Combinatorial and Geometric Properties of Maps on the Sphere

Directed Reading Program, Winter 2026

Jesse Seeligsohn and Claire Driedger

June 1, 2026

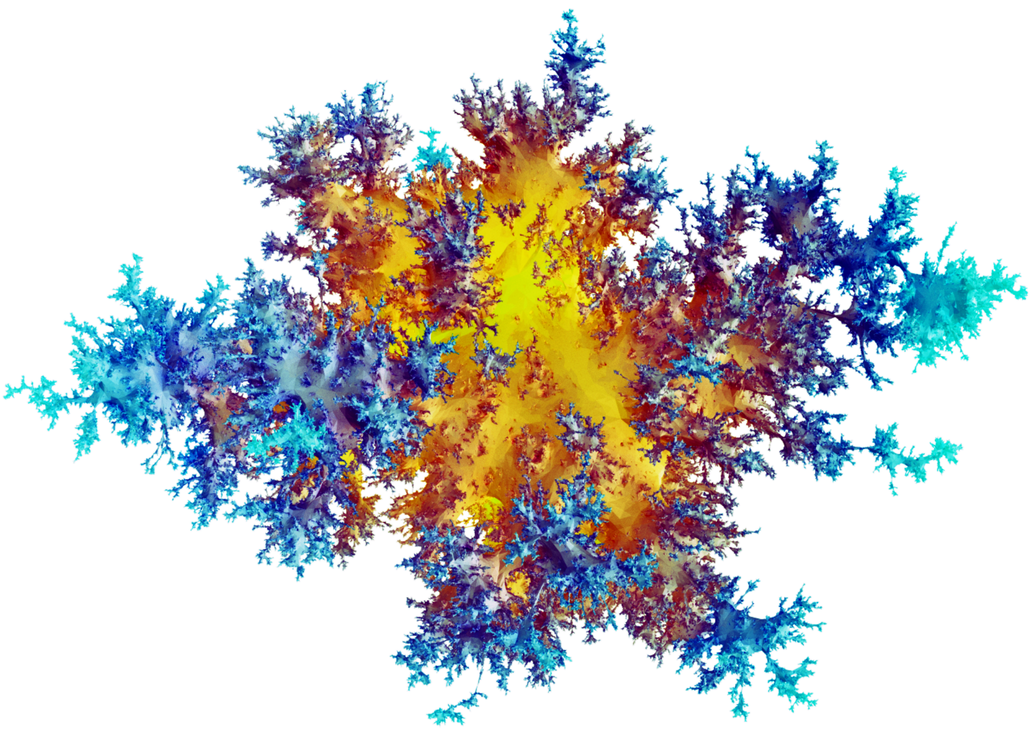

A Large Random Quadrangulation of the Sphere [Stu].

Department of Mathematics and Statistics, McGill University

Contents

Abstract

A planar map is a graph drawn on the sphere with no edge crossings. This report will present the Cori-Vauquelin-Schaeffer bijection, which translates the elaborate geometric and combinatorial properties of quadrangulations into the language of trees, which is better understood. Using this relationship, we gain insights into the number of maps of size n, the distances between vertices in a quadrangulation, and the expected height of a tree.

Introduction

Planar maps have attracted interest from both mathematicians and physicists over the decades. For example, random maps have been used as a discrete model for random geometry in the study of quantum gravity ([JJ97]). In 1963, Tutte derived a formula for the enumeration of planar maps using generating functions; however, we seek a more descriptive approach that illuminates the mechanics of his formula. In this report, we'll address the following two fundamental questions: (i) how can we enumerate planar maps; and (ii) what can we say about distances between vertices in a planar map? We will primarily answer these questions using the Cori-Vauquelin-Schaeffer (CVS) bijection.

Chapter 1 will provide basic definitions and results from graph theory. We will go on to describe maps and two subclasses, quadrangulations and plane trees. The section will culminate with a discussion of the labelling function and contour walk on trees. In Chapter 2, we will explore the Cori-Vauquelin-Schaeffer (CVS) bijection between well-labelled plane trees and pointed quadrangulations. Chapture 3 will use this bijection to recover distances in and the enumeration of quadrangulations, and hence maps. The final chapter (Chapter 4) will briefly discuss random trees.

We will mostly follow the same structure as a set of lecture notes by Grégory Miermont in [Mie14], with the exception of Chapter 4, which will mostly follow [FF70].

1 Graphs on Surfaces

We first recall basis facts and foundational results of graph theory required to introduce planar maps. While much of this section is devoted to building the framework which will allow us to understand the CVS bijection, the study of graphs on surfaces is interesting in its own right and exercises our ability to think about and visualize graphs.

1.1 Graph Basics

Definition 1. An *(undirected) graph* is a triple $G = (V, E, \iota)$ where V is a vertex set, E is an edge set, and

$$\iota : E \to \binom{V}{2} \cup \binom{V}{1}$$

is an incidence function that assigns each edge to its *endpoint(s)*.

Additionally, we define the following terminology. If $\mathbf{e}$ is an edge with $\iota(\mathbf{e}) = \{u, v\}$, then u and v are *adjacent* and they are the *endpoints of* or *vertices incident to* $\mathbf{e}$. Otherwise, if $\iota(\mathbf{e}) = \{u\}$, then $\mathbf{e}$ is a *loop* on u. A graph has *multiedges* if ι is not injective, i.e. there exist distinct edges $\mathbf{e}_i$ and $\mathbf{e}_j$ such that $\iota(\mathbf{e}_i) = \iota(\mathbf{e}_j)$. A *simple graph* is a graph G with no loops nor multiedges, otherwise, G is a *multigraph.* In our discussion, we permit loops and/or multiedges, therefore our usage of the word 'graph' should be interpreted as 'multigraph'. For convenience, we may denote a graph as $G = (V, E)$ with the incidence relation implied and similarly we may use the shorthand $\mathbf{e} = \{u, v\}$ to mean $\iota(\mathbf{e}) = \{u, v\}$. We can also denote the vertex and edge sets of G as $V(G)$ and $E(G)$ for clarity.

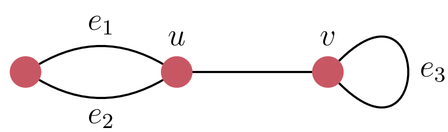

Figure 1: A graph containing multiedges e_1 and e_2 and loop e_3.

The *degree* of a vertex u is the number of edges incident to u. In the graph shown in Figure 1, vertices u and v both have degree 3.

We can now discuss the handshaking lemma, upon which many proofs in graph theory rely, including the interesting corollary that every graph must have an even number of odd-degreed vertices (try it!).

Lemma 1. (Handshaking Lemma) For any undirected graph $G = (V, E)$, the following equality holds:

$$\sum_{v \in V} \deg(v) = 2|E|.$$

Proof. Every edge is incident to two vertices, unless it is a loop, in which case it is incident to the same vertex twice. Therefore, summing over degrees of all the vertices counts every edge twice. $\square$

Definition 2. Let $G = (V, E)$ be a graph. A sequence

$$P = (v_0, \mathbf{e}_1, v_1, \ldots, \mathbf{e}_k, v_k)$$

is a *path of length* k if for each $i \in \{1, \ldots, n\}$,

$$\iota(\mathbf{e}_i) = \{v_{i-1}, v_i\}$$

and all vertices (except possibly the first and last) are distinct. If $v_0 = v_k$, then P is a *cycle*. A graph that contains no cycles is said to be *acyclic*.

Definition 3. Let $G = (V, E)$ be a graph. Two vertices u and v in V are *connected* if there is a path from u to v in G. If every pair of vertices is connected, then G is a *connected* graph. Otherwise, G is *disconnected*.

Definition 4. Let G be a graph. For two vertices u, v in V, the *graph distance* on G is the shortest path (path containing the fewest edges) between u and v, denoted

$$d_G(u, v) = \min\{k : P \text{ is a path of length } k \text{ from } u \text{ to } v\}.$$

If u and v are not connected, set $d_G(u, v) = \infty$.

Definition 5. Let $G = (V, E, \iota)$ and $G' = (V', E', \iota')$ be two graphs. Then G and G' are *isomorphic* if there exists a bijective function $\phi : V \to V'$ such that for all u, v in V,

$$|\{\mathbf{e} \in E : \iota(\mathbf{e}) = \{u, v\}\}| = |\{\mathbf{e}' \in E' : \iota'(\mathbf{e}') = \{\phi(u), \phi(v)\}\}|.$$

That is, ϕ is a *graph isomorphism* if it is a bijective function which preserves the incidence relation between vertices, while also preserving the number of edges between each pair of vertices (because G could have multiedges). See Figure 2 for an example of a graph isomorphism.

1.2 Some Important Graphs

We can categorize graphs based on their structure.

Definition 6. A graph $G = (V, E)$ is *bipartite* if there exists a bipartition of its vertices, denoted $V = U \sqcup W$ such that every edge $\mathbf{e}$ has one endpoint in U and one endpoint in W. That is, there are no edges which have both endpoints in the same side of the bipartition.

Definition 7. A tree $T = (V, E)$ is a graph which is connected and for every u, v in V, there exists a unique path between u and v. The following are equivalent definitions:

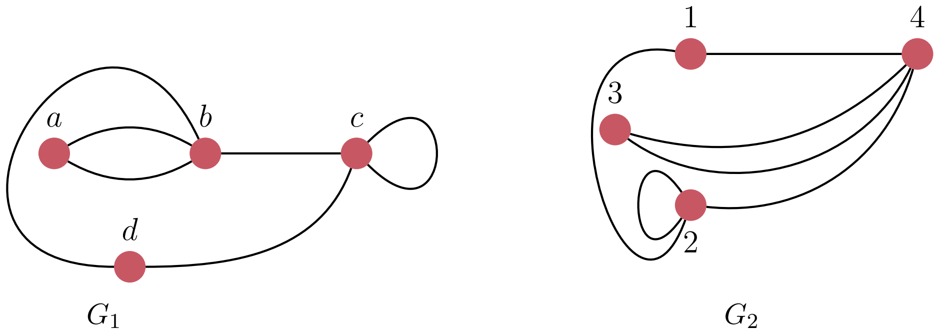

Figure 2: G_1 and G_2 are isomorphic because the function $\phi : \{a, b, c, d\} \to \{1, 2, 3, 4\}$ given by $a \mapsto 3$, $b \mapsto 4$, $c \mapsto 2$, $d \mapsto 1$ is a graph isomorphism.

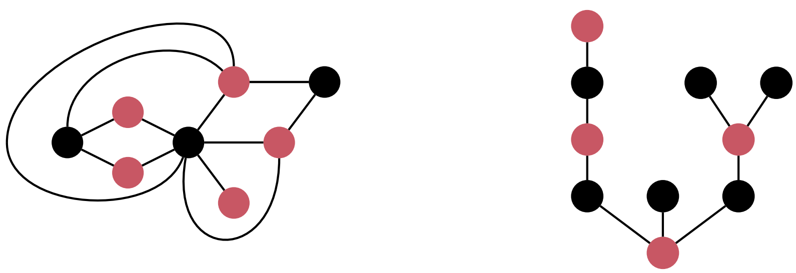

Figure 3: Two bipartite graphs. We use two colours to indicate the bipartition of the vertices.

(*a*) G is connected and is acyclic;
(*b*) G is connected (or, is acyclic) and has $|V| = |E| + 1$ vertices.

A vertex with degree 1 in a tree is called a leaf. A forest is a disjoint union of trees.

For example, the graph on the right in Figure 3 is a tree.

Remark 1. A tree is always bipartite. In fact, being bipartite is equivalent to having no odd cycles.

1.3 Toward a Definition of a Map

We are now ready to extend our understanding of graphs by considering what it means for edges and vertices to be "drawn", or embedded, in a surface. For us, a surface is a 2-dimensional manifold which is oriented, compact, connected and without boundary. The orientability of such surfaces will soon become very convenient as it will allow us to define directions and order, such as 'clockwise' and 'counterclockwise'. The most intuitive surface is the sphere $\mathbb{S}^2 = \{(x, y, z) \in \mathbb{R}^3 : x^2 + y^2 + z^2 = 1\}$.

Now recall that a homeomorphism is bijective, continuous function whose inverse is also continuous. Two surfaces are *homeomorphic* if there exists a homeomorphism which maps one to the other, in other words, if they can be continuously deformed into one another. The Fundamental Theorem of the Classification of Surfaces says that any of our surfaces is homeomorphic to the sphere or the connected sum of g copies of the 2-dimensional torus, i.e. a sphere with g holes in it; therefore it suffices to keep the most convenient surface in mind when thinking about maps. The number g is called the genus of S. The Euler characteristic of an orientable surface is

$$\chi = 2 - 2g.$$

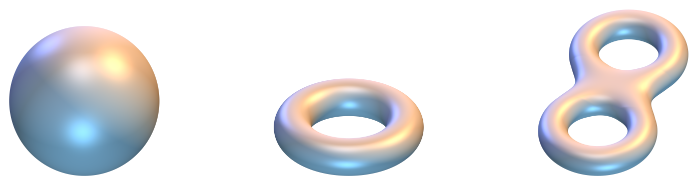

(a) $\chi = 2$ (b) $\chi = 0$ (c) $\chi = -2$

Figure 4: Topological Surfaces with different Euler characteristics.

1.3.1 Graph Embeddings

Definition 8. Let S be a surface. An *oriented edge* on S is the image of a continuous map $e : [0, 1] \to S$ such that e is *almost injective* i.e. e is injective on $[0, 1)$, allowing $e(1) = e(0)$. The *extremeties* of an oriented edge e are the *origin* $e^- = e(0)$ and the *target* $e^+ = e(1)$. The *reversal* of e is $\overline{e} = e(1 - t)$.

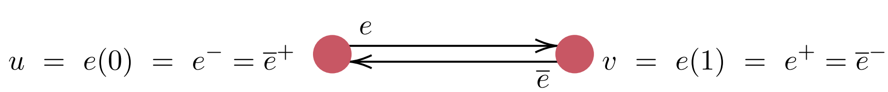

Figure 5: Description of an oriented edge and its reversal.

Almost injectivity ensures oriented edges do not self-intersect except possibly at an endpoint, in which case e is a loop. An *edge* on S is a pair $\mathbf{e} = \{e, \overline{e}\}$, where e is an oriented edge. If E is a set of edges, and $\vec{E}$ is the associated set of oriented edges, then $2|E| = |\vec{E}|$.

Definition 9. An *embedded graph* in S is a graph $G = (V, E)$ where V is a finite subset of S (i.e. vertices are points on the surface) and E is a finite collection of edges such that for every $\mathbf{e} = \{e, \overline{e}\}$ in E,

- the endpoints of $\mathbf{e}$ are exactly the extremeties of e i.e.

$$\iota(\mathbf{e}) = \{e(0), e(1)\};$$

- the interior of $\mathbf{e}$, $\mathrm{int}(\mathbf{e}) = e((0,1))$, does not intersect the interior of any other edge nor the vertex set i.e. for all distinct $\mathbf{e}_i, \mathbf{e}_j$ in E,

$$\mathrm{int}(\mathbf{e}_i) \cap \mathrm{int}(\mathbf{e}_j) = \emptyset \quad \text{and} \quad \mathrm{int}(\mathbf{e}_i) \cap V = \emptyset.$$

That is, an embedding of a graph G involves drawing G on the surface S such that the interior of the edges do not cross one another, themselves, nor the vertex set.

Let G be an embedded graph. The *support* of an embedded graph is the set of all points it inhabits on the surface, denoted by

$$\mathrm{supp}(G) = V \cup \bigcup_{\mathbf{e} \in E} \mathrm{int}(\mathbf{e}).$$

A *face* of G is a connected component support of G. The set of all faces is denoted F. For an oriented edge $e \in \vec{E}$, we take advantage of the fact that our surface is oriented to define $f_e \in F$ as the face to the left of e, when e is traversed in the intuitive direction from e^- to e^+. If $f_e = f_{\bar{e}}$, e is called an *isthmus*. Set the oriented edges $e_1, e_2, \ldots e_d$ bounding a face f to be arranged in counter-clockwise cyclic order, and call this set the *boundary* of f, denoted ∂f. The *degree* of a face is the number of oriented edges which bound the face. Call the region between two consecutive oriented edges a *corner*. We associate a corner with the origin of an oriented edge i.e. denote the corner between e_i and e_{i+1} as e_{i+1} or c_{i+1}, depending on the context. Finally, we impose a cyclic counter-clockwise ordering to the oriented edges extending from a vertex.

Remark 2. Using these new notions of oriented edges and faces, we can extend the Handshaking Lemma to

$$\sum_{v \in V} \deg(v) = 2|E| = |\vec{E}| = \sum_{f \in F} \deg(f).$$

1.3.2 Maps

At last, we arrive at the (topological) definition of a map.

Definition 10. A *map* is an embedded graph whose faces are all homeomorphic to the open unit disk in $\mathbb{R}^2$. When $S = \mathbb{S}^2$, such an embedding is a *plane map*.

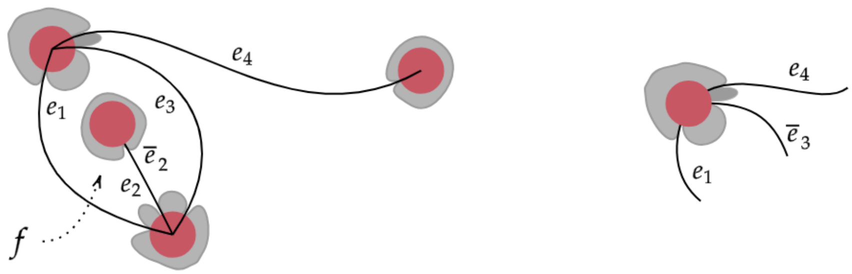

Figure 6: An embedding of a graph in $\mathbb{S}^2$ (left) and a description of the corners belonging to a vertex (right), namely $e_1, \overline{e}_3, e_4$. The face labelled f has degree 4 and has boundary $e_1, e_2, \overline{e}_2, e_3$. Note that $f_{e_2} = f_{\overline{e}_2}$, hence e_2 is an isthmus.

Remark 3. The underlying graph of a map must be connected. Otherwise, the faces of the map will not be homeomorphic to unit disks. For the same reason, the graph embedded in the torus in Figure 7 is not a map because the sole face of the torus is homeomorphic to an annulus in $\mathbb{R}^2$, not a unit disk.

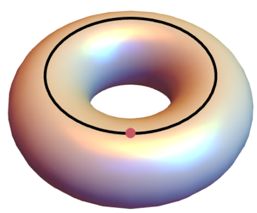

Figure 7: A graph made of a single vertex with a loop embedded around the cirumference of a torus is *not* a map.

Example. The graphs shown in Figures 2 and 3 are plane maps when embedded in the sphere; the outer faces shouldy simply be seen as the surface of the sphere.

Theorem 1. (Jordan's Curve Theorem, on the Sphere)
A Jordan curve c on the sphere is the image of an injective continuous map from the circle to the sphere i.e.

$$c : \mathbb{S}^1 \to \mathbb{S}^2.$$

Then c divides the sphere into exactly two connected regions. The boundary of each component is c.

Proof. See [Jor87]. □

Corollary 1. When $S = \mathbb{S}^2$, maps are exactly the connected embedded graphs.

Proof. The nontrivial direction is $\Longleftarrow$. If G is a connected embedded graph, then the faces are the connected components of $\mathbb{S}^2 \setminus G$, and by Jordan's curve theorem each face is an open region of the sphere, hence is homeomorphic to the unit disk. $\square$

Remark 4. It would be reasonable to ask why maps on the sphere are called plane maps if they are not drawn in the plane. To answer this question, we state the following without proof:

$$G \text{ can be properly embedded in } \mathbb{R}^2 \iff G \text{ can be properly embedded in } \mathbb{S}^2,$$

where G is a connected graph. Therefore, from the perspective of maps, $\mathbb{R}^2$ and $\mathbb{S}^2$ are equivalent.

Proposition 1. *Let G be a connected graph. Then there always exists a g such that G can be embedded as a map into a surface S with genus g.*

Proof. Embed G in the sphere with as few edge crossings as possible. Introduce a new *handle* near every intersection such that one edge can be rerouted over the other. $\square$

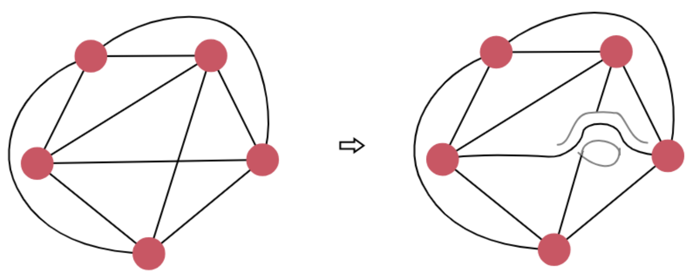

Figure 8: Illustration of the proof of Proposition 1 on the complete graph on 5 vertices, K_5. Therefore, K_5 can be embedded in the torus with $g = 1$.

The problem of finding the smallest such g is an NP-Complete problem known as the graph genus problem [Tho89].

Conversely, for a given surface S with genus g, it is not true that a graph G can be embedded as a map on S. For example, the complete bipartite graph $K_{3,3}$ cannot be embedded in the sphere. We leave the proof for Example 1.3.4. As an aside, this is a direct application of Kuratowski's theorem, which states that a finite graph G can be embedded in the plane if and only if G does not contain K_5 or $K_{3,3}$ as a subdivision. This theorem is beyond the scope of this paper, but illustrates that non-embeddability is witnessed by a finite collection of minimal graphs, which is also true for higher genera (Robertson-Seymour theorem).

1.3.3 Algebraic Description of Maps

Conveniently, there is a simple algebraic analog to our topological point of view of plane maps which involves permutations of the oriented edge set. See the faces of a map $\mathbf{m}$ as a finite set of polygons whose edges are the oriented edges of the faces of the map, again in counter-clockwise cyclic order such that an oriented edge's face lies to its left. Thus, let ϕ be a permutation of $\vec{E}$ whose cycles define the boundaries of these polygons. Then the map can be constructed by gluing together the polygons by pairing up the corresponding oriented edges (Figure 9). To this end, let α be a second permutation of $\vec{E}$ whose cycles are the pairs of oriented edges $\{e, \overline{e}\}$. Finally, let σ be a permutation of $\vec{E}$ whose cycles define the oriented edges extending from the vertices of the map arranged in counter-clockwise order, as we had in the topological definition.

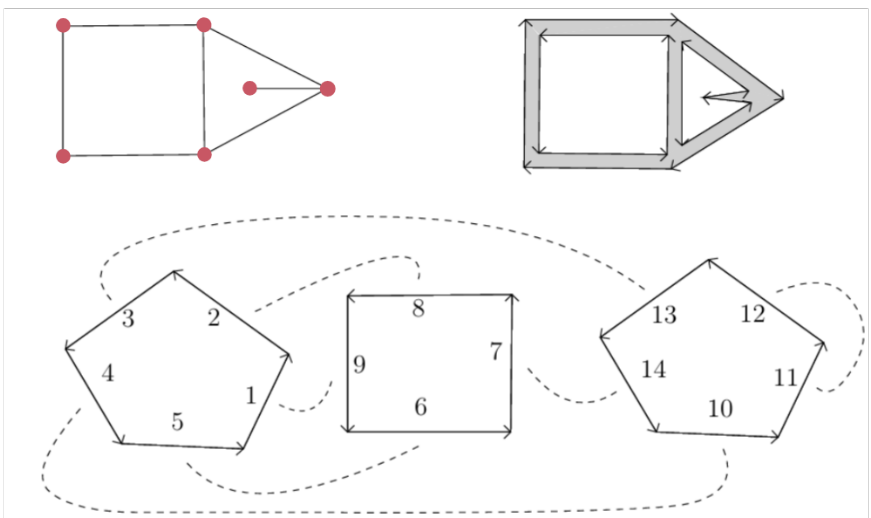

Figure 9: The construction of a map by gluing together its faces. Figure adapted from [Mie14], section 1.2.

Example. The permutations for the map in Figure 9 are

$$\begin{aligned}
\phi &= (1,2,3,4,5)(6,7,8,9)(10,11,12,13,14), \\
\alpha &= (1,9)(2,8)(3,13)(4,10)(5,6)(7,14)(11,12), \\
\sigma &= (1,6)(2,9)(3,8,14)(4,13,11)(5,10,7)(12).
\end{aligned}$$

Remark 5. The connectedness of a map implies that for any two oriented edges e, e' in $\vec{E}$, there exists a word in the letters $\{\sigma, \alpha, \phi\}$ which sends e to e'.

Proposition 2. Any map m can be uniquely defined (up to isomorphism) by any two of the permutations

$$\sigma, \alpha, \phi.$$

Proof. Omit. $\square$

1.3.4 Map Isomorphisms

Which maps should we consider to be the same?

Definition 11. Let m_1 be a map on S_1 and m_2 be a map on S_2. Then m_1 and m_2 are isomorphic if there is an orientation-preserving homeomorphism $\phi : S_1 \to S_2$ such that the restriction of ϕ to m_1 is a graph isomorphism. Represent a set of isomorphic maps using boldface ($\mathbf{m}$, $\mathbf{t}$, etc.).

Intuitively, a map isomorphism not only preserves the incidence relation between vertices (like a map isomophism) but also preserves the incidence relation between edges and faces, or equivalently (by Proposition 2) preserves the cyclic ordering of the edges around each vertex. This definition is subtler than it seems and is best illustrated by an example.

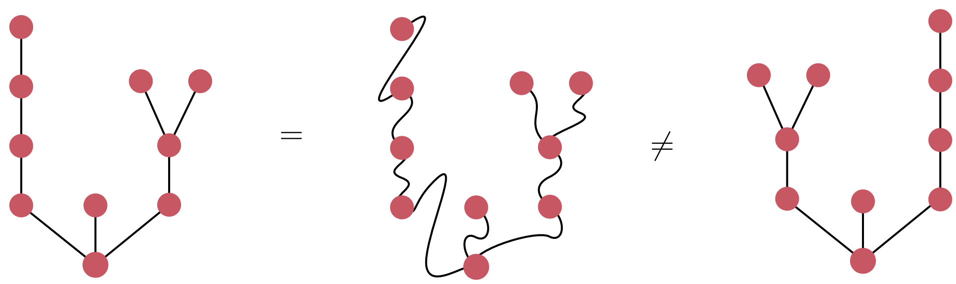

Figure 10: Mirror reflected maps are not isomorphic because the cyclic ordering around at least one vertex has reversed.

We will often root a map along an oriented edge, whose origin is called the root vertex. In doing so, maps with internal symmetry (for example, a square embedded in the sphere has rotational symmetries) become distinguishable under isomophism.

Definition 12. An automorphism of a map is an isomorphism from the map to itself. The trivial isomorphism is the identity map.

We claim that rooting a map removes all nontrivial automorphisms.

Proposition 3. Let m be a map, and let e be its root. Then the only automorphism fixing the root is the identity.

Proof. Let $v = e^-$ be the origin of e. Fixing e also fixes all edges incident to v because an automophism must preserve the cyclic ordering of vertices. By connectivity of the graph, every vertex and edge must be fixed. Hencem the automorphism is the identity. $\square$

Nevertheless, a graph G can give rise to many different rooted maps, depending on the manner in which we embed the graph in the surface. Despite this dependance, the number of regions of a map is always invariant, i.e. is uniquely determined by the number of vertices and edges of the underlying graph and of the topology of the oriented surface. This is the subject of Euler's formula.

Theorem 2. (Euler's formula) Let $G = (V, E)$ be a map embedded in S. Then

$$|V| - |E| + |F| = \chi = 2 - 2g,$$

where χ is the Euler characteristic and g is the genus of S. In particular, when $S = \mathbb{S}^2$,

$$|V| - |E| + |F| = 2,$$

Proof. First, for the particular case of $S = \mathbb{S}^2$, we proceed by induction on $|E| = n$. For the base case, suppose $|E| = 0$. Since G is connected, it consists of only a single vertex, and hence has exactly one face, so the formula holds. Assume the formula holds for a plane map with $n-1$ edges and let G be a plane map with n edges. If G contains no cycles, then G is a tree and $|V| = |E| + 1$. A planar tree has exactly one face, so formula holds. Otherwise, let C be a cycle contained in G. Delete any edge from the cycle. The resulting graph is still connected. Furthermore, the Jordan Curve Theorem asserts that G had at least two faces, and removing this edge merges exactly two faces into one Therefore, we decrease both the number of edges and faces by 1 without changing the number of vertices, so the formula holds by the induction hypothesis.

We omit the proof of Euler's formula for higher genera. $\square$

Example 1.3.4 It is not possible to embed the complete bipartite $K_{3,3}$ in $\mathbb{S}^2$.

Solution. Let G be a planar map isomorphic to $K_{3,3}$ as a graph. $K_{3,3}$ has $3+3=6$ vertices and $3 \times 3 = 9$ edges. By Euler's formula,

$$|F| = 2 - |V| + |E| = 2 - 6 + 9 = 5.$$

Because $K_{3,3}$ has no self-loops or multiple edges, and because all faces are created by crossing pairs (all edges start in one half of the bipartition and end in the other), every face in G must have degree at least 4. Using the extended Handshaking Lemma,

$$18 = 2|E| = |\vec{E}| = \sum_{f \in F} \deg(f) \geq 4|F|.$$

Hence, $|F| \leq 4$, a contradiction.

1.4 Some Important Maps

Definition 13. A quadrangulation $\mathbf{q}$ is a map in which every face has degree 4. The set of all rooted quadrangulations with n faces,distinguished up to isomorphism, is denoted $\mathcal{Q}_n$.

Let $\mathbf{q}$ be a quadrangulation with $|F| = n$ faces. Then because each face is bounded by 4 oriented edges, $\mathbf{q}$ has $4n$ oriented edges hence has $2n$ unoriented edges. By Euler's formula, $\mathbf{q}$ has

$$|F| = 2 - |V| + |E| = 2 - n + 2n = n + 2$$

vertices.

Definition 14. A pointed quadrangulation has a distinguished vertex v^*. The set of all rooted pointed quadrangulations with n faces is $\mathcal{Q}_n^*$.

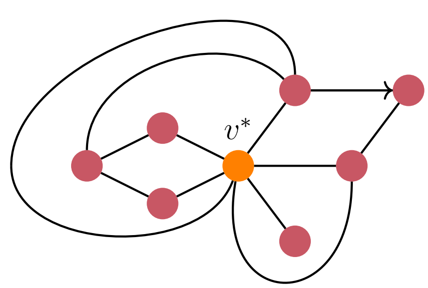

Figure 11: A rooted, pointed quadrangulation with 6 faces.

We now define plane trees, contour walks and the concept of a valid labelling function which assigns to each vertex a number, and the successor function. These tools will be critical in the definition of the CVS bijection.

Definition 15. A plane tree is a map whose underlying graph structure is a tree.

Proposition 4. A plane tree can always be embedded in the sphere and has a unique face.

Proof. The first result follows from the fact that a tree $\mathbf{t}$ has no cycles. The observation that a plane tree has one face follows from Euler's formula and from $|V| = |E| + 1$. □

Remark 6. A map with k oriented edges has $2k$ corners because every oriented edge induces two corners. Therefore, a plane tree on n edges has $2n$ corners.

Among many encodings of trees, including an algebraic fatgraph representation and a labelling of the vertices by integer words (Ulam-Harris encoding), we are interested in encoding trees as a discrete excursion of length $2n$.

Recall that we can traverse the boundary of a face by travelling along the oriented edges in counter-clockwise cyclic order. By Proposition 4, bound the unique face of $\mathbb{S}^2 \setminus \operatorname{supp}(\mathbf{t})$ by the oriented edges of $\mathbf{t}$. Equivalently, call the contour exploration of a tree $\mathbf{t}$ the sequence of oriented edges in *clockwise* cyclic order starting at the root vertex v_0,

$$e_0, e_1, \ldots, e_{2n-1}.$$

Remark 7. Because an oriented edge is uniquely associated with a corner, we will sometimes interpret the elements of the contour exploration as corners, where e_i or c_i is a symbol for the corner associated with the oriented edge e_i.

Definition 16. Let $\mathbf{t}$ be a plane tree and $e_0, e_1, \ldots, e_{2n-1}$ be the contour exploration of $\mathbf{t}$. Denote the i^{th}-visited vertex by $u_i \coloneqq e_i^-$ and set

$$C_{\mathbf{t}}(i) \coloneqq d_{\mathbf{t}}(u_i, v_0), \quad i \in \{0, \ldots, 2n-1\},$$

where $d_{\mathbf{t}}$ is the graph distance. Also, by convention, let $e_{2n} = e_0$ and $u_{2n} = u_0$. Then the *contour function of* $\mathbf{t}$, $C_{\mathbf{t}}(s)$ is the linear interpolation of $C_{\mathbf{t}}(i)$. We call the path traced by the contour function a *discrete excursion of length* $2n$.

Therefore, the contour function can be understood as follows. Imagine an explorer walking along the oriented edges of a tree $\mathbf{t}$. The contour function encodes the explorer's distance to the root at time s as they explore all the vertices of $\mathbf{t}$ in a depth-first-search fashion.

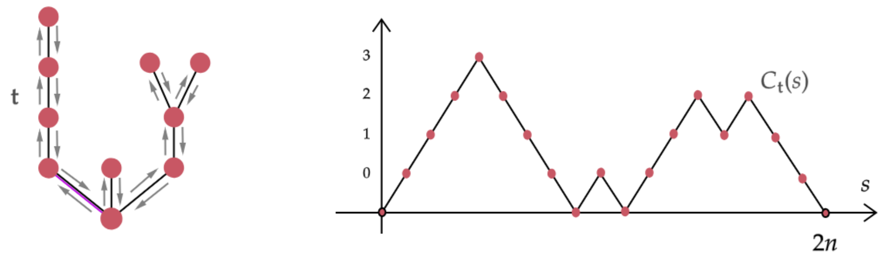

Figure 12: An example of a rooted plane tree and its contour function.

Proposition 5. A rooted plane tree is uniquely encoded by contour function.

Proof. Glue together the discrete excursion in the appropriate fashion. □

Lemma 2. The number of rooted plane trees on n edges is counted by the n^{th} Catalan-number

$$C_n = \frac{1}{n+1}\binom{2n}{n}.$$

Equivalently, the number of rooted plane trees on n edges satisfies the Catalan recurrence

$$|\mathcal{T}_{n+1}| = \sum_{i=0}^{n} |\mathcal{T}_i||\mathcal{T}_{n-i}|, \quad n > 0$$

where $|\mathcal{T}_0| = 1$.

Proof. The Catalan numbers exactly count the number of non-negative paths of length $2n$, starting and ending at 0 with slope ± 1 between integer times, which is exactly a discrete excursion of length $2n$ and hence bijects with the contour function of a plane tree. The formula for the Catalan numbers can be derived by generating functions. Alternatively, see it as a recurrence in which we first decide the leftmost subtree and then decide the remainder. □

Definition 17. Let $\mathbf{t}$ be a rooted plane tree with root vertex u_0. A valid labelling on $\mathbf{t}$ is a function

$$\ell : V(\mathbf{t}) \to \mathbb{Z}$$

such that the root vertex has label 0 i.e. $\ell(v_0) = 0$ and that the labels of adjacent vertices differ by at most 1 i.e.

$$|\ell(u) - \ell(v)| \leq 1, \quad \text{for all adjacent } u, v \in V(\mathbf{t}).$$

Let $\mathbb{T}_n$ be the set of all pairs of rooted plane trees with n edges and valid labellings i.e. of pairs $(\mathbf{t}, \ell)$, where $\mathbf{t} \in \mathcal{T}_n$.

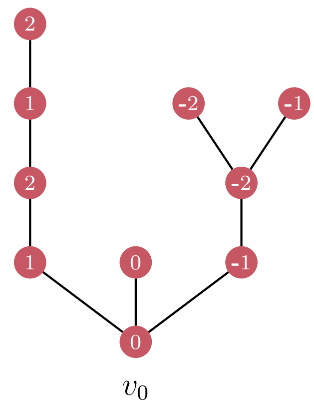

Figure 13: A rooted plane tree with a valid labelling.

That is, a labelled tree is a plane tree whose vertices are adorned with integer labels which differ by $\{-1, 0, 1\}$. The labelling is defined on vertices but we can also see it as a function on the oriented edges such that $\ell(e) := \ell(e^-)$, where e^- the origin of e. Similarly, a corner can be assigned the label of its corresponding vertex.

Definition 18. Let $\mathbf{t}$ be a plane tree with contour exploration $e_0, \ldots e_{2n-1}$. Repeat this sequence of edges with $2n$ periodicity. Then define the successor function on integers as

$$s(i) = \inf\{j > i : \ell(j) = \ell(i) - 1\}, \quad i \in \{0, \ldots, 2n-1\}.$$

Using the direct correspondence between corners and oriented edges, we define the successor of a corner e_i as

$$s(e_i) = e_{s(i).}$$

By convention $\inf \emptyset = \infty$ and therefore

$$s(i) = \infty \iff \ell(e_i) \text{ is minimal }.$$

That is, if we wish for the successor function to be well-defined for the corners of a tree, we will have to introduce a corner with minimum label.

Remark 8. Without loss of generality, we could replace the condition $\ell(j) = \ell(i) - 1$ with $\ell(j) \leq \ell(i) - 1$ because the chain of decreasing labels must decrease by integer steps.

Remark 9. The successor function is not defined for vertices because one vertex can have many corners.

2 The CVS Bijection

Given our interest in planar maps, we seek to encode these combinatorial structures in easier-to-understand objects such as quadrangulations and eventually plane trees. This section is devoted to establishing and understanding the relationships between these classes of maps using bijections.

Proposition 6. A plane map is bipartite if and only if its faces all have even degree.

Proof. We'll say a face f is *contained in* a cycle C when f lies entirely in one of the connected component of $\mathbb{S}^2 \setminus C$. Note that every face is contained in some cycle.

$\Longrightarrow$: Let $\mathbf{m}$ be a bipartite plane map. Then $\mathbf{m}$ contains no odd cycle, so in particular, it contains no odd-degree faces.

$\Longleftarrow$: We prove the contrapositive. Suppose $\mathbf{m}$ is a plane map which is not bipartite, so it contains an odd cycle. We will show it must then contain an odd-degree face. Let C be an odd cycle of length $2k+1$, $k \in \mathbb{N}$. We now induct on the number of faces n contained in C. If $n = 1$, then C is exactly the boundary of a face f, so f has odd degree. Now suppose that the statement holds whenever an odd cycle contains at most $n-1$ faces, and let C be an odd cycle containing $n \geq 2$ faces. Let f be a face contained in C of degree d that shares at least one edge with C. If d is odd, we're done. Otherwise, d is even i.e. $d = 2a$, for some positive integer a, and let ℓ denote the number of oriented edges on the boundary of f that do not lie in C. Then $2a - \ell$ oriented edges in the boundary of f lie in C. Form a new cycle C' by replacing the $2a - \ell$ edges of $C \cap \partial f$ with the ℓ edges of $\partial f \setminus C$. The length of C' is

$$|C'| = |C| - (2a - \ell) + \ell = 2k + 1 - 2a + 2\ell$$

is odd and contains at least one fewer face than C. Therefore, by the induction hypothesis, C' contains an odd-degree face. $\square$

Corollary 2. Quadrangulations are bipartite.

2.1 Encoding of Maps as Quadrangulations & Back

The following bijection was introduced by W.T. Tutte in 1963 [Tut63], and establishes a canonical correspondance between edges of a rooted planar map and faces of a rooted quadrangulation.

Theorem 3. $|\mathcal{M}_n| = |\mathcal{Q}_n|$, where $\mathcal{M}_n$ is the set of rooted plane maps with n edges.

Proof. Let $\mathbf{m} \in \mathcal{M}_n$ be a rooted planar map with n edges. Inside each face of $\mathbf{m}$, add a new vertex, colouring the old vertices black and the new ones pink. Then, connect each pink vertex to every corner of its face with non-intersecting arcs. Finally, erase all original edges of $\mathbf{m}$ and root the quadrangulation at the first outgoing edge from e^- in clockwise order after the root edge e of $\mathbf{m}$. Observe that if a black vertex admits more than one corner belonging to a single face (e.g. isthmus endpoint), the construction produces multiple edges between that vertex and the corresponding white vertex.

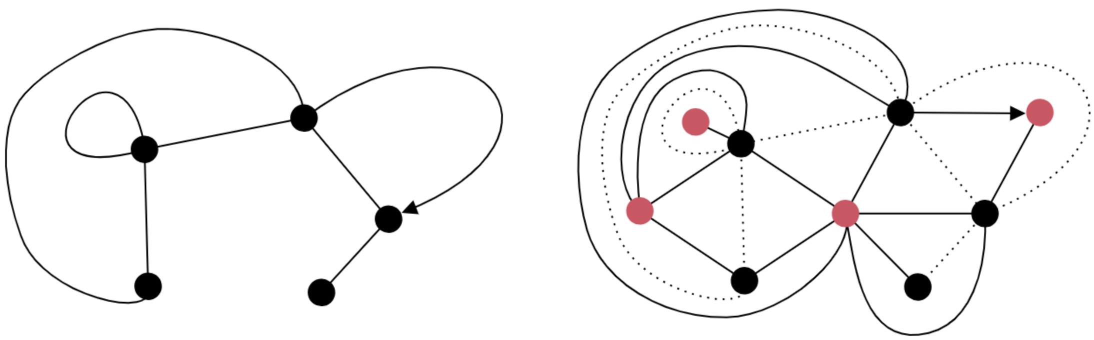

Figure 14: Depiction of the bijection from a rooted plane map to a rooted quadrangulation.

The resulting map is a quadrangulation on n faces; it suffices to show that every edge $\mathbf{e}$ generates exactly one face of $\mathbf{q}$ of degree 4. Write $\mathbf{e} = \{e, \overline{e}\}$ with $e^- = \overline{e}^+ =: u$ and $e^+ = \overline{e}^- =: v$. Since each of the two oriented edges e and $\overline{e}$ is associated to two corners, and each corner produces one new oriented edge, the original edge is envelopped by 4 new oriented edges.

Thus, $\mathbf{q}$ is a quadrangulation on n faces. To root $\mathbf{q}$, declare its root as the first new edge after $\mathbf{e}$ in counter-clockwise order around $\mathbf{e}^-$, where $\mathbf{e}$ is the root of $\mathbf{m}$.

The inverse construction is similar. Given a rooted quadrangulation $\mathbf{q}$, declare a bipartition of its vertices by Corollary 2, colouring the vertices at an even distance from the root edge black and those at an odd distance white. Each face f of $\mathbf{q}$ has degree 4 and adjacent corners have different colours, so there are two black and two white corners incident to each face. Draw an edge between the two black vertices. Finally, erase the original edges of $\mathbf{q}$ to obtain a map $\mathbf{m}$, which we root at the first outgoing edge in counter-clockwise order from e^- after the root edge e of $\mathbf{q}$. $\square$

In particular, this construction provides a correspondence between bipartite plane maps and plane trees.

2.2 Encoding of Quadrangulations as Trees

The CVS bijection, given by Cori and Vauquelin in 1981 in [CV81] defines a function

$$\Phi : \mathbb{T}_n \times \{1, -1\} \to \mathcal{Q}_n^*$$

which sends rooted labelled plane trees (on n edges) to rooted pointed quadrangulations (on n faces) as

$$\Phi((\mathbf{t}, \ell), \varepsilon) \mapsto (\mathbf{q}, v^*).$$

Let $\mathbf{t} \in \mathbb{T}_n$ with root v_0 and label function ℓ. Let $(e_0, \ldots, e_{2n-1})$ be the contour exploration of $\mathbf{t}$. Construct the image of $\mathbf{t}$ as follows:

1. Let the vertex set of $\mathbf{q}$ be the vertex set of $\mathbf{t}$, plus a special vertex v^* placed in the face of $\mathbf{t}$. Define

$$\ell(v^*) = \min\{\ell(v) : v \in V(\mathbf{t})\} - 1,$$

   and denote by e_∞ the corner of v^* i.e. a small area around surrounding v^* which does not intersect any other prexisting corner.

2. Draw non-crossing arcs between the corners e_i and $s(e_i)$, for all $i \in \{0, \ldots, 2n-1\}$. Let s be the successor function with the additional setting of $s(i) \coloneqq \infty$ if $\ell(e_i) = \ell(v^*)$. Arcs should lay entirely in

$$\mathbb{S}^2 \setminus (\mathrm{supp}(\mathbf{t}) \cup v^*),$$

   where $\mathrm{supp}(\mathbf{t})$ is the support of $\mathbf{t}$.

3. Delete the original tree edges.

4. Root the quadrangulation at the arc linking e_0 with $s(e_0)$, oriented from e_0 to $s(e_0)$ if $\varepsilon = 1$ or from $s(e_0)$ to e_0 if $\varepsilon = -1$.

Lemma 3. Φ as described above is well-defined. That is, the image $\Phi(\mathbf{t}, \varepsilon)$ is a pointed quadrangulation.

Proof. We first show that it is possible to draw the arcs in step 2 such that the resulting graph is an embedded graph. That is, adding the new arcs does not necessarily violate the planarity of the graph. To show this, let e_1, e_2, e_3, e_4 be distinct corners arising *in this order* from the contour exploration of $\mathbf{t}$ such that $s(e_1) = e_3$. For contradiction, assume that $s(e_2) = e_4$. Observe that if this were true, the newly-created graph could not be planar. Now, we repeatedly use the definition of successor to

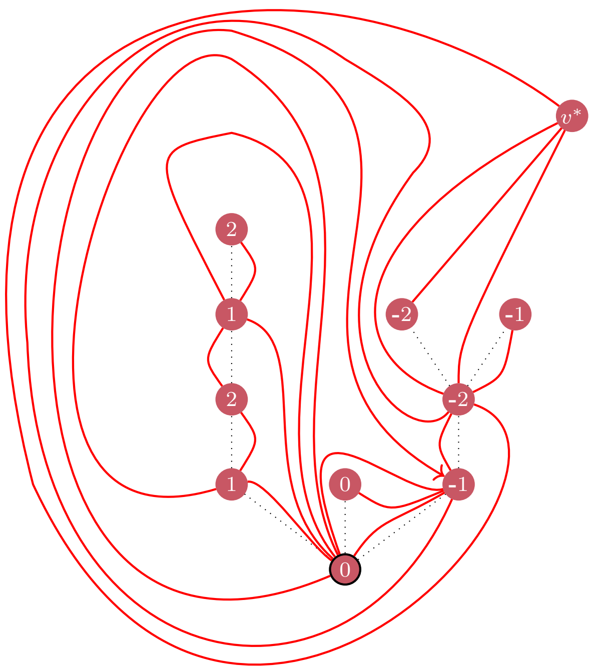

Figure 15: The result of the CVS bijection applied to a rooted labelled tree with $\varepsilon = 1$.

derive a contradiction. First, $s(e_1) = e_3$ implies that the label of e_2 is not smaller than the label of e_1, i.e.

$$\ell(e_2) \geq \ell(e_1).$$

Similarly, $s(e_2) = e_4$ implies that

$$\ell(e_3) \geq \ell(e_2).$$

We know that the successor of e_1 is e_3, therefore

$$\ell(e_3) = \ell(e_1) - 1.$$

Altogether,

$$\ell(e_1) \leq \ell(e_2) \leq \ell(e_3) = \ell(e_1) - 1,$$

a contradiction. In other words, if an arc skips any vertices, the successors of all the enclosed vertices will be found in the enclosure.

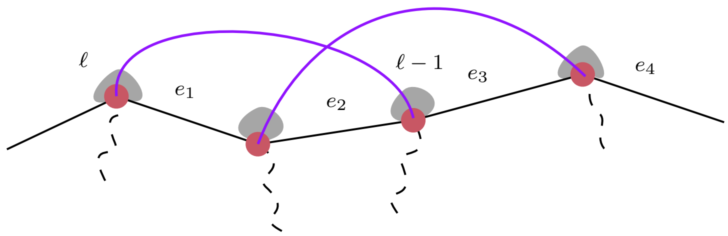

Figure 16: Illustation of the non-planarity of the graph if $s(e_2) = e_4$.

It follows from the previous paragraph that $\mathbf{q}$ is also an embedded graph, since edges are only deleted (not added) during the final steps. Also, $\mathbf{q}$ is connected because there exists a finite sequence of successors from any vertex v to v^*, hence a path between any two vertices u and v via v^*. Therefore, $\mathbf{q}$ is a connected embedded graph, hence is a planar map by Corollary 1.

We claim that every face of $\mathbf{q}$ has degree 4, hence $\mathbf{q}$ is a quadrangulation. We'll show that the CVS bijection encloses every tree edge inside a face of degree 4 by illustrating the four quadrangulate edges case-by-case. Consider an edge $\mathbf{e} = \{e, \overline{e}\}$ of $\mathbf{t}$. Let c be the corner associated with e, c' be the next corner after e in the contour exploration of $\mathbf{t}$, and $\overline{c}$ be the corner associated with $\overline{e}$ and $\overline{c}'$ be the corner associated with the next corner after $\overline{e}$ in the contour exploration.

1. $\ell(\overline{e}) = \ell(e)+1$: First, an arc is drawn from c to $s(c)$ where $s(c)$ is the first corner coming after c with label $\ell(s(c)) = \ell(e) - 1$. Second, an arc is drawn from c' to $s(c')$, where $s(c')$ is the first corner coming after c' with label $\ell(s(c')) = \ell(e)$. Third, because labels can only decrease in integer steps, we know that the vertex associated with $s(c)$ with label $\ell(e) - 1$ must occur after the vertex associated with $s(c')$ with label $\ell(e)$. Therefore, the first corner after $s(c')$ with label one less i.e. $s(s(c'))$ has label $\ell(s(s(c')) = \ell(s(c')) - 1 = \ell(e) - 1$, and hence $s(s(c')) = s(c)$. Finally, the CVS bijection draws an arc from $\overline{c}$ to $\overline{c}'$ because the label of $\ell(\overline{c}') = \ell(\overline{c}) - 1$ by assumption.

2. $\ell(\overline{e}) = \ell(e) - 1$: This case is equivalent to $\ell(e) = \ell(\overline{e}) + 1$, hence is symmetric to Case 1 by interchanging e and $\overline{e}$.

3. $\ell(e) = \ell(\overline{e})$: Let c' be the next corner after e in the contour exploration and $\overline{c}'$ be the next corner after $\overline{e}$, the same as in case 1. Then since

$$\ell(\overline{e}) = \ell(c') = \ell(\overline{c}') = \ell(e),$$

we have $s(c) = s(c')$ and $s(\overline{c}) = s(\overline{c}')$. Therefore, the edges of the quadrangle are the arcs linking c to $s(c)$, c' to $s(c') = s(c)$, $\overline{c}$ to $s(\overline{c})$ and $s(\overline{c}')$ to $s(\overline{c})$.

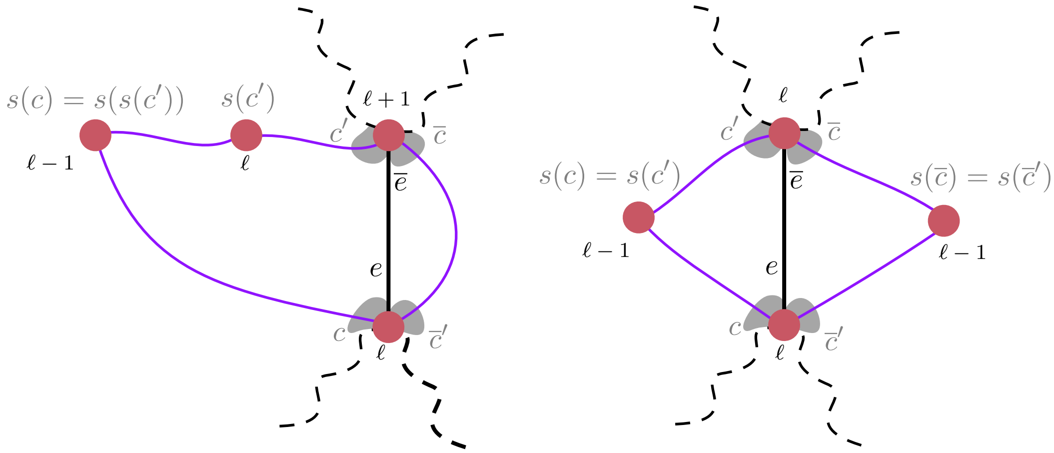

Figure 17: Cases 1 and 3. The arcs drawn by the CVS bijection are purple, and the corners are gray.

Observe that in Case 1, adjacent vertices in the face have labels differing by either 2 or 0 and never 1, hence there are no other quadrangulation edges drawn in the interior of the new face, and similarly in Case 2, the only possible . Therefore, the construction above correctly

It remains to show that every face of $\mathbf{q}$ is of one of the forms above. Observe that $\mathbf{q}$ has $2n$ edges (one per corner) and $n + 2$ vertices, and therefore must have n faces by Euler's formula. The result follows because the original tree had n edges. Therefore, $\mathbf{q}$ is a quadrangulation that is rooted according to ε and is pointed by v^*. □

Theorem 4. Φ is a bijection.

Proof. See [CG04] for the construction of the inverse

$$\Psi : \mathcal{Q}_n^* \to \mathbb{T}_n \times \{1, -1\},$$

and for the proof that $\Phi^{-1} = \Psi$. □

3 Combinatorial and Geometric Applications

3.1 Counting & Sampling Planar Maps

Having now proven the bijection between quadrangulations and planar trees we are ready to enumerate rooted planar maps with n edges, denoted $\mathcal{M}_n$. While counting planar maps is cumbersome (we must account for the possible topologies arising from any number of faces), counting planar trees is straightforward.

Theorem 5.

$$|\mathcal{M}_n| = |\mathcal{Q}_n| = \frac{2}{n+2}3^n \cdot C_n$$

where $C_n = \frac{1}{n+1}\binom{2n}{n}$ is the nth Catalan number.

Proof. Recall our bijection:

$$\Phi : \mathbb{T}_n \times \{1, -1\} \to \mathcal{Q}_n^*$$

$\mathcal{Q}_n^*$ is the set of rooted quadrangulation with a distinguished vertex. We have established that there are $n+2$ vertices in a quadrangulation with n faces, so given a quadrangulation in $\mathcal{Q}_n$, there are $n+2$ ways to pick the distinguished vertex, thus $|\mathcal{Q}_n^*| = (n+2) \cdot |\mathcal{Q}_n|$.

Furthermore, $\mathbb{T}_n$ is the set of planar maps with a valid labelling. Such a labelling requires that neighbouring vertices u, v have labels that differ by no more than 1. Excluding the fixed 0 label of the starting vertex, we have 3 choices for every label $\ell(v) \in \{\ell(u) - 1, \ell(u), \ell(u) + 1\}$. There are n non-rooted vertices in a tree with n edges, so there are 3^n different labellings. Then, by Lemma 2, $|\mathbb{T}_n| = 3^n|\mathcal{T}_n| = 3^n \cdot C_n$, where $\mathcal{T}_n$ is the set of rooted plane trees with n edges. Finally, as our bijection is from $\mathbb{T}_n \times \{1, -1\}$, we have $2 \cdot 3^n \cdot C_n = (n+2)|\mathcal{Q}_n|$ and therefore

$$|\mathcal{M}_n| = |\mathcal{Q}_n| = \frac{2}{n+2} \cdot 3^n \cdot C_n.$$

□

As our goal is to study random planar maps, the following result will be of use to us.

Corollary 3. Choosing a uniformly random rooted planar map is the same (in distribution) as choosing a uniformly random planar tree (along with root edge). That is to say:

Letting $\mathbf{q} \in_u \mathcal{Q}_n$ and choosing $v^* \in_u V(\mathbf{q})$ has the same distribution as $((\mathbf{t}, \ell), \epsilon)$, where $(\mathbf{t}, \ell) \in_u \mathbb{T}_n$ and $\epsilon \in_u \{-1, 1\}$.

Proof. The probability of choosing $(\mathbf{q}, v^*)$ is

$$\frac{1}{(n+2)|\mathcal{Q}_n|} = \frac{1}{|\mathcal{Q}_n^*|}$$

which, by the CVS bijection, is the same probability as choosing our random triplet $((\mathbf{t}, \ell), \epsilon)$, where $(\mathbf{t}, \ell) \in_u \mathbb{T}_n$ and $\epsilon \in_u \{-1, 1\}$.

$\square$

3.2 Exact & Bounded Distances

We now turn our attention to understanding the relationship between the labels in our tree and the distances in its associated quadrangulation.

Theorem 6. Given a pointed, rooted quadrangulation $\mathbf{q} \in \mathcal{Q}_n^*$, for any vertex $v \in V(q)$ the distance between v and v^* in $\mathbf{q}$ is:

$$d_{\mathbf{q}}(v, v^*) = \ell(v) - \min_{v \in V(\mathbf{t})} \ell(v) + 1,$$

where $\mathbf{t}$ is the tree associated to $\mathbf{q}$ by the CVS bijection.

Proof. First, observe that

$$\ell(v) - \min_{v \in V(\mathbf{t})} \ell(v) + 1 = \ell(v) - \left(\min_{v \in V(\mathbf{t})} \ell(v) - 1 \right) = \ell(v) - \ell(v^*)$$

and so we aim to prove that $d_{\mathbf{q}}(v, v^*) = \ell(v) - \ell(v^*)$.

> $\boldsymbol{d_{\mathbf{q}}(v, v^*) \leq \ell(v) - \ell(v^*)}$**:** Let e be a corner in $\mathbf{t}$ incident to v. Then, the sequence $e, s(e), s^{(2)}(e), \ldots, e_\infty$ where $s^{(n)}(e)$ is the successor function applied n times to e, is of length $\ell(v) - \ell(v^*)$ between v and v^* and through the CVS bijection, also a path within $\mathbf{q}$ from v to v^* and so $d_{\mathbf{q}}(v, v^*) \leq \ell(v) - \ell(v^*)$.
>
> $\boldsymbol{d(v, v^*) \geq \ell(v) - \ell(v^*)}$**:** Consider an arbitary path from v to v^* in $\mathbf{q}$ that visits the sequence of vertices $v = v_0, v_1, \ldots, v_n = v^*$. Remember that for two neighbouring vertices in $\mathbf{q}$, $|\ell(v_i) - \ell(v_{i-1})| = 1$ and hence:
>
> $$\begin{aligned} n &= \sum_{i=1}^{n} 1 \\ &= \sum_{i=1}^{n} |\ell(v_i) - \ell(v_{i-1})| \\ &\geq \left| \sum_{i=1}^{n} (\ell(v_i) - \ell(v_{i-1})) \right| \\ &= |\ell(v_n) - \ell(v_0)| \\ &= \ell(v) - \ell(v^*) \end{aligned}$$

Since the sequence $v_0, \dots, v_n$ was arbitrary, the bound holds for all paths, and in particular the shortest path with length $d_{\mathbf{q}}(v, v^*)$ which completes the proof. $\square$

Continuing on to the distance of two arbitrary vertices u, v in $\mathbf{q}$, while we don't have a formula for the exact distance $d_{\mathbf{q}}(u, v)$ in $\mathbf{q}$, we do have a strong upper and lower bound that it exists between.

Before stating the theorem, we prove the following lemma which will be of use to us:

Lemma 4. $d_{\mathbf{q}}(u, v) \geq |\ell(u) - \ell(v)|$.

Proof. The proof follows by the triangle inequality and the previous theorem:

$$\begin{aligned}|\ell(u) - \ell(v)| &= |(\ell(u) - \ell(v^*)) - (\ell(v) - \ell(v^*))| \\ &= |d_{\mathbf{q}}(u, v^*) - d_{\mathbf{q}}(v, v^*)| \\ &\leq d_{\mathbf{q}}(u, v).\end{aligned}$$

$\square$

Theorem 7. Let $\mathbf{t} \in \mathcal{T}_n$ be a rooted plane tree. Denote $[e, e']$ as the sequence of corners encountered in the contour walk on $\mathbf{t}$ from e until we reach e' (including e and e'). Denote $[[u, v]]$ as the set of vertices lying in the shortest path from u to v in $\mathbf{t}$. Let $((\mathbf{t}, \ell), \varepsilon) \in \mathbb{T}_n \times \{-1, 1\}$ and $(\mathbf{q}, v^*)$ be the corresponding tuple from the CVS bijection. Consider two vertices u, v in $V(\mathbf{q}) \setminus \{v^*\}$, and let e, e' be the two corners of $\mathbf{t}$ such that $e^- = u$ and $(e')^- = v$. Then

$$d_{\mathbf{q}}(u, v) \leq \ell(u) + \ell(v) - 2 \min_{e'' \in [e, e']} \ell(e'') + 2$$

and

$$d_{\mathbf{q}}(u, v) \geq \ell(u) + \ell(v) - 2 \min_{w \in [[u, v]]} \ell(w)$$

Proof.

Upper Bound: We construct a path in $\mathbf{q}$ of length $d_{\mathbf{q}}(u, v) \leq \ell(u) + \ell(v) - 2 \min_{e'' \in [e, e']} \ell(e'') + 2$. First, denote $m = \min_{e'' \in [e, e']}$, and let e'' be the first corner in our contour walk that achieves this minimum. Now, since e'' achieves the minimum, its successor must be outside of the $[e, e']$. In particular, it is somewhere past v in the tree's complete contour walk. Consider the two paths in $\mathbf{q}$, one from $u \to s(e) \to s^{(2)}(e) \to \cdots \to s^{(k)}(e) = s(e'')$ of length $\ell(u) - (m - 1) = \ell(u) - m + 1$ and $v \to s(e') \to s^{(2)}(e') \to \cdots \to s^{(j)}(e') =$

$s(e'')$ of length $\ell(v) - m + 1$. Combining these paths gives a path of length $\ell(u) + \ell(v) - 2\min_{e'' \in [e,e']} \ell(e'') + 2$. Since $d_{\mathbf{q}}(u,v)$ is the shortest path, it follows that:

$$d_{\mathbf{q}}(u,v) \leq \ell(u) + \ell(v) - 2 \min_{e'' \in [e,e']} \ell(e'') + 2.$$

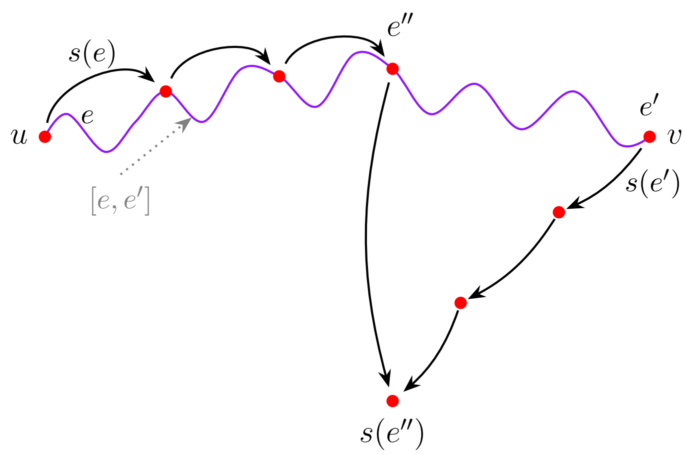

Figure 18: Upper Bound: Construction of the path in $\mathbf{q}$ using successor hops in $\mathbf{t}$.

Lower Bound: Denote the vertex with the smallest label in $[[u,v]]$ as w. If $w = u$ or $w = v$, our bound follows from Lemma 4, so suppose $w \notin \{u,v\}$. Then we can take any partition (there may be more than one) of $\mathbf{t}$ into two connected subtrees $\mathbf{t}_1 \ni u, \mathbf{t}_2 \ni v$ such that $\mathbf{t}_1 \cup \mathbf{t}_2 = \mathbf{t}$ and $\mathbf{t}_1 \cap \mathbf{t}_2 = \{w\}$. Now consider the shortest path from u to v in $\mathbf{q}$. If this path contains v^*, then we have $d_{\mathbf{q}}(u,v) = d_{\mathbf{q}}(u,v^*) + d_{\mathbf{q}}(v^*,u)$ and our bound follows again from Lemma 4. Now suppose v^* is not in the shortest path. We know that the shortest path crosses from $\mathbf{t}_1$ to $\mathbf{t}_2$, so denote the two corners where this crossing happens e_1 and e_2 where $e^-_{(1)} \in \mathbf{t}_1$ and $e^-_{(2)} \in \mathbf{t}_2$, and $e^-_{(1)}, e^-_{(2)}$ are adjacent in $\mathbf{q}$. Since neighbours in $\mathbf{q}$ differ in label by one, without loss of generality, suppose $e_{(2)} = s(e_{(1)})$. Now, w must be visited in the contour walk from $e^-_{(1)}$ to $e^-_{(2)}$. There are two possibilities: either w is the endpoint $e^-_{(2)}$, or else we skip over w when travelling to the successor of $e_{(1)}$ and hence $\ell(w) > \ell(e^-_{(2)})$. Together we have that $\ell(w) \geq \ell(e^-_{(2)})$. Using Lemma 4 again:

$$\begin{aligned} d_{\mathbf{q}}(u,v) &= d_{\mathbf{q}}(u, e^-_{(2)}) + d_{\mathbf{q}}(e^-_{(2)}, v) \\ &\geq \ell(u) - \ell(e^-_{(2)}) + \ell(v) - \ell(e^-_{(2)}) \\ &\geq \ell(u) + \ell(v) - 2\ell(w) \end{aligned}$$

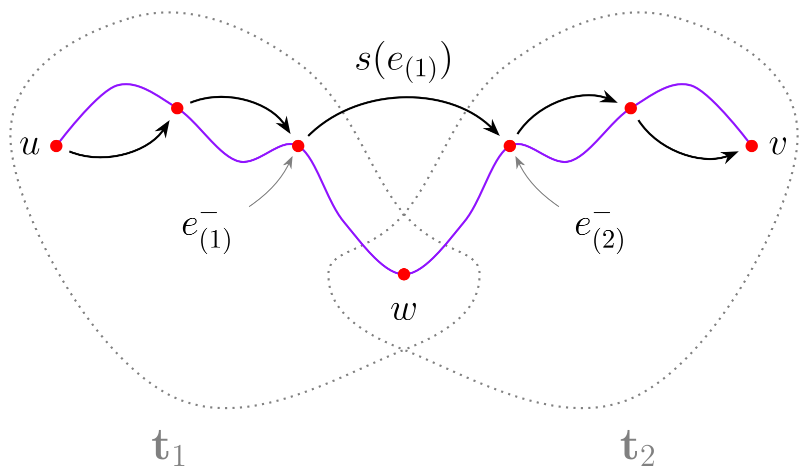

Figure 19: Construction of the path in $\mathbf{q}$ between subtrees $\mathbf{t}_1$ and $\mathbf{t}_2$ jumping over w

$\square$

4 Random Maps

We'd like to study the probabilistic consequences of the CVS bijection. In particular, we will use expected distances in plane trees as a proxy for expected distances in a map. Recall that $\mathcal{Q}_n^*$ is the set of rooted pointed quadrangulations and $\mathcal{Q}_n$ the set of rooted quadrangulations (both on n faces), and $\mathbb{T}_n$ is the set of rooted labelled plane trees and $\mathcal{T}_n$ the set of rooted plane trees (both on n edges).

Given a random rooted plane map $\mathbf{m}$, apply Theorem 3 to obtain a rooted quadrangulation $(\mathbf{q}, v^*)$. Then apply Corollary 3 to obtain a uniformly random rooted labelled plane tree, i.e.

$$\Phi^{-1}(\mathbf{q}, v^*) = ((\mathbf{t}, \ell), \epsilon) \in_u \mathbb{T}_n \times \{-1, 1\},$$

and then forget the root orientation and labels. Indeed, the underlying plane tree is uniformly random because

$$\begin{aligned}\mathbb{P}(T = \mathbf{t}) &= 2 \times 3^n \times \mathbb{P}(((T, L), \varepsilon) = ((\mathbf{t}, \ell), \epsilon)) \\ &= 2 \times 3^n \times \frac{1}{|\mathbb{T}_n \times \{1, -1\}|} \\ &= 2 \times 3^n \times \frac{1}{2 \times 3^n \times |\mathcal{T}_n|} \\ &= \frac{1}{|\mathcal{T}_n|}.\end{aligned}$$

4.1 Foata Fuchs Bijection

Allow us for a second to consider abstract combinatorial trees instead of plane trees. With the goal of quantifying expected distances in trees, we turn to an alternative proof of Cayley's formula, introduced by Foata and Fuchs in 1970 in [FF70] and exposited by Addario-Berry et al. in [Add+22]. The typical proof uses so-called Prüfer codes. Both proofs involve a bijection between a tree and a sequence of the tree's vertices, based on leaf-pruning mechanisms.

Let t be a rooted tree with vertices $\{0, \ldots, n\}$. For a vertex v and a connected subset S of vertices in t, the *path from S to v in t* is the unique shortest path in t which starts at a vertex in S and ends at v. It may also be helpful to keep in mind our prototypical tree, labelled with the indices of the vertices, and its corresponding sequence under the Foata-Fuchs bijection.

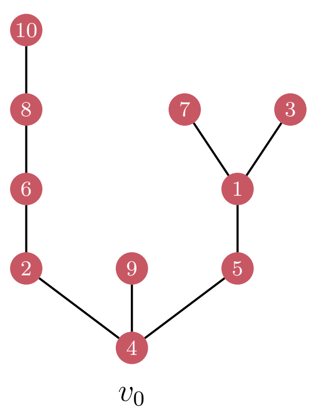

$$(4, 5, 1, 1, 4, 4, 2, 6, 8)$$

Foata Fuchs Bijection

$$\nu : \mathscr{T}_n \to \{0, \ldots, n\}^n$$

For a tree with root v_0:

1. Let $\ell_1 < \ell_2 < \cdots < \ell_k$ be the leaves of the tree in increasing order.
2. Initialize $S_0 = \{v_0\}$ and $\nu = ()$.
3. Recursively, let P_i be the path from S_{i-1} to ℓ_i. Update $S_i = S_{i-1} \cup P_i$, and letting P_i^* be the path without the leaf ℓ_i, append P_i^* to ν.
4. Return ν.

Lemma 5. The Foata Fuchs bijection is well-defined, that is, it maps a tree t on vertices $\{0, \ldots n\}$ to a sequence $\nu \in \{0, \ldots n\}^n$.

Proof. It is trivial that the entries of ν are elements of $\{0, \ldots, n\}$. Thus, we just need to show that when the algorithm terminates, ν has length n. Indeed, for each of the k leaves, the edges in the paths $P_1, \ldots, P_k$ are distinct and since each leaf gets visited, so do all the edges. Moreover, the number of vertices in P_i^* is equal to the length (number of edges) of P_i. Thus, the total length of ν is the sum of lengths of the paths $P_1, \ldots, P_k$, which is equal to the number of edges in t. $\square$

Inverse of the Foata Fuchs Bijection

$$t : \{0, \ldots, n\}^n \to \mathscr{T}_n$$

For a sequence $\nu = (v_0, v_1, \ldots, v_{n-1}) \in \{0, \ldots, n\}^n$:

1. Let $j(0) = 0$ and $j(1), \ldots, j(k-1)$ be the indices of the repeated vertices in ν, and let $j(k) = n$.
2. Let $\ell_1 < \ell_2 < \cdots < \ell_k$ be the elements of $\{0, \ldots, n\}$ not in ν. These will become the leaves of the tree.
3. For $i = 1, \ldots, k$, let P_i be the path

   $$(v_{j(i-1)}, \ldots, v_{j(i)-1}, \ell_i)$$

   with $j(i) - j(i-1)$ edges. That is, P_i is the portion of ν from the ith repeated vertex to just before the $(i+1)$st repeated vertex, appended with ℓ_i.
4. Let $t(\nu)$ be the tree on vertices $\{0, \ldots, n\}$ with root v_0 and edge set

   $$E = \bigcup_{i=1}^{k} E(P_i).$$

Lemma 6. The inverse of the Foata Fuchs bijection is well-defined, that is, it maps a sequence $\nu \in \{0, \ldots, n\}^n$ to a rooted tree t on vertices $\{0, \ldots, n\}$.

Proof. For each i, the path P_i and the vertices

$$\bigcup_{m<i} V(P_m)$$

intersect at a single vertex, namely $v_{j(i-1)}$, the $(i-1)$st repeated vertex. Thus, in particular, $t(\nu)$ is connected and contains no cycles. Additionally, since for each i, ℓ_i appears at the end of P_i and nowhere else (because ℓ_i does not appear in ν), the leaves of $t(\nu)$ are $\ell_1, \ldots, \ell_k$. $\square$

Lemma 7. The Foata Fuchs algorithm is a bijection from $\mathscr{T}_n$ to $\{0, \ldots, n\}^n$.

Proof. We claim that the inverse of the Foata Fuchs bijection is $t(\nu)$, that is,

$$\nu(t(\nu)) = \nu.$$

We are given k; $P_1, \ldots, P_k$ and $v_0, v_{j(1)}, \ldots, v_{j(k-1)}$ defined as they have been previously. By definition, v_0 is the root of the tree $t(\nu)$. Also, P_1 is a path contained in $t(\nu)$ from v_0 to ℓ_1, and from Lemma 6, we know the path is unique. Similarly, for all $1 \leq i \leq k$, P_i is a path from $v_{j(i-1)}$ to ℓ_i and as we saw previously:

$$P_i \cap \bigcup_{m<i} V(P_m) = \{v_{j(i-1)}\}$$

and hence, P_i is the unique path from the union of paths $P_1, \ldots, P_{i-1}$ (the partially constructed tree) to add a unique path to the leaf ℓ_i. Once we apply $\nu(t(\nu))$, we strip each leaf l_i to get $P_i^* = (v_{j(i-1)}, \ldots, v_{j(i)-1})$ and append it to the sequence. This process yields our original ν:

$$P_1^* P_2^* \ldots P_k^* = (v_0, \ldots, v_{n-1}).$$

$\square$

Theorem 8. (Cayley's Formula for Rooted Trees)

$$|\mathscr{T}_n| = (n+1)^n,$$

where $\mathscr{T}_n$ is the set of all rooted trees on n edges.

Proof. Follows from the Foata Fuchs bijection. $\square$

4.2 Expected Height of A Random Leaf

In this section, we seek to quantify the expected height of a uniformly random rooted tree on n edges. In particular, we will show that the distance from the root to a random leaf is $O(\sqrt{n})$. For a tree t, assume the vertices are denoted by $\{0, \ldots, n\}$ and let v_0 be the root. Call the distance from the root to a vertex v the *height of* v, i.e. $\text{height}(v) = d_{\mathbf{t}}(v_0, v)$.

The following lemma demonstrates that to find the distance from the root to a random leaf, it suffices to find the distance from the root to the smallest-indexed leaf.

Lemma 8. Let $\mathbf{t}$ be a uniformly random rooted tree on n edges, let ℓ be the leaf of $\mathbf{t}$ with smallest label, and let ℓ' be a uniformly random leaf of $\mathbf{t}$. Then

$$d_{\mathbf{t}}(v_0, \ell') \stackrel{\mathrm{d}}{=} d_{\mathbf{t}}(v_0, \ell).$$

Proof. Follows from the permutability of the labels of the vertices. $\square$

Definition 19. Let t be a tree on vertices $\{0, \ldots, n\}$. For a vertex $i \in \{0, \ldots, n\}$, the number of children of i is the number of vertices adjacent to i which are further from v_0. The child sequence of a tree t is the sequence $c = (c_0, c_1, \ldots, c_n)$ such that vertex i has c_i children. A tree is said to be 1-free if no vertex has one child i.e. $c_i \neq 1$, for all i.

Lemma 9. Let t be a uniformly random 1-free rooted tree on n edges, and let ℓ be the smallest leaf. Then for all $x > 0$,

$$\mathbb{P}\left(\frac{\text{height}(\ell)}{\sqrt{n}} > x\right) \leq \exp\left(-\frac{x^2}{2}\right).$$

Proof. Under the Foata Fuchs bijection, $\nu(t) = (v_0, v_1, \ldots, v_n)$ is a uniformly random sequence in $\{0, \ldots, n\}^n$ because t was uniformly random to start. Furthermore, by construction, the height of the smallest leaf in t is the index of the first repetition in $\nu(t)$, which is well-defined because t is 1-free and thus has at least one repeated entry. That is, $\text{height}(\ell) = j(1)$, where j is described Lemma 7.

By the chain rule,

$$\begin{aligned}
\mathbb{P}(\mathrm{height}(\ell) > k) &= \mathbb{P}(v_0, \ldots, v_{k-1} \text{ distinct}) \\
&= \prod_{m=1}^{k-1} \mathbb{P}(v_m \notin \{v_0, \ldots, v_{m-1}\} \mid v_0, \ldots, v_{m-1} \text{ distinct}) \\
&\overset{\star}{\leq} \prod_{m=1}^{k-1} \left(1 - \frac{m}{n-m}\right) \\
&\leq \prod_{m=1}^{k-1} \exp\left(-\frac{m}{n-m}\right) && \because\ 1 - x \leq e^{-x} \\
&= \prod_{m=1}^{k-1} \exp\left(-\frac{m}{n}\right) && \because\ n - m \leq n \\
&= \exp\left(-\frac{1}{n}\sum_{m=1}^{k-1} m\right) \\
&= \exp\left(-\frac{1}{n}\frac{k(k-1)}{2}\right) \\
&\leq \exp\left(-\frac{(k-1)^2}{2n}\right)
\end{aligned}$$

with $\star$ holding from the assumption of 1-freeness. Observe that the multiplicity of i in $\nu(t)$ is equal to c_i. Namely, the m previous entries must reappear at least once among the remaining $n - m$ entries, so at most $n - m - m = n - 2m$ of those entries can be new. Thus the probability of seeing a new entry at the mth step is at most ${}^{(n-2m)}\!/\!_{(n-m)} = 1 - {}^{m}\!/\!_{(n-m)}$.

Substituting $k - 1 = x\sqrt{n}$ gives

$$\mathbb{P}\left(\frac{\mathrm{height}(\ell)}{\sqrt{n}} > x\right) \leq \exp\left(-\frac{x^2 n}{2n}\right) = \exp\left(-\frac{x^2}{2}\right).$$

Importantly, our bound only depends on x and not n. □

Remark 10. The assumption of 1-free is not necessary. We omit the proof.

Corollary 4. The height of a uniform random leaf is $O(\sqrt{n})$ in expectation.

Proof. Follows from Lemma 8 and by integrating the bound in Lemma 9. □

5 Acknowledgments

We would like to thank Sasha Bell, Catherine Fontaine, and Hazem Hassan, the organizers of the DRP this year, for giving us the opportunity to learn about random maps. We are especially thankful to our mentor Gabriel Crudele for his patience, knowledge and guidance throughout the semester.

References

[Jor87] Camille Jordan. *Cours d'Analyse*. 1887. URL: `http://www.maths.ed.ac.uk/~aar/jordan/jordan.pdf`.

[Tut63] W.T. Tutte. "A cencus of planar maps". In: *Canadian Journal of Mathematics, 15: 249-271* (1963).

[FF70] Dominique Foata and Aimé Fuchs. "Réarrangements de fonctions et dénombrement". In: *Journal of Combinatorial Theory, 8(4):361-375* (1970).

[CV81] R. Cori and B. Vauquelin. "Planar maps are well labeled trees". In: *Canadian Journal of Mathematics, 33: 1023-1042* (1981).

[Tho89] C. Thomassen. "The graph genus problem is NP-complete". In: *J. Algorithms (10) 568–576* (1989).

[JJ97] B. Durhuus J. Ambjørn and T. Jonsson. "Quantum geometry. A sta tistical field theory approach." In: *Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge* (1997).

[CG04] P. Chassaing and G.Schaeffer. "Random planar lattices and integrated superBrownian excursion". In: *Probab. Theory Related Fields, 128(2):161-212* (2004).

[Mie14] Grégory Miermont. "Aspects of Random Maps". In: *Lecture notes of the 2014 Saint-Flour Probability Summer School* (2014).

[Add+22] Louigi Addario-Berry et al. *The Foata-Fuchs proof of Cayley's formula, and its probabilistic uses*. 2022. arXiv: `2107.09726 [math.CO]`. URL: `https://arxiv.org/abs/2107.09726`.

[Stu] Prof. Dr. Benedikt Stufler. *Simulations: The Brownian sphere approximated by simple quadrangulations*. URL: `https://www.dmg.tuwien.ac.at/stufler/sim/simplequad/` (visited on 04/22/2026).