



The Scaling Limit of Random Two-Connected Series–Parallel Maps

Downloaded from: <https://research.chalmers.se>, 2026-09-14 17:59 UTC

Citation for the original published paper (version of record):

Amankwah, D., Björnberg, J., Stefánsson, S. et al (2026). The Scaling Limit of Random Two-Connected Series–Parallel Maps. *Journal of Theoretical Probability*, 39(4).
<http://dx.doi.org/10.1007/s10959-026-01537-x>

N.B. When citing this work, cite the original published paper.



The Scaling Limit of Random Two-Connected Series–Parallel Maps

Daniel Amankwah¹ · Jakob Björnberg^{2,3} · Sigurdur Örn Stefánsson¹ · Benedikt Stufler⁴ · Joonas Turunen⁵

Received: 20 August 2025 / Revised: 6 July 2026 / Accepted: 16 July 2026
© The Author(s) 2026

Abstract

A finite graph embedded in the plane is called a series–parallel map if it can be obtained from a finite tree by repeatedly subdividing and doubling edges. We study the scaling limit of weighted random two-connected series–parallel maps with n edges and show that under fairly general integrability conditions on these weights, the maps with distances rescaled by a factor $n^{-1/2}$ converge to a constant multiple of Aldous’ continuum random tree (CRT) in the Gromov–Hausdorff sense. The proof relies on a bijection between a set of trees with n leaves and a set of series–parallel maps with n edges, together with a novel *blob decomposition* of the maps.

Keywords Random planar map · Simply generated trees · Scaling limits · Brownian tree

Mathematics Subject Classification 05C80 · 60F17 · 60J80

Dedicated to Svante Janson on the occasion of his 70th birthday.

✉ Sigurdur Örn Stefánsson
sigurdur@hi.is

Daniel Amankwah
dam@tskoli.is

Jakob Björnberg
jakob.bjornberg@gu.se

Benedikt Stufler
benedikt.stufler@tuwien.ac.at

Joonas Turunen
joonas.turunen@alumni.helsinki.fi

- ¹ Mathematics division, University of Iceland, Reykjavik, Iceland
- ² Department of Mathematics, Chalmers University of Technology, Gothenburg, Sweden
- ³ Department of Mathematics, University of Gothenburg, Gothenburg, Sweden
- ⁴ Institute of Discrete Mathematics and Geometry, TU Wien, Vienna, Austria
- ⁵ University of Helsinki, Helsinki, Finland

1 Introduction

A finite connected graph is called a *series-parallel graph* if it can be obtained from a finite tree by repeatedly subdividing or doubling edges. It may equivalently be characterized as a graph that does not contain the complete graph K_4 as a minor. In this work, we consider *series-parallel maps*, which are series-parallel graphs embedded in the plane, viewed up to continuous deformations.

Specifically, we consider series-parallel maps that are *two-connected*, and we are interested in their *scaling limits*. Without the condition of two-connectivity, uniformly sampled series-parallel maps have Aldous' continuum random tree (CRT) as their scaling limit, as easily follows from a general theorem of Stufler [21, Thm. 6.62]. More precisely, uniform samples of such maps with n edges, and distances rescaled by $n^{-1/2}$, converge weakly in the Gromov–Hausdorff sense to a multiple of the CRT. The essence of the proof lies in decomposing the map into two-connected components, or *blocks*, and using a coupling with simply generated trees where each block of the map corresponds to a vertex in the tree. As the tree scales at the order $n^{1/2}$, and since blocks are typically small, after rescaling distances by $n^{-1/2}$ on average each block contributes a constant length to the geodesic. Thus, the global shape of the rescaled map is approximated by the tree, stretched by a constant factor.

In this work, however, we would like to understand the structure of the blocks themselves. Thus, we focus our attention on *two-connected* series-parallel maps. These do not enjoy the same type of coupling to simply generated trees and prove more difficult to handle. Our main result is that, nevertheless, uniformly sampled two-connected series-parallel maps with n edges converge, as $n \rightarrow \infty$ with distances rescaled by $n^{-1/2}$, to a multiple of the CRT. Figure 1 shows simulations of these uniformly sampled maps and suggests that they become tree-like with an increasing number of edges.

Analogous results have been established before for several tree-like random maps, which can often be characterized through excluded minors. Most fundamental is the original result of Aldous, showing that uniform trees with n edges converge to the CRT when rescaled by $n^{-1/2}$ [1–3]. Trees may be characterized by not containing K_3 as a minor. Other notable cases are various classes of outerplanar graphs. A graph is called *outerplanar* if it may be embedded in the plane so that every vertex lies




Fig. 1 A simulation of uniform two-connected series-parallel maps with different numbers of edges. The endpoints of the root edge are represented by $\bullet$ and $\circ$

Table 1 Scaling limits of uniform maps with n edges characterized by exclusion of minors

Map	Excluded minors	Diameter	Scaling limit
Plane trees	K_3	$n^{1/2}$	CRT
Outerplanar maps	$K_4, K_{2,3}$	$n^{1/2}$	CRT
Series–parallel maps	K_4	$n^{1/2}$	CRT
Planar maps	$K_5, K_{3,3}$	$n^{1/4}$	Brownian sphere

on the boundary of the external face; equivalently, they are graphs not containing K_4 or the complete bipartite graph $K_{3,2}$ as minors. In [8], Caraceni established that uniformly sampled outerplanar maps (properly rescaled) converge to the CRT. Stufler [21] generalized this result to a family of face-weighted outerplanar maps, using the aforementioned coupling with simply generated trees. Note that outerplanar maps are series–parallel maps, since series–parallel maps have K_4 as an excluded minor, and outerplanar maps further exclude $K_{3,2}$ as a minor.

If one additionally imposes the condition of two-connectivity on outerplanar maps, one obtains *dissections of a polygon*. Random, face-weighted dissections of polygons were studied by Curien, Haas, and Kortchemski [9] who showed, under a suitable condition on the weights, that the scaling limit is a constant multiple of the CRT. Their proof involves relating the dissections bijectively to trees and using a clever Markov chain argument to compare lengths of geodesics in the dissections and their corresponding trees.

One of the main aims of the above studies is to understand how universally the CRT appears as a scaling limit for different models of random graphs or maps. Note that planar maps themselves, given by excluding K_5 and $K_{3,3}$ as minors, behave vastly differently. Uniform planar maps with n edges, when rescaled by $n^{-1/4}$, admit the so-called *Brownian sphere* as a scaling limit [4], which is completely different from the CRT (see Table 1 for a summary of these results).

Another motivation for understanding the geometry of series–parallel maps better stems from their conjectured connection to the two-dimensional Liouville quantum gravity (LQG), as a part of a more general framework of random surfaces obtained as a mating of trees (see [11]). The exact connection could be figured out via random separable permutations, which is outlined in Sect. 4 at the end of this work. The geometric interpretation of SP maps in the context of LQG has been under active discussion in the random geometry research community, but to our knowledge, nothing has been proven so far. One of the very few ideas is the conjecture that conformally embedded rooted rescaled series–parallel maps converge toward the 2-LQG quantum sphere under a suitable conformal embedding [6].

We use a bijection with certain plane trees (see Sect. 2.3 for details), which enables us to describe the distribution of the random maps with n edges in terms of a Bienaymé–Galton–Watson (BGW) tree conditioned on having n leaves. In order to study the geodesics in the map, we introduce a novel blob decomposition of series–parallel maps. Roughly speaking, such maps get glued together from smaller maps, with the gluing operation identifying a pair of vertices in one blob with a pair of vertices in

another blob. This way, the resulting pair of vertices forms a two-separator of the map. Any geodesics joining vertices from opposite sides are then forced to pass through one of the two gluing points. This paves the way for applying concentration inequalities for Markov-additive processes in order to handle the geodesics. However, the stationary distribution is non-explicit and almost intractable. In order to tackle this, we use a combinatorial trick to carefully define our blobs in a way so that any pair of gluing points is adjacent in at least one of the two blobs. This trick reduces the state space of our Markov-additive process from being infinite to having a finite number of states, thus crucially reducing the complexity for verifying necessary conditions for concentration inequalities.

Moreover, the relation with BGW trees allows us to consider a weighted version of the random maps in a natural way by assigning a general offspring distribution μ to the BGW trees. We will assume that this offspring distribution is critical and that it has finite exponential moments (is light-tailed), which includes the uniform distribution as a special case. The precise definition of the weighted series–parallel maps is given in Definition 2.2 in Sect. 2.6.

The CRT $\mathcal{T}_e$ used in this work is the version constructed from a normalized Brownian excursion e (see [16]). Let SP_n^μ denote the random, weighted, rooted, two-connected series–parallel map with n non-root edges corresponding to a critical offspring distribution μ of a BGW process possessing finite exponential moments (see Definition 2.2). The following is our main result:

Theorem 1.1 *Let $d_{SP_n^\mu}$ denote the usual graph metric on SP_n^μ . There exists a constant $c_\mu > 0$ such that*

$$(SP_n^\mu, c_\mu n^{-1/2} d_{SP_n^\mu}) \xrightarrow{d} (\mathcal{T}_e, d_{\mathcal{T}_e})$$

in the Gromov–Hausdorff sense.

The scaling constant c_μ remains implicit because it is defined in terms of a Markov-additive process with implicitly defined transition probabilities (see Remark 3.4).

We also obtain optimal tail bounds for the diameter $D(SP_n^\mu)$:

Theorem 1.2 *There exist constants $C, c > 0$ such that for all n and all $x > 0$*

$$\mathbb{P}(D(SP_n^\mu) > x) \leq C \exp(-cx^2/n).$$

Remark 1.3 We will consider other closely related classes of series–parallel maps referred to as series–parallel networks. Their definitions are given in Subsection 2.1. The two previous theorems also apply to random weighted series–parallel networks with the same scaling constant c_μ .




Fig. 2 Left: An example of a rooted, two-connected SP map. The directed root edge is indicated by an arrow. Right: Its corresponding labeled plane tree where $\circ$ denotes the label P and $\bullet$ denotes the label S

2 Decomposing and Sampling Series-Parallel Maps

2.1 Rooted, Two-Connected Series-Parallel Maps

A planar map is a drawing of a finite connected graph in the plane so that edges do not cross except possibly at their endpoints, viewed up to orientation preserving homeomorphisms of the plane. A series-parallel (SP) map is a planar map, which does not contain the complete graph K_4 as a minor. The objects of interest in this paper are two-connected SP maps. Another way to describe them is to start from a double edge and repeatedly subdivide or double edges [10, Theorems 1 and 3].

We will assign a direction to one of the edges of the map and refer to it as the *root* of the map. The root will be denoted by $e = (*_0, *_1)$. For concreteness, we draw the maps so that the root edge lies on the boundary of the unbounded face and is directed in a clockwise manner along that face (see Fig. 2 (left) for an example).

Denote the set of rooted two-connected SP maps with $n \geq 1$ non-root edges by $\mathcal{SP}_n$ and the set of all such finite maps by $\mathcal{SP} = \bigcup_{n \geq 1} \mathcal{SP}_n$. We deliberately leave out the map with the single root edge for notational convenience.

In order to be able to describe a decomposition of the maps into series and parallel components, we require two additional definitions. For $M \in \mathcal{SP}_n$, let M' be the map obtained by removing the root edge from M , keeping the endpoints $*_0, *_1$ as marked vertices of M' . Define for all $n \geq 2$

$$\begin{aligned}\mathcal{S}_n &= \{M' \mid M \in \mathcal{SP}_n, M' \text{ is not two-connected}\}, \\ \mathcal{P}_n &= \{M' \mid M \in \mathcal{SP}_n, M' \text{ is two-connected}\}.\end{aligned}$$

Also set

$$\mathcal{S} = \bigcup_{n \geq 2} \mathcal{S}_n, \quad \mathcal{P} = \bigcup_{n \geq 2} \mathcal{P}_n, \quad \text{and} \quad \mathcal{N} = \mathcal{S} \cup \mathcal{P} \cup \{e\}. \quad (2.1)$$

The operation of removing or adding back the root edge defines a bijection

$$\partial_{\text{Map}} : \mathcal{SP} \rightarrow \mathcal{N}, \quad (2.2)$$

see Fig. 3. We refer to the maps in $\mathcal{N}$ as *series-parallel (SP) networks*.




Fig. 3 The root edge is removed from a rooted two-connected SP map. The resulting map is either two-connected (left) or it is not (right)

2.2 Plane Trees and Labels

Recall the standard definition of the Ulam–Harris tree: set $\mathcal{U} = \bigcup_{n=0}^{\infty} \mathbb{N}^n$, where $\mathbb{N} = \{1, 2, 3, \dots\}$ and $\mathbb{N}^0 = \{\emptyset\}$ by convention. If $u = u^1, \dots, u^m$ and $v = v^1, \dots, v^n$ are elements of $\mathcal{U}$, we write the concatenation of u and v as $uv = u^1, \dots, u^m, v^1, \dots, v^n$.

Definition 2.1 A *plane tree* T is a finite or infinite subset of $\mathcal{U}$ such that

1. $\emptyset \in T$,
2. if $v \in T$ and $v = uj$ for some $j \in \mathbb{N}$, then $u \in T$.
3. for every $u \in T$, there exists an integer $\text{out}_T(u) \geq 0$ (the outdegree or number of children of u) such that for every $j \in \mathbb{N}$, $uj \in T$ if and only if $1 \leq j \leq \text{out}_T(u)$.

The element $\emptyset \in \mathcal{U}$ is called the *root vertex*, and the pair $\{\emptyset, 1\}$ is called the *root edge*. One may view plane trees as planar maps with a directed root edge, which do not contain a cycle. A vertex of outdegree 0 is called a *leaf*, and an edge incident to a leaf is called a *leaf edge*. The trivial tree consisting only of the root vertex $\emptyset$, which is then a leaf, will be denoted by $\bullet$.

We denote by $\mathcal{T}$ the set of *finite* plane trees, which have no vertices of outdegree 1, and by $\mathcal{T}_n \subseteq \mathcal{T}$ the set of such trees with $n \geq 1$ leaves. We will apply labels S and P to the trees in $\mathcal{T}$ (related to the series and parallel decompositions of maps). We use the notation $\bar{S} = P$ and $\bar{P} = S$ when we need to swap labels. If $(T, \ell) \in \mathcal{T} \times \{S, P\}$ and $T \neq \bullet$, we say that T is labeled by ℓ and we assign the label ℓ to vertices in every even generation and $\bar{\ell}$ to vertices in every odd generation. In the case, when $T = \bullet$, we take $(\bullet, P)$ and $(\bullet, S)$ to be equivalent and simply denote them by $\bullet$ and say that $\bullet$ is neither labeled by P nor S . Define $\mathcal{TP} = (\mathcal{T} \setminus \{\bullet\}) \times \{P\}$ and $\mathcal{TS} = (\mathcal{T} \setminus \{\bullet\}) \times \{S\}$ as the sets of nontrivial trees labeled by P and S , respectively, and let

$$\mathcal{TN} := \mathcal{TP} \cup \mathcal{TS} \cup \{\bullet\}. \quad (2.3)$$

The corresponding sets of trees with n leaves are denoted using the lower index n .

Next define $\mathcal{T}^* \subseteq \mathcal{T}$ as the set of trees in $\mathcal{T}$ such that the leftmost child of the root is a leaf, and let $\mathcal{T}_n^*$ denote the set of such trees with $n + 1$ leaves (i.e., n leaves not counting the leftmost child of the root). A tree $T \in \mathcal{T}^*$ is commonly called a *planted tree*. We take trees in $\mathcal{T}^*$ to be labeled so that vertices in even generations have label



Fig. 4 The removal of the first child of the root. Here, the vertices $\circ$ correspond to label P and the vertices $\bullet$ to label S

P and vertices in odd generations have label S . There is a useful bijection

$$\partial_{\text{Tree}} : \mathcal{T}^* \rightarrow \mathcal{TN}, \quad (2.4)$$

which is described as follows. Let $T \in \mathcal{T}_n^*$. If the outdegree of the root in T is strictly greater than 2, remove the leftmost child of the root (along with the edge adjacent to it). Note that the resulting tree is still labeled by P , so this yields an element in $\mathcal{TP}_n$. If the outdegree of the root in T equals 2, then remove the root and its leftmost child along with their adjacent edges. The resulting tree is labeled by S and is thus an element in $\mathcal{TS}_n$, see Fig. 4. In the case when T has exactly 1 leaf, let $\partial_{\text{Tree}} T = \bullet$.

2.3 Bijection Between SP Maps and Labeled Trees

Both SP networks $\mathcal{N}$ (2.1) and labeled trees $\mathcal{TN}$ (2.3) may be described in the language of combinatorial species [14]. This provides a natural framework for defining bijections $\varphi : \mathcal{TN} \rightarrow \mathcal{N}$ and $\varphi^* : \mathcal{T}^* \rightarrow \mathcal{SP}$, as we now describe.

First, an element of $\mathcal{N}$ is either the single edge $e = (*_0, *_1)$, or a parallel network or a series network:

$$\mathcal{N} = e + \mathcal{P} + \mathcal{S}. \quad (2.5)$$

The elements of $\mathcal{P}$ are called parallel networks and are constructed from a parallel composition of at least two non-parallel networks:

$$\mathcal{P} = \text{SEQ}_{\geq 2}(\mathcal{N} - \mathcal{P}), \quad (2.6)$$

where the combinatorial species $\text{SEQ}_{\geq 2}$ are the linear orders of length greater than or equal to two. The elements of $\mathcal{S}$ are called series networks and are constructed from a series composition of at least two non-series networks:

$$\mathcal{S} = \text{SEQ}_{\geq 2}(\mathcal{N} - \mathcal{S}). \quad (2.7)$$

The single edge network is treated separately in the above definitions since it is in some sense neither series nor parallel.

Similarly, a tree from $\mathcal{TN}$ is either the unlabeled single leaf $\bullet$ or a tree labeled by P or a tree labeled by S :

$$\mathcal{TN} = \bullet + \mathcal{TP} + \mathcal{TS}. \quad (2.8)$$

$$\begin{array}{ccc}
 \mathcal{T}^* & \xrightarrow{\varphi^*} & \mathcal{SP} \\
 \partial_{\text{Tree}} \downarrow & & \downarrow \partial_{\text{Map}} \\
 \mathcal{TN} & \xrightarrow{\varphi} & \mathcal{N}
 \end{array}$$

Fig. 5 Definition of φ^* from φ

A tree labeled by P is constructed by a composition of at least 2 trees from $\mathcal{T}_{\mathcal{N}}$, which are not labeled by P :

$$\mathcal{TP} = \text{SEQ}_{\geq 2}(\mathcal{TN} - \mathcal{TP}). \quad (2.9)$$

A tree labeled by S is constructed by a composition of at least 2 trees from $\mathcal{T}_{\mathcal{N}}$, which are not labeled by S :

$$\mathcal{TS} = \text{SEQ}_{\geq 2}(\mathcal{TN} - \mathcal{TS}). \quad (2.10)$$

The combinatorial identities (2.5), (2.6), (2.7) and (2.8), (2.9), (2.10) naturally define a bijection

$$\varphi : \mathcal{TN} \rightarrow \mathcal{N}.$$

Moreover, networks with n edges are in bijection with the set of labeled trees with n leaves and networks in $\mathcal{P}$ and $\mathcal{S}$ are in bijection with labeled trees in $\mathcal{TP}$ and $\mathcal{TS}$, respectively. Combining φ with the bijections (2.2) and (2.4) allows us to define a bijection $\varphi^* : \mathcal{T}^* \rightarrow \mathcal{SP}$ as in Fig. 5.

It will be useful to have the following recursive description of φ^* . Given a labeled tree $(T, \ell) \in \mathcal{T} \times \{S, P\}$, first assign a single edge to each of the leaves. Then, start forming networks by combining the edges corresponding to sibling leaves in series or in parallel, as determined by the label of their common parent. Proceeding recursively toward the root of the tree, each vertex corresponds to a series or a parallel network, which is constructed from the subtree of its descendants. Networks corresponding to vertices with a common parent are connected in series if that parent has label S , respectively, in parallel if it has label P . To be definite, parallel connections are made in the same order left-to-right as the order of the corresponding siblings, while series connections are made so that the left-to-right order of the siblings gives the bottom-to-top order of connections. Restricting this construction to $\mathcal{T}^*$ gives the bijection φ^* .

2.4 Blobs

We now define what we call *blobs*. Informally, blobs are parts of M that are separated by bottlenecks, in the sense that a geodesic between two vertices within the same blob has all its vertices inside that blob (see Fig. 6 for an illustration of the description that follows).




Fig. 6 Left: An example of a rooted, two-connected SP map divided into blobs. A geodesic is traced from a to b (in cyan). Right: Its corresponding planted plane tree, labeled by P , divided into segments. Here, $\circ$ and $\bullet$ denote the label P and $\bullet$ denotes the label S . Each segment corresponds to a network, and the red edges depicted in these networks are the places where networks corresponding to fringe subtrees are inserted

Suppose that we are given a rooted two-connected SP map $M \in \mathcal{SP}$ and the corresponding planted tree $T^* = (\varphi^*)^{-1}(M) \in \mathcal{T}^*$ (which we recall is labeled by P). We color a vertex of T^* red if it is marked P and has at least one child that is a leaf. In particular, the root of T^* is colored red. This way, the tree T^* is decomposed into (maximal) subtrees whose roots are red, and where any red non-root vertex is a leaf. We call these subtrees *segments*.

We define the *blobs* of M to be the networks obtained by applying the recursive construction defining φ^* (described at the end of Sect. 2.3) to the segments. We moreover declare the blobs to have an oriented red placeholder edge for each red non-root vertex in the segment, which marks the location where we would insert the network constructed from the corresponding fringe subtree in T^* .

The orientation of each red edge is naturally inherited from the orientation of the edge in the tree from the root to the red vertex, together with the orientation of the root edge of the SP map, where the tree is explored in its usual lexicographic order. We assign placeholder south and north poles to each of the blobs as the start- and endpoints of the oriented edge that corresponds to the red edge assigned to the red leaf of the corresponding parent segment. This way, the map M gets decomposed into blobs, where the placeholder poles of each blob are connected by an edge in M .

We perform a similar decomposition of networks N into blobs, where all definitions are analogous except that we demand the root vertex of the corresponding tree T to be colored red.

2.5 Bienaymé–Galton–Watson Trees and Simply Generated Trees

The main focus of this work is uniformly distributed two-connected SP maps. However, the bijective correspondence with trees that is explained in the previous section allows us easily to consider a more general measure on the maps by assigning weights to the trees in a natural way and pushing these forward to the maps using the bijection.

Let $\mu = (\mu(k))_{k \geq 0}$ be a probability measure with $\mu(1) = 0$. Let BGW^μ denote the distribution of a Bienaymé–Galton–Watson process with offspring distribution μ .

If μ has mean less than or equal to 1, then the tree τ is almost surely finite and, for any finite tree T ,

$$\text{BGW}^\mu(\tau = T) = \prod_{u \in T} \mu(\text{out}_\tau(u)). \quad (2.11)$$

Let BGW_n^μ be the corresponding measure conditioned on the number of leaves being n . The random trees distributed by BGW^μ and BGW_n^μ will be denoted by T^μ and T_n^μ , respectively.

In order to relate the BGW trees to the random maps, it will be useful to consider another (almost) equivalent model of random trees called *simply generated trees with leaves as atoms*. Let $\eta = (\eta(i))_{i \geq 0}$ be a sequence of nonnegative numbers called *weights* and assume that $\eta(1) = 0$. Let $T \in \mathcal{T}_n$ and define the weight of T by

$$\eta(T) = \prod_{u \in T} \eta(\text{out}_T(u)).$$

A *simply generated tree* on n leaves is a random plane tree τ with values in $\mathcal{T}_n$ with a probability distribution

$$\mathbb{P}(\tau = T) = \frac{\eta(T)}{\sum_{T' \in \mathcal{T}_n} \eta(T')}, \quad T \in \mathcal{T}_n. \quad (2.12)$$

The literature contains many results on simply generated trees with a fixed number of vertices, and for these we refer to [13] for details. The current case, where instead the number of leaves is fixed, is considered in, e.g., [21]. Below, we outline the main results which we require.

For $a, b > 0$, let $\eta^{a,b}$ be another weight sequence defined by

$$\eta^{a,b}(k) = \begin{cases} a\eta(0) & k = 0 \\ b^{k-1}\eta(k) & k \geq 2. \end{cases} \quad (2.13)$$

These weights are equivalent to η in the sense that they define the same measure as in (2.12). Define the generating series

$$g^\eta(z) = \sum_{n=0}^{\infty} \eta(n)z^n \quad (2.14)$$

and denote its radius of convergence by ρ_η . We will assume that the weights η are chosen such that $\rho_\eta > 0$. Then, for any $b > 0$ such that $\frac{g^\eta(b) - \eta(0)}{\eta(0)b} < 1$, one may choose

$$a = \frac{1}{\eta(0)} - \frac{g^\eta(b) - \eta(0)}{\eta(0)b}$$

in which case $\eta^{a,b}$ is a probability measure which we will denote by $\eta^{(b)}$. Denote its mean by

$$\mathbf{m}_b = \sum_{k=2}^{\infty} k\eta^{(b)}(k) = \sum_{k=2}^{\infty} k\eta(k)b^{k-1} = (g^\eta)'(b).$$

The mean is a strictly increasing function of b and $\mathbf{m}_0 = 0$. In the current work, we will assume that there is a $b < \rho_\eta$ such that $\mathbf{m}_b = 1$, which is then necessarily unique. The corresponding probability measure will be denoted by μ , and the condition $\mathbf{m}_b = 1$ is referred to as *criticality*. The condition $b < \rho_\eta$ furthermore guarantees that μ has finite exponential moments. It is straightforward to check that the simply generated tree with weights η has the same distribution as T_n^μ .

2.6 Weighted Trees and Maps

Let η be a weight sequence such that there exists a corresponding critical offspring distribution μ . Generate a random planted tree, denoted by $T_n^{\mu,*}$, as follows:

1. Let T_n^μ have distribution BGW_n^μ .
2. Flip a fair coin and label T_n^μ by P or S according to the outcome of the coin flip.
3. Plant the resulting labeled tree, i.e., set $T_n^{\mu,*} := \partial_{\text{Tree}}^{-1} T_n^\mu$.

Definition 2.2 The random, weighted, rooted, two-connected SP map with n non-root edges is defined as:

$$\text{SP}_n^\mu := \varphi^*(T_n^{\mu,*}).$$

The corresponding random network is denoted by:

$$\mathbf{N}_n^\mu = \partial_{\text{Map}} \text{SP}_n^\mu = \varphi(T_n^\mu).$$

The networks obtained by conditioning $\mathbf{N}_n^\mu$ on being a series or a parallel network (i.e., on the outcome of the coin flip above) are denoted by S_n^μ and P_n^μ , respectively.

In order to lighten the notation, we will often suppress the μ and write, e.g., $\text{SP}_n = \text{SP}_n^\mu$, $\mathbf{N}_n = \mathbf{N}_n^\mu$, $\mathbf{T} = \mathbf{T}^\mu$, $\mathbf{T}_n^* = T_n^{\mu,*}$, and $\mathbf{T}_n = T_n^\mu$.

Example 2.3 The uniformly distributed map SP_n is obtained by choosing $\eta(k) = 1$ for all $k \geq 0$, $k \neq 1$. In this case, one finds that

$$\mu(0) = 2 - \sqrt{2}, \quad \mu(1) = 0, \quad \mu(k) = \left(1 - \frac{1}{\sqrt{2}}\right)^{k-1}, \quad k \geq 2. \quad (2.15)$$

Example 2.4 Let $N \geq 2$ be an integer. The map SP_n , which is distributed uniformly among those where each series and each parallel composition in the network $\mathbf{N}_n$ consists of exactly N components, is obtained by choosing the weights $\eta(0) = \eta(N)$

and all other weights equal to 0. In the corresponding labeled tree T_n , every vertex has outdegree 0 or N and the root has outdegree 1 or N . In this case, one finds that $\mu(0) = 1 - \frac{1}{N}$, $\mu(N) = \frac{1}{N}$.

3 Proofs of the Main Results

We now turn to the proofs of Theorems 1.1 and 1.2. The main method is to compare lengths of geodesics between vertices in a series-parallel map and between corresponding vertices in its associated tree. For concreteness, we will work with the random parallel networks P_n . The associated tree is then a BGW tree $T_n = T_n^\mu$ labeled by P . The proof can easily be adapted to the maps S_n , N_n , and SP_n with exactly the same outcome.

Let $T = T^\mu$ be an unconditioned BGW tree and let ξ be a random variable with distribution μ . We will use the *local limit* $\hat{T}$ of T_n (as $n \rightarrow \infty$), which is an infinite random tree with a unique half-infinite path starting from the root, called its *spine*. In $\hat{T}$ non-spine vertices have offspring according to independent copies of ξ , whereas spine vertices have offspring according to independent copies of the size-biased random variable $\hat{\xi}$ with distribution $\mathbb{P}(\hat{\xi} = k) = k\mathbb{P}(\xi = k)$, and the successor spine vertex is chosen uniformly at random among these children. For further details, we refer to Janson's extensive review [13].

For each $\ell \geq 1$, let $T^{(\ell)}$ denote the tree obtained by deleting all descendants of the $(\ell + 1)$ st spine vertex of $\hat{T}$ and identifying the $(\ell + 1)$ st spine vertex with the root of an independent copy of T . We view all the random trees T_n , T , $\hat{T}$, and $\hat{T}^{(\ell)}$ for $\ell \geq 1$ as being colored (i.e., some vertices are red) and labeled by P or S according to the fair coin flip as defined in Sect. 2.

We may define infinite random series-parallel maps/networks SP , N , P and S in a natural way by extending the functions φ and φ^* to infinite trees. The maps/networks may then be viewed as the weak local limits of the corresponding sequences of finite maps/networks. These infinite objects are not strictly needed in the following arguments. However, it is convenient to refer to them occasionally, so for completeness, we provide some details of their construction in the appendix.

We may generate T by starting with a root segment and a list of i.i.d. non-root segments. Likewise, $\hat{T}$ can be generated from a root segment, a list of i.i.d. normal non-root segments, and a list of i.i.d. spine non-root segments. The tree $\hat{T}^{(\ell)}$ additionally has a mixed segment containing the tip of the spine.

Lemma 3.1 *The number of vertices in each type of segment (root, spine, and normal) has finite exponential moments.*

Proof The number of vertices in each random segment may be stochastically bounded by two times the total population of a subcritical branching process (with possibly a modified root degree) whose offspring distribution is light-tailed. The total population of such a process is light-tailed itself by [13, Thm. 18.1].

We detail this argument for the root segment of T^μ . The proofs for all other types of segments are analogous. We construct a BGW tree T^ν with offspring probabilities

$(v_n)_n$ from T^μ as follows. Let T' be the root segment of T^μ . The tree T^ν is defined as the tree with vertex set consisting of the vertices in even generations in T' and vertices in T^ν are connected by an edge if and only if one is the grandchild of the other in T' . Note that since T' has no vertices of outdegree 1, the number of vertices in T' is less than two times the number of its leaves and T^ν has the same number of leaves as T' . Therefore, the number of vertices in T' is bounded by two times the number of leaves in T^ν .

Another way to describe T^ν is the following. Let v be a vertex in T^μ at an even generation whose children we have not yet generated and assume that v is also present in T^ν . If v either has no children in T^μ or it has more than one child and some of its children have no children, then we say that v has no children in T^ν . On the other hand, if v has children in T^μ and all of its children also have children, then we add these grandchildren as the children of v in T^ν . From this description, we find that

$$\begin{aligned} v(0) &= 1 - \sum_{n=2}^{\infty} \mu(n)(1 - \mu(0))^n, \\ v(1) &= 0, \\ v(k) &= \sum_{n=2}^{\infty} \mu(n) \sum_{i_1+\dots+i_n=k} \prod_{j=1}^n \mu(i_j) \mathbb{1}_{\{i_j > 0\}}, \quad k \geq 2. \end{aligned}$$

The generating function g of the probabilities $(v_n)_n$ satisfies

$$g(z) = v(0) + h(h(z)) \quad (3.1)$$

where $h(z) = f(z) - \mu_0$ and f is the generating function of the sequence $(\mu_n)_n$. We then have

$$g'(1) = h'(h(1))h'(1) = h'(1 - \mu(0)) = f'(1 - \mu(0)) < 1$$

since $f'(1) = 1$. Thus, the ν -process is subcritical.

Since μ has finite exponential moments, we find that ν also has finite exponential moments. Thus, the total population of T^ν (and hence also the size of the root segment of T) has finite exponential moments by [13, Thm. 18.1]. $\square$

Corollary 3.2 *There are constants $C, c > 0$ such that the probability for the maximal size of a segment in T_n to be larger than x is bounded by $Cn^{5/2} \exp(-cx)$ uniformly for all $n \geq 2$ and $x > 0$.*

Proof The probability for T to have n leaves satisfies

$$\mathbb{P}(|T| = n) \sim c_{\text{cond}} n^{-3/2} \quad (3.2)$$

for some constant $c_{\text{cond}} > 0$, see, for example, [15]. Let $(X_i)_i$ denote the sizes of the segments in T . There are at most n segments in T_n ; hence, the probability for the

maximal size of a segment in T_n to be larger than $x > 0$ satisfies

$$\begin{aligned} \mathbb{P}\left(\max_i X_i > x \mid |T| = n\right) &= \frac{\mathbb{P}(\cup_i \{X_i > x\}, |T| = n)}{\mathbb{P}(|T| = n)} \\ &\leq c_{\text{cond}}^{-1} n^{3/2} \sum_{i=1}^n \mathbb{P}(X_i > x), \end{aligned}$$

where $|T|$ denotes the number of leaves in T . By Lemma 3.1, each segment of the unconditioned tree has a finite exponential moment and the segments are identically distributed, except possibly the root segment. Thus, each term in the sum is bounded by a constant times e^{-cx} , for some $c > 0$, and the result follows. $\square$

In the infinite map $\mathbf{P}$ corresponding to the tree $\hat{T}$, the blobs $B_0, B_1, \dots$ corresponding to the spine segments are called *spine blobs*. The spine blobs are independent, and the non-root spine blobs $(B_i)_{i \geq 1}$ are identically distributed. For each $\ell \geq 1$, consider the graph distance X_ℓ from the south pole of B_0 to the south pole of B_ℓ .

Lemma 3.3 *There exists a constant $\eta > 0$ such that for each $\epsilon > 0$, there are constants $C, c > 0$ with*

$$\mathbb{P}(|X_\ell - \eta\ell| > \epsilon\ell) \leq C \exp(-c\ell)$$

uniformly for all $\ell \geq 1$.

Proof Consider the graph distance Y_ℓ from the south pole v of B_1 to the south pole of B_ℓ . This way, $|X_\ell - Y_\ell|$ is bounded by the size of B_0 , which has finite exponential moments by Lemma 3.1. It hence suffices to show that there exists a constant $\eta > 0$ such that for each $\epsilon > 0$, there are constants $C, c > 0$ such that for all $\ell \geq 1$,

$$\mathbb{P}(|Y_\ell - \eta\ell| > \epsilon\ell) \leq C \exp(-c\ell). \quad (3.3)$$

Since the poles of B_ℓ are joined by an edge, the graph distances from v to either of them (north or south) differ by at most 1. Hence, the increment $Y_{\ell+1} - Y_\ell$ only depends on this difference $S_\ell \in \{-1, 0, 1\}$ and B_ℓ . Thus, $(Y_\ell, S_\ell)_{\ell \geq 1}$ is a Markov-additive process with a driving chain $(S_\ell)_{\ell \geq 1}$ whose state space has three elements. Inequality (3.3) now immediately follows from a general large deviation result for Markov-additive processes [12, Thm. 5.1] whose conditions are satisfied since the state space of the driving chain under consideration is finite [12, Rem. 3.5, Sec 7 (ii)]. $\square$

Remark 3.4 We leave out details concerning the exact transition probabilities of the Markov-additive process in the above proof but refer to the paper by Curien, Haas, and Kortchemski [9] for a similar situation. In our case, the increments of the additive component $(Y_\ell)_{\ell \geq 1}$ of the chain are not as explicit as in [9] (they are given in terms of distances between certain vertices in a spine blob), and therefore, it is less useful to write them down.

Lemma 3.5 *Let Z_ℓ be the number of segments encountered on the spine from the root of $\hat{T}$ to the $(\ell + 1)$ st spine vertex. There exists a constant $\kappa > 0$ such that for each $\epsilon > 0$, there are constants $C, c > 0$ such that for all $\ell \geq 1$,*

$$\mathbb{P}(|Z_\ell - \kappa\ell| > \epsilon\ell) \leq C \exp(-c\ell).$$

Proof Let E_m denotes the number of edges on the spine we need to cross in order to get from the root vertex to the start of the $(m + 1)$ st segment. Set $E_0 = 0$. Hence, the differences $E_m - E_{m-1}$ are independent for $\ell \geq 1$ and identically distributed for all $m \geq 2$. These increments have a geometric distribution, since they may be identified by a sequence of independent Bernoulli coin flips.

We make use of the following inequality given, for example, in [18, Example 1.4]: Let S_n denotes the sum of n i.i.d. real-valued centered random variables with finite exponential moments. Then, there are constants $\lambda_0, c > 0$ such that for all $n \in \mathbb{N}$, $x > 0$ and $0 \leq \lambda \leq \lambda_0$ it holds that

$$\mathbb{P}(|S_n| \geq x) \leq 2 \exp(cn\lambda^2 - \lambda x). \quad (3.4)$$

It follows that there exists $\kappa^{-1} > 0$ such that for each $\delta > 0$, there exist constants $C, c > 0$ with

$$\mathbb{P}(|E_m - \kappa^{-1}m| > \delta m) \leq C \exp(-cm), \quad \text{for all } m \geq 1.$$

As $\delta > 0$ is arbitrary, it follows by choosing $\delta = \pm \left(1 - \frac{1}{1 \pm \epsilon}\right) \kappa^{-1}$ that for any $\epsilon > 0$, there exist constants $C', c' > 0$ with

$$\mathbb{P}(\ell \notin [E_{\lfloor \kappa\ell(1-\epsilon) \rfloor}, E_{\lfloor \kappa\ell(1+\epsilon) \rfloor}]) \leq C' \exp(-c'\ell).$$

Since $\ell \notin [E_{\lfloor \kappa\ell(1-\epsilon) \rfloor}, E_{\lfloor \kappa\ell(1+\epsilon) \rfloor}]$ is equivalent to $|Z_\ell - \kappa\ell| > \kappa\epsilon\ell$, the proof is complete. $\square$

Before proceeding with the proof, we recall some facts about the Gromov–Hausdorff metric d_{GH} on the space of isometry classes of compact metric spaces. Informally, two compact metric spaces are close in the Gromov–Hausdorff metric d_{GH} if they may be isometrically embedded in the third space Z so that their Hausdorff distance in Z is small. Here, we will use a representation of d_{GH} in terms of distortions of correspondences, see, e.g., [7]. A correspondence $\mathcal{R}$ between two metric spaces (A, d_A) and (B, d_B) is a subset $\mathcal{R} \subseteq A \times B$ such that for any $a \in A$, there is a $b \in B$ such that $(a, b) \in \mathcal{R}$, and for any $b \in B$, there is an $a \in A$ such that $(a, b) \in \mathcal{R}$. A distortion of a correspondence $\mathcal{R}$ is defined by:

$$\text{dis}(\mathcal{R}) = \sup\{|d_A(a_1, a_2) - d_B(b_1, b_2)| : (a_1, b_1), (a_2, b_2) \in \mathcal{R}\}.$$

The Gromov–Hausdorff metric may then be written as:

$$d_{\text{GH}}(A, B) = \frac{1}{2} \inf_{\mathcal{R}} \text{dis}(\mathcal{R}).$$

By the main theorem of [15, 19], there exists a constant $c_{\text{tree}} > 0$ such that

$$(\mathbb{T}_n, c_{\text{tree}} n^{-1/2} d_{\mathbb{T}_n}) \xrightarrow{d} (\mathcal{T}_e, d_{\mathcal{T}_e}) \quad (3.5)$$

in the Gromov–Hausdorff topology, with $(\mathcal{T}_e, d_{\mathcal{T}_e})$ denoting Aldous’ Brownian tree. The rest of the proof involves defining a suitable correspondence between the spaces $\mathbb{P}_n$ and $\mathbb{T}_n$ and showing that the distortion of that correspondence, with properly rescaled metrics, converges to 0 in probability.

We define a correspondence R_n between $\mathbb{P}_n$ and $\mathbb{T}_n$ as follows. Note that there is a bijective relationship between blobs in $\mathbb{P}_n$ and segments in $\mathbb{T}_n$. We let $(u, v) \in R_n$ if and only if there is a blob in $\mathbb{P}_n$, which contains u and a corresponding segment in $\mathbb{T}_n$ which contains v . In that case, we write $u \sim v$.

Let $h_{\mathbb{P}_n}(u)$ denotes the graph distance in the map $\mathbb{P}_n$ from the south pole $*_0$ to u , and $h_{\mathbb{T}_n}(v)$ the height of the vertex v in $\mathbb{T}_n$. The following lemma is the first step in relating distances in $\mathbb{P}_n$ to distances in $\mathbb{T}_n$.

Lemma 3.6 *Let η be the constant in Lemma 3.3, κ the constant in Lemma 3.5, and let $\delta > 0$ be arbitrary. With probability tending to one as $n \rightarrow \infty$, there exist no vertices $u \in \mathbb{P}_n$ and $v \in \mathbb{T}_n$ such that $u \sim v$ and*

$$|h_{\mathbb{P}_n}(u) - \eta \kappa h_{\mathbb{T}_n}(v)| \geq \delta \sqrt{n}. \quad (3.6)$$

Proof Let $\epsilon > 0$ be given. Let $H_n = H\sqrt{n}$ for some $H > 0$. By (3.5), we may choose H sufficiently large such that the height $H(\mathbb{T}_n)$ satisfies

$$\mathbb{P}(H(\mathbb{T}_n) > H_n) \leq \epsilon$$

for all n . Let us call a vertex v in a tree τ *bad* if it corresponds to some vertex u in the associated map P such that Inequality (3.6) holds (with obvious replacements of $\mathbb{P}_n$ and $\mathbb{T}_n$ by P and τ). Let $\mathcal{E}(\tau)$ denote the property that a tree T contains a bad vertex. Let us denote the tip of the spine in $\hat{\mathbb{T}}^{(\ell)}$ by U_ℓ . Given any finite plane tree T with a vertex u having height ℓ , it holds that

$$\mathbb{P}\left(\left(\hat{\mathbb{T}}^{(\ell)}, U_\ell\right) = (T, u)\right) = \mathbb{P}(T = T)$$

(see, e.g., the short derivation in [20, Appendix A3]). Hence, using (3.2), we obtain

$$\begin{aligned}\mathbb{P}(\mathcal{E}) &\leq \epsilon + \mathbb{P}(\mathcal{E}(T_n), H(T_n) \leq H_n) \\ &\leq \epsilon + \frac{2n}{\mathbb{P}(|T| = n)} \sum_{\ell=0}^{\lfloor H_n \rfloor} \mathbb{P}(U_\ell \text{ is bad}) \\ &= \epsilon + O(n^{5/2}) \sum_{\ell=0}^{\lfloor H_n \rfloor} \mathbb{P}(U_\ell \text{ is bad}).\end{aligned}$$

The second inequality is explained as follows. We may write:

$$\begin{aligned}\mathbb{P}(\mathcal{E}(T_n), H(T_n) \leq H\sqrt{n})\mathbb{P}(|T| = n) &= \sum_{\tau \in \mathcal{T}} \mathbb{1}\{\mathcal{E}(\tau), H(\tau) \leq H_n, |\tau| = n\} \mathbb{P}(T = \tau) \\ &\leq \sum_{\tau \in \mathcal{T}, u \in \tau} \mathbb{1}\{u \text{ is bad}, H(\tau) \leq H_n, |\tau| = n\} \mathbb{P}(T = \tau) \\ &= \sum_{\tau \in \mathcal{T}, u \in \tau} \sum_{\ell=0}^{H(\tau)} \mathbb{1}\{u \text{ is bad}, H(\tau) \leq H_n, |\tau| = n, h_\tau(u) = \ell\} \mathbb{P}\left((\hat{T}^{(\ell)}, U_\ell) = (\tau, u)\right) \\ &\leq 2n \sum_{\ell=0}^{\lfloor H_n \rfloor} \sum_{\tau \in \mathcal{T}} \mathbb{P}\left(\hat{T}^{(\ell)} = \tau, U_\ell \text{ is bad}\right) = 2n \sum_{\ell=0}^{\lfloor H_n \rfloor} \mathbb{P}(U_\ell \text{ is bad}).\end{aligned}$$

For the last inequality, we used the fact that if a plane tree has n leaves and no inner vertex with exactly one child, then it has at most $2n$ vertices in total.

If u is a vertex in the map corresponding to $\hat{T}^{(\ell)}$ such that u corresponds to the tip of the spine of $\hat{T}^{(\ell)}$, then the height of u may be bounded by the total number of leaves in the root segments of the subtrees attached to the spine vertices of $\hat{T}^{(\ell)}$. Denote by W_i the total size of all the root segments in the subtrees attached to the spine vertex at distance i from the root, where $0 \leq i \leq \ell$. Note that they are independent and i.i.d for $i < \ell$. At each of these spine vertices, the total number of subtrees is distributed as $\hat{\xi} - 1$ or $\hat{\xi}$ (at the tip of the spine) and is thus light-tailed. Furthermore, the total size of the root segment of each of the subtrees is light-tailed by Lemma 3.1. We conclude that W_i is light-tailed for all $0 \leq i \leq \ell$.

If $\ell < \delta/(2\eta\kappa)\sqrt{n}$, then the event that U_ℓ is bad implies that the height of u is larger than $\delta\sqrt{n}$ and, therefore, that

$$\sum_{i=0}^{\ell} W_i \geq \delta\sqrt{n}.$$

Since the W_i is i.i.d and light-tailed, it follows by the medium deviation Inequality (3.4) that if we take $h > 0$ small enough, then there exist constants $C_1, c_1 > 0$ such that

uniformly for all n and $1 \leq \ell \leq h\sqrt{n}$,

$$\mathbb{P}(U_\ell \text{ is bad}) \leq \mathbb{P}\left(\sum_{i=0}^{h\sqrt{n}} W_i \geq \delta\sqrt{n}\right) \leq C_1 \exp(-c_1\sqrt{n}). \quad (3.7)$$

This ensures

$$O(n^{5/2}) \sum_{\ell=0}^{\lfloor h\sqrt{n} \rfloor} \mathbb{P}(U_\ell \text{ is bad}) = O(n^3) \exp(-c_1\sqrt{n}) = \exp(-\Theta(\sqrt{n})).$$

Furthermore, uniformly for $h\sqrt{n} \leq \ell \leq H\sqrt{n}$, we know by Lemma 3.5 that for any $c > 0$, we have $|Z_\ell - \kappa\ell| \leq c\ell$ with probability at least $1 - \exp(-\Theta(\sqrt{n}))$, which allows us to apply Lemma 3.3 to obtain for any $c > 0$ that $|X_{Z_\ell} - \eta\kappa\ell| \leq c\ell$ again with probability at least $1 - \exp(-\Theta(\sqrt{n}))$. Recall that we constructed $\hat{T}^{(\ell)}$ from $\hat{T}$ by replacing the descendants of the tip of the spine, therefore creating a “mixed” segment that contains the tip of the spine. Hence, in the map corresponding to $\hat{T}^{(\ell)}$, any vertex corresponding to the tip of the spine has a height that differs from X_{Z_ℓ} at most by a summand that by Lemma 3.1 has finite exponential moments. Thus, we obtain:

$$O(n^{5/2}) \sum_{h\sqrt{n} \leq \ell \leq H\sqrt{n}} \mathbb{P}(U_\ell \text{ is bad}) = O(n^3) \exp(-\Theta(\sqrt{n})) = \exp(-\Theta(\sqrt{n})).$$

In summary,

$$\mathbb{P}(\mathcal{E}) \leq \epsilon + \exp(-\Theta(\sqrt{n})).$$

Since $\epsilon > 0$ was arbitrary, this completes the proof. $\square$

We are now ready to prove our main theorems.

Proof of Theorem 1.1 By Corollary 3.2 and Lemma 3.6, we know that, for some constant $C > 0$ and some deterministic sequence $t_n = o(\sqrt{n})$, we have with high probability that the largest blob in P_n has diameter at most $C \log n$, and whenever a vertex u in P_n corresponds to a vertex v in T_n , then

$$|h_{P_n}(u) - \eta\kappa h_{T_n}(v)| \leq t_n. \quad (3.8)$$

Now, let u_1, u_2 be vertices in P_n and let v_1 and v_2 be any vertices in T_n such that v_i corresponds to u_i for $i = 1, 2$. If v_3 denotes the lowest common ancestor of v_1 and v_2 in T_n , then

$$d_{T_n}(v_1, v_2) = h_{T_n}(v_1) + h_{T_n}(v_2) - 2h_{T_n}(v_3).$$

Any geodesic in P_n between u_1 and u_2 must pass through a blob corresponding to v_3 . Letting u_3 denotes some vertex in this blob, it follows that:

$$d_{P_n}(u_1, u_2) = h_{P_n}(u_1) + h_{P_n}(u_2) - 2h_{P_n}(u_3) + O(\log n).$$

By Inequality (3.8), it follows that:

$$d_{P_n}(u_1, u_2) = \eta \kappa d_{T_n}(v_1, v_2) + o(\sqrt{n}).$$

Thus, the correspondence between $\frac{1}{\eta \kappa} d_{P_n}$ and d_{T_n} has distortion $o(\sqrt{n})$. By Equation (3.5), it follows that"

$$\left(P_n, \frac{c_{\text{tree}}}{\eta \kappa} n^{-1/2} d_{P_n} \right) \xrightarrow{d} (\mathcal{T}_e, d_{\mathcal{T}_e}) \quad (3.9)$$

in the Gromov–Hausdorff topology. $\square$

Proof of Theorem 1.2 Since the diameter is bounded by twice the height with respect to the origin of the root edge, it suffices to show that there exists $C, c > 0$ with

$$\mathbb{P}(H(P_n) > x\sqrt{n}) \leq C \exp(-cx^2).$$

By possibly adjusting the constants C, c , it suffices to show this inequality for all $x \geq 1$. Furthermore, the left-hand side equals to zero if $x \geq \sqrt{n}$. Hence, it suffices to treat the case $1 \leq x \leq \sqrt{n}$.

We now use the same notation and argue analogously as in the proof of Lemma 3.6, with the exception that we call a vertex v in T_n *bad* if it corresponds to some vertex u in P_n such that its height in P_n is larger than $x\sqrt{n}$. Analogously, we define for any $\ell \geq 1$ whether the tip of the spine of $\hat{T}^{(\ell)}$ is *bad*, and we let $\mathcal{E}$ denotes the event that T_n contains a bad vertex.

By [21, Lem. 6.61], there exists $C_1, c_1 > 0$ with

$$\mathbb{P}(H(T_n) > y\sqrt{n}) \leq C_1 \exp(-c_1 y^2)$$

for any $y > 0$. Hence, it follows as in the proof of Lemma 3.6 that for any fixed $h > 0$

$$\begin{aligned} \mathbb{P}(\mathcal{E}) &\leq \mathbb{P}(H(T_n) > (x/h)\sqrt{n}) + \mathbb{P}(\mathcal{E}, H(T_n) \leq (x/h)\sqrt{n}) \\ &\leq C_1 \exp(-c_1 x^2/h^2) + O(n^{5/2}) \sum_{\ell=0}^{\lfloor (x/h)\sqrt{n} \rfloor} \mathbb{P}(U_\ell \text{ is bad}). \end{aligned}$$

Furthermore, applying Inequality (3.4) as before yields that there exist constants $\lambda_0, \mu_{\text{blob}}, c_2 > 0$ such that for all $n \in \mathbb{N}$ and $0 \leq \lambda \leq \lambda_0$ it holds that

$$\mathbb{P}(U_\ell \text{ is bad}) \leq 2 \exp(c\ell\lambda^2 - \lambda(x\sqrt{n} - \ell\mu_{\text{blob}})).$$

As $h > 0$ was fixed but arbitrary, we may take h large enough such that $\mu_{\text{blob}}/h < 1/2$. As $x \geq 1$, it follows that uniformly for $\ell \leq (x/h)\sqrt{n}$:

$$c\ell\lambda^2 - \lambda(x\sqrt{n} - \ell\mu_{\text{blob}}) \leq c\ell\lambda^2 - \lambda x\sqrt{n}/2 = -x\Theta(\sqrt{n}).$$

Hence,

$$\mathbb{P}(H(P_n) > x\sqrt{n}) = \mathbb{P}(\mathcal{E}) \leq C_1 \exp(-c_1 x^2/h^2) + \exp(-x\Theta(\sqrt{n})).$$

Since we assumed $1 \leq x \leq \sqrt{n}$, it follows that:

$$\mathbb{P}(H(P_n) > x\sqrt{n}) \leq C_3 \exp(-c_3 x^2)$$

for some constants $C_3, c_3 > 0$ that do not depend on n . $\square$

4 Conclusions and Related Classes

Our results imply, in particular, a scaling limit of uniformly sampled two-connected rooted series-parallel maps. From a different viewpoint, it has been shown in [5] that rooted series-parallel maps are in bijection with permutations avoiding the patterns 2413 and 3142. As a side result, it was also shown that a bipolar planar map admits a unique bipolar orientation precisely when it is series-parallel. These observations link series-parallel maps to the large framework of separable permutations and their connections to the Liouville quantum gravity, which has enjoyed remarkable developments during the last few years. For example, it is conjectured in [6] that conformally embedded rooted series-parallel maps converge toward the 2-LQG quantum sphere. Therefore, we expect interesting research directions following from our work.

Apart from planar maps, we may also consider planar graphs subject to constraints. The terminology for graph classes is a bit at odds with the one used here for planar maps. The class of series-parallel graphs refers to (simple) graphs that do not admit the K_4 as a minor, which may very well be separable (that is, not two-connected). This class is an important example of a subcritical graph class and hence admits the Brownian tree as scaling limit by [17].

Appendix A Local Metrics and Continuity

For completeness, we provide here some details on the constructions of the infinite objects $\mathcal{SP}$, $\mathcal{N}$, $\mathcal{P}$, and $\mathcal{S}$.

Recall the sets $\mathcal{T}$ (finite plane trees with no vertices of outdegree 1) and $\mathcal{T}^*$ (trees in $\mathcal{T}$ where the leftmost child of the root is a leaf) from Sect. 2.2. Denote by $\overline{\mathcal{T}}$ the set of (possibly infinite) plane trees, which have no vertices of outdegree 1, and by $\mathcal{T}_\infty = \overline{\mathcal{T}} \setminus \mathcal{T}$ the infinite trees. Denote by $\overline{\mathcal{T}}^*$ the set of trees in $\overline{\mathcal{T}}$ such that the leftmost child of the root is a leaf and let $\mathcal{T}_\infty^* = \overline{\mathcal{T}}^* \setminus \mathcal{T}^*$. As for $\mathcal{T}^*$, we take trees in $\overline{\mathcal{T}}^*$ to be labeled so that vertices in even generations have label P and vertices in odd generations have label S .

We define *local metrics* and *local topologies* on $\overline{\mathcal{T}}$, $\overline{\mathcal{T}}^*$, and $\mathcal{SP}$ as follows. First, for any graph G with a distinguished vertex v , we let $B_r(G; v)$ denote the subgraph of G spanned by vertices at distance $\leq r$ from the vertex v . For two graphs G and G'




Fig. 7 An infinite tree $T \in \hat{\mathcal{T}}_\infty^*$ with its first two RL pairs

belonging to some set of graphs $\mathcal{G}$, we let

$$d_{\mathcal{G}}(G, G') = \frac{1}{1 + \max\{r \in \mathbb{N} : B_r(G; v) = B_r(G'; v)\}}.$$

The topology generated by this metric is called the *local topology* with respect to v . In the case when $\mathcal{G}$ is one of the sets $\overline{\mathcal{T}}$, $\overline{\mathcal{T}}^*$ we take $v = \emptyset$ and use the notation $d_{\overline{\mathcal{T}}}$ for the metric in both cases. When $\mathcal{G}$ is the set $\mathcal{SP}$, we use $v = *_{\mathcal{O}}$.

Using the local metric, we can define *infinite series-parallel maps*. Namely, we let $\overline{\mathcal{SP}}$ denote the completion of $\mathcal{SP}$ with respect to the metric $d_{\mathcal{SP}}$. Then, $\mathcal{SP}_\infty := \overline{\mathcal{SP}} \setminus \mathcal{SP}$ is the set of infinite series-parallel maps, which are formally equivalence classes of Cauchy sequences of finite series-parallel maps. We now show how to extend the bijection $\varphi^* : \mathcal{T}^* \rightarrow \mathcal{SP}$ of Sect. 2.3 to infinite trees and maps.

A *spine* of an infinite tree is an infinite path $(V(0), V(1), V(2), \dots)$ such that $V(0) = \emptyset$ and $V(i+1)$ is a child of $V(i)$ for all $i \geq 0$. At every vertex $V(i)$ on the spine, there is one subtree rooted on the left of the spine and another one on the right of the spine, which we refer to as the left and right *outgrowths* of the spine. Let $\hat{\mathcal{T}}_\infty^*$ be the subset of $\mathcal{T}_\infty^*$ consisting of trees with a *unique* spine and such that this unique spine has infinitely many outgrowths in odd generations on both the left and the right.

Now let $T \in \hat{\mathcal{T}}_\infty^*$ with unique spine $(V_T(0), V_T(1), V_T(2), \dots)$. We inductively define a sequence $(i_k, j_k)_{k \geq 1}$ of pairs of *odd* integers $i_k \leq j_k$ as follows, see Fig. 7 for an illustration. First, i_1 is the smallest odd integer i so that $V_T(i)$ has a non-empty outgrowth on the right of the spine, and j_1 is the smallest odd integer $j \geq i_1$ so that $V_T(j)$ has a non-empty outgrowth on the left of the spine. Inductively, i_{k+1} is the smallest odd integer $i > j_k$ so that $V_T(i)$ has a non-empty outgrowth on the right of the spine, and j_{k+1} is the smallest odd integer $j \geq i_{k+1}$ so that $V_T(j)$ has a non-empty outgrowth on the left of the spine. We call $(V_T(i_k), V_T(j_k))_{k \geq 1}$ the *RL pairs* of T . Since i_k, j_k are odd, $V_T(i_k), V_T(j_k)$ are labeled by S . Note that $\hat{\mathcal{T}}_\infty^*$ is defined so that T has infinitely many RL pairs.

Proposition A.1 Let $[T]_m := B_m(T; \emptyset)$ denote the first m generations of T .

- If $T \in \hat{\mathcal{T}}_\infty^*$, then $\varphi^*([T]_m)$ is a Cauchy sequence with respect to $d_{\mathcal{SP}}$. Thus, $\varphi^*(T) := \lim_{m \rightarrow \infty} \varphi^*([T]_m)$ gives a well-defined function $\mathcal{T}^* \cup \hat{\mathcal{T}}_\infty^* \rightarrow \overline{\mathcal{SP}}$.
- If a sequence $(T_n)_{n \geq 1}$ in $\mathcal{T}^* \cup \hat{\mathcal{T}}_\infty^*$ converges to $T \in \mathcal{T}^* \cup \hat{\mathcal{T}}_\infty^*$ with respect to the metric $d_{\overline{\mathcal{T}}}$, then $\varphi^*(T_n) \rightarrow \varphi^*(T)$ with respect to $d_{\mathcal{SP}}$.

Proof The key observation is that if $T \in \hat{\mathcal{T}}_\infty^*$ and $v \in T$ is a leaf that is beyond the first r RL pairs, then for any m large enough such that $v \in [T]_m$, the edge e_v corresponding to v in $\varphi^*([T]_m)$ is at least at distance r from both poles $*_0$ and $*_1$. This follows from the recursive description of φ^* , where the non-empty outgrowths of each RL pair become non-empty networks placed in series on either side between e_v and the poles $*_0$ and $*_1$. Each of these networks contributes to the distance from the poles to e_v at least by one and to the total distance at least by r . Due to this observation, the proof can be completed as follows.

- (a) If $[T]_m$ contains the first $r + 1$ RL pairs for all $m > m_0(r)$ then $\forall m, m' > m_0(r)$, $B_r(\varphi^*([T]_m)) = B_r(\varphi^*([T]_{m'}))$ since all the edges beyond the first $r + 1$ RL pairs do not belong to $B_r(\varphi^*([T]_m))$.
- (b) Let $(T_n)_{n \geq 1}$ be a sequence in $\hat{\mathcal{T}}^*$ converging to $T \in \mathcal{T}^* \cup \hat{\mathcal{T}}_\infty^*$. If $T \in \mathcal{T}^*$, then $T_n = T$ for every n large enough so that $\varphi^*(T_n) \rightarrow \varphi^*(T)$ is obvious. Now, we assume that $T \in \hat{\mathcal{T}}_\infty^*$. Let R be such that $[T]_R$ contains the first $r + 1$ RL pairs. There exists a number $n_0(R)$ so that for some $n > n_0(R)$, $[T_n]_R = [T]_R$. For such n , it holds that $B_r(\varphi^*([T_n]_R)) = B_r(\varphi^*([T]_R))$.

□

Since the local limit $\hat{T}$ of the trees T_n almost surely belongs to $\hat{\mathcal{T}}_\infty^*$, it follows that the local limit $SP_\infty = \varphi^*(\hat{T})$ of SP_n is a well-defined element of $\mathcal{SP}_\infty$.

Acknowledgements This research is partially supported by the Icelandic Research Fund, grant number 239736-051, and the Austrian Science Fund, grant numbers 10.55776/F1002 and 10.55776/PAT6732623. The research of JB was supported by Vetenskapsrådet, grant 2023-05002, and by *Ruth och Nils Erik Stenbäcks stiftelse*. The research of JT was partially supported by the Emil Aaltonen Foundation, grant number 240215. JB and SÖS gratefully acknowledge hospitality at the University of Vienna and the Technical University of Vienna at the start of this project. JT acknowledges the Austrian Science Fund (FWF) under grant P33083 (DOI 10.55776/P33083) for supporting their visit. He also thanks William da Silva for introducing the reference [5].

Author contributions All authors worked together on the manuscript, exchanged ideas in meetings, and reviewed the final version.

Funding Open access funding provided by University of Iceland.

Data availability No datasets were generated or analyzed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

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

References

1. Aldous, D.: The Continuum Random Tree. II. An Overview. *Stochastic Analysis* (Durham, 1990). London Math. Soc. Lecture Note Ser., vol. 167, pp. 23–70. Cambridge Univ. Press, Cambridge (1991)
2. Aldous, D.: The continuum random tree. I. *Ann. Probab.* **19**(1), 1–28 (1991)
3. Aldous, D.: The continuum random tree. III. *Ann. Probab.* **21**(1), 248–289 (1993)
4. Bettinelli, J., Jacob, E., Miermont, G.: The scaling limit of uniform random plane maps, via the Ambjørn-Budd bijection. *Electron. J. Probab.* **19**(73), 1–16 (2014)
5. Bonichon, N., Bousquet-Mélou, M., Fusy, E.: Baxter permutations and plane bipolar orientations. *Sém. Lothar. Combin.* **61A**, B61Ah (2010)
6. Borga, J.: The skew Brownian permuton: a new universality class for random constrained permutations. *Proc. Lond. Math. Soc.* **126**(6), 1842–1883 (2023)
7. Burago, D., Burago, Y., Ivanov, S.: *A Course in Metric Geometry*. AMS Graduate Studies in Math., vol. 33. Am. Math. Soc., Providence (2001)
8. Caraceni, A.: The scaling limit of random outerplanar maps. *Ann. Inst. H. Poincaré Probab. Stat.* **52**(4), 1667–1686 (2016)
9. Curien, N., Haas, B., Kortchemski, I.: The CRT is the scaling limit of random Dissections. *Random Struct. Alg.* **47**, 304–327 (2015)
10. Duffin, R.J.: Topology of series-parallel networks. *JMAA* **10**, 303–318 (1965)
11. Duplantier, B., Miller, J., Sheffield, S.: Liouville quantum gravity as a mating of trees. *Astérisque*, 427 (2021)
12. Iscoe, I., Ney, P., Nummelin, E.: Large deviations of uniformly recurrent Markov additive processes. *Adv. Appl. Math.* **6**, 373–412 (1985)
13. Janson, S.: Simply generated trees, conditioned Bienaymé-Galton-Watson trees, random allocations and condensation. *Probab. Surv.* **9**, 103–252 (2012)
14. Joyal, A.: Une théorie combinatoire des series formelles. *Adv. Math.* **42**(1), 1–82 (1981)
15. Kortchemski, I.: Invariance principles for Bienaymé-Galton-Watson trees conditioned on the number leaves. *Stoch. Process. Appl.* **122**, 3126–3172 (2012)
16. Le Gall, J.-F.: Random trees and applications. *Probab. Surv.* **322**(10), 245–311 (1905)
17. Panagiotou, K., Stuffer, B., Weller, K.: Scaling limits of random graphs from subcritical classes. *Ann. Probab.* **44**(5), 3291–3334 (2016)
18. Rassoul-Agha, F., Seppäläinen, T.: *A Course on Large Deviations with an Introduction to Gibbs Measures*. Graduate Studies in Mathematics, vol. 162. American Mathematical Society, Providence, RI (2015)
19. Rizzolo, D.: Scaling limits of Markov branching trees and Galton-Watson trees conditioned on the number of vertices with out-degree in a given set. *Ann. Inst. Henri Poincaré Probab. Stat.* **51**(2), 512–532 (2015)
20. Sénizergues, D., Stefánsson, S.O., Stuffer, B.: Decorated stable trees. *Electron. J. Probab.* **28**, 1–53 (2023)
21. Stuffer, B.: Limits of random tree-like discrete structures. *Probab. Surv.* **17**, 318–477 (2020)

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