Uniform random sampling of simple branched coverings of the sphere by itself

Enrica Duchi, Dominique Poulalhon, and Gilles Schaeffer

LIAFA (Université Paris Diderot) and LIX (École Polytechnique)

Abstract. We present the first polynomial uniform random sampling algorithm for simple branched coverings of the sphere by itself of degree n. More precisely, our algorithm generates in linear time *increasing quadrangulations*, which are equivalent combinatorial structures.
Our result is based on the identification of some canonical *labelled* spanning trees, and yields a constructive proof of a celebrated formula of Hurwitz for the number of some factorizations of permutations in transpositions. The previous approaches were either non constructive or lead to exponential time algorithms for the sampling problem.

Branched coverings of the sphere are 2-dimensional topological structures that have raised a lot of interest ever since the work of Hurwitz at the end of the 19th century. For instance, Okounkov and Pandharipande [17] have used these objects to derive an alternative to Kontsevitch's proof of Witten's celebrated conjecture. More recently, their relations to intersection numbers of moduli spaces and integrable hierachies as studied in mathematical physics have suggested that large random simple branched coverings provide an alternative model of discrete 2-dimensional pure quantum geometry (see *e.g.* [21] for a relatively accessible exposition).
Our aim in the present article is to provide means to effectively sample these alternative random geometries, but since our approach is purely combinatorial we trade the topological definition of branched coverings for Hurwitz fundamental combinatorial representation (see however the appendix, and the complete and elegant treatment in [12]). Define a *factorization in transpositions* of the identity permutation id_n on $\{1,\ldots,n\}$ to be a m-uple of transpositions $\tau_1,\ldots,\tau_m$ such that $\tau_m\cdots\tau_1=\mathrm{id}_n$. It is *transitive* if the graph G_τ on $\{1,\ldots,n\}$ with m edges given by the τ_i is connected, and *minimal* if $m=2n-2$. It can be checked that indeed this is the minimum number of transpositions in a transitive factorization of id_n.

Theorem 1 (Hurwitz (1891)). *Simple branched coverings of the sphere by itself of degree n are encoded up to homeomorphisms of the domain by minimal transitive factorizations in transpositions of the identity of $\mathfrak{S}_n$, and their number, called n-th Hurwitz number, is $n^{n-3}(2n-2)!$.*

The usual model of quantum geometries is the uniform distribution on fixed size unlabelled planar quadrangulations, which was first studied analytically [3] and via Markov chain simulations [2]. Only later has it become possible to perform rigourous exact simulations via efficient (linear time) perfect random sampling [19, 18, 9]. The algorithmic technics underlying these samplers, mainly the identification of carefully chosen canonical spanning plane trees, have in turn triggered important progresses in the comprehension of the intrinsic geometries of random unlabelled quadrangulations [7], culminating with the construction of their continuum limit, the Brownian map [13, 15, 14].
We show here that a similar approach can be undertaken for simple branched coverings of the sphere: starting from a variant of the standard representation of factorizations as graphs embedded on surfaces, we first recast the problem in terms of some *increasing quadrangulations*. We then show that these labelled quadrangulations, which do not fit in the earlier framework, can be decomposed using labelled trees (akin to Cayley trees) instead of plane trees. We so obtain the first constructive proof of Hurwitz formula. Previous proofs were either non constructive (via differential equation hierachies [16], geometric considerations [11], or matrix integrals [5]) or yield exponential generation

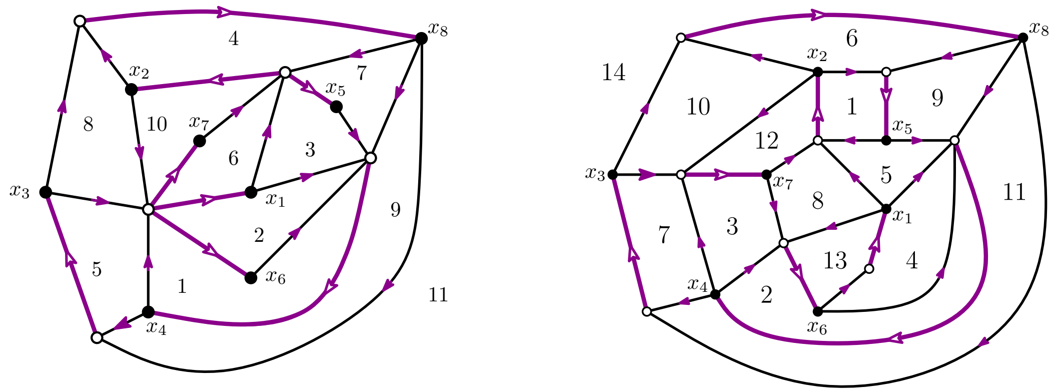

(a) An indexed labelled quadrangulation, (b) an increasing one,

$$\underset{\tau_{14}}{(3,8)}\,\underset{\tau_{13}}{(1,6)}\,\underset{\tau_{12}}{(2,7)}\,\underset{\tau_{11}}{(4,8)}\,\underset{\tau_{10}}{(2,3)}\,\underset{\tau_{9}}{(5,8)}\,\underset{\tau_{8}}{(1,7)}\,\underset{\tau_{7}}{(3,4)}\,\underset{\tau_{6}}{(2,8)}\,\underset{\tau_{5}}{(1,5)}\,\underset{\tau_{4}}{(1,6)}\,\underset{\tau_{3}}{(4,7)}\,\underset{\tau_{2}}{(4,6)}\,\underset{\tau_{1}}{(2,5)} = \mathrm{id}_8$$

(c) and the factorization of id_8 corresponding to quadrangulation (b).

Fig. 1. Labelled quadrangulations and minimal transitive factorizations in transpositions. Quadrangulations in Figures (a) and (b) are endowed with their descent orientation, with descents highlighted.

algorithms (via cut-and-join decompositions [10, 20], or exclusion/inclusion [6]). We then show that the resulting algorithm can be implemented in linear time.

From an algorithmic perspective our contribution is twofold. On the one hand we give a new and unexpected example of the versatility of the canonical spanning trees that derive from minimal α-orientations of plane graphs: these structures appear to underlie a whole chunk of efficient planar algorithmics, from random sampling to graph drawing or optimal coding. On the other hand, we illustrate further the dichotomy between random samplers based on Markov chain simulation and those based on constructive enumeration: while the formers, admittingly much easier to design, are expected to perform at best in quadratic or cubic time, the latters lead to extremely efficient algorithms when they apply. Finally, from the probabilistic and quantum gravity perspective, we believe that our construction, apart from the simulations it allows to perform, could provide a starting point to study the intrinsic geometry of increasing quadrangulations, in the same way as the constructive enumeration of unlabelled quadrangulations has lead to the Brownian map.

1 Preliminaries

Planar maps and labelled quadrangulations. A *planar map* is a proper embedding of a connected graph in the sphere, considered up to orientation preserving homeomorphisms of the sphere. The connected components of the complement of the graph in the sphere are called *faces* and are homeomorphic to discs. A *corner* is an angular sector between two successive edges around a vertex. The *degree* of a vertex or of a face is its number of corners. A *(bicolored) quadrangulation* is a map such that all faces have degree 4 and vertices are bicolored in *black* and *white*, with adjacent vertices having different colors. We require moreover that it be *simple*, that is, without double edges: all faces are real quadrangles with 4 distinct edges and 4 distinct vertices. By Euler's formula, a planar quadrangulation with k black and ℓ white vertices has $m = k + \ell - 2$ faces.

Define a *labelled quadrangulation* as a planar quadrangulation whose m faces have distinct labels $\{1, \ldots, m\}$. It is *indexed* if its n black vertices have distinct labels $\{x_1, \ldots, x_n\}$. We will be interested in the arrangement of face labels around vertices. The *descent orientation* $\mathcal{D}$ of a labelled quadrangulation is such that each oriented edge has its incident face with larger label on its left, see Figure 1. An edge is a *descent* if it is oriented from its white to its black end in $\mathcal{D}$. A labelled quadrangulation is *increasing* if each vertex is incident to exactly one descent – which implies that descent edges provide a perfect matching of black and white vertices. As opposed to the labelled quadrangulation of Figure 1(a), that of Figure 1(b) is increasing. Let $\mathcal{I}_n$ denote the set of indexed increasing quadrangulations with n black and n white vertices.

Graphical representations of transitive factorizations. Let Q be an element of $\mathcal{I}_n$. Then for all $k \leq 2n-2$, let τ_k be the transposition (i, j) given by the labels x_i and x_j of the two black vertices incident to the unique face of Q with label k. This correspondence is illustrated by Figure 1(c).

Proposition 1 (Reformulated folklore). *The above construction is a one-to-one correspondence between indexed increasing quadrangulations in $\mathcal{I}_n$ and minimal transitive factorizations of the identity* id_n.

Increasing quadrangulations can thus be considered as graphical representations of minimal transitive factorizations of the identity in transpositions. From now on we adopt this point of view and concentrate on increasing quadrangulations.

Plane maps and orientations. A *plane map* is the representation of a planar map in the plane, considered up to orientation preserving homeomorphisms of the plane. Plane maps are in one-to-one correspondence with planar maps with a distinguished face, that indicates which face of the planar map is taken as outer (unbounded) face in the plane map.
A *circuit* in an oriented map is an oriented cycle of edges (*i.e.* a cycle that can be traversed following the orientations of the edges). A *simple circuit* is a circuit that does not visit twice the same vertex. In a plane map, each simple circuit divides the plane into two components, the left one and right one (*w.r.t.* the orientation of the circuit), and one of these two components contains the outer face, while the other is bounded. A circuit is *clockwise* if its right hand side component is bounded, and *counterclockwise* otherwise.
Similarly, given a spanning tree T of an oriented map and an edge e not in T, we say that *e turns counterclockwise* around T if the bounded region delimited by e and T lies on the left hand side of e. Observe that this property is independant of the orientation or rooting of T if any.

2 From increasing quadrangulations to Hurwitz trees

Properties of the descent orientation. Given an orientation $\mathcal{O}$ of a map M and a vertex v of M, the *in-degree* of v in $\mathcal{O}$, denoted by $\mathrm{in}_{\mathcal{O}}(v)$, is the number of its incoming edges with respect to $\mathcal{O}$. Its *out-degree* $\mathrm{out}_{\mathcal{O}}(v)$ is defined accordingly. Let us define a *1-1-orientation* of a bipartite map as an orientation $\mathcal{O}$ such that for any black vertex $v_\bullet$ and any white vertex $v_\circ$: $\mathrm{in}_{\mathcal{O}}(v_\bullet) = \mathrm{out}_{\mathcal{O}}(v_\circ) = 1$. Observe that such a 1-1-orientation actually provides a perfect matching of black and white vertices. With this definition, a labelled quadrangulation is increasing if and only if its descent orientation is a 1-1-orientation. Moreover:

Proposition 2. *The descent orientation of any labelled quadrangulation is strongly connected.*

Proof. Otherwise, let v be a vertex that is not accessible from all vertices, and let C_1 and C_2 be the (disjoint) sets of vertices from which v can (*resp.*, cannot) be accessed. All edges between vertices in C_1 and C_2 are oriented from C_1 to C_2. Extract from these a simple co-circuit, that is a sequence of edges $e_1, \ldots, e_k$ such that for any $i \geq 1$, e_i separates faces f_{i-1} and f_i, and $f_k = f_0$. Then for any $i \geq 1$, the label of f_i is strictly larger than that of f_{i-1}, a contradiction. $\square$

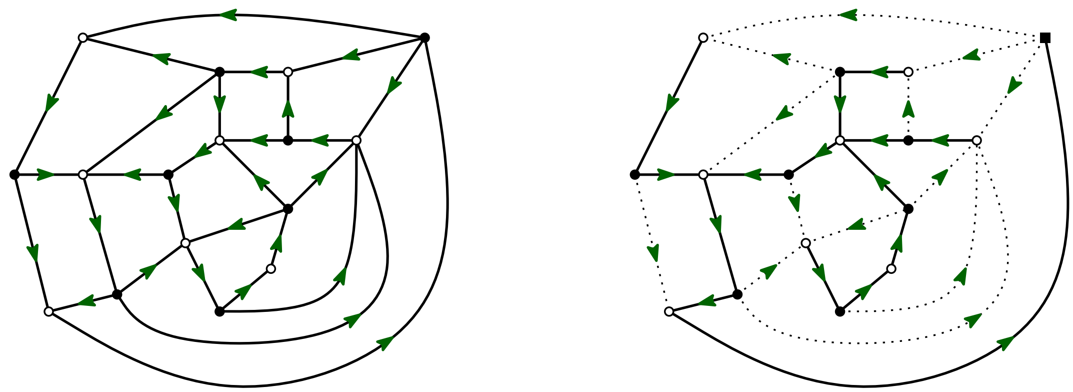

(a) The minimal orientation.

(b) Spanning tree and closure edges.

Fig. 2. Minimal 1-1-orientation of a quadrangulation and bipartition of its edges.

α-orientations of plane maps. 1-1-orientations are actually a special case of so-called α-*orientations* that have been introduced by Felsner [8]. This terminology refers to orientations $\mathcal{O}$ with prescribed $\mathrm{in}_{\mathcal{O}}$ and $\mathrm{out}_{\mathcal{O}}$ functions, α usually denoting the $\mathrm{out}_{\mathcal{O}}$ prescription. The following theorem reveals the remarkable structure of the set of α-orientations of a given graph:

Theorem 2 (Felsner [8]). *Given a plane map M and a feasible mapping α, the set of all α-orientations of M has a lattice structure for the partial order generated by clockwise circuit reversal. In particular, if M admits an α-orientation then it has a unique α-orientation without clockwise circuit, which is the minimum of the lattice.*

Since circuit reversal does not affect the accessibility, this implies moreover that for a given α, either all α-orientations of M are strongly connected, or none of them are. This proves interesting in light of the following theorem:

Theorem 3 (Bernardi [4]). *Let M be a plane map, endowed with an orientation $\mathcal{O}$ without clockwise circuit, in which r is an accessible vertex. Then the set of edges of M can be uniquely partitioned into a spanning tree T, oriented towards its root r, and a set C of edges that turn counterclockwise around T. Moreover, edges in C are in one-to-one correspondence with inner faces of M, each edge corresponding to the face on its left.*

Edges in C are called *closure edges*, since each one closes a bounded face of the plane map.

Application to increasing quadrangulations. Let Q be an increasing quadrangulation of size n, and let us embed Q in the plane by choosing as outer face its face with the largest label among the ones incident to x_n. Let $\mathcal{O}$ be its minimal 1-1-orientation. By Proposition 2 and Theorem 2, $\mathcal{O}$ is strongly connected, hence Theorem 3 may be applied to $(Q, \mathcal{O}, x_n)$, so as to obtain an oriented spanning tree T rooted at x_n, and a set C of closure edges. T is a bipartite tree on n labelled black vertices and n unlabelled white ones, hence it has $2n-1$ edges, while C has cardinality $2n-3$. Now let us transfer to each edge the label of the face on its black-to-white left hand side.

Lemma 1. *One edge in C and one in T have label i, for any i but the label of the outer face.*

Proof. First observe that, since for any white vertex $v_\circ$, $\mathrm{out}_{\mathcal{O}}(v_\circ) = 1$, each white-to-black oriented edge belongs to T – which implies that edges in C are all black-to-white oriented, meaning that

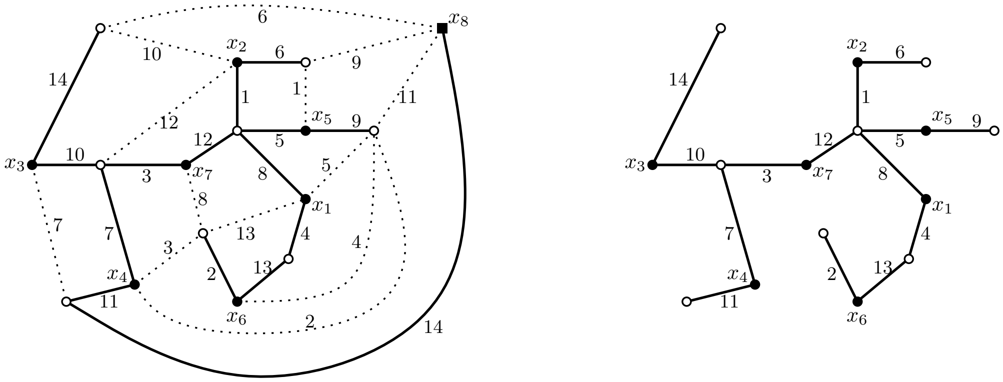

(a) Edge labelling,

(b) and the resulting labelled tree.

Fig. 3. Construction of the Hurwitz tree corresponding to the increasing quadrangulation of Figure 1(b).

their black-to-white left hand side is precisely their left hand side according to $\mathcal{O}$. Since edges in C are in one-to-one correspondence with bounded faces on their left, they do receive distinct labels, and hence edges in T as well, except for the two edges with the outer face on their black-to-white left hand side that both belong to T. $\square$

Let us define a *Hurwitz tree* of size n as any (unrooted) bicolored tree with n unlabelled white vertices, $n-1$ labelled black vertices of degree 2, and $2n-2$ labelled edges, and denote by $\mathcal{H}_n$ the set of such trees. Then the tree H obtained from T after the edge labelling and the removal of the black root vertex x_n is clearly a Hurwitz tree of size n. Let Φ denote the map from $\mathcal{I}_n$ to $\mathcal{H}_n$ that associates with any increasing quadrangulation Q of size n the Hurwitz tree H as above.

Theorem 4. *Φ is a bijection between indexed increasing quadrangulations of size n and Hurwitz trees of size n.*

The proof of this theorem will be given in Section 4, when we have described our sampling algorithm, or equivalently, the inverse of Φ.

3 Sampling trees and mapping them on quadrangulations

Random Hurwitz trees. A *Cayley tree* is a spanning tree of the complete graph with vertices $\{1,\ldots,n\}$. There are n^{n-2} Cayley trees of size n, see *eg* [1].

Proposition 3. *There is a n-to-1 correspondence between pairs (T,π) formed of a Cayley tree t with n vertices and a permutation π in $\mathfrak{S}_{2n-2}$, and Hurwitz trees of size n. In particular the number of Hurwitz trees of size n is the n-th Hurwitz number $n^{n-3}(2n-2)!$.*

Proof. Let T be a Cayley tree on n white vertices, and π a permutation of $\mathfrak{S}_{2n-2}$. Let y_i denote the white vertex with label i. Root T at y_n, and for all $i=1,\ldots,n-1$, insert a black vertex with label x_i on the middle of the first edge on the path from y_i to y_n, and give the labels $\pi(2i-1)$ and $\pi(2i)$ to the two resulting edges. Upon forgetting the (redundant) labels of white vertices we obtain a rooted Hurwitz tree H with n unlabelled white vertices, $n-1$ black vertices of degree 2 with distinct labels in $\{x_1,\ldots,x_{n-1}\}$ and $2n-2$ edges with distinct labels in $\{1,\ldots,2n-2\}$. This construction is clearly bijective. Upon forgetting H's root position, we get a n-to-1 correspondence with (unrooted) Hurwitz trees. $\square$

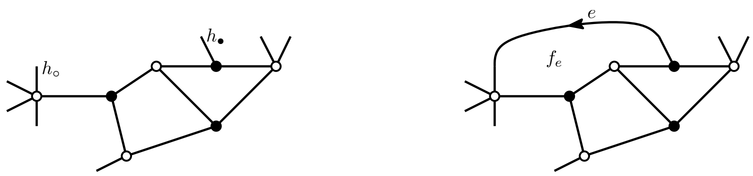

(a) Consecutive half-edges $h_\circ$ and $h_\bullet$, (b) and their local closure.

Fig. 4. The local closure of two half-edges at distance 3.

Corollary 1. *Hurwitz trees of size n can be generated uniformly at random in linear time.*

Proof. Sampling permutations uniformly at random in linear time is a classical textbook exercise. For Cayley trees, it can be done *e.g.* following Joyal's bijective proof of Cayley formula [1].

A general technique to build planar maps out of trees. In order to describe how to construct a quadrangulation out of a tree, we will consider intermediate objects. A *pre-map* is a plane bicolored map with some distinguished pending edges in the outer face called *half-edges* (and whose loose endpoint will not count as vertices). A half-edge is either black or white according to the vertex it is attached to. We say that the half-edges $h_\circ$ and $h_\bullet$ are *consecutive* if, while travelling clockwise from $h_\circ$ on the boundary of the outer face, $h_\bullet$ is the first encountered half-edge; they are *at distance* $p \geq 0$ if moreover $h_\bullet$ is reached after travelling along p sides of edges.
Let us describe our basic operation on pre-maps. Given two consecutive white and black half-edges $h_\circ$ and $h_\bullet$ in the outer face f of a pre-map M, let us merge $h_\bullet$ and $h_\circ$ into a black-to-white oriented edge e in the unique way that preserves planarity, and divides f into a face f_e on the left hand side of e and a new outer face f' on its right hand side. Then immediately:

Proposition 4. *If $h_\circ$ and $h_\bullet$ are consecutive and at distance p, then the bounded face f_e has degree $p+1$ and contains no half-edges.*

The *local closure* of $h_\circ$ and $h_\bullet$ is the resulting pre-map $M' = M \cup \{e\} \setminus \{h_\bullet, h_\circ\}$. In the case of labelled pre-maps, the local closure is compatible with the labelling provided that $h_\circ$ and $h_\bullet$ have the same label. In this case, the label of e is this common label.

The closure of a Hurwitz tree. Given a Hurwitz tree T of size n (with edges labels $\{1, \ldots, m\}$, for $m = 2n-2$), let the associated *hairy tree* $\bar{T}$ be obtained by inserting labelled half-edges at every vertex to complete the cycle of incident labels to be $(1, \ldots, m)$ in clockwise (*resp.* counterclockwise) direction around every white (*resp.* black) vertex.
A pre-map M with labelled edges and half-edges is said *valid* if around every white (*resp.* black) vertex, incident labels form the clockwise (*resp.* counterclockwise) cycle $(1, \ldots, m)$, and if for all $i \in \{1, \ldots, m\}$, at most one edge with label i has its black-to-white left hand side incident to the outer face. In the following, we will actually consider that each edge carries its label (in the face) on its black-to-white left hand side. With this convention, the second condition defining valid pre-maps is that each label occurs at most once in the outer face.

Lemma 2. *In a valid pre-map, any two consecutive half-edges $h_\circ$ and $h_\bullet$ are at distance 1 or 3 and have the same label, and their local closure produces a valid pre-map.*

Proof. If $h_\circ$ has label i, then edges between $h_\circ$ and $h_\bullet$ have alternatively label $i+1$ or i because of the cyclical labelling rules. Hence $h_\bullet$ has label i, and since only one label i may appear in the outer

RandomQuad(n)

uses an initially empty stack S of half-edges, a current pre-map M and a current half-edge h.

1. Generate a uniform random Hurwitz tree T of size n.
2. Let M be the hairy tree $\bar{\mathtt{T}}$ associated to T, and h any of its half-edges.
3. Repeat the following loop: *(loop invariant: M is a valid pre-map with m more white than black half-edges)*
   (a) If h is a white half-edge,
      i. If h is already marked, go to Step 4. *(all black half-edges have been matched)*
      ii. Otherwise, mark h and insert it in S.
   (b) Otherwise, if S is not empty, pop the last half-edge $h_\circ$ from S.
      ($h_\circ$ and h are consecutive, hence at distance 1 or 3, and have equal labels)
      Let M be the local closure of h and $h_\circ$, and give their common label to the new face.
   (c) Let h be the next half-edge around M in clockwise direction.
4. *(S contains m white half-edges, and the successive ones are at distance 0 or 2)* Match the m white half-edges in S to the m black half-edges of a new black vertex of degree m in the outer face to get a new pre-map M.
5. *(M is a valid pre-map without half-edges and with faces of degree 2 and 4)* Contract all faces of degree 2 of M and forget the orientation of closure edges to get an indexed labelled quadrangulation Q.

Fig. 5. The algorithm RandomQuad (assertions leading to the proof of Theorem 5 are emphasized)

face, the distance is at most 3. The local closure of $h_\circ$ and $h_\bullet$ creates an edge that has the outer face on its black-to-white right hand side, so that no new label is created in the outer face and the pre-map remains valid. □

From the definitions, the following lemma is immediate:

Lemma 3. *Hairy trees are valid pre-maps.*

The first algorithm. We can now state and analyse the first algorithm, given in Figure 5. The first steps of an execution are given in Figure 6.

Theorem 5. *Steps 2-5 of the algorithm RandomQuad describe a mapping from $\mathcal{H}_n$ to $\mathcal{I}_n$ which is the inverse of the mapping Φ of Theorem 4. RandomQuad(n) thus generates indexed increasing quadrangulations of size n uniformly at random.*

Proof. Let us first check the emphasized assertions in the algorithm. The first assertion is a clear loop invariant since the half-edges are only modified in Step 3(b) by a local closure, which preserves validity by Lemma 2, and removes simultaneously one black and one white half-edges. Assertion in 3(b) follows from the fact that white half edges are stored in a (last in, first out) stack, so that $h_\circ$ is always the last inserted white half-edge among those that have not yet been matched. Since M is a valid pre-map, Lemma 2 applies. Assertion in 3(a)i follows from the observation that between

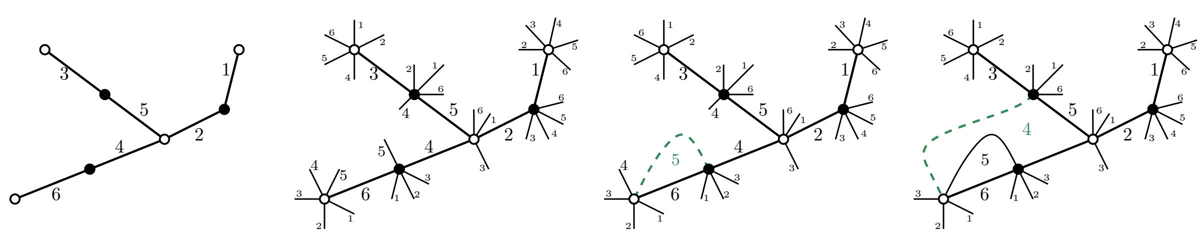

(a) A Hurwitz tree, (b) its hairy tree, (c) a first local closure, (d) and a second one.

Fig. 6. A Hurwitz tree, the associated hairy tree, and two steps of RandomQuad.

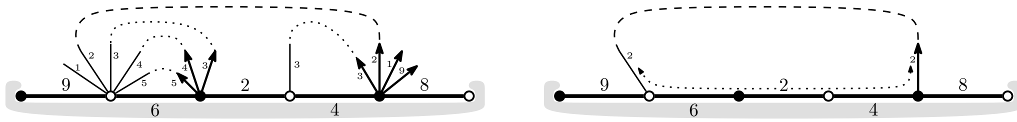

(a) First algorithm, (b) second algorithm

Fig. 7. A partial comparison of the local closures involved in the two algorithms.

two visits to the same white half-edge h, a full turn around the pre-map is performed, and all black half-edges are considered. Since the stack contains (at least) h during this full turn, it is never empty, hence all black half-edges are matched. Assertion in Step 4 follows from the fact that $\mathtt{M}$ is valid. The last assertion follows immediately from the previous ones.
Step 3(b) creates exactly one face of degree 4 for each label i in $\{1, \ldots, m\}$, since the original hairy tree has exactly one edge with label i, and this label disappears from the outer face at the exact step when the face with label i is created. As $\mathtt{M}$ is a valid pre-map, the labels of faces around each white (*resp.* black) vertex in clockwise (*resp.* counterclockwise) direction form a cyclic subsequence of $(1, \ldots, m)$. Hence $\mathtt{Q}$ is an increasing quadrangulation.
This proves the first half of the theorem, namely that the algorithm indeed correctly produces an increasing quadrangulation. The proof that the correspondence is one-to-one and inverse of Φ is delayed to the next section. ☐

Proposition 5. *The algorithm* RandomQuad *can be implemented in linear time and space with respect to the number of edges and half-edges of $\bar{T}$. Since there are n white vertices and $m = 2n-2$ half-edges incident to each white vertex, it has quadratic complexity in n.*

4 The linear complexity algorithm

In this section we give an alternative description of the bijection producing an increasing quadrangulation out of a Hurwitz tree. The idea is to create the half-edges only when they lead to faces of degree 4. In order to do this we analyse more finely the previous algorithm.
Let us consider an edge e with label j, having its white-to-black left hand side in the outer face. Let i and k be the labels of the previous and next edges (not half-edges) e^- and e^+ around the outer face (the relative positions of i, j and k are illustrated by Figure 8). We wish to understand how the first algorithm deals with white half-edges between e^- and e, and black half-edges between e and e^+. Observe that j can be equal to i but not to k because white vertices can be leaves in a Hurwitz tree while black vertices cannot (they all have degree 2). There are two main cases:

- First suppose $i = j$ (*ie* the white endpoint of e has degree 1) or the cycle (i, j, k) is a subcycle of $(m, \ldots, 1)$ (*ie* $i > j > k$, $j > k > i$ or $k > i > j$). Then there are more white half-edges between

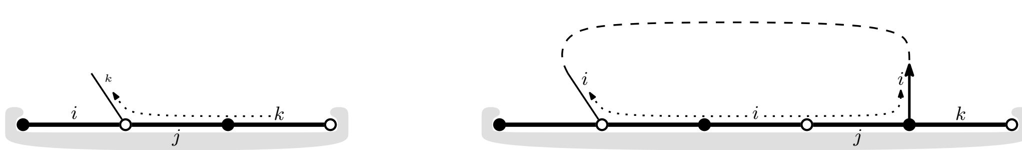

(a) Case $i \geq j > k$: creation of a white half-edge with label k.

(b) Case $i < j < k$: creation of a face with label i.

Fig. 8. The two local rules of FastRandomQuad depending on the (cyclical) order of i, j, k.

FastRandomQuad(n)

uses a stack S of half-edges, a current pre-map M and a current arc $\boldsymbol{e}$.

1. Generate a uniform random Hurwitz tree T of size n; let M $\leftarrow$ T.
2. Let $\boldsymbol{e}$ be an arbitrary edge of M. Orient $\boldsymbol{e}$ from its white to its black endpoint, and repeat the following loop: *(at this point $\boldsymbol{e}$ has the outer face on its left hand side)*
   - If $\boldsymbol{e}$ is a white half-edge then go to Step 3. *(at this point a white half-edge has been encountered twice).*
   - Let $\boldsymbol{e}^-$ and $\boldsymbol{e}^+$ be the previous and next edges around the outer face.
   - Let i, j and k be the labels of $\boldsymbol{e}^-$, $\boldsymbol{e}$ and $\boldsymbol{e}^+$.
   - If $i = j$ or the cycle (i, j, k) is a subcycle of $(m, \ldots, 1)$, create a half-edge $h_\circ$ with label k on the white endpoint of $\boldsymbol{e}$ and insert $h_\circ$ in S. Set $\boldsymbol{e}$ to the edge or half-edge following $\boldsymbol{e}^+$.
   - Otherwise (*i.e.* if the cycle (i, j, k) is a subcycle of $(m, \ldots, 1)$),
     (a) If S is empty then set $\boldsymbol{e}$ to the edge or half-edge following $\boldsymbol{e}^+$.
     (b) Otherwise pop from S a half-edge $h_\circ$ and create a half-edge $h_\bullet$ with label k on the black endpoint of $\boldsymbol{e}$. Match $h_\bullet$ and $h_\circ$ to create a closure edge $\boldsymbol{e}'$ enclosing a face of degree 4 with label k, and let $\boldsymbol{e} \leftarrow \boldsymbol{e}'$.
3. *(at this point S contains at least one white half-edge)* Match the $p \geq 1$ white half-edges in S to the p black half-edges of a new black vertex in the outer face to form p new edges and p new faces of degree 4.

Fig. 9. The algorithm FastRandomQuad (emphasized texts are again assertions)

e^- and e than black half-edges between e and e^+. The white half-edge with label k will be the first of the half-edges between e and e^- to be matched at distance 3.
This case is illustrated with $(i, j, k) = (9, 6, 2)$ in Figure 7(a): half-edges with labels 5, 4, 3 are matched at distance 1 and the half-edge with label 2 is the first to be matched at distance 3.

– Now suppose (i, j, k) is a subcycle of $(1, \ldots, m)$ (*ie* $i < j < k$, $j < k < i$, or $k < i < j$). Then there are more black half-edges between e and e^+ than white ones between e^- and e. The black half-edge with label i will be the first not to match a white half-edge at distance 1: it will either remain unmatched or match a half-edge at distance 3.
This case is illustrated with $(i, j, k) = (2, 4, 8)$ in Figure 7(a): the half-edge with label 3 is matched at distance 1, while the half-edge with label 2 gets matched at distance 3.

Upon iterating this lemma all the closure edges that produce faces of degree 4 can be constructed without constructing those that produce faces of degree 2. Our second algorithm FastRandomQuad, as presented in Figure 9, exactly performs this iteration until all closure edges have been created. The first steps of the execution of this algorithm on the Hurwitz tree of Figure 3(b) are given in Figure 10 in the Appendix.

Proposition 6. *Steps 2-3 of FastRandomQuad are equivalent to Steps 2-5 of RandomQuad. Moreover FastRandomQuad can be implemented to work in linear time and space with respect to the size n of the constructed increasing quadrangulation.*

Proof. The equivalence is a direct consequence of the previous discussion: FastRandomQuad exactly performs the subset of the stack operations performed by RandomQuad that concern half-edges whose closure yield faces of degree 4. This implies that a white half-edge is indeed encountered twice at some point, and that FastRandomQuad stops and produces an increasing quadrangulation. To check that FastRandomQuad works in linear time we observe that less than n closure edges are produced, and that Step (a) is performed at most $2n$ times because RandomQuad visits at most twice each edge side.

End of the proof of Theorem 4 and 5. To conclude the proof we only need to understand why the increasing quadrangulation Q produced by RandomQuad from a tree T is such that $\Phi(Q) = T$. But this follows immediately from the alternative description given by FastRandomQuad. Indeed Step 2 only adds to the tree T closure edges that turn clockwise around T when oriented from their black to their white endpoint: orienting one of the final edges e toward the extra vertex x_n, and all the edges of the tree toward e, we can apply the uniqueness condition of Theorem 3 to conclude. $\square$

Corollary 2. *The numbers of indexed increasing quadrangulations of size n, of minimal transitive factorizations of the identity in $\mathfrak{S}_n$, and of simple branched coverings of degree n of the sphere by itself, are $n^{n-3}(2n-2)!$, and all these objects can be generated uniformly at random in linear time.*

References

1. M. Aigner and G. Ziegler. *Proofs from the Book.* Springer, Berlin, 1998.
2. J. Ambjørn, P. Białas, J. Jurkiewicz, Z. Burda, and B. Petersson. Effective sampling of random surfaces. *Physics Letters B*, 325(3):337–346, 1994.
3. J. Ambjørn, B. Durhuus, and T. Jonsson. *Quantum geometry: a statistical field theory approach.* Cambridge University Press, 1997.
4. O. Bernardi. Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Elec. J. of Combinatorics*, 14(1):R9, 2007.
5. G. Borot, B. Eynard, M. Mulase, and B. Safnuk. A matrix model for simple Hurwitz numbers, and topological recursion. *Journal of Geometry and Physics*, 61(2):522–540, 2011.
6. M. Bousquet-Mélou and G. Schaeffer. Enumeration of planar constellations. *Advances in Applied Mathematics*, 24(4):337–368, 2000.
7. P. Chassaing and G. Schaeffer. Random planar lattices and integrated superbrownian excursion. *Probability Theory and Related Fields*, 128(2):161–212, 2004.
8. S. Felsner. Lattice structures from planar graphs. *Elec. J. Comb.*, 11(1):R15, 2004.
9. É. Fusy, G. Schaeffer, and D. Poulalhon. Dissections, orientations, and trees with applications to optimal mesh encoding and random sampling. *ACM Transactions on Algorithms (TALG)*, 4(2):19, 2008.
10. I. Goulden and D. Jackson. Transitive factorisations into transpositions and holomorphic mappings on the sphere. *Proceedings of the American Mathematical Society*, 125(1):51–60, 1997.
11. S. Lando and D. Zvonkine. Counting ramified coverings and intersection theory on spaces of rational functions. *Moscow Math. J*, 7(1):85–107, 2007.
12. S.K. Lando, A.K. Zvonkin, and D.B. Zagier. *Graphs on Surfaces and Their Applications.* Encyclopaedia of Mathematical Sciences. Springer, 2004.
13. J.-F. Le Gall. The topological structure of scaling limits of large planar maps. *Inventiones mathematicae*, 169(3):621–670, 2007.
14. J.-F. Le Gall. Uniqueness and universality of the Brownian map. *arXiv:1105.4842, to appear in* Ann. Proba., 2011.
15. G. Miermont. The Brownian map is the scaling limit of uniform random plane quadrangulations. *arXiv:1104.1606, to appear in* Acta Math., 2011.
16. A. Okounkov. Toda equations for Hurwitz numbers. *Math. Res. Lett.*, 7(4):447–453, 2000.
17. A. Okounkov and R. Pandharipande. Gromov-Witten theory, Hurwitz numbers, and matrix models, I. *arXiv:math/0101147*, 2001.
18. D. Poulalhon and G. Schaeffer. Optimal coding and sampling of triangulations. *Algorithmica*, 46(3-4):505–527, 2006.
19. G. Schaeffer. Random sampling of large planar maps and convex polyhedra. In *ACM STOC'1999*, pages 760–769. ACM, 1999.
20. V. Strehl. Minimal transitive products of transpositions—the reconstruction of a proof of A. Hurwitz. *Séminaire Lotharingien de Combinatoire, B37c*, 1996.
21. D. Zvonkine. Enumeration of ramified coverings of the sphere and 2-dimensional gravity. *arXiv:math/0506248*, 2005.

A Branched coverings of the sphere by itself

We give here for completeness a definition of branched coverings, but refer again to [12] for a gentle introduction to the topological and combinatorial aspects of their mathematical theory.

A *covering of degree* n of a surface $\mathcal{I}$ by another surface $\mathcal{D}$ is a mapping $\phi : \mathcal{D} \to \mathcal{I}$ such that each value y of $\mathcal{I}$ has n preimages, and each point x of $\mathcal{D}$ has a neiborhood V_x such that ϕ is an homeomorphism from V_x to $\phi(V_x)$. A *branched covering* of degree n of the sphere by itself is a mapping from $\mathbb{S}_2$ to itself such that there is a finite set of values $Y = \{y_1, \ldots, y_m\} \subset \mathbb{S}_2$ such that $\phi|_{\mathbb{S}_2 \setminus \phi^{-1}(Y)}$ is a covering of degree n and for every x in $\phi^{-1}(Y)$ there is an integer k, an open neighborhood V_x of x and two homeomorphisms $h : \mathbb{C} \to V_x$ and $h' : \phi(V_x) \to \mathbb{C}$ such that $h' \circ \phi \circ h$ is the mapping $z \to z^k$ of the complex plane. In this case the preimage x is said to have *order* k. A preimage with order 1 is a *regular point*. By continuity the sum of the orders of all the preimages of a value y by ϕ has to be n, and the multiset of these orders is called the *type* of the critical value y. A critical value y is said to be *simple* if all its preimages but one are regular and the only non regular one has order 2: equivalently a critical value is simple if its type is $1^{n-2}\,2$. A *simple branched covering* is a branched covering whose critical values are all simple.

In order to dispose of symmetry problems we follow the approach of Hurwitz: we fix a regular value and label its preimage with integers 1 to the degree. Finally we consider these coverings up to homeomorphisms of the sphere. The resulting equivalence classes are the simple branched coverings considered by Hurwitz in Theorem 1, for which he gave the quoted formula.

In his work Hurwitz also considered more generally the case where one critical value is non simple, of type $\lambda = 1^{\ell_1} \ldots n^{\ell_n}$ (where ℓ_i denotes the number of preimages of order i). In terms of permutations, these almost simple coverings correspond to minimal transitive factorizations into transpositions of a permutation with cycle type λ. Hurwitz provided also a formula for their number, and our approach extends almost directly to prove this general formula.

B Large random increasing quadrangulations

From the probabilistic and quantum gravity perspective, the main concern is to understand the geometry of natural discrete models of random surfaces.

In order to compare our approach to the existing literature, let X_n (*resp.* Y_n) denote a uniform random increasing (*resp.* planar) quadrangulation with $2n-2$ faces and let $d_{X_n}(.,.)$ be the graph distance on the set of vertices of X_n.

It is known that the expected distance Δ_{Y_n} between two uniform random vertices of Y_n is of order $n^{1/4}$. More precisely, as n goes to infinity the random variable $\Delta_{Y_n} n^{-1/4}$ converges in law to a continuous positive random variable D. The analog question is unsettled for increasing quadrangulations and numerical simulation were out of reach with previous approaches. Our linear time algorithm makes it possible to check experimentally the hypothesis that the distances $\Delta_{X_n} n^{-1/4}$ converge to the same limit.

In the case of Y_n, much more precise results have been obtained in the recent years. In particular upon setting the edge length to $n^{-1/4}$, the random uniform quandrangulation Y_n converges as a metric space to a continuum limit, the Brownian map, which is a random space with the topology of the sphere [13, 15, 14].

Conjecture. The pair $(X_n, n^{-1/4} d_{X_n})$ converges to the Brownian map in the sense of [13, 15, 14].

In other terms we conjecture that large increasing quadrangulations behave very much like large unlabelled quadrangulations. This should be understood as a statement analogous to the well known statement that both random uniform binary trees and uniform random Cayley trees, although quite different at a discrete level, converge upon rescaling edge length to a same continuum limit, the continuum random tree (CRT), when their size go to infinity.

Proving the above convergence would be a remarkable achievement as it would on the one hand give a strong support to the belief of the community that the Brownian map is a new universal limit object, in the same sense as the Brownian motion or the CRT, and on the other hand it would make more precise the connection between the realm of branched coverings and Hurwitz numbers, and that of quantum gravity.

The bijection between Hurwitz trees and increasing quadrangulations that we propose in the present paper can be seen as labelled counterparts to the bijections between plane trees and families of maps that are the basic building blocks of the approach that culminated with [13, 15, 14]. Hopefully they can lead to a proof of the above conjecture.

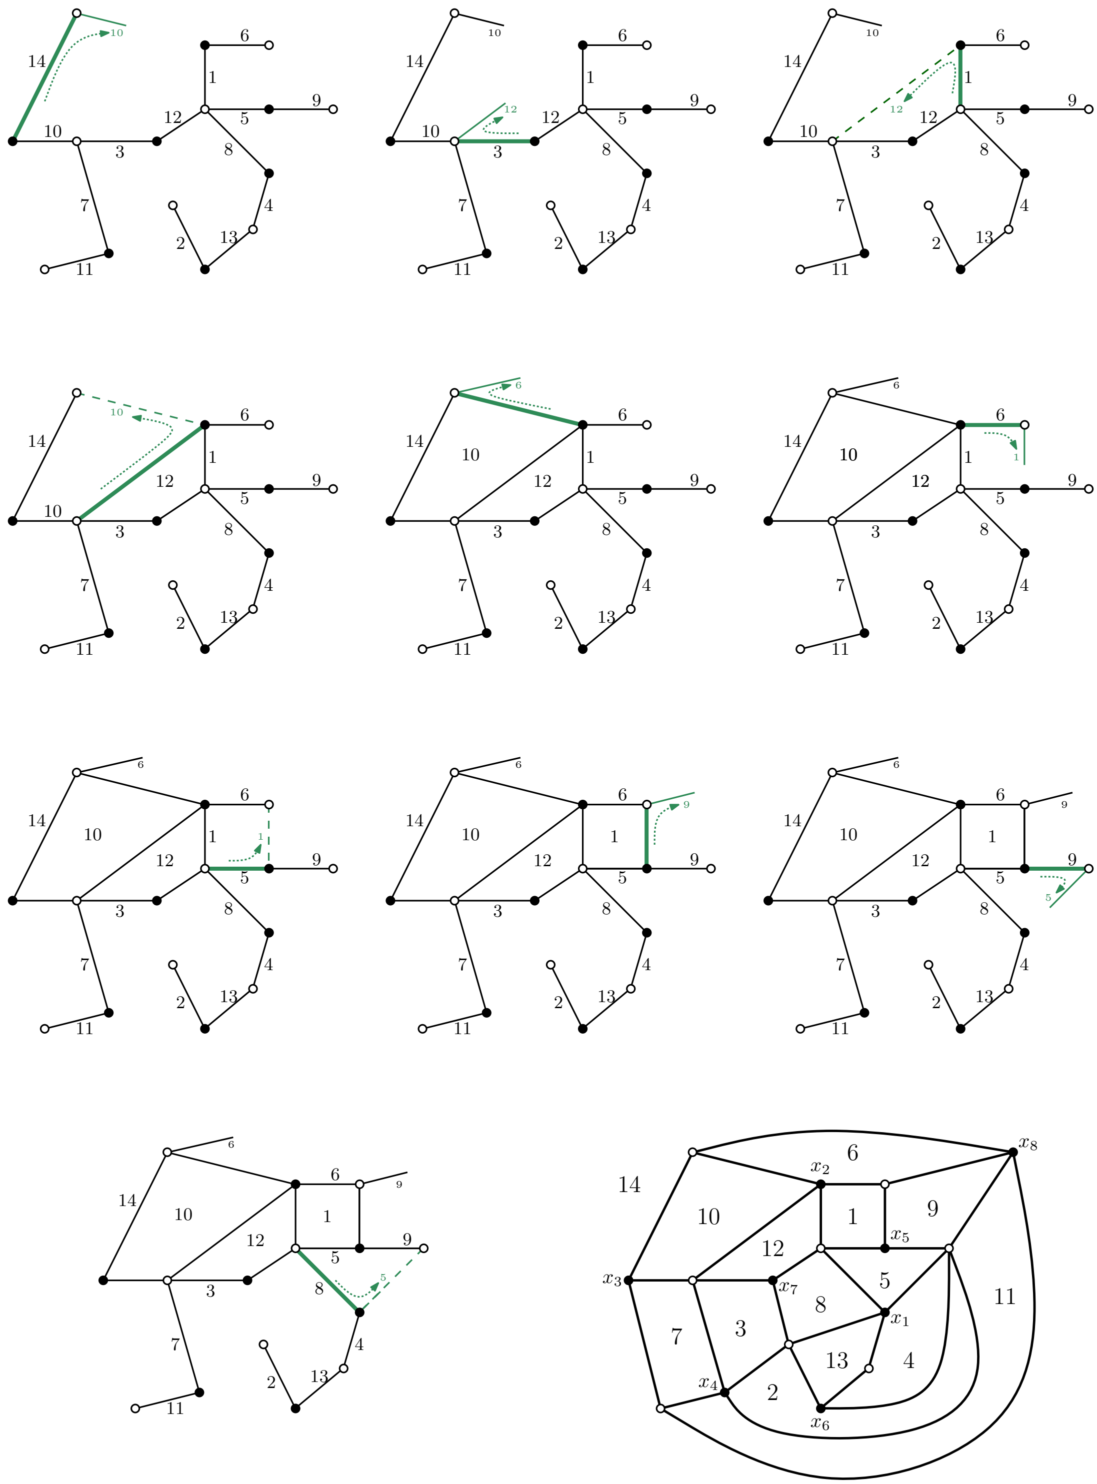

Fig. 10. Execution of FastRandomQuad on the Hurwitz tree of Figure 3(b). At each step, the current edge is the bold green one, and the created (half-)edge is the thin green one.