Central Strips of Sibling Leaves in Laminations of the Unit Disk

by

David J. Cosper, Jeffrey K. Houghton, John C. Mayer,
Luka Mernik, and Joseph W. Olson

Electronically published on April 17, 2015

CENTRAL STRIPS OF SIBLING LEAVES IN LAMINATIONS OF THE UNIT DISK

DAVID J. COSPER, JEFFREY K. HOUGHTON, JOHN C. MAYER,
LUKA MERNIK, AND JOSEPH W. OLSON

ABSTRACT. Quadratic laminations of the unit disk were introduced by William P. Thurston as a vehicle for understanding the (connected) Julia sets of quadratic polynomials and the parameter space of quadratic polynomials. The Central Strip Lemma plays a key role in Thurston's classification of gaps in quadratic laminations and in describing the corresponding parameter space. We generalize the notion of "Central Strip" to laminations of all degrees $d \geq 2$ and prove a Central Strip Lemma for degree $d \geq 2$. We conclude with applications of the Central Strip Lemma to identity return polygons that show for higher degree laminations it may play a role similar to Thurston's lemma.

1. INTRODUCTION

Quadratic laminations of the unit disk were introduced by William P. Thurston as a vehicle for understanding the (connected) Julia sets of quadratic polynomials and the parameter space of quadratic polynomials. The Central Strip Lemma plays a key role in Thurston's classification of gaps in quadratic laminations [12]. It is used to show that there are no wandering polygons for the angle-doubling map σ_2 on the unit circle. Moreover, when a polygon returns to itself, the iteration of σ_2 is transitive on the vertices. For σ_2 it is sufficient to prove these facts for triangles: there are no wandering triangles, and no identity return triangles. From these facts, the classification of types of gaps of a quadratic lamination, and a parameter space for quadratic laminations, follows. Thurston posed

---

a question in his notes on laminations that he deemed important to further progress in the field: "Can there be wandering triangles for σ_3 (and higher degree)?" When Alexander Blokh and Lex Oversteegen [3] [4] showed that the answer was "yes," the need to define and understand central strips for higher degree and their role in "controlling" wandering and identity return polygons of a lamination became imperative. Contributions to this understanding were made by Lisa R. Goldberg and John Milnor [7], Milnor [10] [9], Jan Kiwi [8], Blokh and G. Levin [1], Douglas K. Childers [6], and others.

New results in this paper include definitions of a sibling portrait (Definition 1.9) and of a central strip (Definition 2.2), the statement and proof of the Central Strip Structure (Theorem 2.3), the statement and proof of the generalized Central Strip Lemma (Theorem 2.9), and counting the number of different sibling portraits that could correspond to the preimage of a given non-degenerate leaf (Theorem 3.1). The Central Strip Lemma can be used to provide new proofs of known results that the authors believe to be more transparent, yield more information about the laminations, and yield new results for laminations of degree $d \geq 3$. Initial applications to identity return triangles under σ_3 appear in this paper (§4); other applications, particularly to higher degree, will appear in subsequent papers.

1.1. PRELIMINARIES.

Let $\mathbb{C}$ denote the complex plane, $\mathbb{D} \subset \mathbb{C}$ the open unit disk, $\overline{\mathbb{D}}$ the closed unit disk, and $\mathbb{S}$ the boundary of the unit disk (i.e., the unit circle), parameterized as $\mathbb{R}/\mathbb{Z}$. For $d \geq 2$, define a map $\sigma_d : \mathbb{S} \to \mathbb{S}$ by $\sigma_d(t) = dt$ (mod 1).

Definition 1.1. A *lamination* $\mathcal{L}$ is a collection of chords of $\overline{\mathbb{D}}$, which we call *leaves*, with the property that any two leaves meet, if at all, in a point of $\mathbb{S}$, and such that $\mathcal{L}$ has the property that

$$\mathcal{L}^* := \mathbb{S} \cup \{\cup \mathcal{L}\}$$

is a closed subset of $\overline{\mathbb{D}}$.

It follows that $\mathcal{L}^*$ is a continuum (compact, connected metric space). We allow *degenerate* leaves – all points of $\mathbb{S}$ are degenerate leaves. If $\ell \in \mathcal{L}$ is a leaf, 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 $\overline{\sigma_d(a)\sigma_d(b)}$. If it happens that $\sigma_d(a) = \sigma_d(b)$, then $\sigma_d(\ell)$ is a point, called a *critical value* of $\mathcal{L}$ and we say ℓ is a *critical leaf*.

Proposition 1.2. *Let σ_d^* denote the linear extension of σ_d to leaves of $\mathcal{L}$, so that σ_d^* is defined on $\mathcal{L}^*$. Then σ_d^* is continuous on $\mathcal{L}^*$.*

The proof is left to the reader.

1.2. Leaf length function.

Fix the counterclockwise order $<$ on $\mathbb{S}$ as the preferred (circular) order. Let $|(a,b)|$ denote the length in the parameterization of $\mathbb{S}$ of the arc in $\mathbb{S}$ from a to b counterclockwise. Given a chord $\overline{ab}$, there are two arcs of $\mathbb{S}$ subtended by $\overline{ab}$. Define the *length* of $\overline{ab}$, denoted $|\overline{ab}|$, to be the shorter of $|(a,b)|$ or $|(b,a)|$. The maximum length of a leaf is thus $\frac{1}{2}$. Note that the length of a critical leaf is $\frac{i}{d}$ for some $i \in \{1, 2, \ldots j \mid j = \lfloor \frac{d}{2} \rfloor\}$. See Figure 1.
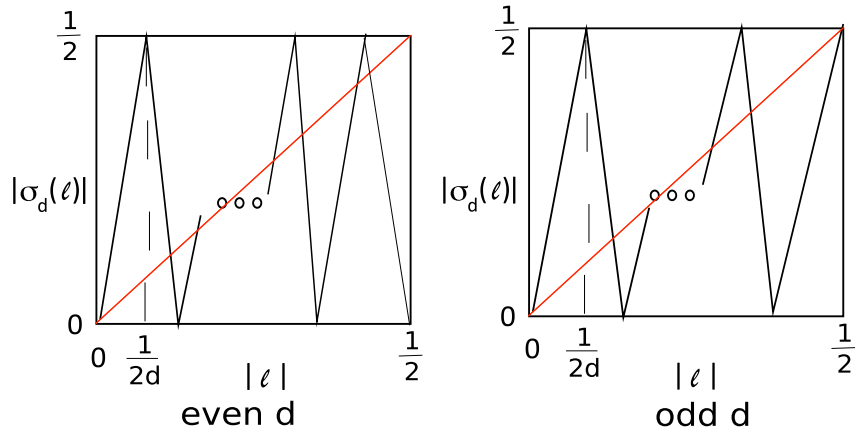

FIGURE 1. Graph of the leaf length function.

Lemma 1.3. *Let $x = |\ell|$ where $\ell \in \mathcal{L}$. Then the length $\tau_d(x)$ of $\sigma_d(\ell)$ is given by the function*

$$\tau_d(x) = \begin{cases} dx \text{ if } 0 \leq x \leq \frac{1}{2d} \\ 1 - dx \text{ if } \frac{1}{2d} \leq x \leq \frac{1}{d} \\ dx - 1 \text{ if } \frac{1}{d} \leq x \leq \frac{3}{2d} \\ 2 - dx \text{ if } \frac{3}{2d} \leq x \leq \frac{2}{d} \\ \vdots \\ (-1)^d dx + (-1)^{d+1} \lfloor \frac{d-1}{2} \rfloor \text{ if } \frac{d-2}{2d} \leq x \leq \frac{d-1}{2d} \\ (-1)^{d+1} dx + (-1)^{d+2} \lfloor \frac{d}{2} \rfloor \text{ if } \frac{d-1}{2d} \leq x \leq \frac{1}{2} \end{cases}$$

defined on the interval $[0, 1/2]$.

The proof is left to the reader.

Proposition 1.4. *The fixed points of τ_d are of the form*

$$\{0, \frac{1}{d+1}, \frac{1}{d-1}, \frac{2}{d+1}, \frac{2}{d-1}, \dots, \frac{j}{d+1}, \frac{j}{d-1}, \frac{j+1}{d+1} \dots, \leq \frac{1}{2}\}.$$

Thus, if $\frac{j}{d+1} < |\ell| < \frac{j}{d-1}$ for some j, then $|\sigma_d(\ell)| < |\ell|$, and if $\frac{j}{d-1} < |\ell| < \frac{j+1}{d+1}$ for some j, then $|\sigma_d(\ell)| > |\ell|$.

The proof is left to the reader.

Lemma 1.5. *A leaf ℓ of length $|\ell| < \frac{1}{d+1}$ will keep increasing in length under iteration of σ_d until the length of some iterate $|\sigma_d^i(\ell)| \geq \frac{1}{d+1}$.*

Proof. Consider τ_d and the identity function. Each monotone interval of the map τ_d and the identity function are linear. The first interval of τ_d, defined on $[0, 1/2d]$, has the equation $\tau_d(x) = dx$; therefore, $(0, 0)$ is the only intersection on the first interval. On the second interval, $[1/2d, 1/d]$, $\tau_d(x) = 1 - dx$. Therefore, the intersection is $1 - dx = x$ or $x = \frac{1}{d+1}$. Since the graph of τ_d is above the identity function on interval $(0, \frac{1}{d+1})$, the length will keep increasing under iteration until it is at least $\frac{1}{d+1}$. $\square$

1.3. Sibling invariant laminations.

Thurston's definition [12] of invariant laminations did not involve sibling leaves. Blokh, et al. showed in [2] that each sibling invariant lamination (defined below) is a Thurston lamination, and each lamination induced by a locally connected Julia set is a sibling invariant lamination. They showed that to understand Julia sets via laminations, it is sufficient to consider sibling laminations. (More precisely, they showed that the closure of the space of quadratic sibling laminations in the Hausdorff metric contains all laminations induced by locally connected Julia sets.)

Definition 1.6 (Sibling Leaves)**.** Let $\ell_1 \in \mathcal{L}$ be a leaf and suppose $\sigma_d(\ell_1) = \ell'$ for some non-degenerate leaf $\ell' \in \mathcal{L}$. A leaf $\ell_2 \in \mathcal{L}$, disjoint from ℓ_1, is called a *sibling* of ℓ_1 provided $\sigma_d(\ell_2) = \ell' = \sigma_d(\ell_1)$. A collection $\mathcal{S} = \{\ell_1, \ell_2, \dots, \ell_d\} \subset \mathcal{L}$ is called a *full sibling collection* provided that for each i, $\sigma_d(\ell_i) = \ell'$, and for all $i \neq j$, $\ell_i \cap \ell_j = \emptyset$.

Definition 1.7. A lamination $\mathcal{L}$ is said to be *sibling d-invariant* (or simply *invariant* if no confusion will result) provided that

    (1) (*Forward Invariant*) for every $\ell \in \mathcal{L}$, $\sigma_d(\ell) \in \mathcal{L}$;

    (2) (*Backward Invariant*) for every non-degenerate $\ell' \in \mathcal{L}$, there is a leaf $\ell \in \mathcal{L}$ such that $\sigma_d(\ell) = \ell'$;

    (3) (*Sibling Invariant*) for every $\ell_1 \in \mathcal{L}$ with $\sigma_d(\ell_1) = \ell'$, a non-degenerate leaf, there is a full sibling collection $\{\ell_1, \ell_2, \dots, \ell_d\} \subset \mathcal{L}$ such that $\sigma_d(\ell_i) = \ell'$.

Definition 1.8. A *gap* in a lamination $\mathcal{L}$ is the closure of a component of $\overline{\mathbb{D}} \setminus \mathcal{L}^*$. A gap is *critical* if and only if two points in its boundary map to the same point. A finite gap is usually called a *polygon*. The leaves bounding a finite gap are called the *sides* of the polygon. A polygon is called *all-critical* if every side is a critical leaf.

Definition 1.9 (Sibling Portrait)**.** The *sibling portrait S* of a full collection of sibling leaves is the collection of regions complementary to the sibling leaves. We call a complementary region a *C-region* provided all of the arcs in which the closure of the region meets the circle are *short* (length $< \frac{1}{2d}$) and call it an *R-region* if all of the arcs are *long* (length $> \frac{1}{2d}$). The *degree* of a complementary region T, denoted $\deg(T)$, is the number of leaves in the boundary of T or, equivalently, the number of circular arcs in the boundary of T.

C-regions which meet the circle in more than one short arc will constitute the components of the central strip, see Definition 2.2. We show in Theorem 2.3, subject to the condition that no sibling maps to a diameter, that each region is either a C-region or an R-region.

Topologically, a *graph* is a finite union of arcs (homeomorphic images of the interval $[0, 1]$) meeting only at endpoints. Endpoints of these arcs are called *vertices* and the arcs themselves are called *edges*. The *degree* of a vertex v is the number of edges that share v as an endpoint. A *tree* is a graph with no closed loops of edges in it.

Definition 1.10. The *dual graph T_S* of the sibling portrait S of a full collection $\mathcal{S}$ of sibling leaves is defined as follows: Let each complementary region correspond to a vertex of T_S and let each sibling leaf on the boundary of two regions correspond to an edge of T_S between the vertices corresponding to the two regions.

Proposition 1.11. *The dual graph of a sibling portrait under σ_d is a connected 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. See Figure 5 (page 76) for an example.

Proposition 1.12. *Let S be the sibling portrait of a full collection $\mathcal{S}$ of sibling leaves under σ_d. Let T denote a complementary region of S and let T' denote the corresponding vertex of the dual graph T_S. Then $deg(T) = deg(T')$.*

The proof is left to the reader.

In Theorem 2.3 below, we show that if the image leaf of a full sibling collection is not a diameter, then each of the complementary regions of

the sibling portrait is either a C-region or an R-region of degree $d \geq 2$, or a *terminal* C- or R-region of degree 1. (By "terminal" region we mean a region corresponding to an endpoint of the dual graph.) Examples of sibling portraits for degrees $d = 2$ and $d = 3$ are in figures 2 and 3.
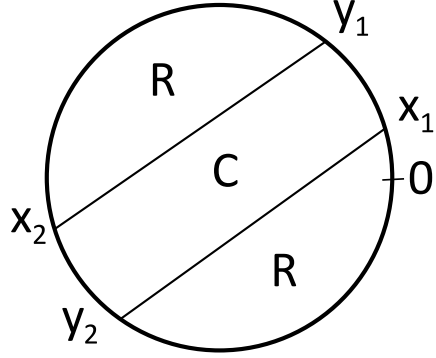

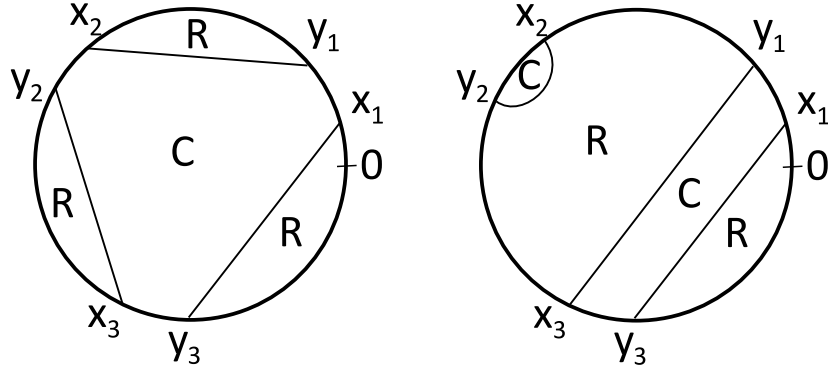FIGURE 2. Example of a sibling portrait with a central strip for σ_2.

FIGURE 3. Examples of all sibling portraits with a central strip for σ_3 with the same spacing of x_i and y_i points, up to rotational symmetry.

Remark 1.13. Note that up to labeling of vertices and rotation, there are only three sibling portraits for $d = 3$. Moreover, the sibling leaves, ℓ_1, ℓ_2, and ℓ_3 in the non-symmetric case are of three different lengths: $\frac{1}{3} < |\ell_1| < \frac{1}{2}$, $\frac{1}{6} < |\ell_2| < \frac{1}{3}$, and $|\ell_3| < \frac{1}{6}$. See Theorem 3.1 for the count for $d > 3$.

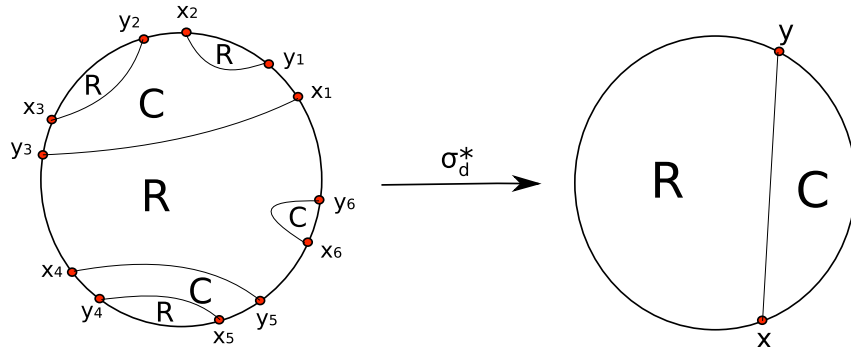Figure 4 shows one of many possibilities for a sibling portrait for σ_6.

FIGURE 4. Example of a sibling portrait for σ_6. (Not to scale; though chords are straight, we sometimes draw them curved to stand out.)

Proposition 1.14. *Suppose S is a sibling portrait. Then the following formula holds.*

$$\sum_{T \in S} (deg(T) - 1) = d - 1.$$

The proof is a direct consequence of the Euler characteristic of a tree.

Remark 1.15. Note that the degree of a C-region T is the number of times the boundary of T wraps, under σ_d, around the analog of the sector of the disk labeled C on the right side of Figure 4. A similar statement holds for R-regions.

2. CENTRAL STRIPS

Let $\mathcal{L}$ be a sibling d-invariant lamination.

Remark 2.1. If $\mathcal{S}$ is a full sibling collection mapping to leaf $\ell = \overline{xy}$, then endpoints x_i and y_i of the preimage leaves alternate counterclockwise around $\mathbb{S}$: $x_1 < y_1 < x_2 < y_2 < \cdots < x_d < y_d < x_1$. (Here we do not suppose that $\ell_i = \overline{x_i y_i}$.) If a leaf is a multiple of $\frac{1}{2d}$ long, then it maps to a leaf of length $\frac{1}{2}$, a diameter. A diameter leaf is either of fixed length or is critical (depending upon whether d is odd or even). As these can be handled as special cases, we consider only full sibling collections not having any leaf mapping to a diameter.

Definition 2.2 (Central Strip). Consider the sibling portrait of a full collection $\mathcal{S}$ of sibling leaves. Then the *central strip* C corresponding to $\mathcal{S}$ is the closure of the union of all C-regions C_i with degree at least 2. The *degree* of the central strip is $deg(C) = \min\{deg(C_i)\}$.

A tree with vertices labeled with two colors, and such that no edge connects vertices of the same color, is said to be *bicolored*. See Figure 5.
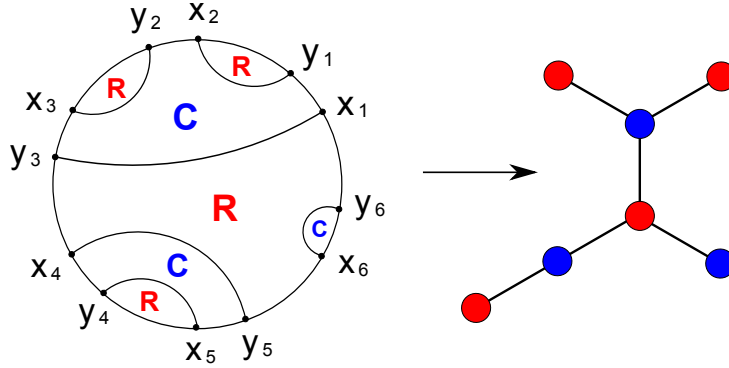

FIGURE 5. Mapping of a sibling portrait to a bicolored tree.

Theorem 2.3 (Central Strip Structure). *Let $\mathcal{S}$ be a full sibling collection of leaf ℓ_1 and its siblings such that $\sigma_d(\ell_1)$ is not a diameter. Then the following hold.*

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

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

Proof. Suppose $\mathcal{S} = \{\ell_1, \ell_2, \ldots, \ell_d\}$. As in Remark 2.1, we label the endpoints of the sibling leaves $x_1 < y_1 < \cdots < x_d < y_d < x_1$ in counterclockwise order so that x_i is an endpoint of ℓ_i. Note that the lengths of circular arcs between successive x and y points alternate between *short* (length $< \frac{1}{2d}$) and *long* (length $> \frac{1}{2d}$), provided the full sibling collection does not map to a diameter. Without loss of generality assume (x_i, y_i) is short.

To prove (1), assume that some leaf is more than $\frac{1}{2d}$ long. Let ℓ_i be a long leaf. We know by assumption that $\ell_i = \overline{x_i y_j}$ for some $j \neq i$. We claim that the region T with arc (x_i, y_i) in its boundary is a C-region and the region on the other side of ℓ_i is an R-region. To see this we traverse

the region T with arc (x_i, y_i) in its boundary counterclockwise. Refer to Figure 4. All sibling leaves map to a single image leaf $\overline{xy}$. All short arcs map to the shorter (counterclockwise) arc (x, y), and all long arcs map to the longer (counterclockwise) arc (y, x). Thus, as we move from x_i to y_i in the domain, we traverse the (shorter) arc (x, y) in the range; then as we move from y_i along a leaf emanating from it, we traverse the leaf $\overline{xy}$ from y to x in the range. Since short and long arcs alternate, we are now again at some x_k in the domain; we cannot be at x_i again, else ℓ_i with endpoint x_i is a short leaf; we thus traverse the arc (x_k, y_k) in the domain while we again traverse the arc (x, y) in the range. Proceeding counterclockwise around the region T, we see that we encounter only short arcs in the domain, each mapping to the counterclockwise arc (x, y) in the range. Note that we traversed at least two short arcs in the domain: (x_i, y_i) and (x_k, y_k). Moreover, we encounter only long leaves in the boundary of T. Thus, T is a C-region with only long leaves and short arcs in its boundary, and $deg(T) \geq 2$.

On the other side of leaf ℓ_i, by a similar argument, the region T' is an R-region bounded only by long arcs, though it may have both long and short leaves in its boundary. However, the short leaves in the boundary of T' can bound only degree 1 C-regions. Thus, T', sharing long boundary leaf ℓ_i with T, is an R-region.

In the above argument, we have shown both that when some leaf has length $> \frac{1}{2d}$, the central strip is nonempty, establishing conclusion (1) of the theorem, for we found at least one C-region of degree ≥ 2, and that any leaf bounds a C-region on one side and an R-region on the other. So, if a pair of complementary regions share a boundary leaf, then one is a C-region and the other is an R-region.

To prove (2), we form the dual graph T_S of the complementary regions of the sibling collection $\mathcal{S}$ as follows: Let each complementary region correspond to a vertex of T_S and each sibling leaf on the common boundary of two regions correspond to an edge of T_S. Since $\overline{\mathbb{D}}$ is connected and each sibling leaf disconnects $\overline{\mathbb{D}}$, T_S is a connected tree. Refer to Figure 5. The proof of part (1) shows that the tree is bicolored, with C-regions being one color and R-regions the other. If no leaves are of length $> \frac{1}{2d}$, then all leaves are short and the regions bounded by them are degree 1 C-regions. The central region bounded by all of them and long arcs of the circle is a degree d R-region. In this case, there is no central strip. This completes part (2) of the proof. $\square$

Proposition 2.4. *Let C_i enumerate the complementary components of the full sibling collection $\mathcal{S}$ meeting the circle in short arcs and let R_i enumerate the complementary components of the full sibling collection*

meeting the circle in long arcs. Let T_S be the corresponding bicolored tree, where C'_i and R'_i denote vertices corresponding to regions C_i and R_i, respectively. Then the number of edges of $T_S =$

$$\sum_i deg(C_i) = \sum_i deg(C'_i) = d = \sum_i deg(R'_i) = \sum_i deg(R_i).$$

The proof is left to the reader.

Proposition 2.5. *The maximum number of disjoint critical chords that can be contained in a component T of a sibling portrait is $deg(T) - 1$.*

Proof. Let T be a component of a sibling portrait under σ_d of degree $k \leq d$. Without loss of generality, let the (closed) arcs of $T \cap \mathbb{S}$ be

$$\mathcal{V} = \{[x_1, y_1], [x_2, y_2], \ldots, [x_k, y_k]\}.$$

Note that a critical chord must join one component of $T \cap \mathbb{S}$ to another going from a point $x_i + \epsilon_i \in [x_i, y_i]$, $0 \leq \epsilon_i \leq |(x_i, y_i)|$, to a point $x_j + \epsilon_i \in [x_j, y_j]$, $j \neq i$. Let $\mathcal{E}$ be a maximal collection of disjoint critical chords in T. It is not hard to see that the cardinality of $\mathcal{E}$ is at least $k - 1$; just draw critical chords from $k - 1$ different points of (x_1, y_1) in succession to points, one in each of $\mathcal{V} - \{[x_1, y_1]\}$, in succession in counterclockwise order.

Let $G(T)$ be a graph whose vertices are the elements of $\mathcal{V}$ and whose edges are elements of $\mathcal{E}$ joining elements of $\mathcal{V}$. The proposition follows from the following claims about $G(T)$:

(1) There are no cycles in $G(T)$.
(2) $G(T)$ is connected.

Given the claims, it follows that $G(T)$ is a tree with k vertices, and thus $k - 1$ edges. Hence, T contains at most $k - 1$ disjoint critical chords.

To prove (1), suppose, by way of contradiction, that there is a cycle in $G(T)$ of length $m \leq k$. Without loss of generality, assume the cycle includes exactly the first m elements of $\mathcal{V}$. Then in T there is a critical chord from a point $x_1 + \epsilon_1 \in (x_1, y_1)$, $\epsilon_1 \geq 0$, to the point $x_2 + \epsilon_1 \in (x_2, y_2)$. Then there is a critical chord from a point $x_2 + \epsilon_2 \in (x_2, y_2)$, $\epsilon_2 > \epsilon_1$ since the critical chords are disjoint, to the point $x_3 + \epsilon_2 \in (x_3, y_3)$. Proceeding in this fashion around the cycle, we finally have a critical chord from a point $x_m + \epsilon_m \in (x_m, y_m)$, $\epsilon_m > \epsilon_{m-1} > \cdots > \epsilon_1$, to a point $x_1 + \epsilon_m \in (x_1, y_1)$. But then the first and last critical chords meet in T since $\epsilon_m > \epsilon_1$, a contradiction.

To prove (2), suppose, by way of contradiction, that G_0 and G_1 are two components of $G(T)$ that are adjacent in counterclockwise order of vertices on $\mathbb{S}$. Suppose that (x_r, y_r) is the last vertex in G_0 and (x_s, y_s) is the last vertex of G_1 in counterclockwise order. Then we can add the

critical chord from y_r to y_s, connecting G_0 to G_1, contradicting maximality of $\mathcal{E}$. It may be necessary to move slightly the endpoints in $\mathbb{S}$ of up to two elements of $\mathcal{E}$ if they happen to have a y_r or y_s as an endpoint. $\quad\square$

2.1. CENTRAL STRIP LEMMA.

The Central Strip Lemma for σ_2, stated below, was used by Thurston [12] to show that there could be no wandering triangle for a lamination invariant under σ_2. A *triangle* in a lamination is a union of three leaves meeting only at endpoints pairwise and forming a triangle inscribed in $\mathbb{S}$. A triangle *wanders* if its forward orbit consists only of triangles (i.e., no side is ever critical), and no two images of the triangle ever meet. This was the first step in Thurston's classification and description of a parameter space for quadratic laminations. The Central Strip Lemma for σ_2 is also used to show that any polygon that returns to itself must return transitive on its vertices; hence, an invariant quadratic lamination cannot have an identity return triangle (see Definition 4.1). In a subsequent paper, we will recover and strengthen Kiwi's theorems that a d-invariant lamination cannot have a wandering $(d+1)$-gon, nor an identity return $(d+1)$-gon.

Theorem 2.6 (Thurston)**.** *Let C be the central strip in a quadratic lamination of leaf ℓ with $|\ell| > \frac{1}{3}$. Then the following hold:*

(1) *The first image $\ell_1 = \sigma_2(\ell)$ cannot reenter C.*
(2) *If an iterate $\ell_j = \sigma_2^j(\ell)$ of ℓ reenters C, for at least $j > 1$, then it must connect the two components of $C \cap \mathbb{S}$.*

Remark 2.7. $|\ell| = \frac{1}{3}$ is a special case: (2) holds with $j = 1$ and $\ell = \overline{\frac{1}{3}\frac{2}{3}}$ maps to itself in reverse order.

In order to state and prove a Central Strip Lemma for $d > 2$, we will need to consider the fact that higher degree laminations can have more than one critical leaf or gap. To discuss the distance between chords, we use a metric on chords defined by Childers [6] which we call the *endpoint metric*.

Definition 2.8. Suppose $\ell_1 = \overline{x_1 y_1}$ and $\ell_2 = \overline{x_2 y_2}$ are chords in $\overline{\mathbb{D}}$, meeting at most in one pair of endpoints. We may suppose the circular order of the endpoints is $x_1 < x_2 < y_2 \leq y_1$. Define the *endpoint distance* between ℓ_1 and ℓ_2 to be

$$d_{\mathrm{E}}(\ell_1, \ell_2) = |(x_1, x_2)| + |(y_2, y_1)|.$$

If $\ell_1 = \ell_2$, we define the distance to be 0.

See Figure 6 for an example. We define the endpoint metric only between non-crossing chords. The reader can check that on such chords it is a metric.
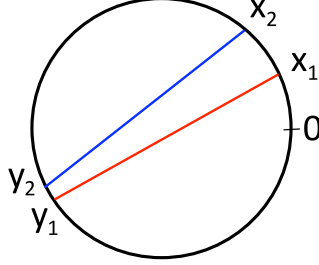
FIGURE 6. Endpoint distance between two disjoint chords: $\mathrm{d_E}(\overline{x_1 y_1}, \overline{x_2 y_2}) = |(x_1, x_2)| + |(y_2, y_1)|$.

See the proof of Theorem 2.3 for the definition of long and short arc lengths in the following theorem and its proof.

Theorem 2.9 (Central Strip Lemma). *Let C be a central strip of leaf ℓ and its siblings with long arc length $> \frac{1}{d+1}$. Let η be the short arc length. Then the following hold.*

(1) *The first image $\ell_1 = \sigma_d(\ell)$ cannot reenter C.*
(2) *The second image $\ell_2 = \sigma_d^2(\ell)$ cannot reenter C with both endpoints in a single component of $C \cap \mathbb{S}$.*
(3) *If an iterate $\ell_j = \sigma_d^j(\ell)$ of ℓ reenters C, for least $j > 1$, and has endpoints lying in one component of $C \cap \mathbb{S}$, then iterate ℓ_k, for some $k \leq j - 1$, gets at least as close in the endpoint metric as $\frac{\eta}{d^{j-k}}$ to a critical chord in $\overline{\mathbb{D}} \setminus C$.*

Proof. Suppose C is a central strip with long arc length $> \frac{1}{d+1}$, so there is leaf $\ell = \overline{xy}$ in its boundary such that $|\ell| > \frac{1}{d+1}$ and no leaf is a multiple of $\frac{1}{2d}$ long. Let the short arc length be η. Then

$$\eta < \frac{1}{d} - \frac{1}{d+1} = \frac{1}{d(d+1)}.$$

Let $\{\ell = \ell_0, \ell_1, \ell_2, \dots\}$ be the orbit of ℓ. Note that the length of ℓ_1 is

$$|\ell_1| = |\sigma_d(\ell_0)| = |(\sigma_d(x), \sigma_d(y)| = d\eta < d(\frac{1}{d(d+1)}) = \frac{1}{d+1}.$$

Since the length of a component of $C \cap \mathbb{S}$ is η, the endpoints of ℓ_1 cannot lie in one component of $C \cap \mathbb{S}$. On the other hand, since the length of ℓ_1 is less than $\frac{1}{d+1}$, it cannot connect two components, because the long arc length is $> \frac{1}{d+1}$. This establishes conclusion (1) of the theorem.

By Lemma 1.5, the length of ℓ_i, for $i > 1$, will grow until it is at least $\frac{1}{d+1}$ long, so will continue not to fit into one component of $C \cap \mathbb{S}$. It

may be that the orbit of ℓ never reenters C, or it may be that it reenters connecting two components of $C \cap \mathbb{S}$.

Suppose now that $\ell_j = \overline{x_j y_j}$ is the first iterate of ℓ that reenters C, and suppose that the endpoints of ℓ_j lie in one component of $C \cap \mathbb{S}$. Then $|\ell_j| \leq \eta$. The only way ℓ_j can get to η or less in length is by approaching a critical chord $D = \overline{ab}$ not contained in C sufficiently close in the endpoint metric. (See Definition 2.8 and Figure 7. In Figure 7 the endpoints of ℓ_{j-1} are denoted e and f.)
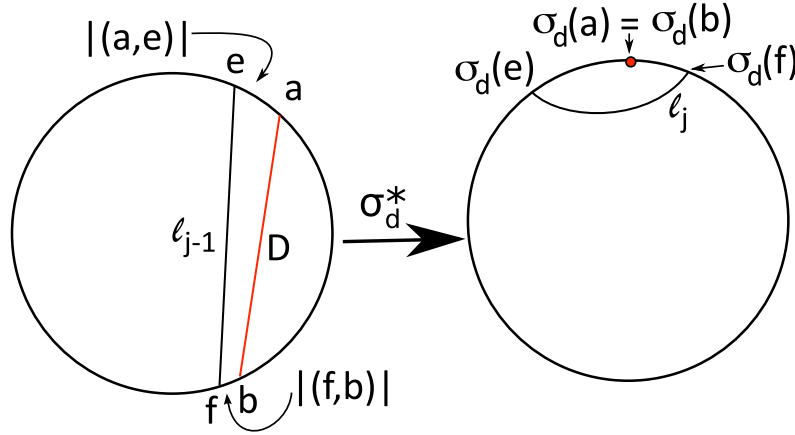

FIGURE 7. An iterate of ℓ approaches a critical chord D.

Suppose

$$d_{\mathrm{E}}(\ell_{j-1}, D) = |(x_{j-1}, a)| + |(b, y_{j-1})| \leq \frac{\eta}{d}.$$

Then, since $\sigma_d(a) = \sigma_d(b)$, we have $|\ell_j| = |(x_j, y_j)| =$

$$= |(\sigma_d(x_{j-1}), \sigma_d(a))| + |(\sigma_d(b), \sigma_d(y_{j-1})|$$

$$= |(\sigma_d(x_{j-1}), \sigma_d(y_{j-1})| \leq d(\frac{\eta}{d}) = \eta.$$

If the last close approach to a critical chord D before entering C were at an iterate $k < j - 1$, it would have to be even closer (by additional factors of $\frac{1}{d}$). This establishes part (3) of the theorem.

To see that the first iterate ℓ_j of ℓ that might enter C with both endpoints in one component of $C \cap \mathbb{S}$ must have $j > 2$, suppose by way of contradiction that ℓ_2 has both endpoints in one component. By the proof

of part (3), we have

$$d\eta = |\overline{xy}| < \frac{1}{d+1} < \frac{1}{d}.$$

Now ℓ_1 must be sufficiently close to a critical chord $D = \overline{ab}$ outside central strip C to shrink ℓ_2, so that

$$d_E(\overline{xy}, D) = |(a, x)| + |(y, b)| \leq \frac{\eta}{d}.$$

But since $|\overline{xy}| < \frac{1}{d}$, ℓ_1 must be under D. Hence,

$$d_E(\overline{xy}, D) + \overline{xy} = \frac{1}{d}.$$

By the above, we can compute that

$$d_E(\overline{xy}, D) = \frac{1 - d^2\eta}{d}.$$

Our supposition that ℓ_2 has both endpoints in one component implies that

$$\frac{1 - d^2\eta}{d} \leq \frac{\eta}{d}.$$

From this and our definition of η, it follows that

$$\frac{1}{d^2 + 1} \leq \eta < \frac{1}{d(d+1)},$$

a contradiction for all $d \geq 2$. This completes the proof of part (2). $\square$

Corollary 2.10 (Unicritical Central Strip Lemma). *Let C be a central strip of degree d of leaf ℓ and its siblings for the map σ_d. Then no image of ℓ can reenter C with both endpoints in a single component of $C \cap \mathbb{S}$.*

Proof. Since C is a central strip of degree d for σ_d, there must be a full sibling collection (d leaves) in its boundary. Since, by definition, it is a central strip, $|\ell| > \frac{1}{2d}$, as is the length of all its siblings. But to have room for an all-critical d-gon inside the strip, $|\ell| < \frac{1}{d}$. If $|\ell| < \frac{1}{d+1}$, it will grow in length under future iterates. So we lose no generality in assuming $|\ell| \geq \frac{1}{d+1}$. The case where $|\ell|$ is fixed at $\frac{1}{d+1}$ is trivial. Since there is no critical chord outside C, it follows from Theorem 2.9(3), that no future image of ℓ can reenter C with both endpoints in a single component of $C \cap \mathbb{S}$. $\square$

3. Counting Sibling Portraits and Central Strips

In the proof of Theorem 2.3, we showed that each sibling portrait corresponds to a bicolored tree. Now we show that the correspondence is one-to-one up to rotational symmetry.

Theorem 3.1. *If ℓ is a non-degenerate leaf, not a diameter, then there are $N(d)$ different full sibling collections which map onto ℓ, distinct up to rotational symmetry, where*

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

and $\phi(x)$ is Euler's totient function, *the number of positive integers less than, and relatively prime to, x.*

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, up to rotation in the plane. The number of bicolored trees with d edges is known to be $N(d)$ [5, Theorem 21]. Thus, we will use this correspondence to show there are $N(d)$ different sibling portraits mapping to the same leaf. Refer to Figure 8 during this proof.
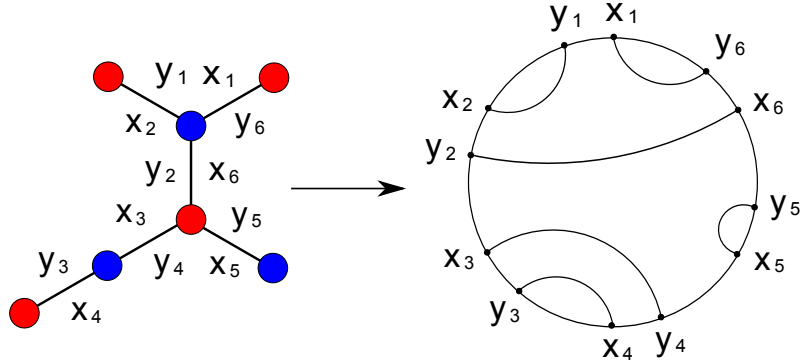

FIGURE 8. Mapping a bicolored tree to a sibling portrait.

The proof of Theorem 2.3 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 in the plane). 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.

Now assume we are given a bicolored tree with d edges like the one in Figure 8. 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 $x_1, y_1, x_2, y_2, \ldots, x_d, y_d$. 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. However, since we generally assume (x_i, y_i) to be a short arc, the vertex of the tree between x_i and y_i 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. $\qquad\square$

The corollary below follows immediately since the only time a sibling portrait does not have a central strip is when all the boundary leaves are short.

Corollary 3.2. *If ℓ is a non-degenerate leaf, not a diameter, then the number of different central strips, distinct up to rotational symmetry, whose boundary leaves map onto ℓ is $N(d) - 1$.*

4. Applications to Identity Return Polygons for σ_d, $d \geq 3$

It will be convenient to be able to refer to points on the circle by their *d-nary* expansions. The pre-images under σ_d of 0 partition the circle into d half-open intervals $[\frac{k-1}{d}, \frac{k}{d})$, for $1 \leq k \leq d$, labeled successively from 0 with symbol $\{0, 1, \ldots, d-1\}$. A point x of the circle is then labeled with its *itinerary*, an infinite sequence $t_0 t_1 t_2 \ldots t_n \ldots$ of symbols selected from $\{0, 1, \ldots, d-1\}$, based upon which labeled interval $\sigma_d^n(x)$ lies in. If a sequence of symbols repeats infinitely a finite sequence of length n, we write $\overline{t_0 t_1 t_2 \ldots t_n}$ to indicate the infinite sequence. For example, the point $\frac{1}{3}$ is of period 2 under σ_2. Thus, its σ_2-itinerary, or binary expansion, is $\overline{01}$. However, the σ_3-itinerary, or ternary expansion, of $\frac{1}{3}$ is $1\overline{0}$ since $\sigma_3(\frac{1}{3}) = 0$ and the point 0 is represented in ternary by $\overline{0}$ since 0 is a fixed point of σ_3. Using the d-nary expansion of a point x, the map σ_d is the *forgetful shift* since the map sends

$$t_0 t_1 t_2 \ldots t_n \ldots \mapsto t_1 t_2 t_3 \ldots t_n \ldots \ .$$

Definition 4.1. A (leaf or) polygon P in a d-invariant lamination is called an *identity return polygon* if and only if P is periodic under iteration of σ_d, the polygons in the orbit of P are pairwise disjoint, and on its first return, each vertex (and thus each side) of P is carried to itself by the identity.

Note that we require the identity return polygon P be in a sibling d-invariant lamination. This is because without this restriction, one can produce examples of identity return polygons that satisfy the other conditions of Definition 4.1, but could not correspond to a Julia set. (See Figure 15 (page 91) in connection with the proof of Lemma 4.7 for an example.) A periodic polygon in a sibling d-invariant lamination corresponds to a periodic branch point in a Julia set. Locally, the circular order of the branches is preserved by the polynomial. It is a non-obvious consequence of the definitions that in a sibling d-invariant lamination, a polygon maps under σ_d preserving the circular order of its vertices [2, Theorem 3.2]. An identity return polygon as we have defined it corresponds to a periodic branch point in a connected Julia set of a polynomial that returns to itself with no rotation around the branch point. Figure 9 illustrates an example of Kiwi [8].
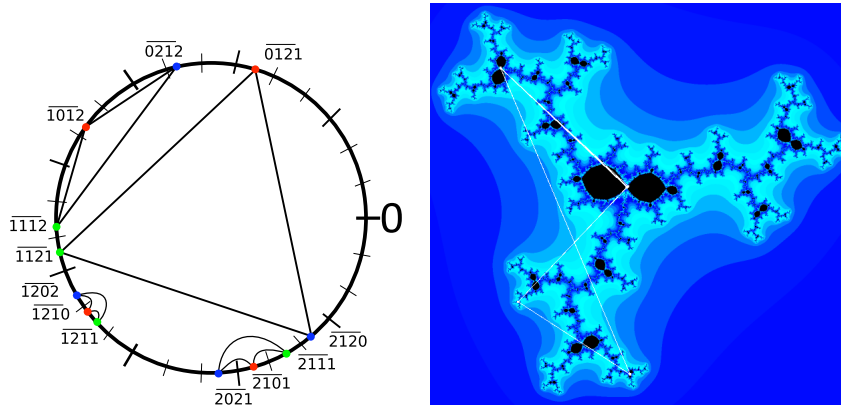

FIGURE 9. A period 4 identity return triangle for σ_3 and a corresponding polynomial Julia set with the orbit of the branch point indicated (after Kiwi).

Example 4.2. There are identity return d-gons of period 3 for all $d \geq 3$. The vertices of an example of such a d-gon for σ_d in d-nary expansion are

$$\{\overline{001}, \overline{002}, \overline{003}, \dots, \overline{00(d-1)}, \overline{(d-1)(d-1)0}\}.$$

In Figure 10 we illustrate an example of an identity return triangle of period 3 for σ_3 and identity return quadrilateral of period 3 for σ_4. As shown by Thurston, invariant quadratic laminations cannot have an identity return triangle [12], though they can have multiple identity return leaves. Generalizing Thurston's result for quadratic laminations, Kiwi [8, Theorem 3.1] proved the next theorem.
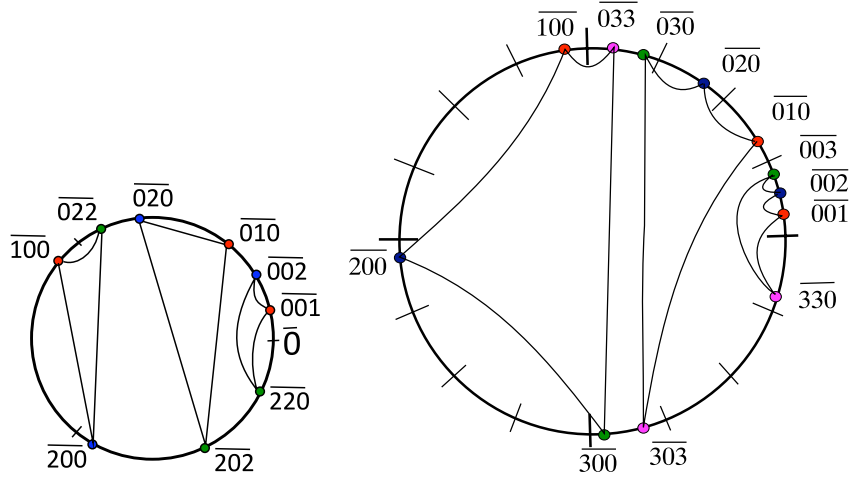
FIGURE 10. Examples of identity return d-gons for $d = 3$ and $d = 4$ under σ_3 and σ_4, respectively; the vertices are labeled by their (periodic) expansions.

Theorem 4.3 (Kiwi). *A d-invariant lamination cannot have an identity return k-gon for any $k > d$.*

4.1. PROPERTIES OF IDENTITY RETURN POLYGONS.

In what follows, we extract some facts about identity return polygons in general. Examples 4.2 and 4.5 show that the following theorem is sharp.

Theorem 4.4. *There can be no period 2 identity return k-gon for σ_d for $k \geq d$.*

Proof. It suffices to show that there does not exist an identity return d-gon of period 2 for σ_d, since an identity return k-gon for $k > d$ would contain an identity return d-gon. We argue by induction on d. It can be easily checked that there does not exist a period 2 identity return triangle for the map σ_3; all possible examples result in circular order reversing (see Figure 15, page 91). Now assume as the induction hypothesis that there does not exist an identity return $(d-1)$-gon of period 2 under σ_{d-1}. By way of contradiction assume that there does exist an identity return d-gon of period 2 under σ_d. Let $a_1 b_1, a_2 b_2, ..., a_d b_d$ denote the itineraries of vertices of the identity return d-gon, with $a_i, b_i \in B = \{0, 1, ..., d-1\}$ for all $i = 1, ..., d$. Consider the list of symbols $A = (a_1, ..., a_d, b_1, ..., b_d)$. Each element of B must appear in at least one entry of A, else we could define an identity return d-gon of period 2 for σ_{d-1}, which would contain

an identity $(d-1)$-gon, contradicting the induction hypothesis. This leads to two cases.

Case 1. Suppose that for some symbol $0 \leq k \leq d - 1$, the symbol k appears more than twice in list A.

By the pigeon-hole principle, this means that some other symbol $m \neq k$ will only appear once in A. Note that any such identity return polygon will be of the form shown in Figure 11 (illustrated for $d = 4$ and $m = 3$, up to some rotational symmetry). By removing the vertex with that symbol from the polygon, one obtains a $(d - 1)$-gon which uses $d - 1$ symbols. This contradicts the induction hypothesis.
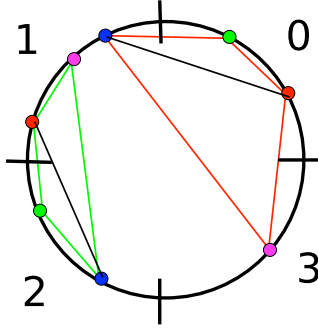

FIGURE 11. Case 1 for σ_4 where $m = 3$.

Case 2. Suppose that each symbol $0 \leq k \leq d - 1$ appears exactly twice in list A.

Then there are two vertices of the polygons in each $1/d$ section of the circle corresponding to the symbols in B. Thus, one of the two polygons in the orbit must have two vertices in a $1/d$ section of the circle corresponding to the symbol 0 in B or to the symbol $d - 1$ in B; otherwise the polygons would cross. First, let us assume two vertices of a polygon are in the "0" section (see Figure 12). The vertex in this section closest to the point $\overline{0} \in \mathbb{S}$, denote it $\overline{0j}$, represents the first vertex (in counterclockwise circular order from $\overline{0} \in \mathbb{S}$) of the identity return polygon on the circle. The next vertex in order will be $\overline{0k}$ for some $k > j \in B$. Since the first coordinates of these vertices are 0, then mapping them forward under σ_d will yield vertices $\overline{j0}$ and $\overline{k0}$. Since $j < k$, these vertices cannot map to the same $1/d$ section.

Since circular order must be preserved, and all symbols are used, the two vertices must map either to vertices in adjacent $\frac{1}{d}$ sections (see Figure 12 on right), or $\overline{j0}$ must be the largest (in circular order) vertex in

the image polygon (see Figure 12 on left). Since both vertices have 0 as their second coordinate, they must each be the first vertex of the polygons in the sections in which they appear. If the vertices map to adjacent $\frac{1}{d}$ sections, this is impossible because $\overline{j0}$ (preceding $\overline{k0}$ in circular order in the image polygon) would not be the smallest in its section. If $\overline{j0}$ is the largest vertex in the image polygon, then preserving circular order requires $\overline{k0}$ to be the smallest vertex in the image polygon, contradicting $j < k$. All options lead to a contradiction, so no such polygon orbit is possible.

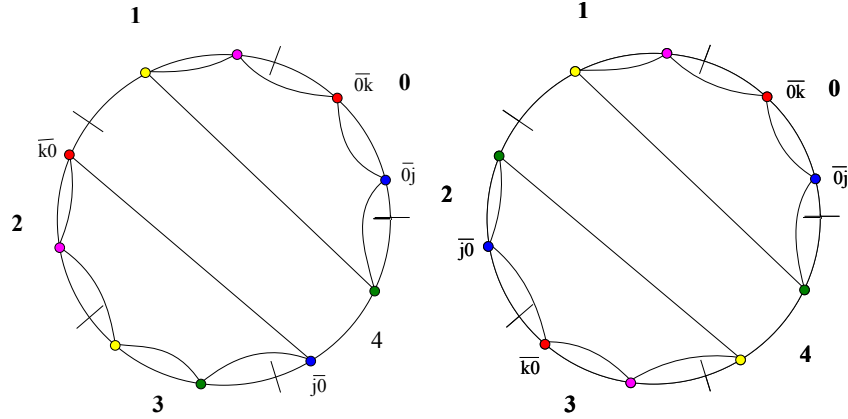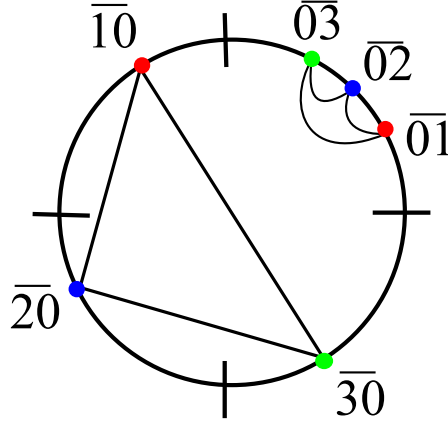The argument is similar if an identity return polygon has two vertices in the $(d-1)$ section.   □

FIGURE 12. Case 2 for σ_5 with two vertices of a polygon in the "0" section.

Example 4.5. Period 2 identity return $(d-1)$-gons exist under σ_d for all $d > 3$. See Figure 13 for $d = 4$. The vertices of an example of such a $(d-1)$-gon for σ_d in d-nary expansion are

$$\{\overline{01}, \overline{02}, \overline{03}, \ldots, \overline{0(d-1)}.$$

FIGURE 13. Period 2 identity return triangle for σ_4.

### 4.2. SOME PROPERTIES OF IDENTITY RETURN TRIANGLES FOR σ_3.

The following propositions present a series of elementary facts about identity return triangles for σ_3.

Proposition 4.6. *Let ℓ be a leaf in a 3-invariant lamination which is not eventually of fixed length. Then ℓ must eventually get within $\frac{1}{12}$ of a critical length.*

Proof. Let ℓ be a leaf in an invariant lamination under σ_3 whose length is not eventually fixed. By Proposition 1.4, the fixed lengths under the leaf length function τ_3 are $\frac{1}{4}$ and $\frac{1}{2}$. By Lemma 1.5, each leaf of length $< \frac{1}{4}$ must eventually iterate to a leaf $\ell_i = \sigma_3^i(\ell)$ of length $\geq \frac{1}{4}$ (see "spider web diagrams" in Figure 14). So assume $|\ell| > \frac{1}{4}$. If $\frac{1}{4} < |\ell| < \frac{5}{12}$, then we can place a critical chord so that in the endpoint metric (Definition 2.8), ℓ is within $\frac{1}{12}$ of the critical chord. So we may assume $\frac{5}{12} < |\ell| < \frac{1}{2}$. We apply the leaf length function τ_3 iteratively, observing that $\tau_3(|\ell|)$ is repelled from $\frac{1}{2}$ because the graph of τ_3 is below the identity on the interval $(\frac{1}{3}, \frac{1}{2})$. Thus, we see that there is a first $k \geq 1$ such that $\frac{1}{4} < \tau_3^k(|\ell|) < \frac{5}{12}$ (see Figure 14 on right). Hence, $\sigma_3^k(\ell)$ is within $\frac{1}{12}$ of a critical chord. $\square$

Proposition 4.7. *An identity return triangle in a 3-invariant lamination cannot have a side of fixed length.*

Proof. By Proposition 1.4, the fixed lengths under the leaf length function τ_3 are $\frac{1}{4}$ and $\frac{1}{2}$. Let ℓ be a side of fixed length of an identity return triangle
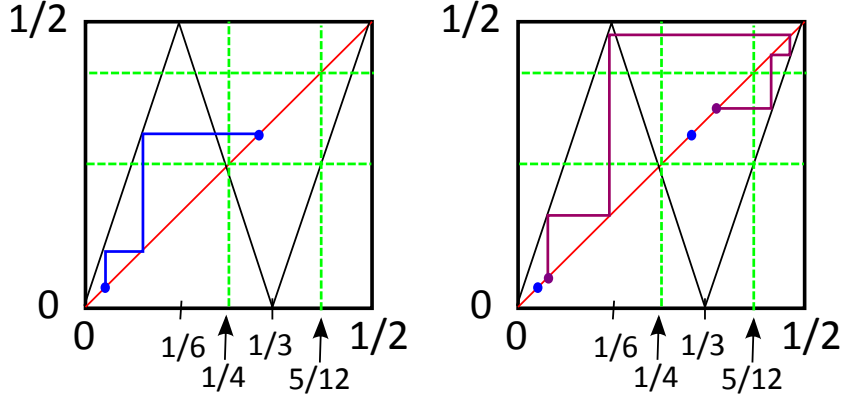
FIGURE 14. Spider web diagrams for leaf length function τ_3.

T. If $|\ell| = \frac{1}{2}$, then ℓ is a diameter, and if ℓ is not fixed, $\sigma_3(\ell)$ will cross ℓ, a contradiction of ℓ being a leaf of an invariant lamination. On the other hand, if ℓ is fixed, then $\sigma_3(T)$ meets T, contradicting T being an identity return triangle. If $|\ell| = \frac{1}{4}$, then the period of ℓ cannot be greater than 3 or it will meet its iterated image. If the period of ℓ were 3, then both endpoints of ℓ would have period 3 repeating ternary expansions (see Figure 10 for an example). But if we add $\frac{1}{4} = \overline{01}$ to a repeating ternary expansion for one endpoint of ℓ, the other endpoint of ℓ will not be period 3 repeating in ternary. It is possible to construct a period 2 triangle which appears to satisfy the definition (see Figure 15) and has a side of fixed length. However, there is no room for critical chords disjoint from the triangles in the orbit, and the triangle maps forward reversing circular order; thus, it cannot be in a sibling d-invariant lamination. $\qquad \square$

Remark 4.8. By way of notation, if T is an identity return triangle of least period n in a 3-invariant lamination, we denote the sides of the triangle by $T = ABC$, and the orbit of the triangle by

$$T = T_0 = A_0 B_0 C_0 \mapsto \sigma_3(T_0) = T_1 = A_1 B_1 C_1 \mapsto \ldots$$
$$\mapsto T_n = A_n B_n C_n = T_0.$$

This notation can be extended to any identity return polygon.

Proposition 4.9. *Let ABC be an identity return triangle of period k. If two sides A and B are of the same length, then there exists a unique $i \in [0, k)$ such that A_i and B_i are within $\frac{1}{12}$ of critical. Further, there is no critical chord c such that A_i and B_i are both within $\frac{1}{12}$ of c in the endpoint metric.*
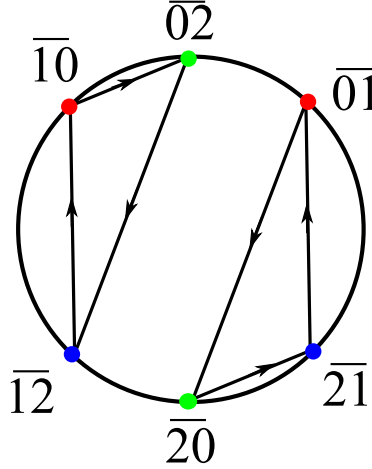
FIGURE 15. Example of an "almost" identity return triangle for σ_3 of period 2; it fails because the triangle maps forward reversing circular order.

Proof. Suppose that $T_0 = A_0 B_0 C_0$ is an identity return triangle and that $|A_0| = |B_0|$. Then for all i, we have $|A_i| = |B_i|$. By propositions 4.6 and 4.7, there is an i such that $\frac{1}{4} < |A_i| = |B_i| < \frac{5}{12}$. If $|A_i| < \frac{1}{3}$, as illustrated on the left in Figure 16, then, similar to the proof of Proposition 4.7, there is no room for two critical chords disjoint from T_i, since A_i and B_i together occupy more than half the circle's arc. Hence, there is no room for siblings of T_i, so no such sibling invariant lamination.

If $|A_i| > \frac{1}{3}$, then there are two regions in $\mathbb{D} \setminus T_i$ where one can place critical chords sufficiently close to A_i and B_i, the region with only A_i on its boundary and the region with only B_i on its boundary. See Figure 16 Right where we include the siblings T_i' and T_i'' of triangle T_i. Since $|A_i| < \frac{5}{12}$, each of A_i and B_i is within $\frac{1}{12}$ of a critical chord. That the ith iterate is unique with this property is clear, since any other such iterate would cross T_i. $\qquad\square$

Proposition 4.10. *Let T be an identity return n-gon ($n \geq 3$) in a 3-invariant lamination. Then, in the orbit of T, two sides will simultaneously approach within $\frac{1}{12}$ of a critical length. In fact, one of the following happens:*

(1) *Two sides of T approach within $\frac{1}{12}$ of two different critical chords at the same iterate.*
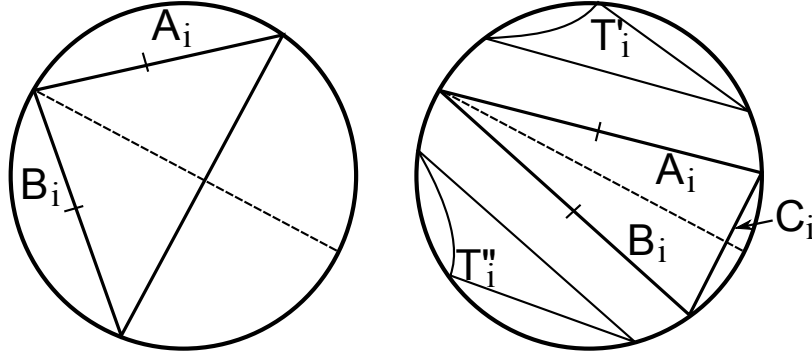
FIGURE 16. **Left:** a triangle with two equal sides of length between $\frac{1}{4}$ and $\frac{1}{3}$. **Right:** the length is between $\frac{1}{3}$ and $\frac{5}{12}$.

(2) *Two sides of T approach within $\frac{1}{12}$ of one side of the same critical chord at the same iterate.*

Furthermore, at that iterate these two sides are longer than any other side.

Remark 4.11. Example 9 shows that the first case can occur. The cubic example in Figure 10 shows that the second case can occur.

Proof. Let $T_0 = A_0 B_0 C_0 \ldots$ be an identity return n-gon in a 3-invariant lamination. By Lemma 4.6 all the sides will eventually get within $\frac{1}{12}$ of a critical chord. In case (1), the polygon must lie between the two critical chords. Hence, those two sides are longer than $\frac{1}{3}$, so all other sides are shorter than $\frac{1}{3}$. Now assume case (1) does not occur.

Without loss of generality, suppose that side A_0 is within $\frac{1}{12}$ of a critical chord ℓ. So there exists an iterate $i \neq 0$ such that side B_i is within $\frac{1}{12}$ of a critical chord, not necessarily ℓ, and there exists an iterate $0 \neq j \neq i$ such that C_j is within $\frac{1}{12}$ of a critical chord, not necessarily ℓ, else we are done.

Suppose that the three sides A_0, B_i, and C_j are close to the same critical chord ℓ and approach no other critical chord closely. Without loss of generality, let A_0 be the side furthest away from ℓ, while still being within $\frac{1}{12}$ of ℓ. See Figure 17 Center. Then T_i is closer to ℓ and by the Central Strip Lemma (Theorem 2.9) sides A_i and B_i are long, since side A has not closely approached a different critical chord before iterate i. Triangle T_j cannot be closer to ℓ since then, by the Central Strip Lemma, the three sides are simultaneously close to $\frac{1}{3}$ in length. Therefore, T_j is

between T_i and T_0 (or siblings). Let T be of period k, so $T_0 = T_k$. Then the iterate T_{k+i} is the first time the central strip of C_j is re-entered, and since side C is never close to a different critical chord, C_{k+i} needs to connect two components of the central strip by the Central Strip Lemma. But $T_{k+i} = T_i$ has C_i inside one component, a contradiction.

Suppose that all three sides are close to three different critical chords at iterates 0, i, and j. Then the three critical chords must form an all-critical triangle. See Figure 17 Left. One of the polygons, say T_0, must be outermost with respect to the all-critical triangle. As argued in the previous paragraph, T_i and T_j are then within the central strip of A_0, contradicting in a similar fashion the Central Strip Lemma.

Therefore, the three sides A_0, B_i, and C_j must approach exactly two different critical chords. By the pigeon hole principle, two sides must approach the same critical chord ℓ. Without loss of generality, let those sides be A_0 and B_i. Let A_0 be the one further away from the critical chord ℓ, but still within $\frac{1}{12}$ of ℓ. Suppose that B_i is within $\frac{1}{12}$ of critical chord at iterate i. Then T_i is inside the central strip of A_0 and its sibling A_0'. Hence, two sides of the polygon T_i are within $\frac{1}{12}$ of a critical chord ℓ at the same iterate and all other sides are within a central strip, so shorter. $\qquad\square$
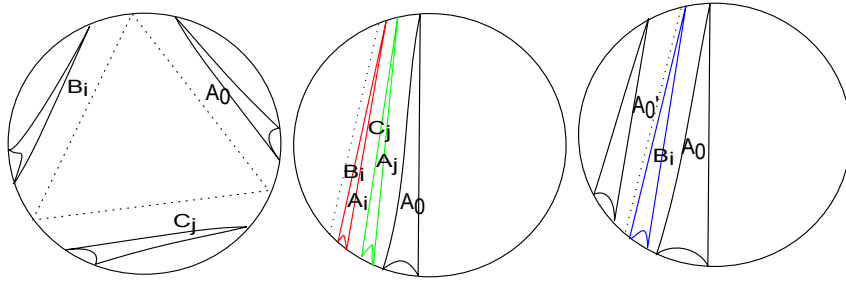

FIGURE 17. **Left:** an identity return polygon approaches three different critical chords. **Center:** an identity return polygon approaches just one critical chord. **Right:** an identity return polygon approaches exactly two critical chords (only one shown).

Proposition 4.12. *Let T be an identity return n-gon in a 3-invariant lamination. Let A and B be two sides that are within $\frac{1}{12}$ of a critical chord and let C be the longest remaining side. If A and B are simultaneously within $\frac{1}{12}$ of*

(1) *two different critical chords at the same iterate, or*
(2) *the same critical chord at the same iterate,*

then the following occur, respectively:

(a) *$\sigma_3(C)$ becomes the longest side, or*
(b) *$\sigma_3(A)$ or $\sigma_3(B)$ becomes the longest side.*

Proof. Let $T = ABCD_4...D_n$.

Suppose (1) holds. Then T lies between two critical chords. Since there is a critical chord between A and its sibling A', there must be a critical value underneath $\sigma_3(A)$. Similarly, there is a critical value underneath $\sigma_3(B)$. Furthermore, those two critical values are different since A and B approach two different critical chords. See Figure 18 on left.
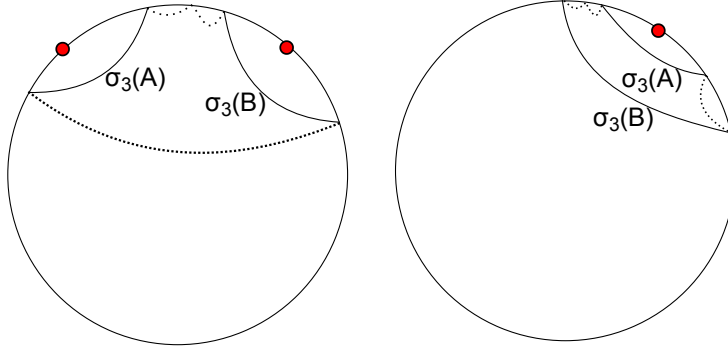

FIGURE 18. Image of a polygon with critical values marked.

Now, since T lies between two critical chords, $\sigma_3(T)$ lies in one half of the circle (else it crosses T). Since $\sigma_3(T)$ is an n-gon, there must be a side C' with $\sigma_3(A)$ and $\sigma_3(B)$ underneath C'. Since $\sigma_3(T)$ lies in one half of the circle, C' is longer than both $\sigma_3(A)$ and $\sigma_3(B)$.

To see that $C' = \sigma(C)$, consider two cases.

Case 1. $|C| \geq \frac{1}{6}$. Let $|A| = \frac{1}{3} + \alpha$, $|B| = \frac{1}{3} + \beta$, and $\gamma = \sum_{k=4}^{n} |D_k|$ Clearly, $\gamma < \frac{1}{6}$ and $|D_i| < \frac{1}{6}$ for all $i = 4, ...n$. Then $|C| = \frac{1}{3} - \alpha - \beta - \gamma$. Then $|\sigma_3(A)| = 3|A| - 1 = 3\alpha$, $|\sigma_3(B)| = 3|B| - 1 = 3\beta$, $|\sigma_3(D_i) = 3|D_i|$

for all $i = 4, ...n$, and $|\sigma_3(C)| = 1 - 3|C| = 3\alpha + 3\beta + 3\gamma = |\sigma_3(A)| + |\sigma_3(B)| + \sum_{k=4}^{n} |\sigma_3(D_k)|$. Therefore, $\sigma_3(C)$ is the longest side.

Case 2. $|C| < \frac{1}{6}$. Then all sides except A and B are shorter than $\frac{1}{6}$. Thus, under the application of σ_3, the side lengths of $\sigma_3(C), \sigma_3(D_4), ...,$ $\sigma_3(D_n)$ triple. Since C was the longest among $C, D_4, ..., D_n$, $\sigma_3(C)$ must be the longest among $\sigma_3(C), \sigma_3(D_4), ..., \sigma_3(D_n)$. However, we have shown above that there exists at least one side that is longer than $\sigma_3(A)$ and $\sigma_3(B)$. Thus, $\sigma_3(C)$ must be the longest.

Cases 1 and 2 establish that (1) implies (a).

Now suppose (2) occurs. See Figure 18 Right. By the hypothesis, both A and B are within $\frac{1}{12}$ of a critical chord ℓ. Then there is the same critical value underneath $\sigma_3(A)$ and $\sigma_3(B)$. Since all remaining sides are between $\sigma_3(A)$ and $\sigma_3(B)$, and since $|\sigma_3(A)|, |\sigma_3(B)| < \frac{1}{4}$, one of those (the outermost from the critical value) will be the longest. Thus, we have that (2) implies (b). $\square$

We conclude this section by proving the special case of Kiwi's Theorem (4.3) for σ_3.

Theorem 4.13. *There are no identity return quadrilaterals under σ_3.*

Proof. Suppose that such a quadrilateral T of period n exists. Let A, B, C, and D be the sides of T, not necessarily in circular order.

Case 1. Suppose, for no iterate i, that two sides are within $\frac{1}{12}$ of two different critical chords.

Then by Proposition 4.10 there exists an iterate T_0 where two sides of T_0 are within $\frac{1}{12}$ of the same critical chord ℓ. Let A_0 and B_0 be those sides. The remaining sides must also eventually approach critical chords within $\frac{1}{12}$ (Proposition 4.6).

CLAIM. We may assume without loss of generality that every approach to a given critical chord is closer than the previous approach to that critical chord.

To see this, suppose that at some iterate we have a quadrilateral T_j further away from a given critical chord than quadrilateral T_i, $i < j$, while at least one of the sides of T_j is different from the sides of T_i that are within $\frac{1}{12}$ of a critical chord. Then, the different "long" side of T_i must approach a different critical chord at some iterate before returning to itself. Otherwise, the Central Strip Lemma is contradicted upon the return to T_i. In that case, we could take T_j as a starting point instead of T_i, establishing the claim.

By Proposition 4.12 let B_1 be the longest side. Since $|B_1| < \frac{1}{4}$, it will grow by Lemma 1.5. Let B_i be the next close approach of side B to a critical chord. If A_i is the second longest side, then we can take this as

a starting point. So, without loss of generality, we can suppose that C_i is the second longest. Then by the Claim and the Central Strip Lemma, the two sides must be close to a different critical chord ℓ'. Now we have a central strip of sides A_0 and B_0 around ℓ and a central strip of B_i and C_i around ℓ'. Denote this state of affairs by $[\ell : A_0, B_0]$ and $[\ell' : B_i, C_i]$. See Figure 19 Left. By Proposition 4.12 and without loss of generality, we may assume that C_{i+1} is the longest side. Let j be the iterate when C_j next approaches a critical chord.
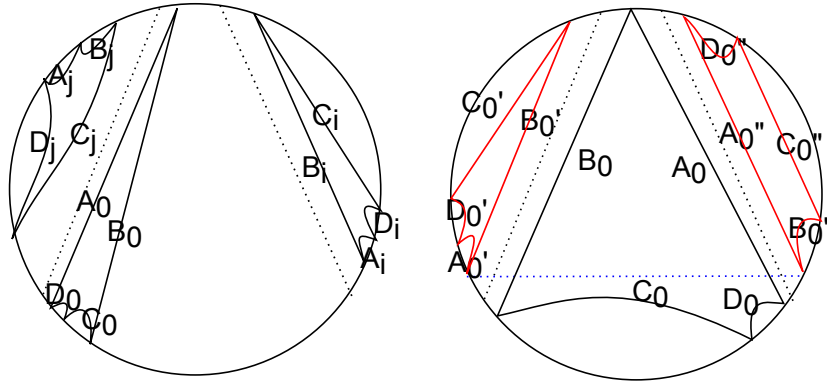

FIGURE 19. Impossible quadrilaterals under σ_3.
Left: Case 1; **Right:** Case 2.

If A_j is second longest then T_j, by the Claim and the Central Strip Lemma, must be close to ℓ. But this gives us $[\ell : A_j, C_j]$, $[\ell' : B_i, C_i]$, which is up to renaming what we had before. If B_j is second longest then by the Claim and the Central Strip Lemma T_j must be close to ℓ'. This gives us $[\ell : A_0, B_0]$, $[\ell' : C_j, B_j]$, which is what we had before but at a later iterate. If D_j is the second longest, then T_j cannot be within $\frac{1}{12}$ of either critical chord, without, by the Claim, violating the Central Strip Lemma. Therefore, no such quadrilateral exists.

Case 2: Now suppose that there is an iterate i where two sides, A_i and B_i, approach two different critical chords.

Since this iterate is unique, if we find a subsequent iterate where both sides are within $\frac{1}{12}$ of the same critical chord, then we can follow the argument in Case 1 to see that no such quadrilateral exists.

Refer to Figure 19 Right. Let $|A_0| = \frac{1}{3} + \alpha$, $|B_0| = \frac{1}{3} + \beta$, and note that $\alpha, \beta < \frac{1}{12}$. Then $|C_0| + |D_0| = \frac{1}{3} - \alpha - \beta > \frac{1}{6}$. First suppose that neither

C_0 nor D_0 is longer than $\frac{1}{6}$. Then by Lemma 1.3 with $d = 3$, $|A_1| = 3\alpha$, $|B_1| = 3\beta$, $|C_1| = 3|C_0|$, and $|D_1| = 3|D_0|$, and so $|A_1|+|B_1|+|C_1|+|D_1| = 1$. This means that T_1 is on both sides of the diameter, so it will intersect T_0. Therefore, we may suppose that $|C_0| > \frac{1}{6}$. Then $|A_1| = 3\alpha$, $|B_1| = 3\beta$, $|C_1| = 1 - 3|C_0|$, and $|D_1| = 3|D_0|$. Hence, $|C_1| = |A_1| + |B_1| + |D_1|$.

If C_0 is not within $\frac{1}{12}$ of a critical chord, then C grows, so C_i must be within a central strip formed by A_0 or B_0 on its first approach to critical length; let it be B_0. Then by the Central Strip Lemma, C_i and B_i must be the two longest sides. Now we can follow the argument in Case 1 to see that no such quadrilateral exists.

Therefore, we may assume that C_0 is within $\frac{1}{12}$ of a critical chord. Since $|C_1| < \frac{1}{4}$, C_1 will continue to grow. So it must approach a critical length at some iterate i. If it approaches a critical chord within the central strip of A_0 or B_0, we can follow the argument in Case 1. So we may suppose that it approaches the same critical chord as at iterate 0. Then it must go underneath C_0, as otherwise side D never gets within $\frac{1}{12}$ of a critical chord. There exists an iterate i when side D_i is within $\frac{1}{12}$ of a critical length. If quadrilateral T_i is inside either of the central strips we can follow the Case 1 argument. So we may suppose that D_i is within $\frac{1}{12}$ of the same critical chord as C_0. Then, by the Central Strip Lemma, sides C_i and D_i must be the two longest sides, so one of them remains the longest. Eventually, one of the sides A or B will become the second longest; otherwise, T cannot return to T_0. Without loss of generality, let C and A be the two longest sides. Then as they approach critical length at iterate j, by the Central Strip Lemma the quadrilateral T_j must be inside the central strip of A_0. Again, we can follow the argument in Case 1. Therefore, there is no such identity return quadrilateral. □

4.3. Questions.

We conclude with a series of questions and remarks.

Question 4.14. What is the appropriate generalization of Proposition 4.6 to σ_d for $d > 3$?

In order to be able to use the full power of the Central Strip Lemma, we want leaves of an identity return polygon for σ_d to get within $\frac{1}{d(d+1)}$ of a critical chord. This need not happen even for a d-gon for σ_d, in general. For example, the period 3 quadrilateral with vertices $\{\overline{132}, \overline{032}, \overline{022}, \overline{200}\}$ under σ_4 has a side, namely $\overline{022}\,\overline{200}$, which never gets within $\frac{1}{20}$ of a critical chord.

Question 4.15. With reference to Proposition 4.7, can an identity return polygon have a side of fixed length?

Question 4.16. With reference to Proposition 4.9, if an identity return polygon has two sides of the same length, must they simultaneously approach two different critical chords?

Question 4.17. What are the appropriate generalizations of propositions 4.10 and 4.12 to σ_d for $d > 3$?

We can show that if the two longest sides of an identity return polygon P in a d-invariant lamination are simultaneously within $\frac{1}{2d}$ of the same critical chord, then one of the two longest sides of $\sigma_d(P)$ remains longest.

One can define "identity return polygon" without assuming it is in a d-invariant lamination.

Definition 4.18 (Alternate Definition of Identity Return Polygon). A polygon P in the closed unit disk is called an *identity return polygon* if and only if P is periodic under iteration of σ_d, the polygons in the orbit of P are pairwise disjoint, circular order of vertices of P is preserved by the action of σ_d on P, and, on its first return, each vertex (and thus each side) of P is carried to itself by the identity.

Question 4.19. Does the existence of an identity return polygon P defined as in Definition 4.18 imply the existence of a d-invariant lamination containing P?

Question 4.20. What polynomials have Julia sets with a vertex that returns to itself in the pattern of the identity return polygons of Figure 10? Of Example 4.2, in general?

Question 4.21. What is the "simplest" 3-invariant lamination that contains a given identity return triangle?

Here "simplest" might mean an invariant lamination with no leaves other than preimages of the triangle or their limit leaves and with each of two sides of the triangle bordering a (different) infinite gap of the lamination.

Question 4.22. Is there any bound on the number of identity return triangle orbits that a 3-invariant lamination can contain?

It is clear from proofs above that a 3-invariant lamination can contain only one identity return triangle orbit where two sides of the triangle approach within $\frac{1}{12}$ of two different critical chords on the same iterate.

Question 4.23. Given $d > 2$ and a period $p > 2$ (necessarily), how many distinct identity return d-gon orbits of period p can be formed under σ_d?

Question 4.24. Given $d > 2$ orbits of period $p > 2$ under σ_d, is it the case that they form at most one identity return d-gon orbit?

That for $d = 2$, the answer is "yes" appears to be intimately connected to the detailed structure of parameter space [12].

REFERENCES

[1] A. Blokh and G. Levin, *An inequality for laminations, Julia sets and "growing trees,"* Ergodic Theory Dynam. Systems **22** (2002), no. 1, 63–97.

[2] Alexander M. Blokh, Debra Mimbs, Lex G. Oversteegen, and Kirsten I. S. Valkenburg, *Laminations in the language of leaves*, Trans. Amer. Math. Soc. **365** (2013), no. 10, 5367–5391.

[3] Alexander Blokh and Lex Oversteegen, *Wandering triangles exist*, C. R. Math. Acad. Sci. Paris **339** (2004), no. 5, 365–370.

[4] _____ , *Wandering gaps for weakly hyperbolic polynomials* in Complex Dynamics: Families and Friends. Ed. Dierk Schleicher. A K Peters, Wellesley, MA, 2009. 139–168.

[5] Miklós Bóna, Michel Bousquet, Gilbert Labelle, and Pierre Leroux, *Enumeration of m-ary cacti*, Adv. in Appl. Math. **24** (2000), no. 1, 22–56.

[6] Douglas K. Childers, *Wandering polygons and recurrent critical leaves*, Ergodic Theory Dynam. Systems **27** (2007), no. 1, 87–107.

[7] Lisa R. Goldberg and John Milnor, *Fixed points of polynomial maps. Part II. Fixed point portraits*, Ann. Sci. École Norm. Sup. (4) **26** (1993), no. 1, 51–98.

[8] Jan Kiwi, *Wandering orbit portraits*, Trans. Amer. Math. Soc. **354** (2002), no. 4, 1473–1485.

[9] John Milnor, *Periodic orbits, externals rays and the Mandelbrot set: An expository account* in Géométrie complexe et systèmes dynamiques (Orsay, 1995). Ed. Marguerite Flexor, Pierrette Sentenac, and Jean-Christophe Yoccoz. Astérisque No. 261 (2000). Paris: Société Mathématique de France, 2000. xiii, 277–333.

[10] _____ , *Dynamics in One Complex Variable.* 3rd edition. Annals of Mathematics Studies, 160. Princeton, NJ: Princeton University Press, 2006.

[11] Sam B. Nadler, Jr. *Continuum Theory. An Introduction.* Monographs and Textbooks in Pure and Applied Mathematics, Vol. 158. New York: Marcel Dekker, Inc., 1992.

[12] William P. Thurston, *On the geometry and dynamics of iterated rational maps* in Complex Dynamics: Families and Friends. Ed. Dierk Schleicher. A K Peters, Wellesley, MA, 2009. 3–137.

(Cosper) Department of Mathematical Sciences; Indiana University–Purdue University at Indianapolis; Indianapolis, Indiana 46202
 E-mail address: `dcosper314@gmail.com`

(Houghton) Department of Mathematics; Metropolitan Community College; Kansas City, MO 64111
 E-mail address: `echo41011@gmail.com`

(Mayer) Department of Mathematics; University of Alabama at Birmingham; Birmingham, AL 35294
 E-mail address: `jcmayer@uab.edu`

(Mernik) Department of Mathematics; The Ohio State University; Columbus, OH 43210
 E-mail address: `mernik.1@osu.edu`

(Olson) Department of Physics; Harvard University; Cambridge, MA 02138
 E-mail address: `joseph.wesley.olson@gmail.com`