Inquiro, the UAB undergraduate science research journal

Volume 2012 | Number 6 Article 23

2012

Counting Sibling Portraits in Laminations of the Unit Disk

Joseph W. Olson

John C. Mayer

Luka Mernik

research paper

Counting Sibling Portraits in Laminations of the Unit Disk

Joseph W. Olson[1], Dr. John C. Mayer[1], Luka Mernik[2]
[1] Department of Mathematics, University of Alabama at Birmingham
[2] Department of Mathematics, California Institute of Technology

Abstract
The study of complex dynamics is currently of great interest in mathematics and related scientific fields. Invariant laminations, invented by William Thurston, are abstract mathematical structures which represent the dynamics of complex polynomials but are easier to understand. We will define a sibling portrait of an invariant lamination which can be thought of as a "snap shot" of the dynamical system. We will show the existence of a one-to-one correspondence between sibling portraits and bicolored trees, mathematical objects from graph theory. This immediately provides a count for the number of different sibling portraits possible.

Introduction
Dynamics is the mathematical study of the evolution of systems which obey some fixed law. The system could be the population of rabbits, the motion of particles in a box, or the motion of planets in a solar system. The dynamical law is simply a formula which describes the system. Typically, the law does not depend on previous states of the system. For example, the population of rabbits tomorrow does not depend on the population of rabbits yesterday; it only depends on the population of rabbits today. However, the law to describe the growth of the population should remain the same. Thus to study the population, we study the answers. Using this method, we can predict long-term behavior of a system.

One simple yet interesting case to study is when the dynamical law is a quadratic polynomial of the form $f(x) = ax^2 + bx + c$. This formula has three parameters: a, b and c. The equation would be easier to analyze if there was only one parameter. By letting the variable be a complex number of the form $z = \alpha + i\beta$ where α and β are real numbers and $i = \sqrt{-1}$ (z typically denotes complex numbers while x denotes real numbers), we can re-write $f(x)$ in terms of one parameter. Note that $f(x) = a(x + \frac{b}{2a})^2 + k$ where $k = c - \frac{b^2}{4a}$. We will let $z^2 = a(x + \frac{b}{2a})^2$ which means $z = \sqrt{a}(x + \frac{b}{2a})$. If a is a negative number, then z is a complex number. It follows that $f(x) = f(z) = z^2 + k$ has only one parameter k. If we think in terms of iterating the formula, we may write the dynamical law as $z_{n+1} = z_n^2 + k$.

The study of quadratic polynomials still poses some challenges. For a given k-value and initial value z_0 we would like to know how the system evolves under many iterations. There are two possibilities. Either the values z_n grow arbitrarily large, in other words $z_n \to \infty$ as $n \to \infty$, or the values never exceed some large number acting as an upper bound, i.e. $z_n < M$ for all n where M is the upper bound. Sometimes, the accuracy in the initial value z_0 is extremely important. For example, let $k = 0$ so $z_{n+1} = z_n^2$. If we know $|z_0| = 1$ with some margin of error, we can have very different results depending on what the value actually is. If $|z_0| = 1$ then $|z_n| = 1$ for all . The magnitude of the numbers never grows nor shrinks. However, if $|z_0| = 0.999$ then $|z_n| \to 0$ as $n \to \infty$ since numbers less than 1 shrink when squared. If $|z_0| = 1.001$ then $|z_n| \to \infty$ as $n \to \infty$. These are drastically different behaviors. This initial value sensitivity is the basis for chaos theory and chaotic dynamics.

A popular choice for z_0 is $z_0 = 0$. Similar to the above example, there are some values of k for which $z_{n+1} = z_n^2 + k$ tends towards ∞ or remains bounded. All the values of k for which $z_{n+1} = z_n^2 + k$ remains bounded when $z_0 = 0$ form the Mandelbrot set. This set is essential in understanding the dynamics of quadratic polynomials as a whole. It is very difficult to understand, however, and there are still important open questions about the properties of this set. In the 1980s, William Thurston invented invariant laminations (described below) as a mathematical tool for understanding the Mandelbrot set [1]. Although more abstract, laminations are easier to study since they have nice topological and combinatorial properties. Just as the Mandelbrot set describes properties of quadratic polynomials, the "space of laminations" describes properties of invariant laminations. There are many similarities between the Mandelbrot set and the space of laminations which motivate this research. The invariant laminations defined by William Thurston, which we will refer to as Thurston invariant laminations, are very well studied in the quadratic polynomial case. However, the definition does not extend well to polynomials of higher dimensions, such as cubic polynomials of degree 3. Recently, Blokh, Mimbs, Oversteegen, and Valkenburg have modified Thurston's definition to one which is called sibling invariant laminations. This definition has demonstrated to be much more powerful in describing the cubic case [2].

This paper introduces sibling invariant laminations including definitions for sibling leaves and sibling portraits. We define central strips of laminations and prove a central strip exists for almost all full sibling families of sibling invariant leaves. We will define a bicolored tree, a mathematical object from graph theory, and show that there is a one-to-one correspondence between sibling portraits and bi-colored trees. This means that there are just as many bi-colored trees as sibling portraits. Thus we will provide counts for both the number of the different possible sibling portraits and consequently the number of different central strips.

Preliminaries
Let $\mathbb{C}$ denote the complex plane, $\mathbb{S}$ denote the unit circle in the complex plane – all the complex numbers z such that $|z| = 1$. The unit circle $\mathbb{S}$ will be labeled from 0 to 1 such that 0

corresponds to the point on the unit circle $\mathbb{S}$ at the $0°$ angle, $\frac{1}{4}$ corresponds to the point on $\mathbb{S}$ at $90°$, $\frac{1}{2}$ corresponds to the point on $\mathbb{S}$ at $180°$, $\frac{3}{4}$ corresponds to the point on $\mathbb{S}$ at $270°$, and 1 corresponds back to the point on $\mathbb{S}$ at $0°$. Fig. 1 depicts this re-labeling.

For $d \geq 2$ and x in $\mathbb{S}$, $0 \leq x \leq 1$, define a function $\sigma_d: \mathbb{S} \to \mathbb{S}$ by $\sigma_d(x) = d \times x \ (mod\ 1)$. For example, for $d = 2$, $\sigma_d\left(\frac{1}{7}\right) = \frac{2}{7}$, $\sigma_d\left(\frac{2}{7}\right) = \frac{4}{7}$, $\sigma_d\left(\frac{4}{7}\right) = \frac{8}{7} - 1 = \frac{1}{7}$. This example is depicted in Fig. 1.

A *lamination* $\mathcal{L}$ is a collection of chords of the unit circle $\mathbb{S}$ with the property that no two chords intersect except for possibly at their endpoint in the unit circle. A lamination $\mathcal{L}$ must also be closed, meaning it contains all of its limit points. If a sequence of points in the lamination converges to a limit, that limit point is also in the lamination. A chord in a lamination which satisfies these properties is called a *leaf* (pl. *leaves*).

If ℓ is a leaf in a lamination $\mathcal{L}$, we write $\ell = \overline{ab}$, where a and b are the endpoints of ℓ in $\mathbb{S}$. We let $\sigma_d(\ell)$ be the chord $\overline{\sigma_d(a)\sigma_d(b)}$. For instance, if $\ell_1 = \overline{\frac{1}{7}\frac{2}{7}}$, $\ell_2 = \overline{\frac{2}{7}\frac{4}{7}}$, and $\ell_3 = \overline{\frac{4}{7}\frac{1}{7}}$, then $\sigma_2(\ell_1) = \overline{\frac{2}{7}\frac{4}{7}}$, $\sigma_2(\ell_2) = \overline{\frac{4}{7}\frac{1}{7}}$, and $\sigma_2(\ell_3) = \overline{\frac{1}{7}\frac{2}{7}}$.

Fig. 1 depicts these three leaves. If it happens that $\sigma_d(a) = \sigma_d(b)$, for example $\sigma_2(0) = \sigma_2\left(\frac{1}{2}\right) = 0$, then $\sigma_2(\ell)$ is a point, called a *critical value* of $\mathcal{L}$ and we say ℓ is a *critical leaf*. In general, as σ_d is applied to ℓ successively, we get iterations of ℓ. The i^{th} iterate of ℓ is denoted $\sigma_d^i(\ell)$. For example, $\sigma_d^2(\ell) = \sigma_d(\sigma_d(\ell))$.

Leaf Length Function

Let (a, b) be an arc segment in the unit circle $\mathbb{S}$. Define the length of (a, b), denoted $|(a, b)|$, to be the length of the shortest path from a to b. For example, $\left|\left(\frac{1}{7}, \frac{4}{7}\right)\right| = \frac{3}{7}$ while $\left|\left(\frac{11}{14}, \frac{1}{14}\right)\right| = \frac{4}{14} = \frac{2}{7}$. For a leaf $\ell = \overline{ab}$, let the length of the leaf, denoted $|\ell|$ or $|\overline{ab}|$, be the length of the arc (a, b). So $|\ell| = |(a, b)|$. Notice that the maximum length of a leaf is $\frac{1}{2}$ since we choose the shortest distance between a and b to be the length. For a fixed $d \geq 2$, a leaf ℓ of length $|\ell| < \frac{1}{2d}$ will grow since $|\sigma_d(\ell)| < d \times \frac{1}{2d} = \frac{1}{2}$. Also note that if two points a and b map to the same point under σ_d, then they are a distance $\frac{n}{d}$ apart for some integer n. Say $b = a + \frac{n}{d}$, then $\sigma_d(b) = \sigma_d\left(a + \frac{n}{d}\right) = d \times \left(a + \frac{n}{d}\right) (mod\ 1) = d \times a + n\ (mod\ 1) = d \times a\ (mod\ 1) = \sigma_d(a)$. So all the points which map to the same point are spaced evenly around the unit circle at $\frac{1}{d}$ intervals.

Sibling Invariant Laminations

***Definition 1.1** (Sibling Leaves)*

Let ℓ and ℓ_1 be leaves in a lamination $\mathcal{L}$ such that $\sigma_d(\ell_1) = \ell$. A leaf ℓ_2 is *disjoint* from ℓ_1 if they do not intersect at all, even at endpoints. If ℓ_2 is disjoint from ℓ_1 and $\sigma_d(\ell_2) = \sigma_d(\ell_1) = \ell$ then ℓ_2 is called a sibling leaf of ℓ_1. A collection of d pair-wise disjoint leaves $\mathcal{S} = \{\ell_1, \ell_2, \ldots, \ell_d\}$ is called a *full sibling family* if for each i, $\sigma_d(\ell_i) = \ell$.

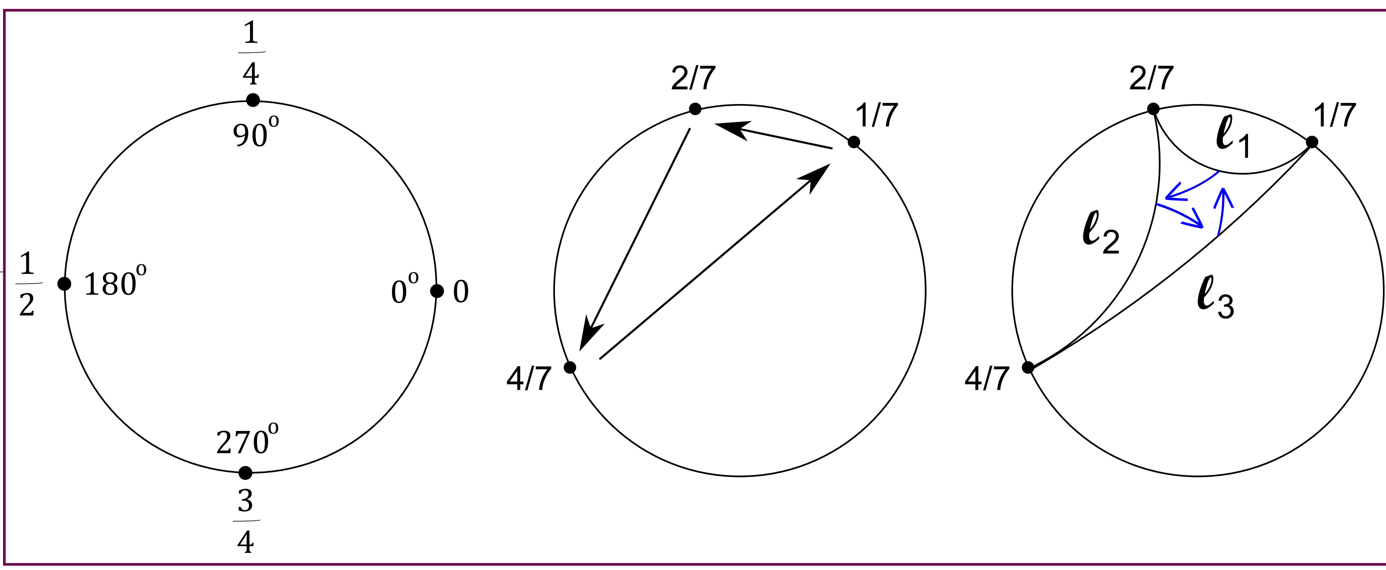

Figure 1. The unit circle re-labeled. Middle: The σ_2 map acting on points in the unit circle. Right: A triplet of leaves mapping to itself by σ_2.

For example, in Fig. 2, the leaves $\ell_1 = \overline{\frac{1}{7}\frac{2}{7}}$ and $\ell_2 = \overline{\frac{9}{14}\frac{11}{14}}$ are sibling leaves. They are disjoint and map to the same leaf. $\sigma_2(\ell_1) = \overline{\frac{2}{7}\frac{4}{7}}$ and $\sigma_2(\ell_2) = \overline{\frac{18}{14}\frac{22}{14}} = \overline{\frac{9}{7}\frac{11}{7}} = \overline{\frac{2}{7}\frac{4}{7}}$. Since $d = 2$, $\{\ell_1, \ell_2\}$ is a full sibling family. Notice that there are always d sibling leaves and their endpoints will be spaced evenly around the unit circle $\mathbb{S}$. In this case where $d = 2$, there are two sibling leaves and they are on opposite sides of $\mathbb{S}$.

Thurston's definition of invariant laminations does not involve sibling leaves. His definition is more broad. Blokh, Mimbs, Oversteegen, and Valkenburg proved that every sibling invariant lamination is also Thurston invariant but not vice-versa [2]. They also proved that sibling invariance does not exclude any of the important laminations. In a sense, sibling invariance keeps the appropriate Thurston invariant laminations and discards the unnecessary ones.

Definition 1.2 (Sibling Invariance)

Recall that chords (called leaves) are in a lamination $\mathcal{L}$ provided that they do not intersect any other leaves except possibly at endpoints. A lamination $\mathcal{L}$ is said to be *sibling d-invariant* provided that:

(1) (Forward Invariance) For every leaf ℓ in $\mathcal{L}$, $\sigma_d(\ell)$ is a leaf in $\mathcal{L}$.
(2) (Backward Invariance) For every leaf ℓ in $\mathcal{L}$, there is a leaf ℓ_1 in $\mathcal{L}$ such that $\sigma_d(\ell_1) = \ell$.
(3) (Sibling Invariance) For every ℓ_1 in $\mathcal{L}$ with $\sigma_d(\ell_1) = \ell$, there is a full sibling family of leaves in $\mathcal{L}$, $\{\ell_1, \ell_2, \ldots, \ell_d\}$, such that $\sigma_d(\ell_i) = \ell$ for all i.

Fig. 2 depicts an example of the first few leaves in a degree 2 sibling invariant lamination. Though chords are straight, we sometimes draw them curved to stand out. The construction begins with the leaves $\ell_1 = \overline{\frac{1}{7}\frac{2}{7}}$, $\ell_2 = \overline{\frac{2}{7}\frac{4}{7}}$, and $\ell_3 = \overline{\frac{4}{7}\frac{1}{7}}$. See that $\sigma_2(\ell_1) = \ell_2$, $\sigma_2(\ell_2) = \ell_3$, and $\sigma_2(\ell_3) = \ell_1$. The leaves ℓ_1, ℓ_2, and ℓ_3 each need a sibling (only one since $d = 2$). The additional leaves $\widetilde{\ell_1}$, $\widetilde{\ell_2}$, and $\widetilde{\ell_3}$ such that $\sigma_2(\widetilde{\ell_1}) = \ell_2$, $\sigma_2(\widetilde{\ell_2}) = \ell_3$, and $\sigma_2(\widetilde{\ell_3}) = \ell_1$ are $\widetilde{\ell_1} = \overline{\frac{9}{14}\frac{11}{14}}$, $\widetilde{\ell_2} = \overline{\frac{11}{14}\frac{1}{14}}$, and $\widetilde{\ell_3} = \overline{\frac{1}{14}\frac{9}{14}}$. Continue to construct the lamination by finding 2 disjoint leaves which map to $\widetilde{\ell_1}$, $\widetilde{\ell_2}$, and $\widetilde{\ell_3}$. Then find 2 disjoint leaves which map to each of those and so forth. Note that this construction method may not produce a unique sibling invariant lamination.

Definition 1.3 (Sibling Portrait)

The *sibling portrait* $\mathcal{S}$ of a full sibling family is the collection of regions complementary to the sibling leaves - the spaces between the leaves. We call a complementary region a C-*region* provided all of the arcs in which the region meets the circle are *short* (length $< \frac{1}{2d}$). It will soon be apparent that the name C-region alludes to the central strip. We call a complementary region an R-*region* if all of the arcs are *long* (length $> \frac{1}{2d}$). The R simply alludes to a not-so-special region. The *degree* of a complementary region T (T is a dummy variable for either C or R), denoted $\deg(T)$, is equal to the number of leaves in the boundary of T or, equivalently, the number of circular arcs in the boundary of T.

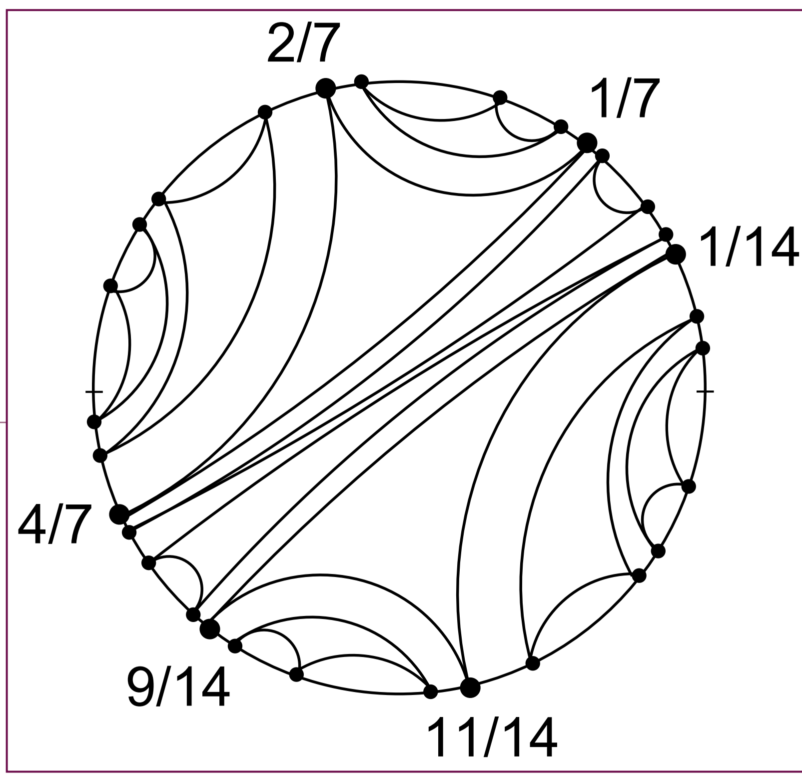

Figure 2. An example of the first few leaves in a degree d=2 *sibling invariant lamination.*

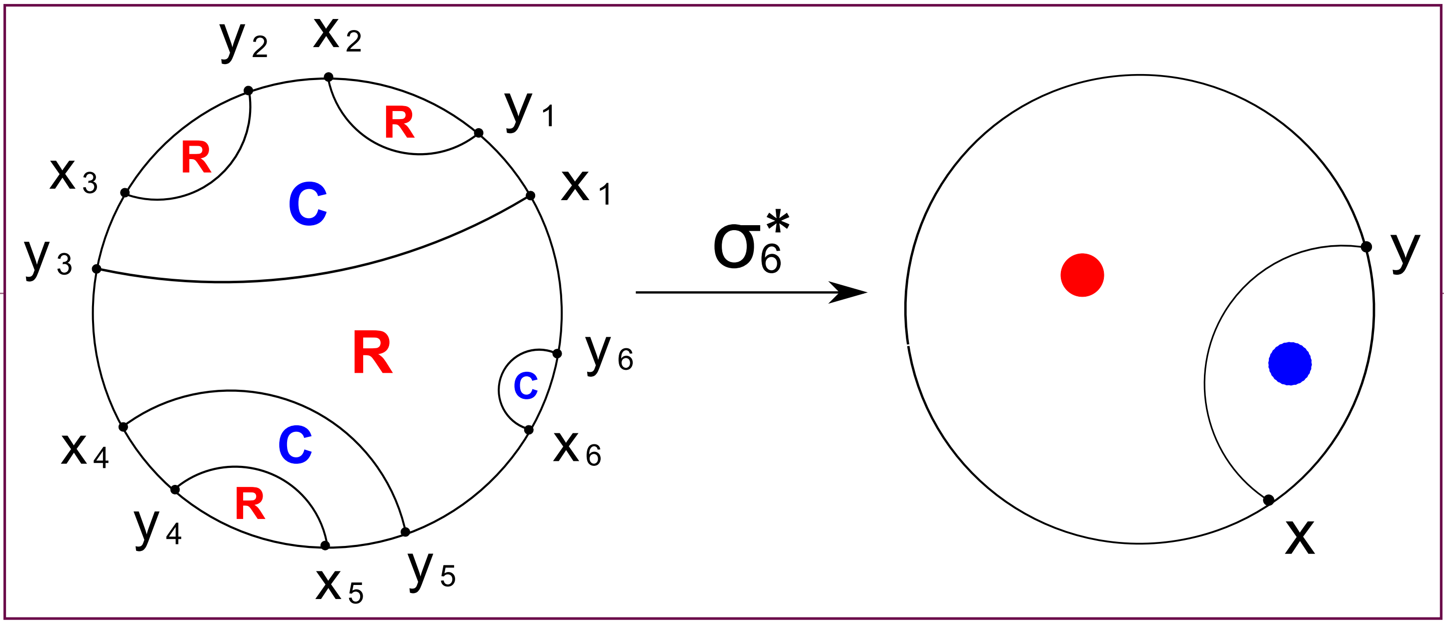

Figure 3. An example of a sibling portrait for degree d=6 *on the left and its image under* σ_6 *on the right.*

Fig. 3 depicts an example a full sibling family of size 6 and the corresponding sibling portrait. All the leaves are mapped under σ_6 to the same leaf depicted on the right. Note that since the C-regions intersect the circle with short arcs and R-regions intersect the circle with long arcs, the C-regions are mapped to the smaller of the two regions depicted on the right.

Definition 1.4 (Graph)
A *graph* is defined to be a finite number of points connected by lines which do not intersect. The points are called *vertices* and the connecting lines are called *edges* [3]. The *degree* of a vertex v is the number of edges that share v as a vertex. A *tree* is a graph with no loops of edges in it so that there is only one path connecting any two vertices [3]. A tree whose vertices are colored with two colors such that no edge connects vertices of the same color is said to be *bicolored*. An example of a bicolored tree is provided in Fig. 4.

Definition 1.5 (Dual Graph)
The *dual graph* $T_{\mathcal{S}}$ of the sibling portrait $\mathcal{S}$ corresponding to a full sibling family is defined as follows: put a dot in each region of the sibling portrait and connect any twos dots exactly when the two regions containing those dots share a leaf in their boundaries. See Fig. 4 for an example.

Proposition 1.6
The dual graph of a sibling portrait under σ_d is a tree consisting of $d+1$ vertices (components of the portrait) and d edges (sibling leaves between components that meet on their boundaries).

The proof is left to the reader. However, by referring to Fig. 4 it is easy to see that there are d edges in the tree since there are d leaves. One can also see that there are $d+1$ vertices of the tree since d leaves partition the circle into $d+1$ regions. Since the leaves in the lamination are required to be disjoint, the graph contains no loops of edges. Therefore the dual graph is indeed a tree. The bicolored property is proven in Theorem 2.2.

Recall that the degree of a complementary region T of a sibling portrait is defined to be the number of leaves in its boundary, and the degree of a vertex in a graph is defined to be the number of edges coming out of it. Since leaves correspond to edges, the degree of a region T is equal to the degree of T's corresponding vertex in the dual graph.

In Theorem 2.2 below, we show that if a full sibling family does not map to a diameter, then all the complementary regions of the sibling portrait are either C-regions or R-regions. If the degree of a region is 1, we will refer to it as a *terminal region* since it corresponds to an endpoint of the dual tree.

Central Strips

During his study of quadratic laminations, Thurston used the idea of a "central strip" to classify types of laminations. This classification helps give structure to the space of laminations which we are trying to correlate to the parameter space for polynomials (the Mandelbrot set in the quadratic case). Since Thurston only worked with quadratic laminations, a rigorous definition for a central strip was not necessary. Below we will define a central strip for any degree $d \geq 2$.

If $\mathcal{S}$ is a full sibling family which maps to the leaf $\ell = \overline{xy}$, then the endpoints x_i, y_i of the sibling leaves may be labeled $x_1, y_1, x_2, y_2, \ldots, x_d, y_d$ in a counterclockwise order around the unit circle $\mathbb{S}$ (see Fig. 4). We do not generally suppose $\ell_i = \overline{x_i y_i}$ unless so stated. Typically we denote $\ell_i = \overline{x_i y_j}$ for some j. We can consider the case where the full sibling family maps to a leaf which is a diameter of the circle as a special case. Therefore, we will assume the sibling leaves map to a non-diametrical leaf.

Definition 2.1 (Central Strip)

Consider the sibling portrait of a full sibling family $\mathcal{S}$. Then the *central strip* C corresponding to $\mathcal{S}$ is the collection of all C-regions C_i with degree at least 2. So in Fig. 4, the central strip would consist of two of the three C-regions.

Theorem 2.2

Let $\mathcal{S} = \{\ell_1, \ell_2, \ldots, \ell_d\}$ be a full sibling family which does not map to a diameter. Then the following hold:

(1) If some leaf in $\mathcal{S}$ is of length $> \frac{1}{2d}$, then there is a central strip C.

(2) The dual graph of the sibling portrait corresponding to $\mathcal{S}$ is a bicolored tree where C-regions are colored one color and R-regions the other.

Proof: Refer to Fig. 4. Since the x_i (and y_i) are evenly spaced, then either the arc (x_i, y_i) or (y_i, x_{i+1}) is short. We may assume that (x_i, y_i) is the short one. Notice that on one side of $\ell = \overline{x_1 y_3}$ in Fig. 4 there is a C-region and on the other an R-region. This is true in general since if the length of arc (x_i, y_i) is short then the length of (y_i, x_{i+1}) must be long. Therefore, C-regions are only adjacent to R-regions and vice-versa. Since the regions correspond to vertices of the dual graph, we naturally see that the vertices must be bicolored in relation to C-regions and R-regions. Proposition 1.6 already established the dual graph is a tree; therefore the dual graph is a bicolored tree.

First, note that if the length of ℓ_1 is $\frac{1}{2d}$ then the length of $\sigma_d(\ell_1)$ is $\frac{1}{2}$ which implies $\sigma_d(\ell_1)$ is a diameter. Since we assumed otherwise, none of the leaves in $\mathcal{S}$ may be of length $\frac{1}{2d}$. Second, if all the leaves in $\mathcal{S}$ are short (length $< \frac{1}{2d}$) then $\ell_i = \overline{x_i y_i}$ for all i. Then each ℓ_i subtends a C-region of degree one. The remaining region in the middle is an R-region of degree d. Since in this case there are no C-regions of degree greater than or equal to two, there is no central strip. Fig. 5 depicts this case.

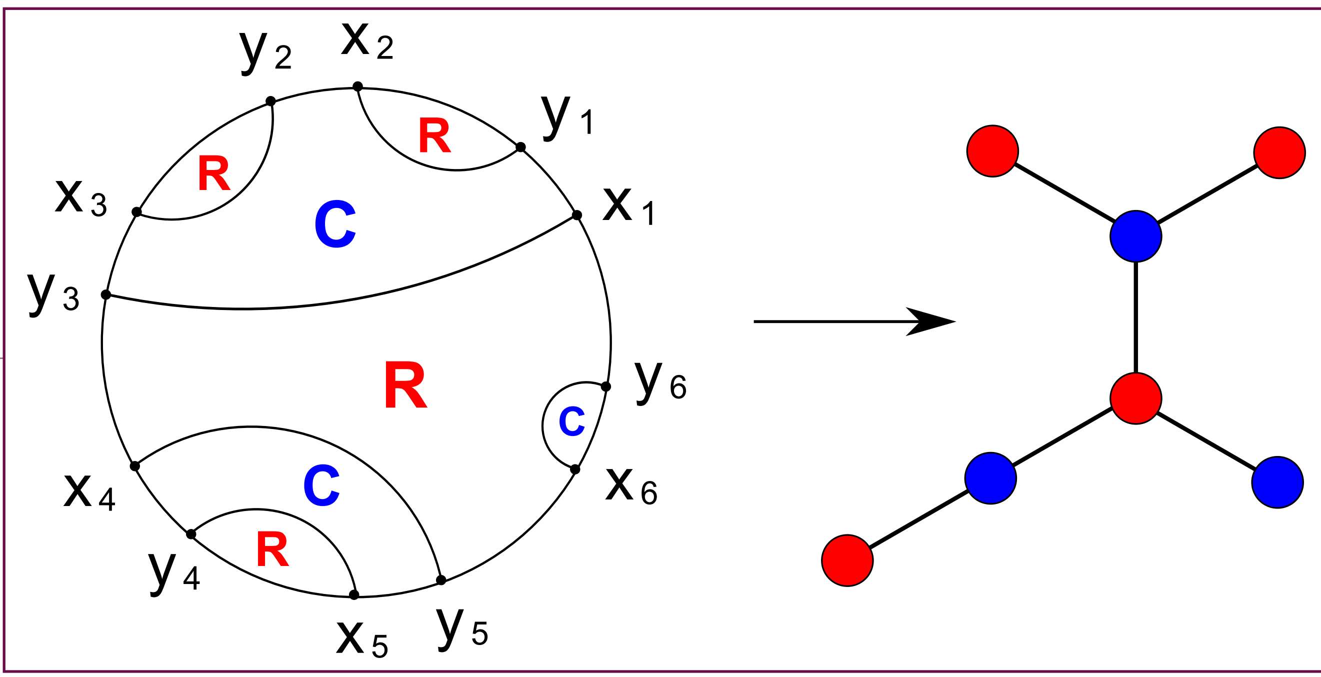

Figure 4. Mapping of a sibling portrait to a bicolored tree.

In the third and last case we may assume the length of at least one leaf in $\mathcal{S}$ is long (length $> \frac{1}{2d}$). Without loss of generality, we may assume the length of ℓ_1 is long. We may also label the generic leaf ℓ_i as the one with endpoint x_i. Since $\ell_1 = \overline{x_1 y_j}$ (in Fig. 4, $\ell_1 = \overline{x_1 y_3}$) is assumed to be long and (x_i, y_i) is assumed to be short, then $j \neq 1$. Thus there is a leaf emanating from y_1, called $\ell_k = \overline{y_1 x_k}$, different from ℓ_1. Then the arc (x_k, y_k) is short and there is another leaf emanating from y_k, $\ell_m = \overline{y_k x_m}$. This continues until a leaf has endpoint x_j (x_3 in Fig. 4). This leaf completes the C-region which has ℓ_1 and ℓ_k as boundary leaves. Since $\ell_1 \neq \ell_k$, this C-region has degree at least two. Therefore, a central strip exists. (End of proof.)

Definition 2.3 (Euler's totient function)
Integers x and y are *relatively prime* if their greatest common divisor is 1. *Euler's totient function*, $\phi(x)$, is defined to be the number of positive integers less than, and relatively prime to, x. For example, $\phi(2) = 1$ since 1 is relatively prime to 2. $\phi(3) = 2$ since both 1 and 2 are relatively prime to 3. $\phi(4) = 2$ since both 1 and 3 are relatively prime to 4. 2 is not relatively prime to 4 since 2 is a divisor of both.

Theorem 2.4
If ℓ is a non-diametrical leaf then there are $N(d)$ different full sibling families which map onto ℓ by σ_d, distinct up to rotations, where

$$N(d) = \frac{1}{d}\left(\frac{1}{d+1}\binom{2d}{d} + \sum_{n|d, n<d} \phi\left(\frac{d}{n}\right)\binom{2n}{n}\right)$$

and $\phi(x)$ is Euler's totient function.

Proof: The goal is to show that there are just as many different full sibling families (equivalently sibling portraits) which map to the same leaf under σ_d as there are different bicolored trees with d edges. The number of bicolored trees with d edges is known to be $N(d)$ [4]. Thus we will use this correspondence to show there are $N(d)$ different sibling portraits mapping to the same leaf. Refer to Fig. 6 during this proof.

The proof of Theorem 2.2 illustrates how to map a sibling portrait to a bicolored tree. It is easy to check that if two sibling portraits map to the same bicolored tree (up to rotation) then those two sibling portraits are the same (up to rotation). Therefore, since for every sibling portrait we can find a unique bicolored tree, there must be at least as many bicolored trees as sibling portraits.

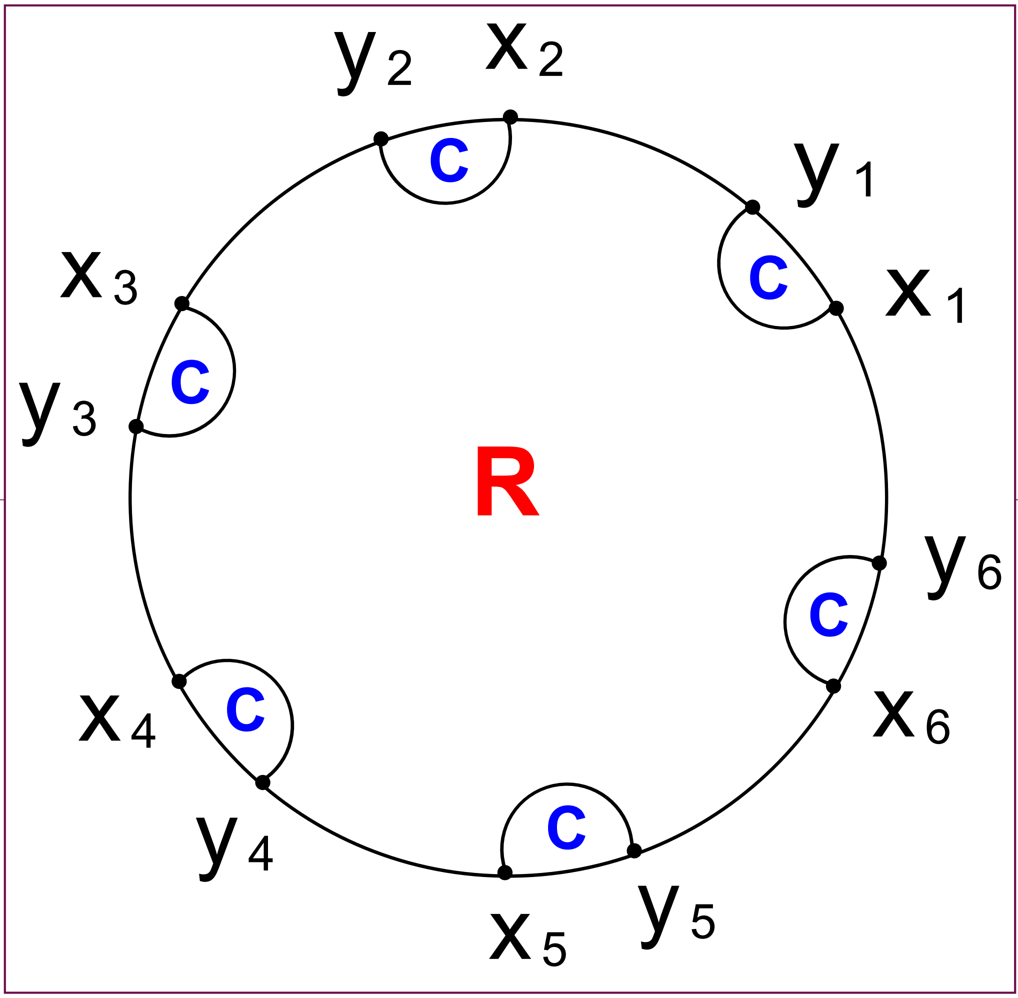

Figure 5. There is no central strip when all the leaves are short.

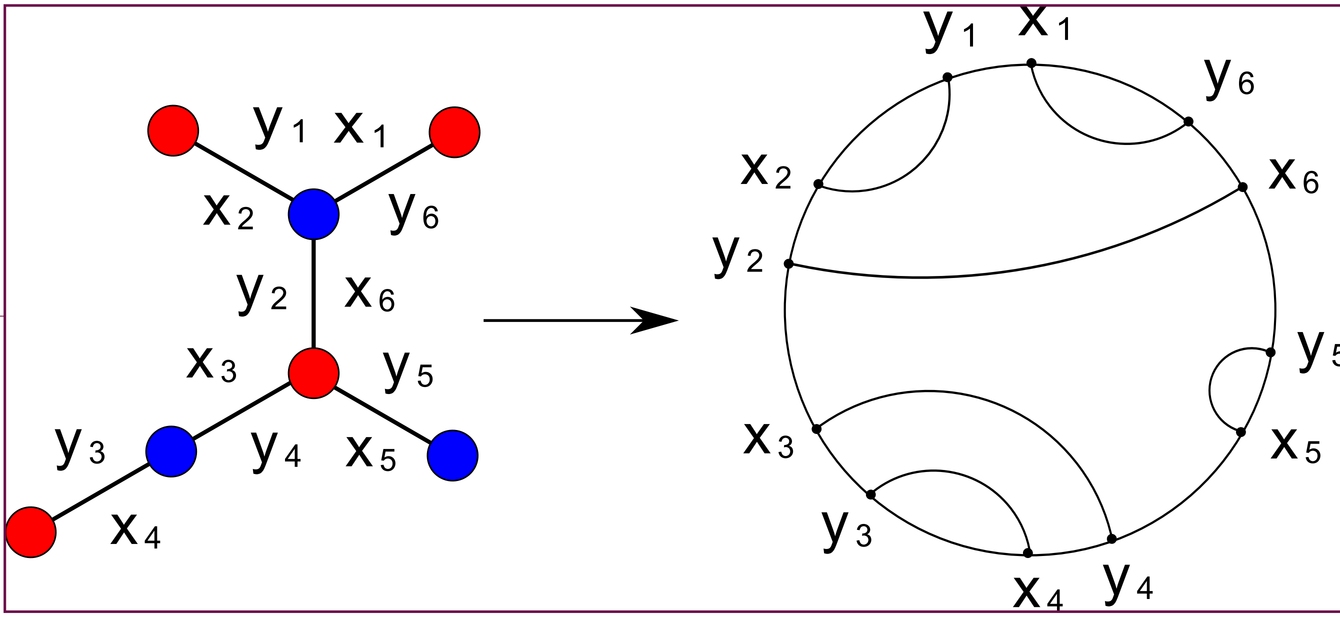

Figure 6. Mapping a bicolored tree to a sibling portrait.

Now assume we are given a bicolored tree with d edges like the one in Fig. 6. Since each edge corresponds to a leaf in the full sibling family and each leaf has two endpoints, we may correlate each side of an edge with an endpoint of the leaf. Label the sides of the edges in the tree in a counterclockwise order $x_1, y_1, x_2, y_2, \dots, x_d, y_d$. Since we consider sibling portraits to be the same if one is a rotation of the other, then it does not matter which edge you choose to label x_1 with. However, since we generally assume (x_i, y_i) to be a short arc, the vertex of the tree between x_i and y_j must correspond to a C-region. Then in the unit circle, connect a leaf between x_i and y_j if they are two sides of the same edge in the tree. This will construct a full sibling family. Similarly, it is easy to check that if two bicolored trees map to the same sibling portrait then they are the same bicolored tree. Therefore for every bicolored tree there is a unique sibling portrait. It then follows that there are exactly as many sibling portraits as bicolored trees. (End of proof.)

Corollary 2.5

If ℓ is a non-diametrical leaf then the number of different central strips, distinct up to rotational symmetry, whose boundary leaves map onto ℓ is $N(d) - 1$.

This corollary follows immediately from Theorem 2.4 since the only time a sibling portrait does not have a central strip is when all the boundary leaves are short, as in Fig. 5.

References

1. W.P. Thurston. On the geometry and dynamics of iterated rational maps. *Complex Dynamics: Families and Friends.* Dierk Schleicher, ed. A K Peters, 2009, pp. 1-108.

2. A. Blokh, D. Mimbs, L. Oversteegen and K. Valkenburg. Laminations in the language of leaves. *Transactions AMS* (to appear). 32 pages. arXiv: 1106.0273.

3. J.R. Munkres. *Topology, 2nd ed.* Pearson Education, Inc. 2000.

4. M. Bona, M. Bousquet, G. Labelle, and P. Leroux. Enumeration of mAry Cacti. *Advances in Applied Mathematics* 24, 2009, 22-56.

Authors' e-mail addresses

John C. Mayer: jcmayer@uab.edu
Luka Mernik: lmernik@caltech.edu
Joseph W. Olson: jwolson@uab.edu

The first, second, and fourth authors thank the UAB Laminations seminar, and in particular Drs. Lex Oversteegen and Alexander Blokh and graduate students Ross Ptacek and David Cosper for useful conversations.