starting points. There are now the following possibilities:

- If $t \in I^{(H,H')}$, then σ_t ends on the path $\{\mathbf{p}(T_h), H \leq h \leq H'\}$. This means that if $\mathbf{p}(t)$ does not belong to the four geodesics $\gamma^{(H,H')}$, $\xi^{(H,H')}$, $\bar{\gamma}^{(H,H')}$, $\bar{\xi}^{(H,H')}$, then we may connect it to, say, the point $\mathbf{p}(T_{(H+H')/2})$ of the bounded set $\text{Qd}^{(H,H')}$, without crossing the four mentioned geodesics.
- If $t \notin I^{(H,H')}$, we distinguish two cases:
 - If $t \notin [\bar{T}_{H'}, T_{H'}]$, then σ_t ends on the unbounded path $\{\mathbf{p}(T_h) : h > H\}$.
 - If $t \in (\bar{T}_H, T_H)$, then σ_t ends on $\{\mathbf{p}(T_h) : 0 \leq h < H\}$.

If $\mathbf{p}(t)$ does not belong to the four geodesics of interest, then it may be joined without crossing the four geodesics either to the unbounded path $\{\mathbf{p}(T_h) : h > H\}$ or to the unbounded path $\{\mathbf{p}(T_h) : 0 \leq h < H\} \cup \gamma_0$, by the assumption that γ_0 does not intersect the four geodesics.

This completes the proof of the claim.

Now fix $H > 0$ and consider another positive number H_0 to be thought of as large. Since we know that $\text{Qd}^{(H_0, H_0+H)}$ under **Plane** has same distribution as $\text{Qd}^{(0, H)}$, we may work with the former space rather than with the latter. For every $\varepsilon > 0$, it holds that there exists some H_0 large enough such that with probability at least

$1 - \varepsilon$, the geodesics $\gamma^{(H_0, H_0+H)}$ and $\bar{\xi}^{(H_0, H_0+H)}$ do not intersect γ_0 . Indeed, this happens whenever $\underline{W}(T_{H_0}, T_{H_0+H}) > \underline{W}(0, T_{H_0})$ or, equivalently,

$$W_{T_{H_0}} - \underline{W}(0, T_{H_0}) > W_{T_{H_0}} - \underline{W}(T_{H_0}, T_{H_0+H}), \tag{5.4}$$

and similarly in negative times. The two sides of (5.4) are independent by the Markov property stated above; the right-hand side has a distribution that depends only on H , while the left-hand side, which has same distribution as $-\underline{W}(0, T_{H_0})$ by a simple time-reversal argument, converges to ∞ in probability as $H_0 \rightarrow \infty$.

By the claim, we obtain that on an event happening with probability at least $1 - \varepsilon$, the set $\mathbf{p}(I^{(H_0, H_0+H)})$ is the closure of the connected component of the complement in BP of the paths

$$\gamma^{(H_0, H_0+H)}, \quad \xi^{(H_0, H_0+H)}, \quad \bar{\gamma}^{(H_0, H_0+H)}, \quad \bar{\xi}^{(H_0, H_0+H)},$$

which all together form a Jordan curve. On this event, the identity on $I^{(H_0, H_0+H)}$ induces, by precomposition with the projection mappings \mathbf{p} and $\mathbf{p}^{(H_0, H_0+H)}$, a bijective mapping ϕ from the compact space $\text{Qd}^{(H_0, H_0+H)}$ to $\mathbf{p}(I^{(H_0, H_0+H)})$, which is 1-Lipschitz since $D_{X,W} \leq \widehat{D}^{(H_0, H_0+H)}$ by Lemma 3.7, (5.3) and Proposition 5.6. This shows that ϕ is a homeomorphism, and therefore, with probability at least $1 - \varepsilon$, the space $\text{Qd}^{(H_0, H_0+H)}$ has the properties claimed in the statement. Using the fact that $\text{Qd}^{(H_0, H_0+H)}$ has same distribution as $\text{Qd}^{(0, H)}$ and that ε was arbitrary, we conclude. ■

The continuum quadrilaterals of the preceding section can be linked to the free quadrilaterals embedded in the Brownian plane by an absolute continuity argument, whose proof is similar to that of Lemma 4.7 and is omitted.

Lemma 5.9. *Fix $0 < K < H$, as well as $A > 0$, $\bar{A} > 0$, and $\Delta \in \mathbb{R}$. Then, for every nonnegative function G that is measurable with respect to the σ -algebra generated by $(X^{(0, K)}, W^{(0, K)})$, one has*

$$\mathbf{Quad}_{A, \bar{A}, H, \Delta}[G] = \mathbf{Plane}[\psi_{A, \bar{A}, H, \Delta}(T_K, -\bar{T}_K, K, W_{T_K}) \cdot G],$$

where

$$\psi_{A, \bar{A}, H, \Delta}(A', \bar{A}', H', \Delta') = \frac{q_{H-H'}(A - A')}{q_H(A)} \frac{q_{H-H'}(\bar{A} - \bar{A}')}{q_H(\bar{A})} \frac{p_{H-H'}(\Delta - \Delta')}{p_H(\Delta)}.$$

This allows us to obtain, as stated in Lemma 3.11, the topology of quadrilaterals.

Proof of Lemma 3.11 for quadrilaterals. The proof is similar to that for slices. We use the fact that the Brownian plane is topologically a plane [40, Proposition 13], as well as [40, Proposition 4] to obtain that it is locally of Hausdorff dimension 4 from the analog result about the Brownian sphere [64]. We deduce from there the desired

properties for a free quadrilateral. To extend this result to quadrilaterals $\text{Qd}_{A,\bar{A},H,\Delta}$, which we view as $\text{Qd}^{(0,H)}$ under the law $\mathbf{Quad}_{A,\bar{A},H,\Delta}$, we use the fact from Proposition 5.1 that it can be seen as the gluing of $\text{Qd}^{(0,H/2)}$ and $\text{Qd}^{(H/2,H)}$ along the boundaries $\xi^{(0,H/2)}$ and $\gamma^{(H/2,H)}$ on the one hand, and $\bar{\gamma}^{(0,H/2)}$ and $\bar{\xi}^{(H/2,H)}$ on the other hand. By the absolute continuity relation stated in Lemma 5.9, we see that the law of $\text{Qd}^{(0,H/2)}$ is absolutely continuous with respect to that of a free quadrilateral with width $H/2$, and the same is true for $\text{Qd}^{(H/2,H)}$. Using Proposition 5.7, we obtain that $\text{Qd}^{(0,H)}$ is obtained by gluing two topological disks, both locally of Hausdorff dimension 4, along part of their boundaries, which allows us to conclude. ■

5.4 The uniform infinite planar quadrangulation

The uniform infinite planar quadrangulation (UIPQ) is the whole plane pendant of the UIHPQ defined in Section 4.4. It is simpler to describe and was introduced earlier [36, 42, 60]. Let $(\mathbf{T}^k, k \in \mathbb{Z})$ be a two-sided sequence of independent Bienaymé–Galton–Watson trees with a geometric offspring distribution of parameter $1/2$. We construct an infinite tree \mathbf{T}_∞ embedded in the plane by mapping the roots of \mathbf{T}^k and of \mathbf{T}^{-k} to the point $\rho^k = (k, 0)$ for every $k \geq 0$, in such a way that, except for these roots, the trees $\mathbf{T}^k, k \geq 0$ are embedded in the open upper half-plane and the trees $\mathbf{T}^k, k < 0$ are embedded in the open lower half-plane, without intersection. Lastly, we link the roots ρ^k, ρ^{k+1} with a horizontal segment for every $k \geq 0$.

Conditionally on \mathbf{T}_∞ , we assign to the edges random numbers, independent and uniformly distributed in $\{-1, 0, 1\}$, and let $\lambda_\infty: V(\mathbf{T}_\infty) \rightarrow \mathbb{Z}$ be the labeling function whose increments along the edges are given by these numbers. Note that this uniquely defines λ_∞ , up to the usual addition of a constant. We call $(\mathbf{T}_\infty, \lambda_\infty)$ the *infinite random well-labeled tree*. We then let $(c_i, i \in \mathbb{Z})$ be the sequence of corners of \mathbf{T}_∞ in contour order, with origin the corner c_0 corresponding to the root of \mathbf{T}^0 . The *uniform infinite planar quadrangulation (UIPQ)* for short) is then the infinite map Q_∞ obtained by applying the CVS construction to $(\mathbf{T}_\infty, \lambda_\infty)$, that is, by linking every corner to its successor as defined in Section 2.1, and removing all edges of the tree afterward. The root of Q_∞ is defined as the corner preceding the arc from c_0 to its successor. As with the UIHPQ, there is no need to add an extra vertex with a corner c_∞ .

As before, we denote by v_i the vertex of \mathbf{T}_∞ incident to c_i and by $\Upsilon(i) \in \mathbb{Z}$ the index of the tree to which v_i belongs. We then define the *contour* and *label processes* on \mathbb{R} by

$$C(i) = d_{\mathbf{T}^{\Upsilon(i)}}(v_i, \rho^{|\Upsilon(i)|}) - |\Upsilon(i)| \quad \text{and} \quad \Lambda(i) = \lambda_\infty(v_i) - \lambda_\infty(v_0), \quad i \in \mathbb{Z},$$

and by linear interpolation between integer values; see Figure 5.2. Observe that, in contrast with the definition of Section 4.4 for an infinite forest, there is an absolute

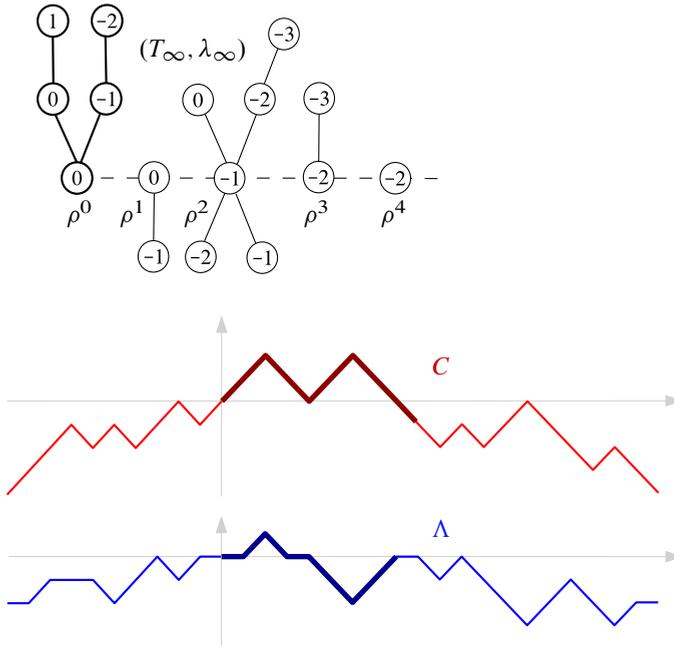


Figure 5.2. Contour and label processes associated with $(\mathbf{T}_\infty, \lambda_\infty)$. The infinite dashed line is the so-called *spine* of the tree. The tree \mathbf{T}^0 and the corresponding encoding processes are highlighted. Similarly to Figure 4.2, one might see the contour process as recording the height of a particle moving at speed one around the infinite tree obtained by now letting ρ^k be located at $(0, -k)$ with \mathbf{T}^k grafted on its right and \mathbf{T}^{-k} on its left (both upright) for $k \geq 0$; see the left of Figure A.1 for an illustration.

value in the definition of C . In fact, changing the $-$ into a $+$ amounts to taking the so-called *Pitman transform*, which is a one-to-one mapping, so this is just a matter of convention. We will come back to this in Appendix A. We can easily check that C is distributed as a two-sided random walk conditioned² on $C(-1) = -1$.

As before, we extend C and Λ to functions on \mathbb{R} by linear interpolation between integer values. For $k \geq 0$, we set

$$\bar{\tau}_k = \max\{i \leq 0 : C(i) = -k\} \quad \text{and} \quad \tau_k = \min\{i \geq 0 : C(i) = -k\}.$$

Note that, for a fixed $k \geq 0$, the process $(k + C(s + \tau_k), 0 \leq s \leq \tau_{k+1} - \tau_k)$ is the contour process of \mathbf{T}^k , while, for $k \geq 1$, $(k + C(s + \bar{\tau}_{k+1} + 1), 0 \leq s \leq \bar{\tau}_k - \bar{\tau}_{k+1} - 1)$ is the contour process of \mathbf{T}^{-k} without the last descending step. Therefore, in this notation, the forest composed of the k leftmost trees in the upper half-plane is coded

²See Remark 5.10 for the explanation of this conditioning.

by the interval $[0, \tau_k]$, while the forest composed of the k leftmost trees in the lower half-plane is coded by the interval $[\bar{\tau}_{k+1} + 1, 0]$. This slightly annoying shift will appear later on, in particular in the statement of Lemma 5.18.

Remark 5.10. As with the UIHPQ, the above definition gives a slight variant of the usual UIPQ, which is similarly defined by adding a further tree rooted at ρ^0 embedded in the lower half-plane, or equivalently, by removing the conditioning by $\{C(-1) = -1\}$. This bias is similar to the one we had for the UIHPQ. Here again, the reason for using this definition is that it will give the natural semigroup property for the discrete quadrilaterals.

We set

$$D_\infty(i, j) = d_{Q_\infty}(v_i, v_j), \quad i, j \in \mathbb{Z},$$

and extend it to a function on \mathbb{R}^2 by bilinear interpolation between integer values, as in (4.12). We define the renormalized versions $C_{(n)}, \Lambda_{(n)}, D_{(n)}$ of C, Λ, D_∞ by (4.13). The following proposition builds on the convergence obtained in [40] of the UIPQ to the Brownian plane. As was the case for the UIHPQ, it does not appear in this exact form in [40] and calls for a proof, which is postponed to Appendix A. Recall from Section 5.3 the definition of the distribution **Plane**.

Proposition 5.11. *The following convergence in distribution holds on $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)}$:*

$$(C_{(n)}, \Lambda_{(n)}, D_{(n)}) \xrightarrow[n \rightarrow \infty]{(d)} (X, W, D_{X,W}),$$

where the limiting triple is understood under **Plane**.

5.5 Discrete quadrilaterals in the uniform infinite planar quadrangulation

We proceed as in the last paragraph of Section 4.4. But, here, the lack of an analog of Corollary 4.2 makes matter substantially more intricate. We consider a sequence $(h_n) \in \mathbb{N}^{\mathbb{N}}$ such that

$$\frac{h_n}{\sqrt{2n}} \xrightarrow[n \rightarrow \infty]{} H > 0.$$

For each n , we let F_n be the random forest consisting of the h_n trees $\mathbf{T}^0, \mathbf{T}^1, \dots, \mathbf{T}^{h_n-1}$, and ρ^{h_n} , as well as \bar{F}_n be the random forest consisting of the h_n trees $\mathbf{T}^{-h_n}, \mathbf{T}^{-h_n+1}, \dots, \mathbf{T}^{-1}$ and ρ^0 . The pair (F_n, \bar{F}_n) is a double forest in the terminology of Section 2.4 and the map $F_n \cup \bar{F}_n$ is well labeled by the restriction of λ_∞ . We denote by Q_n the corresponding quadrilateral and by v_*, \bar{v}_* its apexes; similarly to the previous section, we see it as part of the UIPQ Q_∞ .

For each $i \in \mathbb{Z}$, the vertex v_i of \mathbf{T}_∞ incident to c_i can still be seen as a vertex of Q_n when $\bar{\tau}_{h_n+1} + 1 \leq i \leq \tau_{h_n}$. We set

$$\hat{D}_n(i, j) = d_{Q_n}(v_i, v_j), \quad \bar{\tau}_{h_n+1} + 1 \leq i, j \leq \tau_{h_n},$$

extend it to a function on $[\bar{\tau}_{h_n+1} + 1, \tau_{h_n}]^2$ by bilinear interpolation between integer values as in (4.12), and define its renormalized version

$$\hat{D}_{(n)}(s, t) = \frac{\hat{D}_n(2ns, 2nt)}{(8n/9)^{1/4}}, \quad \frac{\bar{\tau}_{h_n+1} + 1}{2n} \leq s, t \leq \frac{\tau_{h_n}}{2n}. \quad (5.5)$$

This section is devoted to the proof of the following result, which essentially amounts to stating that, jointly with the convergence of $\Omega_n(Q_\infty)$ to the Brownian plane, the properly rescaled quadrilateral $\Omega_n(Q_n)$ converges to $\text{Qd}^{(0,H)}$.

Theorem 5.12. *The following convergence in distribution holds in $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)} \times \mathcal{C}^{(2)}$:*

$$(C_{(n)}, \Lambda_{(n)}, D_{(n)}, \hat{D}_{(n)}) \xrightarrow[n \rightarrow \infty]{(d)} (X, W, D_{X,W}, \hat{D}^{(0,H)}),$$

where the limiting quadruple is understood under **Plane**.

The first step in the proof is the following tightness statement.

Lemma 5.13. *From every increasing sequence of integers, one may extract a subsequence along which the following convergence holds in $\mathcal{C}^{(2)}$, jointly with the convergence of Proposition 5.11:*

$$\hat{D}_{(n)} \xrightarrow[n \rightarrow \infty]{(d)} \tilde{D}, \quad (5.6)$$

where \tilde{D} is a random pseudometric on $[\bar{T}_H, T_H]$.

Proof. The classical tightness argument from [65, Proposition 3.2] implies that the laws of $\hat{D}_{(n)}$, $n \geq 1$, are tight in $\mathcal{C}^{(2)}$. Together with Proposition 5.11, this yields the tightness of the laws of the sequence of the quadruples $(C_{(n)}, \Lambda_{(n)}, D_{(n)}, \hat{D}_{(n)})$, and therefore, by Prokhorov's theorem, their joint convergence in distribution, at least along some subsequence, to a limiting process $(X, W, D_{X,W}, \tilde{D})$, where the law of the first three components is determined by Proposition 5.11. Since $\hat{D}_{(n)}$ is a pseudometric on $[(\bar{\tau}_{h_n+1} + 1)/2n, \tau_{h_n}/2n]$, and because of the convergence of $C_{(n)}$ to X implying the joint convergence of the bounds of this interval to \bar{T}_H, T_H , it is straightforward to check that all subsequential limits of these laws are carried by functions that are pseudometrics on the interval $[\bar{T}_H, T_H]$. ■

From now on, we fix a subsequence along which (5.6) holds, and only consider for the time being values of n that belong to this particular subsequence. By the Skorokhod representation theorem, we may and will assume that this convergence furthermore holds almost surely.

We define $\widetilde{\text{Qd}}$ as the set $[\bar{T}_H, T_H]/\{\widetilde{D} = 0\}$, endowed with the metric \widetilde{D} . Beware that it is not clear at all that $\widetilde{\text{Qd}} = \text{Qd}^{(0,H)}$, and this is precisely what we want to prove. More precisely, we aim at showing that, almost surely, for every $s, t \in [\bar{T}_H, T_H]$, it holds that $\widetilde{D}(s, t) = \widehat{D}^{(0,H)}(s, t)$, which will entail Theorem 5.12.

Since the real number H is fixed once and for all, we will use in the remainder of this section the shorthand pieces of notation

$$\widehat{D} = \widehat{D}^{(0,H)} \quad \text{as well as } D = D_{X,W}.$$

We let $\mathbf{p}: \mathbb{R} \rightarrow \text{BP}$ and $\widetilde{\mathbf{p}}: [\bar{T}_H, T_H] \rightarrow \widetilde{\text{Qd}}$ be the canonical projections, which are continuous since D and \widetilde{D} are continuous functions. Note that, clearly, for every n , it holds that $D_\infty \leq \widehat{D}_n$ on $[\bar{\tau}_{h_{n+1}} + 1, \tau_{h_n}]^2$, so that $D \leq \widetilde{D}$ on $[\bar{T}_H, T_H]$. As a result, there exists a unique continuous (even 1-Lipschitz) projection $\pi: \widetilde{\text{Qd}} \rightarrow \mathbf{p}([\bar{T}_H, T_H])$ such that $\mathbf{p} = \pi \circ \widetilde{\mathbf{p}}$ on $[\bar{T}_H, T_H]$.

The inequality $\widetilde{D} \leq \widehat{D}$ follows from the usual following arguments. First we come back to discrete maps and observe that, for integers $i, j \in [\bar{\tau}_{h_{n+1}} + 1, \tau_{h_n}]$, we have $d_C(i, j) = 0$ if and only if v_i and v_j are the same vertex of $F_n \cup \bar{F}_n$, which implies that $\widehat{D}_n(i, j) = 0$. Next, by considering the so-called *maximal wedge path* consisting of the concatenation of the two geodesics from c_i and from c_j obtained by following subsequent successors up to the point where they coalesce, we obtain the classical upper bound similar to (2.1):

$$\widehat{D}_n(i, j) \leq d_\Lambda(i, j) + 2, \quad i, j \in [\bar{\tau}_{h_{n+1}} + 1, \tau_{h_n}] \text{ with } ij \geq 0. \quad (5.7)$$

Passing to the limit yields that $\{d_X = 0\} \subseteq \{\widetilde{D} = 0\}$ and that $\widetilde{D} \leq \widehat{d}_W$, which imply the inequality $\widetilde{D} \leq \widehat{D}$. The converse inequality is harder and is the focus of what follows.

Let us start with some key properties of the pseudometrics D , \widetilde{D} , and \widehat{D} . The following lemma is proved in the exact same way as [25, Lemma 14].

Lemma 5.14. *The spaces $\widetilde{\text{Qd}}$ and $\text{Qd}^{(0,H)}$ are compact geodesic metric spaces.*

We will need the following identification of the set $\{\widetilde{D} = 0\}$, analog to Lemma 5.8.

Lemma 5.15. *The following holds almost surely. For every $s, t \in [\bar{T}_H, T_H]$ with $s \neq t$, it holds that $\widetilde{D}(s, t) = 0$ if and only if either $d_X(s, t) = 0$ or $\widehat{d}_W(s, t) = 0$, these two cases being mutually exclusive.*

Proof. It follows very similar lines to that of [67, Proposition 3.1], and we will only sketch the main arguments. The fact that $d_X(s, t) = 0$ or $\widehat{d}_W(s, t) = 0$ implies $\widetilde{D}(s, t) = 0$ is immediate from the inequality $\widetilde{D} \leq \widehat{D}$. Conversely, assume that $\widetilde{D}(s, t) = 0$ for some $s \neq t$ in $[\bar{T}_H, T_H]$. Then, in particular, since $D \leq \widetilde{D}$, it holds that $D(s, t) = 0$, so that either $d_W(s, t) = 0$ or $d_X(s, t) = 0$, and these two cases are

exclusive. If we are in the case that s, t are of the same sign and that $d_W(s, t) = 0$, this trivially implies $\widehat{d}_W(s, t) = 0$, as wanted. And since $\widehat{d}_W \geq d_W$, it cannot hold that $\widehat{d}_W(s, t) = d_X(s, t) = 0$ at the same time. Hence, the proof will be complete if we can show that the situation where $d_W(s, t) = 0$ necessarily implies that s and t are of the same sign.

For this, we argue by contradiction, assuming that $t < 0 < s$ and $d_W(s, t) = 0$. Note that this implies in particular that s, t lie on some point of the geodesics Γ_0 and Ξ_0 , respectively, meaning that $W_s = \inf_{u \in [0, s]} W_u$ and $W_t = \inf_{u \in [t, 0]} W_u$. Then, by the convergence of $\Lambda_{(n)}$ to W , there exist $i_n \in [0, \tau_{h_n}]$ and $j_n \in [\bar{\tau}_{h_n+1} + 1, 0]$ such that $i_n/2n \rightarrow s$ and $j_n/2n \rightarrow t$, with the property that $\Lambda_n(i_n) = \min_{k \in [0, i_n]} \Lambda_n(k)$ and $\Lambda_n(j_n) = \min_{k \in [j_n, 0]} \Lambda_n(k)$. This means that v_{i_n} lies on the maximal geodesic γ_n of Q_n , and v_{j_n} lies at distance 1 from the shuttle $\bar{\xi}_n$ of Q_n .

Now any geodesic path in Q_n from v_{i_n} to v_{j_n} will necessarily intersect the spine of the tree \mathbf{T}_∞ at some tree root ρ^{l_n} with $0 \leq l_n \leq h_n$. Let $k_n \in [\bar{\tau}_{h_n+1} + 1, \tau_{h_n}]$ be an integer such that $v_{k_n} = \rho^{l_n}$. In terms of the contour process C_n , this means that $C_n(k_n) \leq C_n(l)$ for every $l \in [0 \wedge k_n, 0 \vee k_n]$. Up to extracting along a further subsequence, we may assume that $k_n/2n \rightarrow u \in [\bar{T}_H, T_H]$ as $n \rightarrow \infty$, and we observe that u must be such that $X_u \leq X_t$ for every $t \in [0 \wedge u, 0 \vee u]$, and in particular, we observe that $d_X(u, T_{H'}) = d_X(u, \bar{T}_{H'}) = 0$, where $H' = -X_u$. We may exclude the case where $H' = 0$ by noting that, necessarily, $W_s = W_t = W_u < 0$.

On the other hand, since v_{k_n} lies on a geodesic path from v_{i_n} to v_{j_n} , which has length $o(n^{1/4})$ because of our assumption that $\widetilde{D}(s, t) = 0$, it holds that $\widetilde{D}(s, u) = \widetilde{D}(u, t) = 0$. We arrive at the wanted contradiction since we have found four points $s \neq t, T_{H'} \neq \bar{T}_{H'}$ that are all identified by D but such that $d_W(s, t) = 0$ and $d_X(T_{H'}, \bar{T}_{H'}) = 0$. ■

As $\widetilde{D} \leq \widehat{D} \leq \widehat{d}_W$ and $\{d_X = 0\} \subseteq \{\widehat{D} = 0\}$, Lemma 5.15 implies that the equivalence relations $\{\widetilde{D} = 0\}$ and $\{\widehat{D} = 0\}$ coincide, and that $\tilde{\mathbf{p}} = \widehat{\mathbf{p}}^{(0, H)}$. For this reason, we may, and will, systematically identify points of $\widetilde{\text{Qd}}$ with points of $\text{Qd}^{(0, H)}$. Moreover, the identity mapping $\text{Qd}^{(0, H)} \rightarrow \widetilde{\text{Qd}}$ is continuous, and by compactness of these spaces, we conclude that $\widetilde{\text{Qd}}$ is homeomorphic to $\text{Qd}^{(0, H)}$.

Theorem 5.12 will be obtained by compactness and continuity arguments from the following local version, stating that, locally and away either from both maximal geodesics or from both shuttles, the three distances under consideration are equal. The proof of the following lemma can straightforwardly be adapted from [25, Lemma 15], so that we only sketch it and refer the reader to the latter reference for the details. In an arbitrary pseudometric space (M, d) , we denote by

$$d(x, A) = \inf\{d(x, y) : y \in A\}$$

the distance from a point $x \in M$ to a subset $A \subseteq M$.

Lemma 5.16. *The following holds almost surely. Fix $\varepsilon > 0$, and let $s, t \in [\bar{T}_H, T_H]$ be such that $\tilde{D}(s, t) < \varepsilon$ and*

- *either $\tilde{D}(s, \Gamma_0 \cup \Gamma_{\bar{T}_H}) \wedge \tilde{D}(t, \Gamma_0 \cup \Gamma_{\bar{T}_H}) > \varepsilon$,*
- *or $\tilde{D}(s, \Xi_0 \cup \Xi_{T_H}) \wedge \tilde{D}(t, \Xi_0 \cup \Xi_{T_H}) > \varepsilon$.*

Then, it holds that $D(s, t) = \tilde{D}(s, t) = \hat{D}(s, t)$.

Proof. Let i_n, j_n be integers in $[\bar{\tau}_{h_n+1} + 1, \tau_{h_n}]$ such that $i_n/2n \rightarrow s$ and $j_n/2n \rightarrow t$ as $n \rightarrow \infty$. From the assumption that $\tilde{D}(s, t) < \varepsilon$ and the convergence of $D_{(n)}$ toward \tilde{D} , we deduce that $d_{Q_n}(v_{i_n}, v_{j_n}) < \varepsilon(8n/9)^{1/4}$ for every n large enough.

Next, keeping the same notation, assume that we are in the first alternative of the statement. Then we claim that for every n large enough, v_{i_n} and v_{j_n} must be at d_{Q_n} -distance at least $\varepsilon(8n/9)^{1/4}$ from the maximal geodesics γ_n and $\bar{\gamma}_n$ of Q_n . Indeed, if we assume otherwise, then up to taking an extraction along a further subsequence, we would find a point $k_n \in [\bar{\tau}_{h_n+1} + 1, \tau_{h_n}]$ such that for every n , v_{k_n} belongs to (the same) one of these maximal geodesics, and is at d_{Q_n} -distance at most $(8n/9)^{1/4}$ from (the same) one of two points v_{i_n} or v_{j_n} . To fix the ideas, assume that v_{k_n} is on γ_n and is close to v_{i_n} in the latter sense, the discussion being similar in the other cases. Up to taking yet another subsequence if necessary, we may assume that $k_n/2n$ converges to some $u \in [\bar{T}_H, T_H]$. Note that k_n , being a time of visit of the maximal geodesic γ_n , must be a left-minimum for the label process Λ_n restricted to nonnegative times, and, by passing to the limit, u must be a left-minimum of W restricted to nonnegative times, entailing that $\tilde{D}(u, \Gamma_0) = 0$. Therefore, by passing to the limit in the inequality $D_{(n)}(i_n/2n, k_n/2n) \leq \varepsilon$, we would obtain that $\tilde{D}(s, \Gamma_0) \leq \varepsilon$, a contradiction with our assumption.

Now observe that Q_∞ is obtained by the following two gluing operations, from Q_n and the infinite quadrangulation Q_n^c , encoded by the labeled double forest with trees grafted above ρ^{h_n+i} , $i \geq 0$, and below ρ^{h_n+i} , $i \geq 1$.

- First, by gluing the geodesic sides ξ_n and $\bar{\gamma}_n$ of Q_n to the (unique) maximal geodesic and shuttle of Q_n^c . Note that the resulting infinite quadrangulation is also obtained by performing the interval CVS construction on \mathbf{T}_∞ with the intervals $\{c_i, i \leq 0\}$ and $\{c_i, i \geq 0\}$.
- Second, by gluing together the (unique) maximal geodesic and shuttle of the infinite quadrangulation obtained at the first step. Note that the geodesic sides of this infinite map are prolongations of γ_n and $\bar{\xi}_n$.

Therefore, Lemma 3.7 (ii) applied twice (once for each gluing operation) shows that if $v, w \in V(Q_n)$ are such that $d_{Q_n}(v, w) < K$ and

- either $d_{Q_n}(v, \gamma_n) \wedge d_{Q_n}(w, \gamma_n) > K$,
- or $d_{Q_n}(v, \bar{\gamma}_n) \wedge d_{Q_n}(w, \bar{\gamma}_n) > K$,

then $d_{Q_n}(v, w) = d_{Q_\infty}(v, w)$. Applying this to $v = v_{i_n}$, $w = v_{j_n}$, and $K = (8n/9)^{1/4}\varepsilon$ yields, after passing to the limit, that $\tilde{D}(s, t) = D(s, t)$. Since $\tilde{D} \leq \hat{D} \leq D$, this yields the result in the first alternative of the statement. The second case, with shuttles instead of maximal geodesics, is similar. ■

We may finally prove Theorem 5.12.

Proof of Theorem 5.12. We follow the same lines as in the proof of [25, Theorem 11]. As we observed before, the metric spaces $\tilde{\text{Qd}}$ and $\text{Qd}^{(0,H)}$ are homeomorphic. Therefore, since the geodesics γ and $\bar{\gamma}$ do not intersect in $\text{Qd}^{(0,H)}$, the same is true in $\tilde{\text{Qd}}$, and similarly, the geodesics ξ and $\bar{\xi}$ do not intersect in these spaces. Moreover, as we know, these four geodesics intersect only at $\gamma(0) = \bar{\xi}(0)$, $\xi(0) = \bar{\gamma}(0)$, $x_*^{(0,H)}$ and $\bar{x}_*^{(0,H)}$. Therefore, for every $x \in \tilde{\text{Qd}} \setminus \{\gamma(0), \bar{\gamma}(0), x_*^{(0,H)}, \bar{x}_*^{(0,H)}\}$, there exists $\varepsilon > 0$ such that the open ball $B_{\tilde{D}}(x, \varepsilon)$ of radius ε around x for the metric \tilde{D} intersects neither $\gamma \cup \bar{\gamma}$ nor $\xi \cup \bar{\xi}$. By Lemma 5.16, this implies that the balls $B_{\tilde{D}}(x, \varepsilon)$, and $B_{\hat{D}}(x, \varepsilon)$ are isometric. Hence, $\tilde{\text{Qd}}$ and $\text{Qd}^{(0,H)}$ are two compact geodesic metric spaces that are locally isometric except possibly around four points. Therefore, the lengths of paths that do not go through these four points must be the same in both spaces. It is then easy to see that the same is true for all paths that visit each of these four points at most once, by splitting into subpaths, and by standard properties of lengths of paths. One concludes by observing that, given a path in $\tilde{\text{Qd}}$, one may construct another path of length smaller than or equal to that of the initial path, and that visits each of the four distinguished points at most once. Since a geodesic space is a length space [32], the distance between two points is given by the infimum of length of paths between these points. Therefore, $\tilde{\text{Qd}}$ and $\text{Qd}^{(0,H)}$ are isometric. ■

5.6 Scaling limit of conditioned quadrilaterals

In this section, we finally prove Theorem 2.8. As a preliminary result, we will need a simple estimate on distances in quadrilaterals. We invite the reader to recall the combinatorial setting of Section 2.4 and to consult Figure 5.3. Let $((\mathbf{f}, \bar{\mathbf{f}}), \lambda)$ be a well-labeled double forest and let \mathbf{q} be the corresponding quadrilateral. For $k \in \{1, 2, \dots, h-1\}$, keeping only the *first* k trees in \mathbf{f} and the *last* k trees in $\bar{\mathbf{f}}$ yields a submap of $\mathbf{f} \cup \bar{\mathbf{f}}$, well labeled by the restriction of λ . We let \mathbf{q}_k be the corresponding quadrilateral, which we naturally see as a submap of \mathbf{q} . We will need the following coarse comparison between distances in \mathbf{q} and \mathbf{q}_k .

Lemma 5.17. *Let $\varpi = 2 + \max\{\lambda(u) : u \in V(\mathbf{q}) \setminus V(\mathbf{q}_k)\} - \min\{\lambda(u) : u \in V(\mathbf{q}) \setminus V(\mathbf{q}_k)\}$. Then, for any $v, w \in V(\mathbf{q}_k)$, one has*

$$d_{\mathbf{q}}(v, w) \leq d_{\mathbf{q}_k}(v, w) \leq d_{\mathbf{q}}(v, w) + \varpi.$$

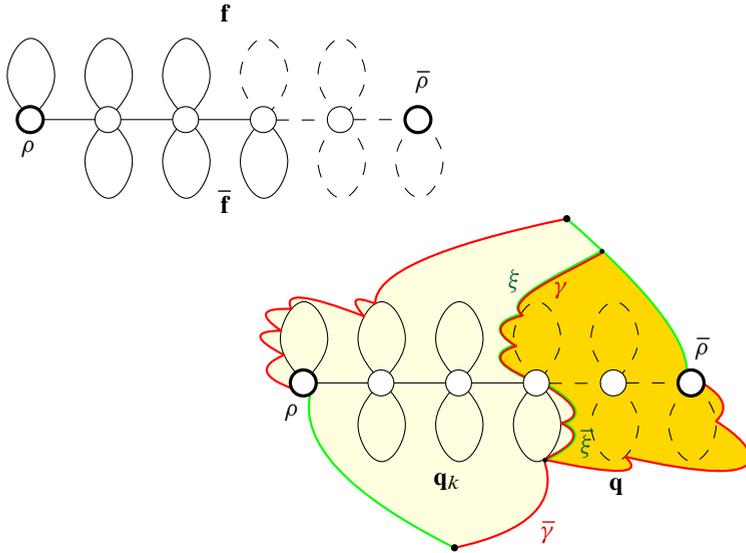


Figure 5.3. Here, $h = 5$ and $k = 3$. On the top left, a schematic picture of a double forest $(\mathbf{f}, \bar{\mathbf{f}})$, assumed to be well labeled, and its “truncation” obtained after removing the dashed elements. On the bottom right, a schematic picture of the corresponding quadrilaterals: the quadrilateral \mathbf{q} is obtained by gluing \mathbf{q}_k (in light yellow) along its sides ξ and $\bar{\gamma}$ with the quadrilateral (in dark yellow) coded by the dashed elements along its sides γ and $\bar{\xi}$ (only these four geodesic sides of interest are named in the picture).

Proof. Observe that \mathbf{q} may be obtained by gluing \mathbf{q}_k along its sides ξ and $\bar{\gamma}$ with the quadrilateral coded by the double forest obtained by taking the last $h - k$ trees in \mathbf{f} and the first $h - k$ trees in $\bar{\mathbf{f}}$, well labeled by the restriction of λ , along its sides γ and $\bar{\xi}$. The lemma is then a straightforward consequence of Lemma 3.7 (i) since the lengths of the glued geodesics are bounded by the quantity ϖ . ■

We now prove Theorem 2.8 by proceeding similarly to Section 4.5. Recall the notation $(C(t), \Lambda(t), t \in \mathbb{R})$, $\tau_k, \bar{\tau}_k$ from Section 5.4. Let \mathbb{P}_∞ be the law of (C, Λ) and assume without loss of generality that the latter is the canonical process. Although we use the same notation \mathbb{P}_∞ as in Section 4.5, we believe that there is little risk of confusion. For $j \geq 1$, let \mathcal{F}_j be the σ -algebra generated by $(C(i), \Lambda(i), 0 \leq i \leq j)$, and let \mathcal{G}_j be the one generated by $(C(i), \Lambda(i + 1), -j \leq i \leq -1)$. Note that \mathcal{F}_{τ_h} is the σ -algebra generated by the h leftmost trees of \mathbf{T}_∞ in the upper half-plane, together with their labels, as well as the label of the root ρ^h . Similarly, $\mathcal{G}_{-\bar{\tau}_{h+1}}$ is the σ -algebra generated by the h leftmost trees of \mathbf{T}_∞ in the lower half-plane, together with their labels, as well as the label of the root ρ^0 : the meaning of the shift by $+1$ in the process Λ is that we do not want to incorporate the information of the label of the root ρ^{h+1} in $\mathcal{G}_{-\bar{\tau}_{h+1}}$.

Next, for $a, \bar{a}, h \in \mathbb{N}$ and $\delta \in \mathbb{Z}$, we denote by $\mathbb{P}_{a, \bar{a}, h, \delta}$ the distribution of

$$(C|_{[-2\bar{a}-h, 2a+h]}, \Lambda|_{[-2\bar{a}-h, 2a+h]})$$

under

$$\mathbb{P}_\infty[\cdot | \tau_h = 2a + h, \bar{\tau}_{h+1} + 1 = -2\bar{a} - h, \Lambda(\tau_h) = \delta].$$

The corresponding double forest encoded by this process is thus composed of a spine ρ^0, \dots, ρ^h of length h , onto which are grafted $2h$ Bienaymé–Galton–Watson trees with Geometric(1/2) offspring distribution and uniform admissible labels, conditioned on the fact that the total number of edges in the upper half-plane h trees is a , the number of edges in the lower half-plane h trees is \bar{a} , and the label of the last root ρ^h is δ . The following lemma gives an absolute continuity relation between the laws $\mathbb{P}_{a, \bar{a}, h, \delta}$ and \mathbb{P}_∞ . Its proof, which we omit, is similar to that of Lemma 4.11, using the enumeration results of Proposition 2.5. Recall from Propositions 2.3 and 2.5 the definitions of Q_ℓ and M_ℓ .

Lemma 5.18. *Fix the integers $0 < k < h$, as well as positive integers $a, \bar{a} \in \mathbb{N}$ and $\delta \in \mathbb{Z}$. For every nonnegative functional G that is $\mathcal{F}_{\tau_k} \vee \mathcal{G}_{-\bar{\tau}_{k+1}}$ -measurable, we have*

$$\mathbb{E}_{a, \bar{a}, h, \delta}[G] = \mathbb{E}_\infty[\Psi_{a, \bar{a}, h, \delta}(\tau_k, -(\bar{\tau}_{k+1} + 1), k, \Lambda(\tau_k)) \cdot G],$$

where

$$\Psi_{a, \bar{a}, h, \delta}(s, t, h', j) = \frac{Q_{h-h'}(2a + h - s)}{Q_h(2a + h)} \frac{Q_{h-h'}(2\bar{a} + h - t)}{Q_h(2\bar{a} + h)} \frac{M_{h-h'}(\delta - j)}{M_h(\delta)}.$$

From now on, in addition to the sequence (h_n) , we fix three sequences (a_n) , (\bar{a}_n) , (δ_n) as in (2.5). The following is a tedious but straightforward consequence of the local limit theorem [26, Theorem 8.4.1].

Lemma 5.19. *If the integer-valued sequence (h'_n) satisfies $h'_n/\sqrt{2n} \rightarrow H' \in (0, H)$, then*

$$\sup_{\substack{0 \leq s \leq a_n \\ 0 \leq t \leq \bar{a}_n \\ j \in \mathbb{Z}}} \left| \Psi_{a_n, \bar{a}_n, h_n, \delta_n}(s, t, h'_n, j) - \psi_{A, \bar{A}, H', \Delta} \left(\frac{s}{n}, \frac{t}{n}, H', \left(\frac{9}{8n} \right)^{1/4} j \right) \right| \xrightarrow{n \rightarrow \infty} 0.$$

We proceed to the conditioned version of Theorem 5.12. Recall the definition of $\hat{D}_{(n)}$ given in (5.5).

Proposition 5.20. *On $\mathcal{C} \times \mathcal{C} \times \mathcal{C}^{(2)}$, the triple $(C_{(n)}, \Lambda_{(n)}, \hat{D}_{(n)})$ considered under the distribution $\mathbb{P}_{a_n, \bar{a}_n, h_n, \delta_n}$ converges in distribution to $(X, W, \hat{D}_{X, W} = \hat{D}^{(0, H)})$, considered under **Quad** $_{A, \bar{A}, H, \Delta}$.*

Proof. The arguments are very close to those used in the proof of Proposition 4.13 in Section 4.5, adding Lemma 5.17 and Proposition 5.2 to cover the additional difficulty. The joint convergence of the first two coordinates is also standard. Then, fix $\varepsilon \in (0, H)$ and set

$$h_n^\varepsilon = h_n - \lfloor \varepsilon \sqrt{2n} \rfloor.$$

Let $\widehat{D}_n^\varepsilon$ and $\widehat{D}_{(n)}^\varepsilon$ be defined as in (5.5) and above, but with h_n^ε instead of h_n . For simplicity, for every $i \in \mathbb{R}$, let

$$i^\varepsilon = (\bar{\tau}_{h_n^\varepsilon+1} + 1) \vee i \wedge \tau_{h_n^\varepsilon}$$

and define j^ε similarly for any $j \in \mathbb{R}$. Define also

$$\kappa_n^\varepsilon = (\tau_{h_n} - \tau_{h_n^\varepsilon}) + (\bar{\tau}_{h_n^\varepsilon+1} - \bar{\tau}_{h_n+1}).$$

From (5.7), we obtain

$$\begin{aligned} |\widehat{D}_n(i, j) - \widehat{D}_n(i^\varepsilon, j^\varepsilon)| &\leq \widehat{D}_n(i, i^\varepsilon) + \widehat{D}_n(j, j^\varepsilon) \\ &\leq 4(\omega(\Lambda_n; \kappa_n^\varepsilon) + 1). \end{aligned}$$

Using Lemma 5.17, we have for every $i, j \in [\bar{\tau}_{h_n^\varepsilon+1} + 1, \tau_{h_n^\varepsilon}]$,

$$|\widehat{D}_n(i, j) - \widehat{D}_n^\varepsilon(i, j)| \leq \omega(\Lambda_n; \kappa_n^\varepsilon) + 2.$$

These two facts together then imply that

$$\text{dist}_{\mathcal{C}(2)}(\widehat{D}_{(n)}^\varepsilon, \widehat{D}_{(n)}) \leq \frac{\kappa_n^\varepsilon}{2n} + 5\omega\left(\Lambda_{(n)}, \frac{\kappa_n^\varepsilon}{2n}\right) + \mathcal{O}(n^{-1/4}).$$

We now use the convergence of the first two coordinates, implying, for every $\eta > 0$,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathbb{P}_{a_n, \bar{a}_n, h_n, \delta_n}(\text{dist}_{\mathcal{C}(2)}(\widehat{D}_{(n)}^\varepsilon, \widehat{D}_{(n)}) \geq \eta) \\ \leq \mathbf{Quad}_{A, \bar{A}, H, \Delta}(\kappa^\varepsilon + 5\omega(W; \kappa^\varepsilon) \geq \eta), \end{aligned} \quad (5.8)$$

where

$$\kappa^\varepsilon = A - T_{H-\varepsilon} + \bar{T}_{H-\varepsilon} + \bar{A}.$$

Since almost surely under $\mathbf{Quad}_{A, \bar{A}, H, \Delta}$, the quantity κ^ε tends to 0 as $\varepsilon \rightarrow 0$, we deduce that the left-hand side in (5.8) also converges to 0. It remains to show that $\widehat{D}_{(n)}^\varepsilon$ under $\mathbb{P}_{a_n, \bar{a}_n, h_n, \delta_n}$ converges toward $\widehat{D}^{(0, H-\varepsilon)}$ under $\mathbf{Quad}_{A, \bar{A}, H, \Delta}$ to conclude, by the principle of accompanying laws, that $\widehat{D}_{(n)}$ converges to the distributional limit of $\widehat{D}^{(0, H-\varepsilon)}$ as $\varepsilon \rightarrow 0$, which is $\widehat{D}^{(0, H)}$ by Proposition 5.2. To this end, we consider the restrictions $C_{(n)}^\varepsilon, \Lambda_{(n)}^\varepsilon$ of $C_{(n)}, \Lambda_{(n)}$ to the intervals $[(\bar{\tau}_{h_n^\varepsilon+1} + 1)/2n, \tau_{h_n^\varepsilon}/2n]$ and, letting F be a nonnegative bounded continuous function, we observe that, using

Lemma 5.18, then Theorem 5.12 (for the choice of $H - \varepsilon$ instead of H) and Lemma 5.19, and finally Lemma 5.9, we have

$$\begin{aligned}
 & \mathbb{E}_{a_n, \bar{a}_n, h_n, \delta_n} [F(C_{(n)}^\varepsilon, \Lambda_{(n)}^\varepsilon, D_{(n)}^\varepsilon)] \\
 &= \mathbb{E}_\infty [\Psi_{a_n, \bar{a}_n, h_n, \delta_n}(\tau_{h_n^\varepsilon}, -1 - \bar{\tau}_{h_n^\varepsilon+1}, \Lambda(\tau_{h_n^\varepsilon})) G(C_{(n)}^\varepsilon, \Lambda_{(n)}^\varepsilon, D_{(n)}^\varepsilon)] \\
 & \xrightarrow{n \rightarrow \infty} \mathbf{Plane}[\psi_{A, \bar{A}, H, \Delta}(T_{H-\varepsilon}, -\bar{T}_{H-\varepsilon}, H - \varepsilon, W_{T_{H-\varepsilon}}) \\
 & \quad \times G(X^{(0, H-\varepsilon)}, W^{(0, H-\varepsilon)}, \hat{D}^{(0, H-\varepsilon)})] \\
 &= \mathbf{Quad}_{A, \bar{A}, H, \Delta}[G(X^{(0, H-\varepsilon)}, W^{(0, H-\varepsilon)}, \hat{D}^{(0, H-\varepsilon)})].
 \end{aligned}$$

This concludes the proof. ■

From there, we easily obtain the wanted GHP convergence by arguments similar to those developed in the proof of Theorem 2.6 at the end of Section 4.5.

Chapter 6

Construction from a continuous unicellular map

Our proof of Theorem 1.1 gives a description of the limiting Brownian surfaces as gluings of elementary pieces, which appear either in the Brownian plane or in the Brownian half-plane. Although this construction has a clear geometric content, it can be arguably cumbersome to work with, having in mind, for instance, the universal character that the spaces $\mathbb{S}_L^{[g]}$ are expected to bear.

Indeed, we believe that Brownian surfaces arise as universal limits for many more classes of maps satisfying mild conditions (for instance, uniformly distributed maps) and a more direct description seems to be useful in order to show such results. In particular, we believe that the Brownian torus is the scaling limit of essentially simple triangulations, as considered in [15]. In fact, most of the known results of convergence toward the Brownian sphere use a re-rooting technique due to Le Gall [67], which, very roughly speaking, says that if maps in a given class are properly encoded by discrete objects converging to the random snake driven by a normalized Brownian excursion and if these maps and the limiting object exhibit a property of invariance under uniform re-rooting, then the limiting space is the Brownian sphere. We expect this approach to be generalizable to our context and we now give a description of Brownian surfaces that is a direct generalization of the classical definition of the Brownian sphere. This can be thought of a continuum version of the Cori–Vauquelin–Schaeffer bijection, building on a continuum version of a unicellular map (a map with only one internal face).

For a function $f \in \mathcal{C}$ and $s, t \in I(f)$ with $t < s$, we extend (4.1) by setting

$$\underline{f}(s, t) = \inf_{I(f) \setminus [t, s]} f$$

and we set, for $s, t \in I(f)$,

$$\tilde{d}_f(s, t) = f(s) + f(t) - 2 \max\{\underline{f}(s, t), \underline{f}(t, s)\}. \quad (6.1)$$

The difference with (4.2) is that we now take into account the minimum of f on the “interval” from $s \vee t$ to $s \wedge t$ on the “circle” $I(f)/\{\bar{\tau}(f) = \tau(f)\}$.

The Brownian sphere. As a warm-up, let us first recall the definition of the Brownian sphere. It is the metric space $\mathbb{S}_\emptyset^{[0]} = ([0, 1], \tilde{d}_Z)/\{d_e = 0\}$, where Z is the random snake driven by a normalized Brownian excursion \mathbf{e} .

Recall that the continuum random tree (CRT) introduced by Aldous [7, 8] is the \mathbb{R} -tree¹ $\mathcal{T}_e = ([0, 1]/\{d_e = 0\}, d_e)$, hence the Brownian sphere $\mathbb{S}_\emptyset^{[0]}$ may actually be

¹See Section 4.1.1.

seen as a quotient of the CRT. In fact, Le Gall [65] showed that the pseudometric $\tilde{d}_Z/\{d_e = 0\}(s, t) = 0$ if and only if $\tilde{d}_Z(s, t) = 0$ or $d_e(s, t) = 0$, so the topological space $S_{\emptyset}^{[0]}$ is obtained by a continuous analog to the Cori–Vauquelin–Schaeffer bijection.

The Brownian disk. Let us turn to the Brownian disk with perimeter $L \in (0, \infty)$. It is the metric space $S_{(L)}^{[0]} = ([0, 1], \tilde{d}_W)/\{d_X = 0\}$, where (X, W) is the pair encoding a slice with area 1, width L and tilt 0, that is, distributed according to **Slice** $_{1,L,0}$ (defined in Section 4.2).

The most natural continuous object generalizing the CRT in the case of the disk is the gluing

$$\mathcal{M}_{(L)}^{[0]} = ([0, 1], d_X)/R,$$

where R is the coarsest equivalence relation containing $\{d_X = 0\}$ and $\{(0, 1)\}$. Since $\tilde{d}_W(0, 1) = 0$, the Brownian disk is also $([0, 1], \tilde{d}_W)/R$ and can be seen as a quotient of $\mathcal{M}_{(L)}^{[0]}$. Visually, $\mathcal{M}_{(L)}^{[0]}$ is obtained by taking a circle of length L and gluing a Brownian forest of mass 1 and length L on it. The random snake W then assigns Brownian labels to it (with a Brownian bridge multiplied by $\sqrt{3}$ on the circle and standard Brownian motions everywhere else).

The general case. The CRT and the structure $\mathcal{M}_{(L)}^{[0]}$ are the continuous equivalent to the encoding objects of Section 2.2 in the particular cases of the sphere and the disk. In general, we have a similar yet even more intricate construction, which we now describe. Let $g \geq 0$ be fixed and $\mathbf{L} = (L^1, \dots, L^b)$ be a b -tuple of positive real numbers. Let then $(S, (A^e)_{e \in \vec{E}(S)}, (H^e)_{e \in \vec{I}(S)}, (L^e)_{e \in \vec{B}(S)}, (\Lambda^v)_{v \in V(S)})$ be a random vector distributed according to the distribution $\text{Param}_{\mathbf{L}}$, defined around (3.8). Conditionally given this vector, we consider the following collection of processes. For each $e \in \vec{E}(S)$,

- the process X^e is a first-passage bridge of standard Brownian motion from 0 to $-H^e$ with duration A^e ,
- the process Z^e is a random snake driven by the reflected process $X^e - \underline{X}^e$,

the processes (X^e, Z^e) , $e \in \vec{E}(S)$, being independent. Independently, the process ζ^e is a Brownian bridge

- of duration H^e from Λ^{e^-} to Λ^{e^+} , with variance 1 if $e \in \vec{I}(S)$,
- of duration L^e from Λ^{e^-} to Λ^{e^+} , with variance 3 if $e \in \vec{B}(S)$.

Furthermore, for $e \in \vec{I}(S)$, the bridges are linked through the relation

$$\zeta^{\bar{e}}(s) = \zeta^e(H^e - s), \quad 0 \leq s \leq H^e,$$

and, except for these relations, are independent. We then set, for each $e \in \vec{E}(S)$,

$$W_t^e = Z_t^e + \zeta_{-\underline{X}_t^e}^e, \quad 0 \leq t \leq A^e.$$

In the end, we obtain a collection of processes (X^e, W^e) , $e \in \vec{E}(S)$, which are linked through the relations linking ζ^e with $\zeta^{\bar{e}}$, translating the fact that the labels of the floors of forests grafted on both sides of the same internal edge of the scheme should correspond.

We arrange the half-edges e_1, \dots, e_κ incident to the internal face of S according to the contour order, starting from the root, and we define the concatenation

$$(W_s)_{0 \leq s \leq 1} = W^{e_1} \bullet \dots \bullet W^{e_\kappa},$$

which is a continuous process. We define \tilde{d}_W by (6.1) as above and now define the equivalence relation along which to glue.

Roughly speaking, we glue together Brownian forests coded by the X^e 's according to the scheme structure. For $s \in [0, 1)$, we denote by $[s]$ the integer in $\{1, \dots, \kappa\}$ such that

$$\sum_{i=1}^{[s]-1} A^{e_i} \leq s < \sum_{i=1}^{[s]} A^{e_i} \quad \text{and} \quad \langle s \rangle = s - \sum_{i=1}^{[s]-1} A^{e_i} \in [0, A^{e_{[s]}}).$$

By convention, we also set $[1] = 1$ and $\langle 1 \rangle = 0$. We define the relation R on $[0, 1]$ as the coarsest equivalence relation for which $s R t$ if one of the following occurs:

$$[s] = [t], \quad d_{X^{e_{[s]}}}(\langle s \rangle, \langle t \rangle) = 0, \tag{6.2a}$$

$$\begin{aligned} e_{[s]} = \bar{e}_{[t]}, \quad X^{[s]} \langle s \rangle = \underline{X}^{[s]} \langle s \rangle, \quad X^{[t]} \langle t \rangle = \underline{X}^{[t]} \langle t \rangle, \\ X^{[s]} \langle s \rangle = H^{e_{[t]}} - X^{[t]} \langle t \rangle, \end{aligned} \tag{6.2b}$$

where we wrote $X^{[s]} \langle s \rangle$ instead of $X^{e_{[s]}}(\langle s \rangle)$ for short. Formulas (6.2a) identify numbers coding the same point in one of the Brownian forests, while formulas (6.2b) identify the floors of forests “facing each other”: the numbers s and t should code floor points (second and third equalities) of forests facing each other (first equality) and correspond to the same point (fourth equality).

Proposition 6.1. *The Brownian surface $S_{\mathcal{L}}^{[g]}$ has same distribution as $([0, 1], \tilde{d}_W)/R$.*

Let us give a similar interpretation as in the case of the disk. Let first $(X_s)_{0 \leq s \leq 1}$ be the continuous process obtained by shifting and concatenating $X^{e_1}, \dots, X^{e_\kappa}$. Then $S_{\mathcal{L}}^{[g]}$ may be seen as a quotient of

$$\mathcal{M}_{\mathcal{L}}^{[g]} = ([0, 1], d_X)/R,$$

which can be pictured as follows. Starting from the random vector $(S, (A^e)_{e \in \vec{E}(S)}, (H^e)_{e \in \vec{I}(S)}, (L^e)_{e \in \vec{B}(S)})$, we first construct the metric graph obtained from S by assigning either the length H^e or L^e to the edge corresponding to e . For every half-edge e incident to the internal face of S , we then glue a Brownian forest of mass A^e

and length H^e or L^e on e . We equip this space $\mathfrak{M}_L^{[g]}$ with Brownian labels (with variance $\sqrt{3}$ on the boundary edges) and define $\mathfrak{S}_L^{[g]}$ from there by the same process as in the case of the Brownian disk.

Proof of Proposition 6.1. First of all, recall from Section 3.6 that the Brownian surface $\mathfrak{S}_L^{[g]}$ is defined as the gluing along geodesic sides of a collection of continuum elementary pieces distributed as follows. Conditionally given

$$(S, (A^e)_{e \in \vec{E}(S)}, (H^e)_{e \in \vec{I}(S)}, (L^e)_{e \in \vec{B}(S)}, (\Lambda^v)_{v \in V(S)}),$$

the elementary pieces EP^e , $e \in \vec{E}(S)$, are only dependent through the relation linking EP^e with $\text{EP}^{\bar{e}}$ and, setting $\Delta^e = \Lambda^{e^+} - \Lambda^{e^-}$,

- if $e \in \vec{B}(S)$, then EP^e is a slice with area A^e , width L^e and tilt Δ^e ,
- if $e \in \vec{I}(S)$, then EP^e is a quadrilateral with half-areas A^e and $A^{\bar{e}}$, width H^e and tilt Δ^e .

Furthermore, it is straightforward from the definition of the pairs (X^e, W^e) , $e \in \vec{E}(S)$, that if $e \in \vec{B}(S)$, then the pair $(X^e, W^e - \Lambda^{e^-})$ is distributed as $\mathbf{Slice}_{A^e, L^e, \Delta^e}$. When $e \in \vec{I}(S)$, we denote by

$$(\bar{X}^e, \bar{W}^e) = (X_{s+A^e}^e - 2\underline{X}_{s+A^e}^e - H^e, W_{s+A^e}^e)_{-A^e \leq s \leq 0}$$

the process obtained by shifting the Pitman transform of X^e in order to obtain a process from $-H^e$ to 0, as well as changing the time range to $[-A^e, 0]$. By standard results on Brownian motion and random snakes, the pair obtained by concatenating $(\bar{X}^e, \bar{W}^e - \Lambda^{\bar{e}^-})$ with $(X^e, W^e - \Lambda^{e^-})$ has the law of a process distributed as $\mathbf{Quad}_{A^e, A^{\bar{e}}, H^e, \Delta^e}$.

As a result, we may assume that the elementary piece EP^e is encoded by

- the pair $(X^e, W^e - \Lambda^{e^-})$ if $e \in \vec{B}(S)$,
- the concatenation of $(\bar{X}^e, \bar{W}^e - \Lambda^{\bar{e}^-})$ with $(X^e, W^e - \Lambda^{e^-})$ if $e \in \vec{I}(S)$.

This yields a collection of elementary pieces with the proper laws and dependencies; the fact that, for $e \in \vec{I}(S)$, EP^e and $\text{EP}^{\bar{e}}$ are the same with exchanged shuttles and maximal geodesics is a simple application of the Pitman transform.

For $s \in [0, 1]$, we denote by $\pi(s)$ the projection in the gluing $\mathfrak{S}_L^{[g]}$ of the point $\langle s \rangle$ of the elementary piece $\text{EP}^{e[s]}$. We claim that $\pi: [0, 1] \rightarrow \mathfrak{S}_L^{[g]}$ is onto. Indeed, for each half-edge $e \in \vec{E}(S)$, recall that the elementary piece EP^e is defined as a quotient of $[0, A^\epsilon]$ and observe that $\{\langle s \rangle : s \text{ such that } e_{[s]} = \epsilon\} = [0, A^\epsilon]$; furthermore, the ‘‘missing point’’ A^ϵ of EP^e is glued to a point 0 of some elementary piece, which is $\pi(s)$ for some s satisfying $\langle s \rangle = 0$. Writing d_S the distance in $\mathfrak{S}_L^{[g]}$ and $d_R = \tilde{d}_W/R$, it is sufficient to show that, for $s, t \in [0, 1]$,

$$d_R(s, t) = d_S(\pi(s), \pi(t)).$$

As the pseudometric d_f defined in (4.2) is unchanged by the addition of an additive constant, setting

$$d_\epsilon = \begin{cases} d_W^\epsilon & \text{if } \epsilon \in \vec{B}(S), \\ \widehat{d}_{\vec{W}^\epsilon \bullet_W \epsilon} & \text{if } \epsilon \in \vec{I}(S), \end{cases}$$

the quantity $d_S(\boldsymbol{\pi}(s), \boldsymbol{\pi}(t))$ is the infimum of sums of the form $\sum_{i=1}^\ell d_{\epsilon_i}(s_i, t_i)$, where

- $\epsilon_1 = e_{[s]}$, $s_1 = \langle s \rangle$, $\epsilon_\ell = e_{[t]}$, $s_\ell = \langle t \rangle$,
- for all i , it holds that

$$s_i, t_i \in \begin{cases} [0, A^{\epsilon_i}] & \text{if } \epsilon_i \in \vec{B}(S), \\ [-A^{\bar{\epsilon}_i}, A^{\epsilon_i}] & \text{if } \epsilon_i \in \vec{I}(S), \end{cases}$$

- for all i ,
 - (a) either $\epsilon_i = \epsilon_{i+1} \in \vec{B}(S)$ and $d_{X^{\epsilon_i}}(t_i, s_{i+1}) = 0$,
 - (b) or $\epsilon_i = \epsilon_{i+1} \in \vec{I}(S)$ and $d_{\vec{X}^{\bar{\epsilon}_i} \bullet_X \epsilon_i}(t_i, s_{i+1}) = 0$,
 - (c) or the point t_i of EP^{ϵ_i} is glued to the point s_{i+1} of $\text{EP}^{\epsilon_{i+1}}$.

As $d_\epsilon(u, v) = \infty$ whenever $uv < 0$, we may furthermore assume that, for all i , $s_i t_i \geq 0$. Now, for each i , we set

$$\tilde{s}_i = \sum_{j=1}^{[\epsilon_i]-1} A^{e_j} + s_i \text{ if } s_i \geq 0, \quad \tilde{s}_i = \sum_{j=1}^{[\epsilon_i]} A^{e_j} + s_i \text{ if } s_i < 0,$$

where we wrote $[\epsilon]$ the index of the half-edge ϵ in the ordering e_1, \dots, e_κ of $\vec{E}(S)$. We define \tilde{t}_i similarly. It is easy to check that $\tilde{s}_1 = s$, $\tilde{t}_\ell = t$ and that, for each i , we have $d_{\epsilon_i}(s_i, t_i) = d_W(\tilde{s}_i, \tilde{t}_i)$. Furthermore, for each i , we have the following:

- (a) If $\epsilon_i = \epsilon_{i+1} \in \vec{B}(S)$ and $d_{X^{\epsilon_i}}(t_i, s_{i+1}) = 0$, then, unless $t_i = s_{i+1} = A^{\epsilon_i}$ (in which case $\tilde{t}_i = \tilde{s}_{i+1}$), it holds that $s_i < A^{\epsilon_i}$ and $t_i < A^{\epsilon_i}$, which yields that $\tilde{t}_i R \tilde{s}_{i+1}$ by (6.2a).
- (b) If $\epsilon_i = \epsilon_{i+1} \in \vec{I}(S)$ and $d_{\vec{X}^{\bar{\epsilon}_i} \bullet_X \epsilon_i}(t_i, s_{i+1}) = 0$, then
 - if $t_i s_{i+1} \geq 0$, then $\tilde{t}_i R \tilde{s}_{i+1}$ by (6.2a) as above,
 - if $t_i s_{i+1} < 0$, then $\tilde{t}_i R \tilde{s}_{i+1}$ by (6.2b).
- (c) If the point t_i of EP^{ϵ_i} is glued to the point s_{i+1} of $\text{EP}^{\epsilon_{i+1}}$, then it implies that $\tilde{d}_W(\tilde{t}_i, \tilde{s}_{i+1}) = 0$ (recall the situation depicted in Figure 3.4).

As a result, since $\tilde{d}_W \leq d_W$, it holds that

$$d_R(s, t) \leq d_S(\boldsymbol{\pi}(s), \boldsymbol{\pi}(t)).$$

The converse inequality is very similar, noting that R identifies points in the elementary piece EP^ϵ as does $d_{X^\epsilon} = 0$ or $d_{\vec{X}^{\bar{\epsilon}} \bullet_X \epsilon} = 0$, and that \tilde{d}_W encodes all the

functions d_ϵ , together with the gluings of the elementary pieces. The use of \tilde{d}_W and not d_W takes into account the gluings of shuttles with maximal geodesics of elementary pieces “overflying” the root, as, for instance, in Figure 3.4, the shuttle of $EP^{e_{14}}$ with part of the maximal geodesic of EP^{e_1} , or part of the shuttle of $EP^{e_{12}}$ with part of the maximal geodesic of EP^{e_7} . The details are left to the reader. ■

Appendix A

Technical lemmas on the Brownian plane

We now recall the definition of the Brownian plane from [40], then show that it is equivalent to the one we gave in Section 5.3, and we finally prove Proposition 5.11.

A.1 Equivalence of definitions of the Brownian plane

The original definition goes as follows. Let $(\mathfrak{X}_t, t \in \mathbb{R})$ be such that $(\mathfrak{X}_t, t \geq 0)$ and $(\mathfrak{X}_{-t}, t \geq 0)$ are two independent three-dimensional Bessel processes. Since the overall minimum of \mathfrak{X} is reached at 0, the maximum in the definition of $\tilde{d}_{\mathfrak{X}}$ – given in (6.1) – is equal to

$$\max(\underline{\mathfrak{X}}(s, t), \underline{\mathfrak{X}}(t, s)) = \begin{cases} \inf\{\mathfrak{X}_u, u \in [s \wedge t, s \vee t]\} & \text{if } st \geq 0, \\ \inf\{\mathfrak{X}_u, u \notin [s \wedge t, s \vee t]\} & \text{if } st < 0. \end{cases}$$

Next, define $(\mathfrak{W}_t, t \in \mathbb{R})$ to be a centered Gaussian process conditionally given \mathfrak{X} , with covariance function specified by

$$\mathbb{E}[|\mathfrak{W}_s - \mathfrak{W}_t|^2 \mid \mathfrak{X}] = \tilde{d}_{\mathfrak{X}}(s, t).$$

The Brownian plane was defined in [40] as

$$(\tilde{M}_{\mathfrak{X}, \mathfrak{W}}, \tilde{D}_{\mathfrak{X}, \mathfrak{W}}) = (\mathbb{R} / \{\tilde{D}_{\mathfrak{X}, \mathfrak{W}} = 0\}, \tilde{D}_{\mathfrak{X}, \mathfrak{W}}), \quad \text{where } \tilde{D}_{\mathfrak{X}, \mathfrak{W}} = d_{\mathfrak{W}} / \{\tilde{d}_{\mathfrak{X}} = 0\}.$$

The following proposition shows that the definition given in Section 5.3 is equivalent to the one above. Recall the piece of notation $\underline{X}_t = \underline{X}(0 \wedge t, 0 \vee t)$ and define the process $(\Pi_t = X_t - 2\underline{X}_t, t \in \mathbb{R})$ by taking the Pitman transform of X on $\mathbb{R}_{\geq 0}$ and on $\mathbb{R}_{\leq 0}$.

Proposition A.1. *The process (Π, W) considered under **Plane** has same distribution as $(\mathfrak{X}, \mathfrak{W})$ defined above. Moreover, as metric spaces, $(\tilde{M}_{\Pi, W}, \tilde{D}_{\Pi, W} = d_W / \{\tilde{d}_{\Pi} = 0\})$ and $(M_{X, W}, D_{X, W})$ are almost surely equal.*

Proof. We claim that $\tilde{d}_{\Pi} = d_X$. This entails that, conditionally given X , the process W is also Gaussian with $\mathbb{E}[(W_s - W_t)^2 \mid X] = \tilde{d}_{\Pi}(s, t)$ and, since Π has same distribution as \mathfrak{X} by Pitman's $2M - X$ theorem [87, Theorem 1.3], we see that (Π, W) and $(\mathfrak{X}, \mathfrak{W})$ have same distribution. We then have

$$\tilde{D}_{\Pi, W} = d_W / \{\tilde{d}_{\Pi} = 0\} = d_W / \{d_X = 0\} = D_{X, W}.$$

Checking that $\tilde{d}_\Pi = d_X$ is a classical exercise, based on the fact that, for $0 \leq s < t$,

$$\underline{\Pi}(s, t) = \underline{X}(s, t) - \underline{X}_s - \underline{X}_t \quad \text{and} \quad \inf_{u \geq s} \Pi_u = -\underline{X}_s. \quad (\text{A.1})$$

The right equation is obtained from the left one by letting $t \rightarrow \infty$, noting that, for t large enough, $\underline{X}(s, t) = \underline{X}_t$. The left equation comes from a straightforward case analysis. If $\underline{X}_s = \underline{X}_t$, then, for all $u \in [s, t]$, $\underline{X}_u = \underline{X}_s = \underline{X}_t$ and so $\Pi_u = X_u - \underline{X}_s - \underline{X}_t$; taking the infimum on $u \in [s, t]$ gives the result. If $\underline{X}_s > \underline{X}_t$, then $\underline{X}(s, t) = \underline{X}_t$ so the right-hand side is $-\underline{X}_s$. Let $r \in [s, t]$ be such that $X_r = \underline{X}_r = \underline{X}_s$. We have

$$\Pi_r = X_r - 2\underline{X}_r = -\underline{X}_s$$

and, for $u \geq s$,

$$\Pi_u = X_u - 2\underline{X}_u \geq -\underline{X}_u \geq -\underline{X}_s.$$

For $0 \leq s < t$, the left equation of (A.1) entails

$$\begin{aligned} \tilde{d}_\Pi(t, s) &= \Pi_s + \Pi_t - 2\underline{\Pi}(s, t) \\ &= X_s - 2\underline{X}_s + X_t - 2\underline{X}_t - 2(\underline{X}(s, t) - \underline{X}_s - \underline{X}_t) = d_X(s, t). \end{aligned}$$

For $s < 0 < t$, we have that

$$\begin{aligned} \underline{\Pi}(t, s) &= \inf_{u \geq t} \Pi_u \wedge \inf_{u \leq s} \Pi_u = -(\underline{X}(0, t) \vee \underline{X}(s, 0)) \\ &= \underline{X}(s, t) - \underline{X}(s, 0) - \underline{X}(0, t) \end{aligned}$$

and we conclude as above. The remaining case $s < t < 0$ is treated similarly. ■

A.2 Convergence of the uniform infinite planar quadrangulation to the Brownian plane

We use here the setting of Section 5.4. The proof of Proposition 5.11 will follow similar lines as that of Proposition 4.9, using the coupling results of [40]. As the law of \mathcal{X} is obtained from that of X by taking the Pitman transform on $\mathbb{R}_{\geq 0}$ and on $\mathbb{R}_{\leq 0}$, the same should be done for the contour process C of the tree \mathbf{T}_∞ . We thus define the process $(\mathfrak{C}(t) = C(t) - 2\underline{C}(t), t \in \mathbb{R})$.

Note that this gives an alternate natural contour process since, for $i \in \mathbb{Z}$, it holds that

$$\mathfrak{C}(i) = d_{\mathbf{T}^{\Upsilon(i)}}(v_i, \rho^{|\Upsilon(i)|}) + |\Upsilon(i)| = d_{\mathbf{T}_\infty}(v_i, \rho^0).$$

In this setting of discrete trees, the Pitman transform on the contour process is very visual: it merely consists of going from reading the trees while moving *down* between trees to reading them while moving *up* between trees; see Figure A.1 for an illustration. We may now proceed to the proof of the convergence of the UIPQ to the Brownian plane.

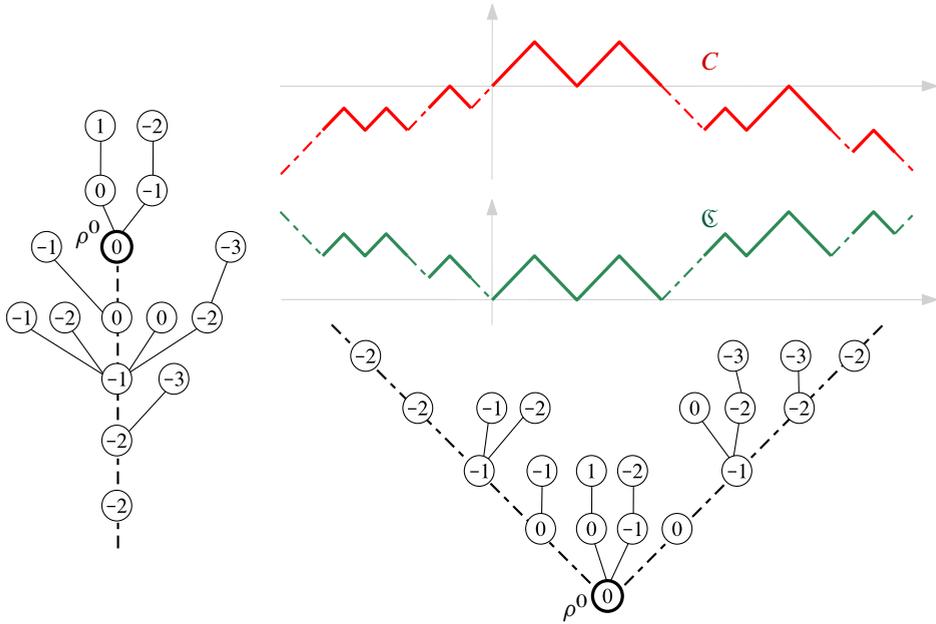


Figure A.1. *Left.* Representation of the infinite tree from Figure 5.2 after moving the trees in such a way that, for $k \geq 0$, ρ^k is located at $(0, -k)$ with \mathbf{T}^k grafted on its right and \mathbf{T}^{-k} on its left. *Top right.* Taking the Pitman transform of the contour process on $\mathbb{R}_{\geq 0}$ and on $\mathbb{R}_{\leq 0}$ yields an alternate contour process. The differing parts are represented with dot-dashed lines; they correspond to the edges of the infinite spine of the tree. Visually, the process C records the height of a particle moving at speed one around the tree when represented as on the left. *Bottom right.* Representation of the tree from Figure 5.2 where the root of \mathbf{T}^k is now located at $(k, |k|)$ for each $k \in \mathbb{Z}$. Note that, in this representation, the roots of \mathbf{T}^{-k} and \mathbf{T}^k differ so that the spine is duplicated. The process \mathbb{C} records the height of a particle moving at speed one around this bi-infinite tree.

Proof of Proposition 5.11. Similarly to the proof of Proposition 4.9, we fix some number $K > 0$ and will sample a large plane quadrangulation such that its properly scaled version and its limit, the Brownian sphere, are indistinguishable from the rescaled UIPQ and the Brownian plane, in a neighborhood of 0 of amplitude K . We use again a superscript prime symbol $'$ for the objects related to the plane quadrangulation and its limit. Here, some care will also be needed when taking an inverse Pitman transform, since this operation a priori involves more than just a neighborhood of 0.

We fix $L > 0$ and $n \geq 1$, and consider a uniform random element (M'_n, λ'_n) of $\vec{\mathbf{M}}_{a_n, \emptyset}^{[0]}$, where $a_n = \lfloor nL \rfloor$, that is, M'_n is a uniform rooted plane tree with a_n edges, which we view as a map with a unique face f_* , and λ'_n is a labeling function uniformly distributed among those yielding a well-labeled tree.

We let $(\mathfrak{C}'_n, \mathfrak{L}'_n)$ be the contour and label function of this tree, we let

$$Q'_n = \text{CVS}(M'_n, \lambda'_n; f_*)$$

be the quadrangulation encoded by (M'_n, λ'_n) , and we set

$$\mathfrak{D}'_n(i, j) = d_{Q'_n}(v_i, v_j)$$

for $0 \leq i, j \leq 2a_n$, where v_i is the i -th visited vertex in M'_n in contour order, starting from the root corner, and viewed as a vertex of Q'_n . We extend \mathfrak{D}'_n into a continuous function on $[0, 2a_n]^2$ by bilinearity, and all processes $\mathfrak{C}'_n, \mathfrak{L}'_n, \mathfrak{D}'_n$ to $[-2a_n, 2a_n]$ by the same formulas as (4.15) and (4.16) but with $l_n = 0$. We also define the rescaled versions $\mathfrak{C}'_{(n)}, \mathfrak{L}'_{(n)}, \mathfrak{D}'_{(n)}$ exactly as in (4.17). The joint convergence

$$(\mathfrak{C}'_{(n)}, \mathfrak{L}'_{(n)}, \mathfrak{D}'_{(n)}) \xrightarrow[n \rightarrow \infty]{(d)} (\mathfrak{X}', \mathfrak{W}', \mathfrak{D}')$$

on the space $\mathcal{C}([-L, L]) \times \mathcal{C}([-L, L]) \times \mathcal{C}([-L, L]^2)$ is then a consequence of [80, Theorem 3], where the limit is as follows. Restricted to $[0, L]$, the process \mathfrak{X}' is a Brownian excursion of duration L and \mathfrak{W}' is the random snake driven by \mathfrak{X}' , while \mathfrak{D}' is a random pseudometric, which is an explicit function of $(\mathfrak{X}', \mathfrak{W}')$. Moreover, all these processes are extended to $[-L, L]$ by a simple translation of their argument by L .

Let us now recall the relevant aspects of the coupling results of [40], between the pairs $(\mathfrak{X}, \mathfrak{W})$ and $(\mathfrak{X}', \mathfrak{W}')$. It will be convenient to let

$$\bar{\mathfrak{T}}_x = \inf\{t \leq 0 : \mathfrak{X}_t = x\}, \quad \mathfrak{T}_x = \sup\{t \geq 0 : \mathfrak{X}_t = x\}.$$

Fix $r > 0$ and $\varepsilon > 0$. Then by [40, Lemmas 5 and 6], it is possible to find $A > 1$ and then $\alpha > 0$ and $L_0 > 0$ large, such that for $L > L_0$ the two processes $(\mathfrak{X}, \mathfrak{W})$ and $(\mathfrak{X}', \mathfrak{W}')$ can be coupled in such a way that on some event \mathcal{F} of probability $\mathbb{P}(\mathcal{F}) \geq 1 - \varepsilon$, the following properties hold:

- For every $s, t \in [-\alpha, \alpha]$, one has

$$\mathfrak{X}_t = \mathfrak{X}'_t, \quad \mathfrak{W}_t = \mathfrak{W}'_t. \tag{A.2}$$

- It holds that

$$-\alpha < \bar{\mathfrak{T}}_{A^4} \quad \text{and} \quad \mathfrak{T}_{A^4} < \alpha. \tag{A.3}$$

- For every $s, t \in [\bar{\mathfrak{T}}_A, \mathfrak{T}_A]$, the two conditions $\max(\bar{D}_{\mathfrak{X}, \mathfrak{W}}(0, t), \bar{D}_{\mathfrak{X}, \mathfrak{W}}(0, s)) \leq r$ and $\max(\mathfrak{D}'(0, t), \mathfrak{D}'(0, s)) \leq r$ are equivalent, and, if these are satisfied, one has

$$\mathfrak{D}'(s, t) = \bar{D}_{\mathfrak{X}, \mathfrak{W}}(s, t).$$

This choice of coupling being fixed, let us now define $(X_t, t \in \mathbb{R})$ as the unique process such that $\mathfrak{X}_t = X_t - 2\underline{X}_t$ is the Pitman transform of X on $\mathbb{R}_{\geq 0}$ and on $\mathbb{R}_{\leq 0}$; more explicitly,

$$X_t = \begin{cases} \mathfrak{X}_t - 2 \inf_{s \geq t} \mathfrak{X}_s & \text{if } t \geq 0, \\ \mathfrak{X}_t - 2 \inf_{s \leq t} \mathfrak{X}_s & \text{if } t < 0. \end{cases}$$

Let us also define $W = \mathfrak{W}$. Then X indeed has the law of a two-sided Brownian motion, and W is the random snake driven by X , so that (X, W) has law **Plane**. Note that, in this particular coupling, we have $\bar{T}_x = \bar{\mathfrak{T}}_x$ and $T_x = \mathfrak{T}_x$ for every $x \geq 0$, and also $D_{X,W} = \bar{D}_{\mathfrak{X},\mathfrak{W}}$, by the observation in the proof of Proposition A.1. Moreover, on the event \mathcal{F} , the restriction $X|_{[\bar{T}_{A^2}, T_{A^2}]}$ is actually a function of $\mathfrak{X}'|_{[-\alpha, \alpha]}$. Indeed, by (A.2) and (A.3),

$$X_t = \begin{cases} \mathfrak{X}'_t - 2 \inf_{t \leq s \leq \alpha} \mathfrak{X}'_s & \text{if } 0 \leq t \leq T_{A^2}, \\ \mathfrak{X}'_t - 2 \inf_{-\alpha \leq s \leq t} \mathfrak{X}'_s & \text{if } \bar{T}_{A^2} \leq t < 0 \end{cases}$$

since, for $0 \leq t \leq T_{A^2}$, one has $\underline{\mathfrak{X}}(t, \infty) = \underline{\mathfrak{X}}(t, T_{A^2}) = \underline{\mathfrak{X}}(t, \alpha)$, and similarly in negative times.

By choosing appropriately the values of r , and enlarging the values of A and α if necessary, then, similarly to the proof of Proposition 4.9, we obtain that (A.2) holds on $[-K, K]$, and that the restrictions to $[-K, K]^2$ of \mathfrak{D}' and $D_{X,W}$ coincide with probability at least $1 - \varepsilon$.

Next, keeping K, ε fixed, and possibly up to choosing L even larger, we need to couple the processes $(C(n), \Lambda(n), D(n))$ and $(\mathfrak{C}'_{(n)}, \mathfrak{L}'_{(n)}, \mathfrak{D}'_{(n)})$ appropriately. To this end, we use the techniques of [40, Proposition 9]. The latter states that for $\varepsilon > 0$, there exists $\alpha > 0$ (independent of the choice of L arising in the definition of the scaling constant a_n) such that for every n large enough, one may couple the quadrangulations Q'_n and Q_∞ in such a way that, with probability at least $1 - \varepsilon$, the balls of radius $\alpha a_n^{1/4}$ around the root of Q'_n and Q_∞ are isometric. The proof proceeds by coupling the encoding labeled trees (M'_n, λ'_n) and $(\mathbf{T}_\infty, \lambda_\infty)$ in such a way that, with even larger probability, the first $\lfloor \delta a_n^{1/2} \rfloor$ generations of these trees coincide for some $\delta > 0$, and the minimal value of λ_∞ taken on the vertices $\rho^0, \rho^1, \dots, \rho^{\lfloor \delta a_n^{1/2} \rfloor}$ of \mathbf{T}_∞ is less than $-4\alpha a_n^{1/4}$. By choosing R and then L large enough in the first place, for our choice of K , we may also require that with probability at least $1 - \varepsilon$,

- the contour and label processes $\mathfrak{C}'_n, \mathfrak{L}'_n$ of (M'_n, λ'_n) and $\mathfrak{C}, \mathfrak{L}$ of $(\mathbf{T}_\infty, \lambda_\infty)$ on the interval $[-2nK, 2nK]$ involve only vertices of generations less than $\lfloor Rn^{1/2} \rfloor$, and
- the most recent common ancestor of the vertices at generation $\lfloor \delta a_n^{1/2} \rfloor$ has generation at least $\lfloor Rn^{1/2} \rfloor$.

In particular, on this event, the restriction of the process C'_n to $[-2nK, 2nK]$ is equal to the restriction of the process C on this same event – in words, the second itemized event means that the spine of \mathbf{T}_∞ is determined by the data of M'_n up to generation $Rn^{1/2}$. Since the process C is the inverse Pitman transform of \mathfrak{C} , it is then a simple exercise to conclude that $(C'_{(n)}, \Lambda'_{(n)}, D'_{(n)})$, which coincides with $(C_{(n)}, \Lambda_{(n)}, D_{(n)})$ on $[-K, K]$ with high probability, converges to some (X', W', D') , which coincides with $(X, W, D_{X,W})$ on $[-K, K]$ with high probability. ■

Appendix B

Scaling limit of size parameters in labeled maps

B.1 Preliminaries

In this appendix, we prove Proposition 3.10, following the method of [19, Proposition 7] and [21, Proposition 7]. In the meantime, we obtain an asymptotic enumeration result for $\vec{\mathcal{Q}}_{n, l_n}^{[g]}$ in Proposition B.1 below, which will also allow us to deduce Theorem 1.5 and Corollary 1.6 from Theorem 1.1.

Recall that $(g, k) \notin \{(0, 0), (0, 1)\}$, that $\mathbf{L} = (L^1, \dots, L^k)$ is a fixed k -tuple such that $L^1, \dots, L^b > 0$, while $L^{b+1}, \dots, L^k = 0$, and that we consider a fixed sequence of k -tuples $\mathbf{l}_n = (l_n^1, \dots, l_n^k) \in (\mathbb{Z}_{\geq 0})^k$, $n \geq 1$, such that $l_n^i / \sqrt{2n} \rightarrow L^i$ as $n \rightarrow \infty$ for $1 \leq i \leq k$.

We furthermore assume that n is sufficiently large so that $l_n^i \geq 1$ for each $i \leq b$. We denote by $\vec{\mathcal{S}}$ the set of rooted genus g schemes with k holes, such that h_1, \dots, h_b are faces. Note that our assumption on n ensures that $S_n \in \vec{\mathcal{S}}$.

“Free” parameters and notation. For every scheme $\mathbf{s} \in \vec{\mathcal{S}}$, not necessarily dominant, we arbitrarily fix, once and for all, half-edges $\epsilon_0 \in \vec{I}(\mathbf{s})$ and $\epsilon_i \in \vec{B}_i(\mathbf{s})$ for $1 \leq i \leq b$. We fix an orientation $I(\mathbf{s})$ of $\vec{I}(\mathbf{s})$ that contains ϵ_0 , and we set $I'(\mathbf{s}) = I(\mathbf{s}) \setminus \{\epsilon_0\}$. We also let v_0 be the root vertex of \mathbf{s} , and $V'(\mathbf{s}) = V(\mathbf{s}) \setminus \{v_0\}$, as in Section 3.6. Finally, we set

$$\begin{aligned} \vec{B}_0(\mathbf{s}) &= \bigsqcup_{b+1 \leq i \leq k} \vec{B}_i(\mathbf{s}), & \vec{B}_+(\mathbf{s}) &= \bigsqcup_{1 \leq i \leq b} \vec{B}_i(\mathbf{s}), \\ \vec{B}'_i(\mathbf{s}) &= \vec{B}_i(\mathbf{s}) \setminus \{\epsilon_i\} \text{ for } 1 \leq i \leq b, & \vec{B}'_+(\mathbf{s}) &= \bigsqcup_{1 \leq i \leq b} \vec{B}'_i(\mathbf{s}). \end{aligned}$$

The motivation for introducing \vec{B}_+ and \vec{B}_0 is that we need a different treatment depending on whether the hole perimeters are in the scale \sqrt{n} or $o(\sqrt{n})$. The sets with a prime symbol should be thought of as the sets containing the parameters on which there is a “degree of freedom.” (*The reason for removing one element from I will become clear in a moment. We will not need a \vec{B}'_0 since the corresponding perimeters are all asymptotically null in the scale \sqrt{n} of interest.*)

From now on, we use the shorthand piece of notation $\mathbf{x}^{\mathcal{E}}$ for a family $(x^j)_{j \in \mathcal{E}}$ indexed by a set \mathcal{E} . For any subset $\mathcal{F} \subseteq \mathcal{E}$, we also denote by $\mathbf{x}^{\mathcal{F}} = (x^j)_{j \in \mathcal{F}}$ the subfamily indexed by \mathcal{F} , and, in the case of real nonnegative numbers, by $\|\mathbf{x}\|_{\mathcal{F}} = \sum_{j \in \mathcal{F}} x^j$ (note in particular that $\|\mathbf{x}\|_{\emptyset} = 0$).

Counting scheme-rooted labeled maps with given size parameters. For the time being, we do not take the areas parameters into account. We fix a rooted scheme $\mathbf{s} \in \vec{\mathcal{S}}$,

and size parameters $\mathbf{h}^{\vec{I}(\mathbf{s})}$, $\mathbf{l}^{\vec{B}(\mathbf{s})}$ and $\lambda^{V(\mathbf{s})}$. We say that a labeled map is *scheme-rooted on \mathbf{s}* if its scheme carries an extra root and the scheme rooted at this extra root is \mathbf{s} . We consider the elements of $\mathbf{M}_{n,I}^{[g]}$ scheme-rooted on \mathbf{s} whose size parameters are $\mathbf{h}^{\vec{I}(\mathbf{s})}$, $\mathbf{l}^{\vec{B}(\mathbf{s})}$ and $\lambda^{V(\mathbf{s})}$. Reasoning as in Propositions 2.3 and 2.5, we can express the number of such elements as

$$12^{n-\|\mathbf{h}\|/2} 2^{\|\mathbf{h}\|+\|I\|} Q_{\|\mathbf{h}\|+\|I\|}(2n+\|I\|) \prod_{e \in I(\mathbf{s})} 3^{h^e} M_{h^e}(\delta\lambda^e) \prod_{e \in \vec{B}(\mathbf{s})} 2^{2l^e+\delta\lambda^e} P_{l^e}(\delta\lambda^e),$$

the products over $I(\mathbf{s})$ and $\vec{B}(\mathbf{s})$ respectively counting the number of ways to label the vertices along the edges of $I(\mathbf{s})$ and $\vec{B}(\mathbf{s})$, and the remaining term counting the labeled forests, which can be seen as one big labeled forest obtained by concatenating all the labeled forests indexed by the half-edges of $\vec{E}(\mathbf{s})$. After recalling that $\sum_{e \in \vec{B}_i(\mathbf{s})} \delta\lambda^e = 0$ and that $\|I\|_{\vec{B}_i(\mathbf{s})} = l^i$ for every $i \in \{1, 2, \dots, k\}$ corresponding to an external face of \mathbf{s} , we may recast this quantity as

$$12^n 8^{\|I\|} Q_{\|\mathbf{h}\|+\|I\|}(2n+\|I\|) \prod_{e \in I(\mathbf{s})} M_{h^e}(\delta\lambda^e) \prod_{e \in \vec{B}(\mathbf{s})} P_{l^e}(\delta\lambda^e).$$

Consequently, the number of elements of $\vec{\mathbf{M}}_{n,I}^{[g]}$ scheme-rooted on \mathbf{s} (these labeled maps are thus rooted twice) whose size parameters are $\mathbf{h}^{\vec{I}(\mathbf{s})}$, $\mathbf{l}^{\vec{B}(\mathbf{s})}$ and $\lambda^{V(\mathbf{s})}$ is equal to

$$\begin{aligned} \vec{\mathcal{S}}_n^{\mathbf{s}}(\mathbf{h}, \mathbf{l}, \lambda) &= (2n+\|I\|) 12^n 8^{\|I\|} Q_{\|\mathbf{h}\|+\|I\|}(2n+\|I\|) \\ &\quad \times \prod_{e \in I(\mathbf{s})} M_{h^e}(\delta\lambda^e) \prod_{e \in \vec{B}(\mathbf{s})} P_{l^e}(\delta\lambda^e), \end{aligned} \quad (\text{B.1})$$

since there are $2n+\|I\|$ possible rootings of the map.

Counting rooted labeled maps. Next, for $n, h \in \mathbb{N}$, $\mathbf{s} \in \vec{\mathcal{S}}$, we set

$$\mathcal{Z}_1^{\mathbf{s}}(h, n) = \sum_{\mathcal{T}_{\mathbf{s}}(h, n)} \prod_{e \in I(\mathbf{s})} M_{h^e}(\delta\lambda^e) \prod_{e \in \vec{B}(\mathbf{s})} P_{l^e}(\delta\lambda^e), \quad (\text{B.2})$$

where the sum is taken over the set $\mathcal{T}_{\mathbf{s}}(h, n)$ of all size parameters from labeled maps in $\mathbf{M}_{n, I_n}^{[g]}$ scheme-rooted on \mathbf{s} , having h edges in total on the internal edges of \mathbf{s} . More precisely, it is the set of tuples

$$(\mathbf{h}^{\vec{I}(\mathbf{s})}, \mathbf{l}^{\vec{B}(\mathbf{s})}, \lambda^{V(\mathbf{s})}) \in \mathbb{N}^{\vec{I}(\mathbf{s})} \times \mathbb{N}^{\vec{B}(\mathbf{s})} \times \mathbb{Z}^{V(\mathbf{s})}$$

such that $\|\mathbf{h}\| = 2h$, $\|I\|_{\vec{B}_i(\mathbf{s})} = l_n^i$ for $1 \leq i \leq k$, $h^{\bar{e}} = h^e$ for all $e \in \vec{I}(\mathbf{s})$, $\lambda^{v_0} = 0$. Note that the conditions

$$\|I\|_{\vec{B}_i(\mathbf{s})} = l_n^i$$

may only be satisfied if I_n is *compatible* with \mathbf{s} in the sense that $l_n^i = 0 \Leftrightarrow \vec{B}_i(\mathbf{s}) = \emptyset$ for all i . As a result, $\mathcal{T}_{\mathbf{s}}(h, n) = \emptyset$ and thus $\mathcal{Z}_1^{\mathbf{s}}(h, n) = 0$ whenever I_n is not compatible with \mathbf{s} .

By double counting the elements of $\vec{\mathbf{M}}_{n, I_n}^{[g]}$ scheme-rooted on \mathbf{s} , we thus obtain that

$$|\vec{\mathbf{M}}_{n, I_n}^{[g]}| = 12^n 8^{\|I_n\|} \sum_{\mathbf{s} \in \vec{\mathbf{S}}} \frac{2n + \|I_n\|}{2|E(\mathbf{s})|} \sum_{h \in \mathbb{N}} Q_{2h + \|I_n\|} (2n + \|I_n\|) \mathcal{Z}_1^{\mathbf{s}}(h, n), \quad (\text{B.3})$$

since we sum over $\bigcup_{\mathbf{s} \in \vec{\mathbf{S}}, h \in \mathbb{N}} \{\mathbf{s}\} \times \mathcal{T}_{\mathbf{s}}(h, n)$ the number $\vec{\mathbf{S}}_n^{\mathbf{s}}(\mathbf{h}, I_n, \boldsymbol{\lambda})$ given by (B.1), divided by the number $2|E(\mathbf{s})|$ of possible extra rootings on the scheme.

B.2 Asymptotics of the scheme

When we work with a fixed scheme, which will be the case in all but the fourth paragraph, we drop the argument from the sets in the notation in order to ease the reading, thus writing I instead of $I(\mathbf{s})$, for instance.

Law of the scheme. Recall that the triple (M_n, λ_n, S_n) is a rooted, scheme-rooted, labeled map, where (M_n, λ_n) is uniformly distributed over $\vec{\mathbf{M}}_{n, I_n}^{[g]}$, while, conditionally given it, S_n is rooted by uniformly choosing its root among $\{e, \bar{e} : e \in \vec{I}(S_n) \cup \vec{B}_+(S_n)\}$. Let us fix a rooted scheme $\mathbf{s} \in \vec{\mathbf{S}}$ whose root or its reverse belongs to $\vec{I} \cup \vec{B}_+$. Writing $\underline{\mathbf{s}}$ the nonrooted scheme corresponding to \mathbf{s} , observe that the set of rooted labeled maps in $\vec{\mathbf{M}}_{n, I_n}^{[g]}$ with scheme $\underline{\mathbf{s}}$ is in bijection with the set of rooted labeled maps in $\vec{\mathbf{M}}_{n, I_n}^{[g]}$ scheme-rooted on \mathbf{s} . Then,

$$\begin{aligned} \mathbb{P}(S_n = \mathbf{s}) &= \sum_{\substack{(\mathbf{m}, \lambda) \in \vec{\mathbf{M}}_{n, I_n}^{[g]} \\ \text{with scheme } \underline{\mathbf{s}}}} \mathbb{P}((M_n, \lambda_n) = (\mathbf{m}, \lambda), S_n = \mathbf{s}) \\ &= \sum_{\substack{(\mathbf{m}, \lambda) \in \vec{\mathbf{M}}_{n, I_n}^{[g]} \\ \text{scheme-rooted on } \mathbf{s}}} \mathbb{P}((M_n, \lambda_n) = (\mathbf{m}, \lambda)) \mathbb{P}(S_n = \mathbf{s} \mid (M_n, \lambda_n) = (\mathbf{m}, \lambda)) \\ &= \sum_{h \in \mathbb{N}} \sum_{\mathcal{T}_{\underline{\mathbf{s}}}(h, n)} \vec{\mathbf{S}}_n^{\mathbf{s}}(\mathbf{h}, I_n, \boldsymbol{\lambda}) \frac{1}{|\vec{\mathbf{M}}_{n, I_n}^{[g]}|} \frac{1}{|\vec{I}| + 2|\vec{B}_+|} = \frac{\mathcal{Z}_1^{\mathbf{s}}(n)}{\mathcal{Z}_1(n)}, \end{aligned}$$

where

$$\mathcal{Z}_1^{\mathbf{s}}(n) = \frac{1}{|\vec{I}(\mathbf{s})| + 2|\vec{B}_+(\mathbf{s})|} \sum_{h \in \mathbb{N}} Q_{2h + \|I_n\|} (2n + \|I_n\|) \mathcal{Z}_1^{\mathbf{s}}(h, n) \quad (\text{B.4})$$

and $\mathcal{Z}_1(n) = \sum_{\mathbf{s} \in \vec{\mathbf{S}}} \mathcal{Z}_1^{\mathbf{s}}(n)$ is the proper normalization constant.

Schemes with tadpoles. Here, we fix a scheme \mathbf{s} whose external faces among h_i , $b + 1 \leq i \leq k$, are all tadpoles. Equivalently, each \vec{B}_i , $b + 1 \leq i \leq k$, is either empty or a singleton. In this case, by the Euler characteristic formula,

$$|V'| - |I| - |\vec{B}'_+| = -2g,$$

since \mathbf{s} has $|I| + |\vec{B}'_+| + b + |\vec{B}_0|$ edges and $1 + b + |\vec{B}_0|$ faces.

Assuming that I_n is compatible with \mathbf{s} , we write the sum over $\mathcal{T}_s(h, n)$ in (B.2) as an integral under the Lebesgue measure

$$dL_s = d\mathbf{h}^{I'} \otimes d\mathbf{l}^{\vec{B}'_+} \otimes d\lambda^{V'} \quad \text{over } (\mathbb{R}_{\geq 0})^{I'} \times (\mathbb{R}_{\geq 0})^{\vec{B}'_+} \times \mathbb{R}^{V'}$$

and obtain

$$\mathcal{Z}_1^s(h, n) = \prod_{e \in \vec{B}_0} P_{l^e}(0) \int dL_s \prod_{e \in I} M_{\underline{h}^e}(\delta \underline{\lambda}^e) \prod_{e \in \vec{B}_+} P_{\underline{l}^e}(\delta \underline{\lambda}^e),$$

where $l^e = l_n^i$ if e is the unique element of \vec{B}_i for $i > b$, and

$$\begin{aligned} \underline{h}^e &= \lceil h^e \rceil \text{ for } e \in I', & \underline{h}^{\epsilon_0} &= h - \sum_{e \in I'} \underline{h}^e, \\ \underline{l}^e &= \lceil l^e \rceil \text{ for } e \in \vec{B}'_+, & \underline{l}^{\epsilon_i} &= l_n^i - \sum_{e \in \vec{B}'_i} \underline{l}^e \text{ for } 1 \leq i \leq b, \\ \underline{\lambda}^v &= \lceil \lambda^v \rceil \text{ for } v \in V', & \underline{\lambda}^{v_0} &= 0. \end{aligned}$$

(Note that the ceiling function is superfluous for integer parameters; we kept it for notational simplicity.) In order to deal with the cases where $\underline{h}^{\epsilon_0} \leq 0$ or $\underline{l}^{\epsilon_i} \leq 0$, we simply declare¹

$$M_\ell(j) = P_\ell(j) = 0$$

whenever $\ell \leq 0$.

Observe that I_n compatible with \mathbf{s} means that \vec{B}_0 corresponds to $\{i > b : l_n^i \geq 1\}$. We then make the changes of variables in the natural scales to obtain

$$\begin{aligned} \mathcal{Z}_1^s(h, n) &= 3^{b/2-g} 2^{|V'|/2-g/2-3b/4} n^{|V'|/2+g/2-b/4-1/2} \prod_{i>b:l_n^i \geq 1} P_{l_n^i}(0) \\ &\times \int dL_s \prod_{e \in I} \left(\frac{8n}{9}\right)^{1/4} M_{\underline{h}^e}(\delta \underline{\lambda}^e) \prod_{e \in \vec{B}_+} \left(\frac{8n}{9}\right)^{1/4} P_{\underline{l}^e}(\delta \underline{\lambda}^e), \end{aligned} \quad (\text{B.5})$$

where

$$\begin{aligned} \underline{\underline{h}}^e &= \lceil \sqrt{2nh^e} \rceil \text{ for } e \in I', & \underline{\underline{h}}^{\epsilon_0} &= h - \sum_{e \in I'} \underline{\underline{h}}^e, \\ \underline{\underline{l}}^e &= \lceil \sqrt{2nl^e} \rceil \text{ for } e \in \vec{B}'_+, & \underline{\underline{l}}^{\epsilon_i} &= l_n^i - \sum_{e \in \vec{B}'_i} \underline{\underline{l}}^e \text{ for } 1 \leq i \leq b, \\ \underline{\underline{\lambda}}^v &= \left[\left(\frac{8n}{9}\right)^{1/4} \lambda^v \right] \text{ for } v \in V', & \underline{\underline{\lambda}}^{v_0} &= 0. \end{aligned}$$

¹This is just a convenience. Note that we set $M_0(0) = P_0(0) = 0$ here, although it would be more natural from a combinatorial point of view to set both these quantities to 1.

We finally use the same method to treat the summation over $h \in \mathbb{N}$ in (B.4), that is, we see it as an integral and do the proper change of variables. We write $\mathbf{l}_n \bowtie \mathbf{s}$ to mean “ \mathbf{l}_n compatible with \mathbf{s} ”:

$$\mathcal{Z}_1^{\mathbf{s}}(n) = \frac{\mathbf{1}_{\mathbf{l}_n \bowtie \mathbf{s}}}{|\vec{I}| + |\vec{B}_+|} \sqrt{\frac{2}{n}} \int_{\mathbb{R}_{\geq 0}} dh n Q_{2\lceil\sqrt{2nh}\rceil + \|\mathbf{l}_n\|} (2n + \|\mathbf{l}_n\|) \mathcal{Z}_1^{\mathbf{s}}(\lceil\sqrt{2nh}\rceil, n). \tag{B.6}$$

Setting $h^{\epsilon_0} = h - \sum_{e \in I'} h^e$, $l^{\epsilon_i} = L^i - \sum_{e \in \vec{B}_i} l^e$ for $1 \leq i \leq b$, and $\lambda^{v_0} = 0$, by the local limit theorem [86, Theorem VII.1.6], it holds that, when $h \sim \sqrt{2nh}$,

$$\left(\frac{8n}{9}\right)^{1/4} M_{\underline{h}^e}(\delta \underline{\lambda}^e) \xrightarrow{n \rightarrow \infty} p_{h^e}(\delta \lambda^e), \quad \left(\frac{8n}{9}\right)^{1/4} P_{\underline{l}^e}(\delta \underline{\lambda}^e) \xrightarrow{n \rightarrow \infty} p_{3l^e}(\delta \lambda^e), \tag{B.7}$$

and

$$n Q_{2\lceil\sqrt{2nh}\rceil + \|\mathbf{l}_n\|} (2n + \|\mathbf{l}_n\|) \xrightarrow{n \rightarrow \infty} q_{2h + \|\mathbf{L}\|}(1). \tag{B.8}$$

Consequently, provided the domination hypothesis obtained in the following paragraph, we get the following equivalent:

$$\mathcal{Z}_1^{\mathbf{s}}(n) \underset{n \rightarrow \infty}{\sim} c_1^{\mathbf{s}}(\mathbf{L}) \mathbf{1}_{\mathbf{l}_n \bowtie \mathbf{s}} n^{|V|/2 + g/2 - b/4 - 1} \prod_{i > b: l_i^j \geq 1} P_{l_i^j}(0), \tag{B.9}$$

where the constant $c_1^{\mathbf{s}}(\mathbf{L})$ is given by

$$\begin{aligned} c_1^{\mathbf{s}}(\mathbf{L}) &= \frac{1}{|\vec{I}| + 2|\vec{B}_+|} 3^{b/2 - g} 2^{|V|/2 - g/2 - 3b/4 + 1/2} \\ &\times \int_{\mathbb{R}_{\geq 0}} dh q_{2h + \|\mathbf{L}\|}(1) \int dL_{\mathbf{s}} \prod_{e \in I} p_{h^e}(\delta \lambda^e) \prod_{e \in \vec{B}_+} p_{3l^e}(\delta \lambda^e). \end{aligned}$$

Domination hypothesis. In order to show that the convergence is dominated, we use the bounds of Petrov [86, Theorem VII.3.16], stating that there exists a constant C such that, for any $\ell \in \mathbb{N}$, $j \in \mathbb{Z}$, $i \in \mathbb{N}$, and $r \in \mathbb{N}$,

$$M_{\ell}(j) \vee P_{\ell}(j) \leq C \frac{1}{\sqrt{\ell}} \quad \text{and} \quad Q_i(\ell) \leq C \frac{i}{\ell^{3/2}} \frac{1}{1 + (i^2/\ell)^r}. \tag{B.10}$$

We fix an arbitrary spanning tree of \mathbf{s} , that is, a tree with vertex-set V and edge-set a subset of E . We associate with any vertex $v \neq v_0$ the first edge of the unique path in the tree from v to v_0 and we denote by e_v the unique half-edge of $I \cup \vec{B}$ that corresponds to this edge.

We bound the integrand in (B.5) as follows. First, by (B.10), we have, for $e \in I'$,

$$\left(\frac{8n}{9}\right)^{1/4} M_{\underline{h}^e}(\delta \underline{\lambda}^e) \leq \left(\frac{8n}{9}\right)^{1/4} \frac{C}{\sqrt{\underline{h}^e}} \leq \left(\frac{8n}{9}\right)^{1/4} \frac{C}{\sqrt{\sqrt{2nh^e}}} = \sqrt{\frac{2}{3}} \frac{C}{\sqrt{h^e}} \leq \frac{C}{\sqrt{h^e}}.$$

For $h = \lceil \sqrt{2nh} \rceil$ and $h^{\epsilon_0} = h - \sum_{e \in I'} h^e$, a similar bound holds for $e = \epsilon_0$, up to possibly enlarging the constant. Indeed, it suffices to show that $\underline{h}^{\epsilon_0}$ is bounded from below by a constant times $\sqrt{2nh^{\epsilon_0}}$ in order to complete the computation. We may assume that $\underline{h}^{\epsilon_0} \geq 1$ as otherwise the left-hand side is null. Then, if $\sqrt{2nh^{\epsilon_0}} \leq 2|I'|$, it immediately holds that $\underline{h}^{\epsilon_0} \geq \frac{1}{2|I'|} \sqrt{2nh^{\epsilon_0}}$. Otherwise,

$$\underline{h}^{\epsilon_0} = h - \sum_{e \in I'} \underline{h}^e \geq \sqrt{2nh^{\epsilon_0}} - |I'| \geq \frac{1}{2} \sqrt{2nh^{\epsilon_0}}.$$

In conclusion, up to changing the constant C , it holds that, for all $e \in I$,

$$\left(\frac{8n}{9}\right)^{1/4} M_{\underline{h}^e}(\delta \underline{\lambda}^e) \leq \frac{C}{\sqrt{h^e}}.$$

Similarly, up to enlarging the constant C even more, setting $l^{\epsilon_i} = L^i - \sum_{e \in \vec{B}_i} l^e$ for $1 \leq i \leq b$, it holds that, for $e \in \vec{B}_+$,

$$\left(\frac{8n}{9}\right)^{1/4} P_{\underline{l}^e}(\delta \underline{l}^e) \leq \frac{C}{\sqrt{l^e}}.$$

We use these bounds whenever $e \notin \vec{E}_V = \{e_v : v \in V \setminus \{v_0\}\}$ and then, we operate the integral with respect to $d\lambda^{V'}$ vertices by vertices, starting from a leaf of the fixed spanning tree, then from a leaf of the tree remaining after removing the first vertex, and so on until only v_0 remains. Since for any $\ell \in \mathbb{N}$,

$$\int dx \left(\frac{8n}{9}\right)^{1/4} M_\ell \left(\left[\left(\frac{8n}{9}\right)^{1/4} x \right] \right) = 1,$$

and similarly with P_ℓ instead of M_ℓ , we obtain that, for n sufficiently large and after integration with respect to $d\lambda^{V'}$, the integrand in (B.5) is bounded by

$$\mathbf{1}_{\{\|h\|_I \leq 2h\}} \mathbf{1}_{\{\|l\|_{\vec{B}'_+} \leq 2\|L\|\}} \prod_{e \in I \setminus \vec{E}_V} \frac{C}{\sqrt{h^e}} \prod_{e \in \vec{B}_+ \setminus \vec{E}_V} \frac{C}{\sqrt{l^e}}. \tag{B.11}$$

This is integrable with respect to $d\mathbf{h}^{I'} \otimes d\mathbf{l}^{\vec{B}'_+}$ and is bounded, after integration, by some constant times some power of h . Taking r sufficiently large in (B.10) yields that this quantity multiplied by $n Q_{2\lceil \sqrt{2nh} \rceil + \|I_n\|} (2n + \|I_n\|)$ is integrable with respect to dh . The claimed dominated convergence follows.

Dominant schemes. We will now see which schemes are such that $Z_1^s(n)$ has the highest possible order in n . The exponent of n in the equivalent (B.9) is maximal when $|V(\mathbf{s})|$ is the largest; in this case,

$$|V(\mathbf{s})| = 2(2g + k - 1).$$

This equality is obtained as in the proof of Lemma 2.1, since $|V(\mathbf{s})|$ being the largest means that the vertices have the lowest possible degrees, namely 3 for the internal vertices and 1 for the external vertices. More precisely, denoting by v, e, f the numbers of vertices, edges and faces of \mathbf{s} , as well as t the number of tadpoles among h_{b+1}, \dots, h_k , we obtain

$$f = b + t + 1, \quad 2e = 3(v - k + b + t) + k - b - t,$$

and the result from the Euler characteristic formula

$$v - e + f = 2 - 2g.$$

Next, the local limit theorem [86, Theorem VII.1.6] yields the existence of a compact set $K \subset (0, \infty)$ such that, for all $\ell \in \mathbb{N}$, $\sqrt{\ell}P_\ell(0) \in K$. Finally, for any $\mathbf{s} \in \vec{\mathcal{S}}$, we denote by \mathbf{s}° the scheme obtained by shrinking every tadpole among h_{b+1}, \dots, h_k into a vertex. For any fixed dominant scheme $\mathbf{d} \in \vec{\mathcal{S}}^*$, observe that there exists exactly one scheme among $\{\mathbf{s} \in \vec{\mathcal{S}} : \mathbf{s}^\circ = \mathbf{d}\}$ that is compatible with I_n , namely the one whose tadpoles among h_{b+1}, \dots, h_k are the holes indexed by $\{i > b : l_n^i \geq 1\}$. Furthermore, if $\mathbf{s} \in \vec{\mathcal{S}}$ is such that $\mathbf{s}^\circ \in \vec{\mathcal{S}}^*$, then the external faces among $h_i, b + 1 \leq i \leq k$, of \mathbf{s} are all tadpoles. We may thus use the equivalent (B.9) for these schemes. Consequently, as $n \rightarrow \infty$,

$$\sum_{\substack{\mathbf{s} \in \vec{\mathcal{S}} \\ \mathbf{s}^\circ = \mathbf{d}}} \mathcal{Z}_1^{\mathbf{s}}(n) = \Theta\left(n^{5(g-1)/2+k-b/4} \prod_{i>b:l_n^i \geq 1} (l_n^i)^{-1/2}\right). \quad (\text{B.12})$$

In particular, if \mathbf{s} has only tadpoles among its external faces indexed by $b + 1 \leq i \leq k$ but is such that \mathbf{s}° is not dominant, then $\mathcal{Z}_1^{\mathbf{s}}(n)$ is negligible with respect to this sum.

Nondominant schemes. We will now see that the above is the highest order in n and that it is only obtained for the schemes that are dominant after the tadpoles shrinkage. To this end, we fix an arbitrary scheme $\mathbf{s} \in \vec{\mathcal{S}}$. As above, we consider an arbitrary spanning tree of \mathbf{s} and still denote by $e_v \in I \cup \vec{B}$ the half-edge corresponding to $v \in V'$, as well as $\vec{E}_V = \{e_v : v \in V'\}$.

In (B.2), we bound $M_{h^e}(\delta^e)$ or $P_{l^e}(\delta^e)$ thanks to (B.10) if $e \notin \vec{E}_V$, and we operate the sum over $\lambda^{V'}$ leaf by leaf as we did in the previous paragraph. Since for any $\ell \in \mathbb{N}$,

$$\sum_{j \in \mathbb{Z}} M_\ell(j) = \sum_{j \in \mathbb{Z}} P_\ell(j) = 1,$$

we obtain the bound

$$\mathcal{Z}_1^{\mathbf{s}}(h, n) \leq \sum_{\mathbf{h}^I, \mathbf{l}^{\vec{B}}} \prod_{e \in I \setminus \vec{E}_V} \frac{C}{\sqrt{h^e}} \prod_{e \in \vec{B} \setminus \vec{E}_V} \frac{C}{\sqrt{l^e}},$$

where the sum is over the tuples $\mathbf{h}^{\vec{I}}, \mathbf{l}^{\vec{B}}$ satisfying the conditions of $\mathcal{T}_s(h, n)$. Seeing the sums as integrals under the simplex Lebesgue measures $\Delta_I^h \Delta_{\vec{B}_i}^{l_n^i}$ whenever $\vec{B}_i \neq \emptyset$ yields integrals of Dirichlet distributions (with parameter vectors containing only $1/2$'s and 1 's), after renormalization by h or l_n^i . As a result,

$$\begin{aligned} \mathcal{Z}_1^s(h, n) &\lesssim h^{|I|-1-|I \setminus \vec{E}_V|/2} \prod_{\substack{1 \leq i \leq k \\ \vec{B}_i \neq \emptyset}} (l_n^i)^{|\vec{B}_i|-1-|\vec{B}_i \setminus \vec{E}_V|/2} \\ &= h^{|I|/2+|I \cap \vec{E}_V|/2-1} \prod_{\substack{1 \leq i \leq k \\ \vec{B}_i \neq \emptyset}} (l_n^i)^{|\vec{B}_i|/2+|\vec{B}_i \cap \vec{E}_V|/2-1}, \end{aligned}$$

where we used the symbol \lesssim to mean bounded up to a constant independent² of \mathbf{s}, h , and n . Since l_n^i is in the scale \sqrt{n} for $i \leq b$, the part of the product concerning these indices is bounded by a constant times

$$n^{|\vec{B}'_+|/4+|\vec{B}_+ \cap \vec{E}_V|/4-b/4}.$$

Recall that $B_i \neq \emptyset \Leftrightarrow l_n^i \geq 1$ when \mathbf{l}_n is compatible with \mathbf{s} . Using (B.6), which is valid for any scheme, as well as the bound (B.10) as above to get integrability, we obtain

$$\begin{aligned} \mathcal{Z}_1^s(n) &\lesssim \mathbf{1}_{\mathbf{l}_n \triangleright \mathbf{s}} n^{(|I|+|\vec{B}'_+|)/4+(I \cup \vec{B}_+) \cap \vec{E}_V / 4 - b/4 - 1} \prod_{i > b : l_n^i \geq 1} (l_n^i)^{|\vec{B}_i|/2+|\vec{B}_i \cap \vec{E}_V|/2-1} \\ &\lesssim \mathbf{1}_{\mathbf{l}_n \triangleright \mathbf{s}} n^{(|I|+|\vec{B}'_+|+|V'|)/4-b/4-1} \prod_{i > b : l_n^i \geq 1} (l_n^i)^{|\vec{B}_i|/2-1}, \end{aligned}$$

since $l_n^i = \mathcal{O}(\sqrt{n})$ for all i , and $|(I \cup \vec{B}_+ \cup \vec{B}_0) \cap \vec{E}_V| = |V'|$. Using again the Euler characteristic formula, as well as the bound $|V| \leq 2(2g + k - 1)$, we obtain

$$\mathcal{Z}_1^s(n) \lesssim \mathbf{1}_{\mathbf{l}_n \triangleright \mathbf{s}} n^{5(g-1)/2+k-b/4} \prod_{i > b : l_n^i \geq 1} (l_n^i)^{-1/2} \left(\frac{l_n^i}{\sqrt{n}} \right)^{(|\vec{B}_i|-1)/2},$$

which gives an order lower than that of (B.12) as soon as there exists $i > b$ such that $|\vec{B}_i| \geq 2$ since $l_n^i = o(\sqrt{n})$. As a result, the normalization constant $\mathcal{Z}_1(n)$ is of the order appearing in (B.12) and $\mathcal{Z}_1^s(n) = o(\mathcal{Z}_1(n))$ whenever $\mathbf{s}^\circ \notin \vec{\mathbf{S}}^*$. In particular, $\mathbb{P}(S_n^\circ \in \vec{\mathbf{S}}^*) \rightarrow 1$ as $n \rightarrow \infty$ and we obtain the first statement of Proposition 3.10: with asymptotic probability 1, every vanishing face of M_n induces a tadpole in S_n .

²Recall from Lemma 2.1 that the number of edges in the schemes from $\vec{\mathbf{S}}$ is bounded.

B.3 Asymptotics of the size parameters

Limiting distribution of the size parameters. Given a bounded continuous function ϕ on the set

$$\bigcup_{\mathbf{s} \in \vec{\mathcal{S}}: \mathbf{s}^\circ = \mathbf{s}} \{\mathbf{s}\} \times (\mathbb{R}_{\geq 0})^{\vec{I}(\mathbf{s})} \times (\mathbb{R}_{\geq 0})^{\vec{B}(\mathbf{s})} \times \mathbb{R}^{V(\mathbf{s})},$$

we set, for $n, h \in \mathbb{N}, \mathbf{s} \in \vec{\mathcal{S}}$,

$$\mathcal{Z}_\phi^{\mathbf{s}}(h, n) = \sum_{\mathcal{I}_s(h, n)} \phi\left(\mathbf{s}^\circ, \frac{\mathbf{h}^{\vec{I}(\mathbf{s}^\circ)}}{\sqrt{2n}}, \frac{\mathbf{l}^{\vec{B}(\mathbf{s}^\circ)}}{\sqrt{2n}}, \frac{\boldsymbol{\lambda}^{V(\mathbf{s}^\circ)}}{(8n/9)^{1/4}}\right) \prod_{e \in I(\mathbf{s})} M_{he}(\delta\lambda^e) \prod_{e \in \vec{B}(\mathbf{s})} P_{1e}(\delta\lambda^e),$$

and

$$\mathcal{Z}_\phi^{\mathbf{s}}(n) = \frac{1}{|\vec{I}(\mathbf{s})| + 2|\vec{B}_+(\mathbf{s})|} \sum_{h \in \mathbb{N}} Q_{2h + \|\mathbf{l}_n\|} (2n + \|\mathbf{l}_n\|) \mathcal{Z}_\phi^{\mathbf{s}}(h, n),$$

so that

$$\mathbb{E}\left[\phi\left(\mathcal{S}_n^\circ, \frac{\mathbf{H}_n^{\vec{I}(\mathcal{S}_n^\circ)}}{\sqrt{2n}}, \frac{\mathbf{L}_n^{\vec{B}(\mathcal{S}_n^\circ)}}{\sqrt{2n}}, \frac{\boldsymbol{\Lambda}_n^{V(\mathcal{S}_n^\circ)}}{(8n/9)^{1/4}}\right)\right] = \frac{1}{\mathcal{Z}_1(n)} \sum_{\mathbf{s} \in \vec{\mathcal{S}}} \mathcal{Z}_\phi^{\mathbf{s}}(n).$$

Conducting with $\mathcal{Z}_\phi^{\mathbf{s}}(n)$ exactly the same computations as the ones we did with $\mathcal{Z}_1^{\mathbf{s}}(n)$, we obtain the same domination (up to $\sup |\phi|$) when \mathbf{s}° is not dominant and a similar equivalent when \mathbf{s}° is dominant, namely (B.9), where $c_1^{\mathbf{s}}(\mathbf{L})$ is replaced by

$$\begin{aligned} c_\phi^{\mathbf{s}}(\mathbf{L}) &= \frac{1}{|\vec{I}(\mathbf{s})| + 2|\vec{B}_+(\mathbf{s})|} 3^{b/2-g} 2^{|V'(\mathbf{s})|/2-g/2-3b/4+1/2} \int_{\mathbb{R}_{\geq 0}} dh q_{2h + \|\mathbf{L}\|}(1) \\ &\times \int d\mathbf{L}_s \phi\left(\mathbf{s}^\circ, \mathbf{h}^{\vec{I}(\mathbf{s}^\circ)}, \mathbf{l}^{\vec{B}(\mathbf{s}^\circ)}, \boldsymbol{\lambda}^{V(\mathbf{s}^\circ)}\right) \prod_{e \in I(\mathbf{s})} p_{he}(\delta\lambda^e) \prod_{e \in \vec{B}_+(\mathbf{s})} p_{31e}(\delta\lambda^e). \end{aligned}$$

From the Euler characteristic formula, we obtain $|\vec{I}(\mathbf{s})| + 2|\vec{B}_+(\mathbf{s})| = 2|E(\mathbf{s}^\circ)| = 2(6g + 2p + b - 3)$ does not depend on \mathbf{s} , and we remind that $|V'(\mathbf{s})| = 4g + 2p - 3$ does not either. Let us consider a dominant scheme $\mathbf{d} \in \vec{\mathcal{S}}^*$ and an integer $n \in \mathbb{N}$. We let $\mathbf{d}_n \in \vec{\mathcal{S}}$ be the unique scheme compatible with \mathbf{l}_n and such that $\mathbf{d}_n^\circ = \mathbf{d}$. Recall that this is the scheme obtained from \mathbf{d} by making into tadpoles the external vertices indexed by $\{i > b : l_n^i \geq 1\}$. Since the above integral only involves \mathbf{s}° , we have $c_\phi^{\mathbf{d}_n}(\mathbf{L}) = c_\phi^{\mathbf{d}}(\mathbf{L})$, and then

$$\mathcal{Z}_1(n) = \sum_{\mathbf{s} \in \vec{\mathcal{S}}} \mathcal{Z}_1^{\mathbf{s}}(n) \sim \sum_{\mathbf{d} \in \vec{\mathcal{S}}^*} \mathcal{Z}_1^{\mathbf{d}_n}(n) \sim n^{5(g-1)/2+k-b/4} \prod_{i>b: l_n^i \geq 1} P_{i_h}(0) \sum_{\mathbf{d} \in \vec{\mathcal{S}}^*} c_1^{\mathbf{d}}(\mathbf{L}),$$

and

$$\mathbb{E}\left[\phi\left(S_n^\circ, \frac{\mathbf{H}_n^{\vec{I}(S_n^\circ)}}{\sqrt{2n}}, \frac{\vec{B}_n(S_n^\circ)}{\sqrt{2n}}, \frac{\Lambda_n^{V(S_n^\circ)}}{(8n/9)^{1/4}}\right)\right] \xrightarrow{n \rightarrow \infty} \frac{1}{\sum_{\mathbf{d} \in \vec{\mathcal{S}}^\star} c_1^{\mathbf{d}}(\mathbf{L})} \sum_{\mathbf{d} \in \vec{\mathcal{S}}^\star} c_\phi^{\mathbf{d}}(\mathbf{L}). \quad (\text{B.13})$$

In passing, we obtain the following asymptotic formula for the cardinality of $\vec{\mathcal{Q}}_{n, l_n}^{[g]}$, which readily yields Proposition 1.4 (corresponding to the case $k = b$), the unrooting giving a factor $1/4n$ coming from (1.2). We define the continuous function in $\mathbf{L} \in (0, \infty)^b$,

$$t_g(\mathbf{L}) = \sum_{\mathbf{d} \in \vec{\mathcal{S}}^\star} c_1^{\mathbf{d}}(\mathbf{L}),$$

and we add the excluded cases $(g, k) = (0, 0)$ and $(g, k, b) = (0, 1, 1)$, which are needed in Section 1.5: we set

$$t_0(\emptyset) = \frac{1}{2\sqrt{\pi}} \quad \text{and} \quad t_0(L) = \frac{2^{-9/4}}{\pi\sqrt{L}} \exp\left(-\frac{L^2}{2}\right).$$

Proposition B.1. *As $n \rightarrow \infty$, it holds that*

$$|\vec{\mathcal{Q}}_{n, l_n}^{[g]}| \sim 4t_g(\mathbf{L})12^n 8^{\|\mathbf{l}_n\|} n^{5(g-1)/2+k-b/4} \frac{e^\star}{e^\star + p_n^\diamond} \prod_{i>b: l_n^i \geq 1} P_{l_n^i}(0),$$

where $e^\star = 6g + 2p + b - 3$ is the common number of edges of all dominant schemes, and $p_n^\diamond = |\{i > b : l_n^i \geq 1\}|$ is the number of external faces among $\mathfrak{h}_{b+1}, \dots, \mathfrak{h}_k$ in the maps of $\mathbf{M}_{n, \mathbf{l}_n}^{[g]}$.

In the excluded cases $(g, k) = (0, 0)$ and $(g, k, b) = (0, 1, 1)$, a similar formula holds:

$$|\vec{\mathcal{Q}}_{n, \emptyset}^{[0]}| \sim 4t_0(\emptyset)12^n n^{-5/2} \quad \text{and} \quad |\vec{\mathcal{Q}}_{n, (l_n)}^{[0]}| \sim 4t_0(L)12^n 8^{l_n} n^{-7/4}$$

for $L > 0$ and $l_n \sim \sqrt{2n}L$ as $n \rightarrow \infty$.

Proof. Recall from Section 2.2 that $\vec{\mathbf{M}}_{n, \mathbf{l}_n}^{[g]}$ is in 1-to-2 correspondence with $\vec{\mathcal{Q}}_{n, \mathbf{l}_n, 0}^{[g]}$, and that (B.3) gives its cardinality. Using (1.3) then (B.3) and (B.4), we obtain that

$$\begin{aligned} |\vec{\mathcal{Q}}_{n, \mathbf{l}_n}^{[g]}| &= \frac{2}{n + \|\mathbf{l}_n\| + 2 - 2g - k} |\vec{\mathbf{M}}_{n, \mathbf{l}_n}^{[g]}| \\ &= 2 \frac{2n + \|\mathbf{l}_n\|}{n + \|\mathbf{l}_n\| + 2 - 2g - k} 12^n 8^{\|\mathbf{l}_n\|} \sum_{\mathbf{s} \in \vec{\mathcal{S}}} \frac{|\vec{I}(\mathbf{s})| + 2|\vec{B}_+(\mathbf{s})|}{2|E(\mathbf{s})|} \mathcal{Z}_1^{\mathbf{s}}(n) \\ &\sim 4 \times 12^n 8^{\|\mathbf{l}_n\|} \sum_{\mathbf{d} \in \vec{\mathcal{S}}^\star} \frac{|E(\mathbf{d})|}{|E(\mathbf{d}_n)|} \mathcal{Z}_1^{\mathbf{d}_n}(n) \\ &\sim 4 \times 12^n 8^{\|\mathbf{l}_n\|} n^{5(g-1)/2+k-b/4} \frac{e^\star}{e^\star + p_n^\diamond} \prod_{i>b: l_n^i \geq 1} P_{l_n^i}(0) \sum_{\mathbf{d} \in \vec{\mathcal{S}}^\star} c_1^{\mathbf{d}}(\mathbf{L}), \end{aligned}$$

which gives the desired first statement.

The excluded cases $(g, k) = (0, 0)$ and $(g, k, b) = (0, 1, 1)$ are standard; they are obtained similarly, by computing $|\vec{\mathbf{M}}_{n, \emptyset}^{[0]}|$ and $|\vec{\mathbf{M}}_{n, (l_n)}^{[0]}|$. More precisely, it is well known that

$$|\vec{\mathbf{M}}_{n, \emptyset}^{[0]}| = 3^n \frac{(2n)!}{n!(n+1)!} \sim \frac{12^n}{\sqrt{\pi}} n^{-3/2}.$$

In order to compute the remaining cardinality, we proceed as in Appendix B.1 and obtain

$$|\vec{\mathbf{M}}_{n, (l_n)}^{[0]}| = \frac{2n + l_n}{l_n} 12^n 8^{l_n} Q_{l_n}(2n + l_n) P_{l_n}(0),$$

the division by l_n taking into account the fact that seeing an element of $\vec{\mathbf{M}}_{n, (l_n)}^{[0]}$ as a forest amounts to choose a first tree among l_n . From equivalents (B.7) and (B.8), this yields

$$|\vec{\mathbf{Q}}_{n, (l_n)}^{[0]}| \sim \frac{2}{n} \frac{2n}{\sqrt{2nL}} 12^n 8^{l_n} \frac{1}{n} q_L(1) \left(\frac{8n}{9}\right)^{-1/4} p_{3L}(0),$$

which gives the desired result. \blacksquare

Limiting distribution of the areas. We finally take into account the areas. To this end, observe that, conditionally given

$$(S_n, \mathbf{H}_n^{\vec{I}(S_n)}, \mathbf{L}_n^{\vec{B}(S_n)}),$$

the area vector $\mathbf{A}_n^{\vec{E}(S_n)}$ is distributed as follows. We arrange the half-edges e_1, \dots, e_κ incident to the internal face of S_n according to the contour order, starting arbitrarily, and let $x_i = \sum_{j=1}^i \ell_j$, where $\ell_j = H_n^{e_j}$ if $e_j \in \vec{I}(S_n)$ or $\ell_j = L_n^{e_j}$ if $e_j \in \vec{B}(S_n)$. Then, $A_n^{e_1}, A_n^{e_1} + A_n^{e_2}, A_n^{e_1} + A_n^{e_2} + A_n^{e_3}, \dots$ are distributed as the hitting times of the successive levels $-x_1, -x_1 - x_2, -x_1 - x_2 - x_3, \dots$ by a simple random walk conditioned on hitting the final level $-\sum_{j=1}^\kappa x_j = -\|\mathbf{H}_n\| - \|\mathbf{L}_n\|$ at time $2n + \|\mathbf{L}_n\|$. The desired convergence (3.9) easily follows from this together with (B.13), as well as the fact that, for every $e \in \vec{B}_0(S_n)$, we have $A_n^e + L_n^e = \Theta((L_n^e)^2)$ in probability.

B.4 Boltzmann quadrangulations

We finally prove Theorem 1.5; in its setting,

$$\begin{aligned} & \mathcal{W}(F(\Omega_{a^{-1}}(Q)) \mathbf{1}_{\mathbf{Q}_{I_a \mathbf{0}^\rho}^{[g]}}) \\ &= \sum_{n \in [a^{-1}/K, a^{-1}K] \cap \mathbb{Z}_{\geq 0}} \mathcal{W}(\mathbf{Q}_{n, I_a \mathbf{0}^\rho}^{[g]}) \mathcal{W}[F(\Omega_{a^{-1}}(Q)) \mid \mathbf{Q}_{n, I_a \mathbf{0}^\rho}^{[g]}] \\ &= a^{-1} \int_{a^{-1}/K}^{a^{-1}K} dA \mathcal{W}(\mathbf{Q}_{[A/a], I_a \mathbf{0}^\rho}^{[g]}) \mathcal{W}[F(\Omega_{a^{-1}}(Q)) \mid \mathbf{Q}_{[A/a], I_a \mathbf{0}^\rho}^{[g]}]. \end{aligned}$$

By Theorem 1.1 and the definition of $\mathbb{S}_{A,\mathbf{L}}^{[g]}$, it holds that

$$\mathcal{W}[F(\Omega_{a^{-1}}(Q)) \mid \mathbf{Q}_{[A/a], t_a \mathbf{0}^p}^{[g]}] \xrightarrow{a \downarrow 0} \mathbb{E}[F(\mathbb{S}_{A,\mathbf{L}}^{[g]} \mathbf{0}^p)],$$

while Proposition 1.4 yields that

$$\mathcal{W}(\mathbf{Q}_{[A/a], t_a \mathbf{0}^p}^{[g]}) \underset{a \downarrow 0}{\sim} t_g \left(\frac{\mathbf{L}}{\sqrt{A}} \right) \left(\frac{A}{a} \right)^{(5g-7)/2 + 3b/4 + p}.$$

Hence, Theorem 1.5 will be proved if we can show that the convergence in the last integral expression is dominated. However, this is a direct consequence of the discussion of the domination hypothesis around (B.11). Corollary 1.6 is proved in a very similar way, this time summing over all possible values of the perimeters, which results in the integral with respect to $d\mathbf{L}$ on $(0, \infty)^b$.

References

- [1] C. Abraham, [Rescaled bipartite planar maps converge to the Brownian map](#). *Ann. Inst. Henri Poincaré Probab. Stat.* **52** (2016), no. 2, 575–595
- [2] R. Abraham, J.-F. Delmas, and P. Hoscheit, [A note on the Gromov–Hausdorff–Prokhorov distance between \(locally\) compact metric measure spaces](#). *Electron. J. Probab.* **18** (2013), article no. 14
- [3] L. Addario-Berry and M. Albenque, [The scaling limit of random simple triangulations and random simple quadrangulations](#). *Ann. Probab.* **45** (2017), no. 5, 2767–2825
- [4] L. Addario-Berry and M. Albenque, [Convergence of non-bipartite maps via symmetrization of labeled trees](#). *Ann. H. Lebesgue* **4** (2021), 653–683
- [5] L. Addario-Berry and Y. Wen, [Joint convergence of random quadrangulations and their cores](#). *Ann. Inst. Henri Poincaré Probab. Stat.* **53** (2017), no. 4, 1890–1920
- [6] M. Albenque and D. Poulalhon, [A generic method for bijections between blossoming trees and planar maps](#). *Electron. J. Combin.* **22** (2015), no. 2, article no. 2.38
- [7] D. Aldous, [The continuum random tree. I](#). *Ann. Probab.* **19** (1991), no. 1, 1–28
- [8] D. Aldous, [The continuum random tree. III](#). *Ann. Probab.* **21** (1993), no. 1, 248–289
- [9] M. Ang, N. Holden, and X. Sun, [The SLE loop via conformal welding of quantum disks](#). *Electron. J. Probab.* **28** (2023), article no. 30
- [10] M. Ang, G. Remy, and X. Sun, [The moduli of annuli in random conformal geometry](#). *Ann. Sci. Éc. Norm. Supér. (4)* **58** (2025), no. 4, 1037–1087
- [11] O. Angel, B. Kolesnik, and G. Miermont, [Stability of geodesics in the Brownian map](#). *Ann. Probab.* **45** (2017), no. 5, 3451–3479
- [12] O. Angel and O. Schramm, [Uniform infinite planar triangulations](#). *Comm. Math. Phys.* **241** (2003), no. 2–3, 191–213
- [13] E. Baur, G. Miermont, and G. Ray, [Classification of scaling limits of uniform quadrangulations with a boundary](#). *Ann. Probab.* **47** (2019), no. 6, 3397–3477
- [14] E. Baur and L. Richier, [Uniform infinite half-planar quadrangulations with skewness](#). *Electron. J. Probab.* **23** (2018), article no. 54
- [15] V. Beffara, C. B. Huynh, and B. Lévêque, [Scaling limits for random triangulations on the torus](#). 2019, arXiv:1905.01873v1
- [16] J. Beltran and J.-F. Le Gall, [Quadrangulations with no pendant vertices](#). *Bernoulli* **19** (2013), no. 4, 1150–1175
- [17] E. A. Bender and E. R. Canfield, [The asymptotic number of rooted maps on a surface](#). *J. Combin. Theory Ser. A* **43** (1986), no. 2, 244–257
- [18] J. Bertoin, N. Curien, and I. Kortchemski, [Random planar maps and growth-fragmentations](#). *Ann. Probab.* **46** (2018), no. 1, 207–260
- [19] J. Bettinelli, [Scaling limits for random quadrangulations of positive genus](#). *Electron. J. Probab.* **15** (2010), 1594–1644

- [20] J. Bettinelli, [The topology of scaling limits of positive genus random quadrangulations](#). *Ann. Probab.* **40** (2012), no. 5, 1897–1944
- [21] J. Bettinelli, [Scaling limit of random planar quadrangulations with a boundary](#). *Ann. Inst. Henri Poincaré Probab. Stat.* **51** (2015), no. 2, 432–477
- [22] J. Bettinelli, [Geodesics in Brownian surfaces \(Brownian maps\)](#). *Ann. Inst. Henri Poincaré Probab. Stat.* **52** (2016), no. 2, 612–646
- [23] J. Bettinelli, [A bijection for nonorientable general maps](#). *Ann. Inst. Henri Poincaré D* **9** (2022), no. 4, 733–791
- [24] J. Bettinelli, E. Jacob, and G. Miermont, [The scaling limit of uniform random plane maps, via the Ambjørn–Budd bijection](#). *Electron. J. Probab.* **19** (2014), article no. 74
- [25] J. Bettinelli and G. Miermont, [Compact Brownian surfaces I: Brownian disks](#). *Probab. Theory Related Fields* **167** (2017), no. 3–4, 555–614
- [26] N. H. Bingham, C. M. Goldie, and J. L. Teugels, *Regular variation*. Encyclopedia Math. Appl. 27, Cambridge University Press, Cambridge, 1989
- [27] J. Bouttier, Planar maps and random partitions. 2019, Habilitation thesis, Université Paris-Sud, <https://tel.archives-ouvertes.fr/tel-02417269>, visited on 10 March 2026
- [28] J. Bouttier, P. Di Francesco, and E. Guitter, [Planar maps as labeled mobiles](#). *Electron. J. Combin.* **11** (2004), no. 1, article no. 69
- [29] J. Bouttier and E. Guitter, [Distance statistics in quadrangulations with a boundary, or with a self-avoiding loop](#). *J. Phys. A* **42** (2009), no. 46, article no. 465208
- [30] J. Bouttier and E. Guitter, [Planar maps and continued fractions](#). *Comm. Math. Phys.* **309** (2012), no. 3, 623–662
- [31] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer, Berlin, 1999
- [32] D. Burago, Y. Burago, and S. Ivanov, *A course in metric geometry*. Grad. Stud. Math. 33, American Mathematical Society, Providence, RI, 2001
- [33] A. Caraceni and N. Curien, [Geometry of the uniform infinite half-planar quadrangulation](#). *Random Structures Algorithms* **52** (2018), no. 3, 454–494
- [34] G. Chapuy and M. Dołęga, [A bijection for rooted maps on general surfaces](#). *J. Combin. Theory Ser. A* **145** (2017), 252–307
- [35] G. Chapuy, M. Marcus, and G. Schaeffer, [A bijection for rooted maps on orientable surfaces](#). *SIAM J. Discrete Math.* **23** (2009), no. 3, 1587–1611
- [36] P. Chassaing and B. Durhuus, [Local limit of labeled trees and expected volume growth in a random quadrangulation](#). *Ann. Probab.* **34** (2006), no. 3, 879–917
- [37] P. Chassaing and G. Schaeffer, [Random planar lattices and integrated superBrownian excursion](#). *Probab. Theory Related Fields* **128** (2004), no. 2, 161–212
- [38] R. Cori and B. Vauquelin, [Planar maps are well labeled trees](#). *Canad. J. Math.* **33** (1981), no. 5, 1023–1042
- [39] N. Curien, *Peeling random planar maps*. *École d’Été de Probabilités de Saint-Flour XLIX – 2019*. Lecture Notes in Math. 2335, Springer, Cham, 2023

- [40] N. Curien and J.-F. Le Gall, [The Brownian plane](#). *J. Theoret. Probab.* **27** (2014), no. 4, 1249–1291
- [41] N. Curien and J.-F. Le Gall, [First-passage percolation and local modifications of distances in random triangulations](#). *Ann. Sci. Éc. Norm. Supér. (4)* **52** (2019), no. 3, 631–701
- [42] N. Curien, L. Ménard, and G. Miermont, A view from infinity of the uniform infinite planar quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.* **10** (2013), no. 1, 45–88
- [43] N. Curien and G. Miermont, [Uniform infinite planar quadrangulations with a boundary](#). *Random Structures Algorithms* **47** (2015), no. 1, 30–58
- [44] F. David, [Planar diagrams, two-dimensional lattice gravity and surface models](#). *Nuclear Phys. B* **257** (1985), no. 1, 45–58
- [45] F. David, A. Kupiainen, R. Rhodes, and V. Vargas, [Liouville quantum gravity on the Riemann sphere](#). *Comm. Math. Phys.* **342** (2016), no. 3, 869–907
- [46] J. Ding, J. Dubédat, A. Dunlap, and H. Falconet, [Tightness of Liouville first passage percolation for \$\gamma \in \(0, 2\)\$](#) . *Publ. Math. Inst. Hautes Études Sci.* **132** (2020), 353–403
- [47] J. Ding, J. Dubédat, and E. Gwynne, [Introduction to the Liouville quantum gravity metric](#). In *ICM—International Congress of Mathematicians. Vol. 6. Sections 12–14*, pp. 4212–4244, European Mathematical Society, Berlin, 2023
- [48] B. Eynard, [Counting surfaces](#). *CRM Aisenstadt Chair lectures*. Prog. Math. Phys. 70, Birkhäuser, Cham, 2016
- [49] C. Guillarmou, A. Kupiainen, R. Rhodes, and V. Vargas, Segal’s axioms and bootstrap for Liouville theory. 2025, arXiv:2112.14859v2
- [50] C. Guillarmou, R. Rhodes, and V. Vargas, [Polyakov’s formulation of 2d bosonic string theory](#). *Publ. Math. Inst. Hautes Études Sci.* **130** (2019), 111–185
- [51] E. Gwynne, N. Holden, and X. Sun, Mating of trees for random planar maps and Liouville quantum gravity: a survey. In *Topics in statistical mechanics*, pp. 41–120, Panor. Synthèses 59, Société Mathématique de France, Paris, 2023
- [52] E. Gwynne and J. Miller, [Scaling limit of the uniform infinite half-plane quadrangulation in the Gromov–Hausdorff–Prokhorov-uniform topology](#). *Electron. J. Probab.* **22** (2017), article no. 84
- [53] E. Gwynne and J. Miller, [Metric gluing of Brownian and \$\sqrt{8/3}\$ -Liouville quantum gravity surfaces](#). *Ann. Probab.* **47** (2019), no. 4, 2303–2358
- [54] E. Gwynne and J. Miller, [Convergence of the self-avoiding walk on random quadrangulations to \$SLE_{8/3}\$ on \$\sqrt{8/3}\$ -Liouville quantum gravity](#). *Ann. Sci. Éc. Norm. Supér. (4)* **54** (2021), no. 2, 305–405
- [55] E. Gwynne and J. Miller, [Existence and uniqueness of the Liouville quantum gravity metric for \$\gamma \in \(0, 2\)\$](#) . *Invent. Math.* **223** (2021), no. 1, 213–333
- [56] E. Gwynne, J. Miller, and S. Sheffield, [The Tutte embedding of the Poisson–Voronoi tessellation of the Brownian disk converges to \$\sqrt{8/3}\$ -Liouville quantum gravity](#). *Comm. Math. Phys.* **374** (2020), no. 2, 735–784
- [57] E. Gwynne, J. Miller, and S. Sheffield, [An invariance principle for ergodic scale-free random environments](#). *Acta Math.* **228** (2022), no. 2, 303–384

- [58] N. Holden and X. Sun, [Convergence of uniform triangulations under the Cardy embedding](#). *Acta Math.* **230** (2023), no. 1, 93–203
- [59] V. G. Knizhnik, A. M. Polyakov, and A. B. Zamolodchikov, [Fractal structure of 2D-quantum gravity](#). *Modern Phys. Lett. A* **3** (1988), no. 8, 819–826
- [60] M. Krikun, [Local structure of random quadrangulations](#). 2006, arXiv:math/0512304v2
- [61] A. Kupiainen, R. Rhodes, and V. Vargas, [Integrability of Liouville theory: proof of the DOZZ formula](#). *Ann. of Math. (2)* **191** (2020), no. 1, 81–166
- [62] S. K. Lando and A. K. Zvonkin, [Graphs on surfaces and their applications \(with an appendix by Don B. Zagier\)](#). Encyclopaedia Math. Sci. 141, Springer, Berlin, 2004
- [63] J.-F. Le Gall, [Spatial branching processes, random snakes and partial differential equations](#). Lect. Math. ETH Zürich, Birkhäuser, Basel, 1999
- [64] J.-F. Le Gall, [A conditional limit theorem for tree-indexed random walk](#). *Stochastic Process. Appl.* **116** (2006), no. 4, 539–567
- [65] J.-F. Le Gall, [The topological structure of scaling limits of large planar maps](#). *Invent. Math.* **169** (2007), no. 3, 621–670
- [66] J.-F. Le Gall, [Geodesics in large planar maps and in the Brownian map](#). *Acta Math.* **205** (2010), no. 2, 287–360
- [67] J.-F. Le Gall, [Uniqueness and universality of the Brownian map](#). *Ann. Probab.* **41** (2013), no. 4, 2880–2960
- [68] J.-F. Le Gall, [Brownian disks and the Brownian snake](#). *Ann. Inst. Henri Poincaré Probab. Stat.* **55** (2019), no. 1, 237–313
- [69] J.-F. Le Gall, [Brownian geometry](#). *Jpn. J. Math.* **14** (2019), no. 2, 135–174
- [70] J.-F. Le Gall, [The Brownian disk viewed from a boundary point](#). *Ann. Inst. Henri Poincaré Probab. Stat.* **58** (2022), no. 2, 1091–1119
- [71] J.-F. Le Gall, [Geodesic stars in random geometry](#). *Ann. Probab.* **50** (2022), no. 3, 1013–1058
- [72] J.-F. Le Gall and G. Miermont, [Scaling limits of random planar maps with large faces](#). *Ann. Probab.* **39** (2011), no. 1, 1–69
- [73] J.-F. Le Gall and F. Paulin, [Scaling limits of bipartite planar maps are homeomorphic to the 2-sphere](#). *Geom. Funct. Anal.* **18** (2008), no. 3, 893–918
- [74] J.-F. Le Gall and A. Riera, [Growth-fragmentation processes in Brownian motion indexed by the Brownian tree](#). *Ann. Probab.* **48** (2020), no. 4, 1742–1784
- [75] J.-F. Le Gall and A. Riera, [Spine representations for non-compact models of random geometry](#). *Probab. Theory Related Fields* **181** (2021), no. 1–3, 571–645
- [76] J.-F. Marckert and A. Mokkadem, [The depth first processes of Galton–Watson trees converge to the same Brownian excursion](#). *Ann. Probab.* **31** (2003), no. 3, 1655–1678
- [77] C. Marzouk, [On scaling limits of random trees and maps with a prescribed degree sequence](#). *Ann. H. Lebesgue* **5** (2022), 317–386
- [78] G. Miermont, [On the sphericity of scaling limits of random planar quadrangulations](#). *Electron. Commun. Probab.* **13** (2008), 248–257

- [79] G. Miermont, [Tessellations of random maps of arbitrary genus](#). *Ann. Sci. Éc. Norm. Supér. (4)* **42** (2009), no. 5, 725–781
- [80] G. Miermont, [The Brownian map is the scaling limit of uniform random plane quadrangulations](#). *Acta Math.* **210** (2013), no. 2, 319–401
- [81] J. Miller and W. Qian, [Geodesics in the Brownian map: Strong confluence and geometric structure](#). *Mem. Amer. Math. Soc.* **315** (2025), no. 1602, vi+117 pp.
- [82] J. Miller and S. Sheffield, [Liouville quantum gravity and the Brownian map I: the QLE\(8/3, 0\) metric](#). *Invent. Math.* **219** (2020), no. 1, 75–152
- [83] J. Miller and S. Sheffield, [An axiomatic characterization of the Brownian map](#). *J. Éc. Polytech. Math.* **8** (2021), 609–731
- [84] J. Miller and S. Sheffield, [Liouville quantum gravity and the Brownian map II: Geodesics and continuity of the embedding](#). *Ann. Probab.* **49** (2021), no. 6, 2732–2829
- [85] J. Miller and S. Sheffield, [Liouville quantum gravity and the Brownian map III: the conformal structure is determined](#). *Probab. Theory Related Fields* **179** (2021), no. 3–4, 1183–1211
- [86] V. V. Petrov, *Sums of independent random variables*. *Ergeb. Math. Grenzgeb.* 82, Springer, New York, 1975
- [87] J. W. Pitman, [One-dimensional Brownian motion and the three-dimensional Bessel process](#). *Adv. in Appl. Probab.* **7** (1975), no. 3, 511–526
- [88] A. M. Polyakov, [Quantum geometry of bosonic strings](#). *Phys. Lett. B* **103** (1981), no. 3, 207–210
- [89] L. B. Richmond and N. C. Wormald, [Almost all maps are asymmetric](#). *J. Combin. Theory Ser. B* **63** (1995), no. 1, 1–7
- [90] G. Schaeffer, *Conjugaison d'arbres et cartes combinatoires aléatoires*. Ph.D. thesis, 1998, Université Bordeaux
- [91] S. Sheffield, [What is a random surface?](#) In *ICM—International Congress of Mathematicians. Vol. 2. Plenary lectures*, pp. 1202–1258, European Mathematical Society, Berlin, 2023
- [92] D. W. Stroock, *Probability theory. An analytic view*. 2nd edn., Cambridge University Press, Cambridge, 2011
- [93] C. Villani, *Optimal transport: Old and new*. Grundlehren Math. Wiss. 338, Springer, Berlin, 2009
- [94] B. Wu, [Conformal bootstrap on the annulus in Liouville CFT](#). 2024, arXiv:2203.11830v5

Jérémie Bettinelli, Grégory Miermont

Compact Brownian Surfaces II

Orientable Surfaces

Fix an arbitrary compact orientable surface with a boundary and consider a uniform bipartite random quadrangulation of this surface with n faces and boundary component lengths of order \sqrt{n} or of lower order. Endow this quadrangulation with the usual graph metric renormalized by $n^{-1/4}$, mark it on each boundary component, and endow it with the counting measure on its vertex set renormalized by n^{-1} , as well as the counting measure on each boundary component renormalized by $n^{-1/2}$. We show that, as n goes to infinity, this random marked measured metric space converges in distribution for the Gromov–Hausdorff–Prokhorov topology, toward a random limiting marked measured metric space called a *Brownian surface*.

This extends known convergence results of uniform random planar quadrangulations with at most one boundary component toward the *Brownian sphere* and toward the *Brownian disk*, by considering the case of quadrangulations on general compact orientable surfaces. Our approach consists in cutting a Brownian surface into elementary pieces that are naturally related to the Brownian sphere and the Brownian disk and their noncompact analogs, the Brownian plane and the Brownian half-plane, and to prove convergence results for these elementary pieces, which are of independent interest.

<https://ems.press>

ISSN 2747-9080

ISBN 978-3-98547-102-7



EM
S 
PRESS