Realization of Hyperbolic Laminations and Properties of Limit Laminations

Adam Carty
University Of Alabama At Birmingham

REALIZATION OF HYPERBOLIC LAMINATIONS AND PROPERTIES OF LIMIT LAMINATIONS

by

ADAM CARTY

LEX OVERSTEEGEN, COMMITTEE CHAIR
RAGIB HASAN
JOHN MAYER
LAWRENCE ROBERTS
NIKITA SELINGER
DONGSHENG WU

A DISSERTATION

Submitted to the graduate faculty of the University of Alabama at Birmingham,
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

BIRMINGHAM, ALABAMA

2023

REALIZATION OF HYPERBOLIC LAMINATIONS AND PROPERTIES OF LIMIT LAMINATIONS

ADAM CARTY

APPLIED MATHEMATICS

ABSTRACT

The dynamics of polynomials is a topic of great interest for those studying complex dynamics. Polynomial dynamics is concerned with how collections of points behave under iteration of a polynomial P. The *Julia set* of P, which is the boundary of the collection of points which iterate off to ∞, is an interesting collection of points to look at because P is sensitive to small pertubations of points in the Julia set. Thanks to the work of William Thurston, we have a tool known as a *lamination* (a closed collection of chords in the unit disk $\overline{\mathbb{D}}$ which do not cross) to aid us in studying the behavior of P on the Julia set. Given a polynomial P with a locally connected Julia set, Thurston showed us how to construct a *q-lamination* corresponding to P. In this thesis, we go the other direction by finding a polynomial that corresponds to a given q-lamination (with additional assumptions). We then explore the properties of such laminations which can be realized by a polynomial. We conclude by looking at an appropriate space of laminations. The space that we choose is known as the space of *limit laminations* which is the closure of the set of q-laminations. The space of limit laminations is used to understand the space of complex polynomials by giving us a way to assign a lamination to a polynomial without a locally connected Julia set. We focus on properties of limit laminations as well as the q-laminations which converge to limit laminations.

DEDICATION

I dedicate this dissertation to my wife, Jessica, who has supported me immensely on this journey.

ACKNOWLEDGEMENTS

Many people in my life are responsible for pushing me and encouraging me in my mathematical journey. Mr. Curtis Dishner is largely responsible for sparking an interest in learning difficult mathematics as a high school student, and for that, I am grateful. Dr. Debra Gladden and Dr. Laura Singletary graciously took me under their wings as an undergraduate student and helped guide my scholarly life and invested in my personal life. I also thank Dr. Lex Oversteegen for his countless aid, advice, and patience with me as I navigated the research for this dissertation. Without him, I certainly would have been lost. Lastly, I wish to thank the many members of the Laminations Seminar (including Dr. Mayer, Dr. Blokh, and Dr. Selinger) who have gone before me and will continue after me. I will miss my time in the seminar and am thankful for the several years I was a part of it.

LIST OF FIGURES

TABLE OF CONTENTS

CHAPTER 1
INTRODUCTION

In the field of complex dynamics, a major topic of study is that of complex polynomials. On one hand, we are interested in the behavior of points (as well as collections of points) in the complex plane $\mathbb{C}$ under continued iteration of a complex polynomial P. On the other hand, we are also curious about how the set of complex polynomials of degree d can be viewed as a topological space. Thanks to the work of Thurston [11], we have a tool which can be used to help us understand the dynamics of polynomials as well as the space of degree d polynomials.

In order to understand the dynamics of a polynomial P, we split $\mathbb{C}$ up in the following way: 1) those points which run off to ∞ under continued iteration of P (the *basin of attraction to* ∞) and 2) those points whose forward orbit under P remains bounded (the *filled-in Julia set*, F_P). The common boundary of these two sets is called the *Julia set*, J_P, and the complement of the Julia set is known as the *Fatou set*. Breaking up $\mathbb{C}$ in this way allows us to isolate our study of points which behave in a *chaotic* fashion. The Julia set is chaotic in the sense that any open set around a point in the Julia set contains points with very different dynamics. In particular, some points in the open set will iterate off to ∞, some points may have orbits which converge to a *periodic orbit*, some points may themselves have a periodic orbit, and some points may have orbits with other phenomena.

The helpful tool that we use to study the dynamics of a polynomial is called a *lamination*, a closed collection of chords of the closed unit disk $\overline{\mathbb{D}}$ which do not intersect in $\mathbb{D}$. Special types of laminations on which we put a dynamical structure allow us to simplify our study of points in $\mathbb{C}$ which have more chaotic dynamics under a polynomial P. The construction of a lamination which models the dynamics of a degree d polynomial P first requires that P has a connected Julia set. This allows us to use the Riemann Mapping Theorem to find a conformal mapping $\Psi : \widehat{\mathbb{C}} \setminus \overline{\mathbb{D}} \to \widehat{\mathbb{C}} \setminus F_P$ where $\widehat{\mathbb{C}}$ is the Riemann sphere. Under the additional assumption that the Julia set of P is locally connected, we can extend Ψ to the boundary of $\widehat{\mathbb{C}} \setminus \overline{\mathbb{D}}$, i.e. the unit circle $\mathbb{S}$, and we get a map $\psi : \mathbb{S} \to J_P$ which, as it turns out, semi-conjugates the d-tupling map σ_d on $\mathbb{S}$ and $P|_{J_P}$. This defines a natural equivalence relation on $\mathbb{S}$ where x and y are related if $\psi(x) = \psi(y)$. We use the equivalence classes to determine which chords of $\overline{\mathbb{D}}$ will be included in the *q-lamination* $\mathcal{L}$ corresponding to P. In particular, for a given equivalence class, we take the boundary edges of the convex hull of that class and include those chords (now called *leaves*) in $\mathcal{L}$. We are able to extend σ_d to leaves of $\mathcal{L}$ as well as components of $\overline{\mathbb{D}} \setminus \bigcup \mathcal{L}$, which we call *gaps*. $\mathcal{L}$ is useful because if we collapse all the leaves of $\mathcal{L}$ to points, we obtain a combinatorial model for the Julia set of P in the sense that this quotient space is homeomorphic to J_P. The details of the construction of a lamination as well as the its dynamical properties are presented in the next chapter.

We see that q-laminations are constructed as a way to understand the dynamics of a specific polynomial with a locally connected Julia set. However, laminations (in particular laminations with dynamical properties which we call *invariant laminations*) have been greatly studied in order to understand their structure. One may ask if there is a benefit to studying invariant laminations in general. This thesis argues that there is. In particular, the first problem we solve is to understand when we can reverse the connection between a polynomial and a lamination. Instead of starting with a polynomial and constructing a q-lamination for that polynomial, can we start with a q-lamination $\mathcal{L}$ and find a polynomial which corresponds to $\mathcal{L}$ in the same sense as before? If this is possible, we say that we can realize $\mathcal{L}$ as a complex polynomial. We answer this in the affirmative by observing a special type of lamination that we call a *hyperbolic lamination.* A hyperbolic lamination is a lamination which imposes certain conditions on the *critical sets* (leaves which collapse to a point under σ_d or gaps which map forward under σ_d with degree greater than 1) of $\mathcal{L}$. In particular, we insist that all the critical sets must be *(pre)-periodic Fatou gaps* which eventually map to a critical set under iterations of σ_d. We make note that the realization of laminations as complex polynomials has been completed in a more general setting in [7] (in particular for laminations not containing any *Siegel gaps*). However, we present the argument in greater detail and completeness for the hyperbolic lamination case.

The motivation for looking at hyperbolic laminations is because we are able to apply *Thurston's Theorem*. Thurston's Theorem is a powerful existence theorem that states when a given *topological polynomial* (a branched covering map f from the Riemann sphere to itself such that $f^{-1}(\infty) = \{\infty\}$)) is equivalent (in a special sense of the word) to a complex polynomial g. The equivalence here is known as *Thurston equivalence*. It is weaker than topological conjugacy in that it gives us two homeomorphisms which are isotopic to one another and are both needed in order to commute f and g. Thurston's Theorem is helpful for us because for a given hyperbolic lamination $\mathcal{L}$, it has been shown in [1] that the map σ_d can be continuously extended over $\mathcal{L}$ and the rest of the Riemann sphere as a branched covering and hence as a topological polynomial. Such a topological polynomial given by a hyperbolic lamination satisfies the conditions of Thurston's Theorem.

However, we cannot simply stop after applying Thurston's Theorem. This is because at this stage, we have the given hyperbolic lamination $\mathcal{L}$, and we can find the corresponding q-lamination $\mathcal{L}_g$ generated by the complex polynomial g. The issue that needs to be resolved is showing that $\mathcal{L}$ and $\mathcal{L}_g$ are equal up to certain rotations of one another. It is not necessarily true that $\mathcal{L} = \mathcal{L}_g$, but there is only a small number (the number of σ_d fixed points on $\mathbb{S}$ to be precise) of ways they can differ up to rotation. In order to show this, we first prove that not only are f and g Thurston equivalent, but rather they are topologically conjugate in this case. From there, we go about showing that the Julia sets of f and g are homeomorphic which is sufficient for proving the claim regarding $\mathcal{L}$ and $\mathcal{L}_g$.

Once we know that we can realize hyperbolic laminations as complex polynomials, we are justified in studying such laminations and their properties. This is because we know there is a polynomial which can be modeled by the given lamination. Therefore, we explore periodic critical sets of invariant laminations to observe a few phenomena that occur since we know hyperbolic laminations must contain a periodic critical set. We focus our efforts, in this study, on degree 3 invariant laminations instead of the general degree d invariant laminations case due to the increase in complexity of the structure of critical sets for higher degree laminations.

Given a 3-invariant lamination $\mathcal{L}$ and a periodic critical gap G of $\mathcal{L}$ with period n, there are two important types of leaves on the boundary of G which can determine the gap completely. The first type is a *major* M of G which is a leaf closest (in length) to a critical chord. The second is a *refixed* leaf R which is a leaf such that it returns to itself under n iterates of σ_3. With these two types of leaves in mind, we will show that the major M is either periodic or eventually maps to the refixed leaf R under certain conditions. Under these same conditions, we also see that the leaves on the boundary of G consist of either one or two orbits of leaves. In particular, each leaf on the boundary of G maps eventually to either M or R, and we can have two different orbits of leaves if M and R themselves are not in each other's orbit. The key theorem used in making these conclusions is the Central Strip Lemma from [6] which determines the behavior of leaves as they grow and shrink in length as they map forward under σ_3

In the final chapter, we shift topics to that of *limit laminations*. This work is inspired by the work in [5]. In [5], the goal was to understand the space of degree 2 polynomials by understanding an appropriate topological space of degree 2 invariant laminations. In the general degree d case, the space used is the space of limits of d-invariant q-laminations, $\overline{\mathbb{L}_d^q}$, where the convergence of a sequence of q-laminations is determined by an appropriate Hausdorff metric. This space is natural in the sense that polynomials with locally connected Julia sets can be modeled with q-laminations.

Therefore, to understand the space of all degree d polynomials, we approximate a polynomial P with non-locally connected Julia set with polynomials P_i with locally connected Julia sets. We then observe the q-laminations $\mathcal{L}_i$ generated by the P_is and see what they converge to. It is known that the limit is an invariant lamination, but it is not necessarily a q-lamination.

The goal of the final chapter is to study some properties of limit laminations as a way to understand the space of invariant laminations. In particular, we study two properties: 1) stability of *rotational sets* and 2) existence of invariant objects. Rotational sets of $\mathbb{S}$ under σ_d were well studied and categorized in [8]. A rotational set A is a closed, invariant subset of $\mathbb{S}$ such that σ_d is order-preserving on A. For each rotational set, we have a unique rotation number that we can assign it. In a lamination, we have the corresponding notion of rotational gaps. We prove that a limit lamination $\mathcal{L}$ is *stable* in the sense that if $\mathcal{L}$ contains a rotational gap, then for all $\mathcal{L}_i$ sufficiently close to $\mathcal{L}$ $\mathcal{L}_i$ also contains a rotational gap. Furthermore, the rotation numbers of the $\mathcal{L}_i$s converge to the rotation number of $\mathcal{L}$.

Finally, we study the existence of invariant objects (either leaves or gaps) in 3-invariant laminations as well as in 3-invariant limit laminations. We are able to prove that a 3-invariant lamination $\mathcal{L}$ contains an invariant gap or an invariant leaf so long as $\mathcal{L}$ does not contain a critical leaf with a fixed endpoint. We then observe the case when $\mathcal{L}$ is a limit lamination and allow for a critical leaf with a fixed endpoint. In this case, we are able to determine when an invariant object exists.

CHAPTER 2

LAMINATIONS

In this chapter, we provide the basic definitions of laminations. We begin with laminations in general and then narrow our focus to invariant laminations. We then further refine our focus to that of q-laminations, specific types of invariant laminations, which will be what we use to model the dynamics of polynomials with locally connected Julia sets. The notions and definitions in this section are adapted from [9], [12], [10] and [4].

1. Invariant Laminations

Definition 2.1. A *lamination* $\mathcal{L}$ is a collection of chords in the closed unit disk $\overline{\mathbb{D}}$ such that the following hold:

(L1) no two chords in $\mathcal{L}$ intersect in $\mathbb{D}$

(L2) $\bigcup \mathcal{L}$ is closed.

If $\ell \in \mathcal{L}$ then we call ℓ a *leaf* of $\mathcal{L}$. If ℓ has endpoints a and b, then we denote ℓ by $\ell = \overline{ab}$.

Note that this definition allows for a leaf to intersect another leaf at one of its endpoints on the unit circle $\mathbb{S}$. If (L1) is not met by two chords, i.e. they intersect inside the open unit disk $\mathbb{D}$, then we say that those chords *cross*. (L2) ensures us that if there exists a sequence of leaves $\ell_i = \overline{a_i b_i} \in \mathcal{L}$ such that $a_i \to a$ and $b_i \to b$ for some $a, b \in \mathbb{S}$, then the chord $\ell = \overline{ab}$ must also be a leaf in $\mathcal{L}$. In this instance, we see that the leaves ℓ_i converge to ℓ. Another consequence of (L2) is that we must allow for single points in $\mathbb{S}$ to be considered as leaves. For instance, if a sequence of leaves

$\ell_i = \overline{a_i b_i}$ is such that $a_i \to x$ and $b_i \to x$ for some $x \in \mathbb{S}$, then x must be a leaf in $\mathcal{L}$, even though it is a singleton. If $x \in \mathcal{L}$ is a point, then we call x a *degenerate leaf.*

At this point in understanding laminations, we see that they are static objects. We now wish to incorporate dynamics with laminations. To do so, we introduce a simple map on $\mathbb{S}$, σ_d. Note that, for simplicity's sake, we identify $\mathbb{S}$ with $\mathbb{R}/\mathbb{Z}$. Hence, we view a point $z \in \mathbb{S}$ as being in the interval $[0, 1)$.

DEFINITION 2.2. Let $d \geq 2$ be an integer. The d-tupling map $\sigma_d : \mathbb{S} \to \mathbb{S}$, is defined as follows: $\sigma_d(z) = dz \pmod 1$.

In [1], it was carefully shown how to extend σ_d to a lamination $\mathcal{L}$. In particular, we can extend σ_d to a leaf $\ell = \overline{ab}$ in the following way.

DEFINITION 2.3. Let $\ell = \overline{ab}$ be a chord in $\overline{\mathbb{D}}$. Then, define $\sigma_d(\ell) = \overline{\sigma_d(a)\sigma_d(b)}$.

Note that if $\ell = \overline{ab}$ is such that $\sigma_d(a) = \sigma_d(b)$ (i.e. ℓ *collapses* to a single point under σ_d), then we call ℓ a *critical chord/leaf.* We are now ready to put dynamics on laminations.

DEFINITION 2.4. A lamination $\mathcal{L}$ is said to be a *sibling d-invariant lamination*, or simply a *d-invariant lamination*, if all of the following hold:

(D1) $\mathcal{L}$ is forward invariant; i.e., if $\ell \in \mathcal{L}$, then $\sigma_d(\ell) \in \mathcal{L}$

(D2) $\mathcal{L}$ is backward invariant; i.e., if $\ell \in \mathcal{L}$, then there is a leaf $\ell' \in \mathcal{L}$ such that $\sigma_d(\ell') \in \mathcal{L}$.

(D3) $\mathcal{L}$ is sibling invariant; i.e., if $\ell \in \mathcal{L}$ is not a critical leaf, then there exists d disjoint leaves $\ell_1, ..., \ell_d \in \mathcal{L}$ such that $\ell_1 = \ell$ and $\sigma_d(\ell_i) = \sigma_d(\ell)$ for all $i = 1, ..., d$.

We note here that (D1) and (D2) are not dependent upon ℓ being a nondegenerate leaf. These two properties still hold for single points of $\mathbb{S}$ in $\mathcal{L}$.

Definition 2.5. The collection of leaves $\{\ell_1, ..., \ell_d\}$ from (D3) above is called a *full sibling collection*. Given two distinct leaves ℓ_i and ℓ_j in this collection, we say that ℓ_i and ℓ_j are *siblings* or *sibling leaves*.

Another important notion regarding laminations in general is that of gaps.

Definition 2.6. Given a lamination $\mathcal{L}$, we define a *gap*, G, of $\mathcal{L}$ to be the closure of a component of $\overline{\mathbb{D}} \setminus \bigcup \mathcal{L}$. If $\ell \in \mathcal{L}$ is on the boundary of G, then we call ℓ an *edge* of G.

We can classify a gap G as finite or infinite by considering the set $\partial G \coloneqq G \cap \mathbb{S}$. If ∂G is infinite, then we say that G is an *infinite gap*. Otherwise, if ∂G is finite, we say that G is a *finite gap* or *polygon*. An infinite gap can be further described as either a *Fatou gap* (if ∂G is uncountable) or a *caterpillar gap* (if ∂G is countable). With ∂G in mind, we can now extend our notion of σ_d to a gap G. It is clear that G is equal to the convex hull, CH(∂G), of ∂G. Therefore, we define $\sigma_d(G) \coloneqq \text{CH}(\sigma_d(\partial G))$.

Definition 2.7. A lamination $\mathcal{L}$ is said to be *gap invariant* if for each gap G of $\mathcal{L}$, $\sigma_d(G)$ is either a gap, a leaf, or a point of $\mathcal{L}$. Moreover, if $\sigma_d(G)$ is a gap, then $\sigma_d|_{\text{Bd}(G)} : \text{Bd}(G) \to \text{Bd}(\sigma_d(G))$ maps as the composition of a monotone map and a covering map to the boundary of the image gap with positive orientation. The *degree* of σ_d on G is given by the number of preimages each leaf on the boundary of $\sigma_d(G)$ has on the boundary of G. If the degree of G is greater than 1, then we call G a *critical gap*.

The following result (Theorem 3.2 from [4]) demonstrates that sibling invariant laminations exhibit gap invariance.

Theorem 2.8. *Every d-invariant lamination is gap invariant.*

Thus, we often exploit gap invariance of sibling invariant laminations to help us understand their structure better.

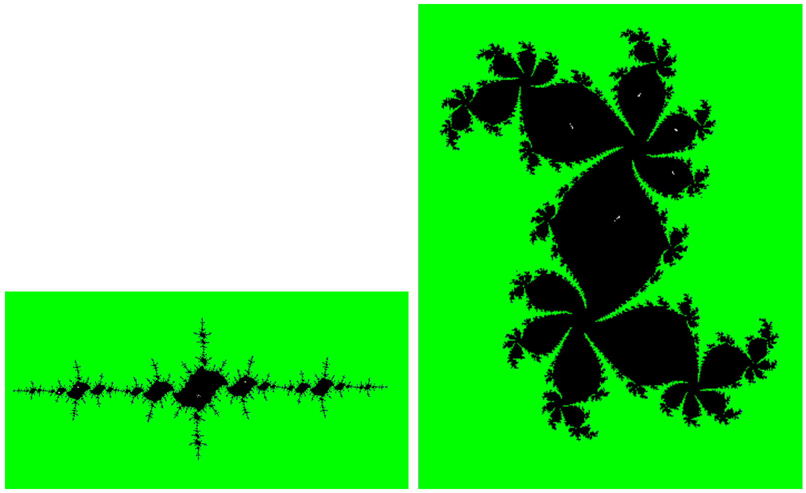

FIGURE 2.1. The Julia sets for two quadratic polynomials $P_1(z) = z^2 - 1.315 - .04i$ (left) and $P_2(z) = z^2 + .365 + .335i$. We see that altering the coefficents of a quadratic polynomial has a large impact on the structure of its Julia set.

2. q-laminations

We now wish to make the connection between polynomials and laminations. Let P be a complex polynomial of degree $d \geq 2$. Each polynomial P of degree d is affinely conjugate to a monic centered polynomial, i.e. of the type $z^d + a_{d-2}z^{d-2} + \cdots + a_0$. So, it is sufficient to study the space of monic centered polynomials. So, we consider P to be a monic centered polynomial. Our goal is to associate to P a special type of lamination called a q-lamination. We must first define a few notions with respect to P.

DEFINITION 2.9. The *basin of attraction to* ∞ is the set $B_\infty(P) = \{z \in \mathbb{C} \mid P^n(z) \to \infty \text{ as } n \to \infty\}$. Here we remark that the notation $P^n(z)$ represents iteration of P, i.e. $P^n(z) = P(P^{n-1}(z))$. The *Julia set* of P is the topological boundary of $B_\infty(P)$, and we denote it by J_P. The *filled-in Julia set* F_P of P is the collection of points in $\mathbb{C}$ whose forward orbits remain bounded under P, i.e. $F_P = \mathbb{C} \setminus B_\infty(P)$. The *Fatou set* of P is defined as $\mathbb{C} \setminus J_P$.

In Figure 2.1, the black region is the filled-in Julia set. The boundary of the black region is the Julia set. The green region on the outside of the black region is the basin of attraction to ∞.

There are several properties of Julia sets which motivate our definition and properties of laminations. Some of these important properties are as follows:

(1) J_P is fully invariant under P, i.e. $J_P = P(J_P) = P^{-1}(J_P)$.

(2) J_P is a perfect set.

(3) J_P is the topological boundary of F_P.

Before we can obtain a lamination that is associated with P, we first make the assumption that J_P is connected. Furthermore, we make the assumption that J_P is locally connected. Also, we denote the Riemann sphere by $\widehat{\mathbb{C}}$. Under these assumptions, we can apply the Riemann Mapping Theorem to obtain a conformal map $\Psi : \widehat{\mathbb{C}} \setminus \overline{\mathbb{D}} \to \widehat{\mathbb{C}} \setminus F_P$. Ψ can be chosen so that $\Psi(\infty) = \infty$, $\Psi'(\infty) > 0$, and Ψ conjugates $P|_{\widehat{\mathbb{C}} \setminus F_P}$ and the map $\theta_d(z) = z^d$ restricted to $\widehat{\mathbb{C}} \setminus \overline{\mathbb{D}}$, where d is the degree of P. The following commuting diagram depicts the setup thus far.

$$\begin{array}{ccc} \widehat{\mathbb{C}} \setminus \overline{\mathbb{D}} & \xrightarrow{\theta_d} & \widehat{\mathbb{C}} \setminus \overline{\mathbb{D}} \\ \downarrow \Psi & & \downarrow \Psi \\ \widehat{\mathbb{C}} \setminus F_P & \xrightarrow{P} & \widehat{\mathbb{C}} \setminus F_P \end{array}$$

Now, since J_P is assumed to be locally connected, a theorem of Carathéodory allows us to continually extend Ψ to the boundary of $\widehat{\mathbb{C}} \setminus \overline{\mathbb{D}}$, i.e. the unit circle $\mathbb{S}$. This gives rise to a map $\overline{\Psi}$ defined on $\widehat{\mathbb{C}} \setminus \mathbb{D}$. We now define $\psi : \mathbb{S} \to J_p$ as $\overline{\Psi}|_{\mathbb{S}}$. Define an equivalence relation $\sim_P$ on $\mathbb{S}$ by $x \sim_P y$ if and only if $\psi(x) = \psi(y)$. We see that ψ semi-conjugates the maps σ_d and $P|_{J_P}$ since Ψ conjugates θ_d and P. Thus $\sim_P$ is invariant. It is well known that $\sim_P$ classes are finite and have pairwise disjoint convex hulls. Define the *topological Julia set* to be $J_{\sim_P} = \mathbb{S}/\sim_P$, which is homeomorphic to J_P. σ_d induces a *topological polynomial* $f_\sim : J_{\sim_P} \to J_{\sim_P}$ which is topologically conjugate to $P|_{J_P}$. The following commutative diagram depicts this relationship (h is a homeomorphism between $J_{\sim P}$ and J_P).

$$\begin{array}{ccc} J_{\sim_P} & \xrightarrow{f_{\sim_P}} & J_{\sim_P} \\ \downarrow h & & \downarrow h \\ J_P & \xrightarrow{P} & J_P \end{array}$$

We now extract the properties of $\sim_P$ in order to study an equivalence relation on $\mathbb{S}$ in a more general sense without the need of starting with a given complex polynomial.

Definition 2.10. An equivalence relation $\sim$ on $\mathbb{S}$ is a *laminational equivalence relation* if all of the following hold:

(E1) the graph of $\sim$ is a closed subset of $\mathbb{S} \times \mathbb{S}$

(E2) convex hulls of distinct equivalence classes are disjoint

(E3) all equivalence classes are finite.

We now put dynamics on laminational equivalence relations.

Definition 2.11. A laminational equivalence relation $\sim$ is *(σ_d)-invariant* if all of the following hold:

(I1) $\sim$ is *forward invariant*: for a class $\mathbf{g}$, the set $\sigma_d(\mathbf{g})$ is a class too

(I2) $\sim$ is *backward invariant*: for a class $\mathbf{g}$, its preimage $\sigma_d^{-1}(\mathbf{g})$ is a union of classes

(I3) for any $\sim$-class $\mathbf{g}$, $\sigma_d(\mathbf{g})$ is either a point or a class with more than two points such that the map $\sigma_d|_{\mathbf{g}} : \mathbf{g} \to \sigma_d(\mathbf{g})$ is a *covering map with positive orientation*, i.e., for every connected component (s, t) of $\mathbb{S} \setminus \mathbf{g}$, the arc of the circle $(\sigma_d(s), \sigma_d(t))$ is a connected component of $\mathbb{S} \setminus \sigma_d(\mathbf{g})$.

Laminational equivalence relations give rise to a special type of laminations called *q-laminations*.

Definition 2.12. A lamination $\mathcal{L}$ is called a *q-lamination* if $\sim_{\mathcal{L}}$ is a laminational equivalence relation and $\mathcal{L}$ consists exactly of boundary edges of the convex hulls of $\sim_{\mathcal{L}}$-classes.

Thus, we can construct q-laminations from laminational equivalence classes. In particular, in our previous discussion with a complex polynomial P, we saw that ψ induced an equivalence relation on $\mathbb{S}$ having the properties of an invariant laminational equivalence relation. Therefore, given a complex polynomial with locally connected Julia set, we can associate to it a q-lamination.

We can also think of q-laminations in a different way. In particular, how can we verify if a given lamination $\mathcal{L}$ is a q-lamination? We do so by defining an equivalence relation $\sim_{\mathcal{L}}$ on $\mathbb{S}$ by $x \sim_{\mathcal{L}} y$ if and only if there exists a finite concatenation of leaves of $\mathcal{L}$ joining x and y. If $\sim_{\mathcal{L}}$ is an invariant laminational equivalence relation and $\mathcal{L}$ is exactly consisting of boundary edges of convex hulls of $\sim_{\mathcal{L}}$, then $\mathcal{L}$ is a q-lamination.

Another helpful tool that helps us translate the locally connected Julia set of a polynomial into a q-lamination is that of external rays.

Definition 2.13. Let P be a polynomial of degree d with locally connected Julia set and Ψ the associated Riemann map. Then, for a given angle θ, we define the *external ray* as $R_\theta = \{\Psi(re^{2\pi i\theta}) : r > 1\}$.

External rays are the Ψ images of radial rays outside of $\overline{\mathbb{D}}$. If $\lim\limits_{r\to 1^+} \Psi(re^{2\pi i\theta})$ exists, then we say that R_θ *lands* at the limit point. This landing point is known to be in the Julia set J_P. From this viewpoint, we can construct a lamination equivalence relation on $\mathbb{S}$ by declaring $x \sim y$ if x and y correspond to two angles whose corresponding external rays land at the same point. In turn, this gives rise to a q-lamination as argued before.

In Figure 2.2, we see that three external rays (corresponding to angles 1/7, 2/7, and 4/7) land at the same point. Therefore, in the corresponding lamination $\mathcal{L}$, the points at those three angles must be joined together by leaves. If we consider quotienting $\mathcal{L}$ by collapsing all leaves to points, we obtain a model for the Julia set.

The following theorem (Lemma 3.18 of [4] makes the connection between sibling invariant laminations and q-laminations.

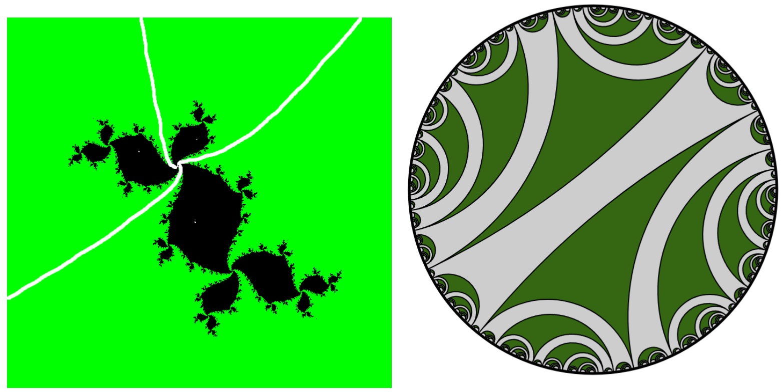

FIGURE 2.2. The Julia set for the quadratic polynomial $P(z) = z^2 - .08 + .73i$ (left) and the corresponding q-lamination (right). There are three external rays landing at the same point in the Julia set of P.

THEOREM 2.14. *d-invariant q-laminations are sibling invariant.*

CHAPTER 3

REALIZATION OF HYPERBOLIC LAMINATIONS AS COMPLEX POLYNOMIALS

It is well understood how to construct an invariant q-lamination given a complex polynomial with locally connected Julia set. In this section, we wish to go the other direction. That is, given a q-lamination, can we find a complex polynomial whose corresponding lamination is the given lamination? As it turns out, we can put assumptions on the given lamination to ensure this. We note that the given lamination (with these assumptions) and the polynomial's lamination are equal up to rotation by a fixed point of $\mathbb{S}$ under σ_d. To begin this study, we first define a few general terms. In this section, we will denote the Riemann sphere by S^2.

1. Thurston Equivalence

Definition 3.1. A branched covering map $f : S^2 \to S^2$ is called a *topological polynomial* if ∞ is a critical point and $f^{-1}(\infty) = \{\infty\}$.

Note that, while a topological polynomial f does not necessarily have any analytic structure, we can still observe its dynamical behavior. In particular, for a topological polynomial f, we can define the basin of attraction to ∞ and therefore the Julia set of f.

Definition 3.2. Let f be a topological polynomial. The *Julia set* of f is the boundary of the basin of attraction to ∞, $B_\infty(f) = \{z \in S^2 \mid f^n(z) \to \infty \text{ as } n \to \infty\}$.

The key property we use to ensure the realization of a given lamination is based upon knowing how the critical sets look like and how they map. In particular, from here on we will assume that a given lamination $\mathcal{L}$ is such that all of its criticality is located inside of Fatou gaps, i.e there are no critical chords in $\mathcal{L}$ and there are no finite critical gaps in $\mathcal{L}$. Moreover, we assume that for each critical gap G of $\mathcal{L}$, G is preperiodic and contains a critical gap in its forward orbit. Such invariant laminations are called *hyperbolic laminations*, and we record this in the following definition.

DEFINITION 3.3. Suppose $\mathcal{L}$ is a d-invariant q-lamination which does not contain any critical chords. Furthermore, suppose all compatible critical chords are interior to (pre)-periodic Fatou gaps. If each critical gap G of $\mathcal{L}$ is such that $\sigma_d^n(G)$ is a critical gap for some n, then we call $\mathcal{L}$ a *hyperbolic lamination*.

We state an important theorem from [1] that relates laminations and topological polynomials.

THEOREM 3.4. *Let $\mathcal{L}$ be a hyperbolic lamination. Then there is a branched covering map $\sigma_d^\# : \overline{\mathbb{D}} \to \overline{\mathbb{D}}$, a continuous extension of σ_d, that maps $\mathcal{L}$ to $\mathcal{L}$.*

This map $\sigma_d^\#$ can be further extended to all of S^2 by defining the map to be equal to $z \mapsto z^d$ on $S^2 \setminus \overline{\mathbb{D}}$. We call this map $f : S^2 \to S^2$. Clearly, $f^{-1}(\infty) = \{\infty\}$ so that f is a topological polynomial. In [1], the branch points of f are shown to be the centers of critical gaps of $\mathcal{L}$. Moreover, f maps centers of gaps to centers of gaps. These facts motivate us to utilize known results regarding special types of topological polynomials with properties f satisfies.

DEFINITION 3.5. Let $f : S^2 \to S^2$ be a branched covering map with critical set C_f. The *postcritical set of* f is the set $P_f = \bigcup_{n=1}^{\infty} f^n(C_f)$. If P_f is a finite collection of points, then we say that f is *postcritically finite*.

DEFINITION 3.6. A *Thurston map* $f : S^2 \to S^2$ is a postcritically finite branched covering map. We say that two Thurston maps f and g are *Thurston equivalent* if there exists homeomorphisms $\theta_1, \theta_2 : S^2 \to S^2$ such that the following diagram commutes:

$$\begin{array}{ccc} (S^2, P_f) & \xrightarrow{\theta_2} & (S^2, P_g) \\ \big\downarrow f & & g \big\downarrow \\ (S^2, P_f) & \xrightarrow{\theta_1} & (S^2, P_g) \end{array}$$

where P_f, P_g are the postcritical sets of f and g, respectively, $\theta_1(P_f) = \theta_2(P_f) = P_g$, and θ_1 is isotopic to θ_2 relative P_f (denoted by $\theta_1 \sim \theta_2$).

Thurston equivalence between a topological polynomial and complex polynomial is the important relationship that will allow us to realize a hyperbolic lamination as a complex polynomial. However, given a topological polynomial f which corresponds to a hyperbolic lamination, we need to know if there exists a complex polynomial g which satisfies Thurston equivalence with f. The work of Bielefeld, Fisher, and Hubbard in [2] as well as Thurston provides remarkable results which help us. We provide definitions and results without proof from [2].

Definition 3.7. Let f be a Thurston map. Suppose there exists a collection of curves $\Lambda = \{\gamma_0, ..., \gamma_k = \gamma_0\}$ such that for each $i = 0, ..., k-1, \gamma_i$ is homotopic rel P_f to exactly one component of $f^{-1}(\gamma_{i+1})$ and $f : \gamma' \to \gamma_{i+1}$ has degree 1. Then Λ is called a *Levy cycle.*

Levy cycles are a subset of a larger class of multicurves with respect to a topological polynomial called *Thurston obstructions* which we do not describe.

Lemma 3.8. *If $f : S^2 \to S^2$ is a rational function, then f cannot have a Levy cycle.*

Lemma 3.9. *If a topological polynomial f has a Thurston obstruction, then f has a Levy cycle.*

We now state Thurston's Theorem (located in [2]).

Theorem 3.10. *A postcritically finite topological polynomial $f : S^2 \to S^2$ is Thurston equivalent to a complex polynomial if and only if there are no Thurston obstructions.*

Putting these lemmas together with Thurston's theorem, we have the following theorem.

THEOREM 3.11. *Let f be a topological polynomial corresponding to a hyperbolic lamination $\mathcal{L}$. Then, f is Thurston equivalent to a complex polynomial g.*

PROOF. Since $\mathcal{L}$ is hyperbolic, f is postcritically finite. Furthermore, f cannot have a Levy cycle by Lemma 3.8, and hence by (the contrapositive of) Lemma 3.9 f has no Thurston obstructions. Therefore, by Theorem 3.10, there exists a complex polynomial g such that f and g are Thurston equivalent. □

2. Conjugating f and g

While it is remarkable to guarantee the existence of a complex polynomial g that we can associate with a topological polynomial f, Thurston equivalence alone is not sufficient in helping us relate the dynamics of f and g. The following theorem strengthens the relationship between f and g. Not only are they Thurston equivalent, but they are also topologically conjugate.

THEOREM 3.12. *Suppose f and g are Thurston equivalent where f is a topological polynomial corresponding to a hyperbolic lamination $\mathcal{L}$ and g is a complex polynomial guaranteed by Theorem 3.11. Suppose there exists a metric on $S^2 \setminus (P_f \cup P_g)$ such that f is uniformly expanding on $S^2 \setminus P_f$ and g is uniformly expanding on $S^2 \setminus P_g$. Then, f and g are topologically conjugate.*

PROOF. Note that a desired metric exists by [9]. An intuitive understanding of this metric is that, near each critical point, distances between points close to the critical point grow as they approach the critical point. Our goal is to extend the above commuting diagram upward as to obtain a sequence of homeomorphisms which converges to a conjugating map between f and g. Before doing so, we need appropriate compact subsets, K_g and K_f, of S^2 to work with. The purpose of these compact spaces is that we can bound the diameters of paths in these spaces. After constructing a conjugating homeomorphism between these smaller spaces, we will then be able to extend the homeomorphism to all of S^2. To construct K_g, remove pairwise disjoint disks around each point in the grand orbit of the critical points of g from S^2. In order for θ_1 and θ_2 to remain defined on this new space, we must remove the θ_1-preimages and θ_2-preimages of these disks from S^2, giving us the compact space K_f. We will now consider the maps $f|_{K_f}$, $g|_{K_g}$, $\theta_1|_{K_f}$, and $\theta_2|_{K_f}$. We will construct the conjugating homeomorphism $\theta_\infty : K_f \to K_g$. The following diagram is helpful in understanding this construction.

$$\begin{array}{ccc}
K_f & \xrightarrow{\theta_\infty} & K_g \\
\downarrow{\scriptstyle f} & & {\scriptstyle g}\downarrow \\
K_f & \xrightarrow{\theta_\infty} & K_g \\
 & \vdots & \\
K_f & \xrightarrow{\theta_3 \sim \theta_4} & K_g \\
\downarrow{\scriptstyle f} & & {\scriptstyle g}\downarrow \\
K_f & \xrightarrow{\theta_2 \sim \theta_3} & K_g \\
\downarrow{\scriptstyle f} & & {\scriptstyle g}\downarrow \\
K_f & \xrightarrow{\theta_1 \sim \theta_2} & K_g
\end{array}$$

Note that the bottom portion of this diagram is the same as the diagram given by Thurston equivalence but restricted to K_f and K_g (with the exception of the presence of θ_3). We begin by constructing the map $\theta_3 : K_f \to K_g$. Our goal is to ensure that $\theta_2 \sim \theta_3$, i.e. θ_2 and θ_3 are isotopic.

For each $t \in [1, 2]$, let $\theta_t^1 : S^2 \to S^2$ denote a member the family of homeomorphisms forming an isotopy between θ_1 and θ_2 where $\theta_1^1 = \theta_1$ and $\theta_2^1 = \theta_2$. Now, let $s \in K_f$. Run $f(s)$ through the isotopy maps θ_t^1 to form a path I_1 with endpoints $\theta_1(f(s))$ and $\theta_2(f(s))$. That is, for each $t \in [1, 2]$, we have $\theta_1^t(f(s)) \in I_1$. We now wish to uniquely lift this path, I_1, by pulling it back under an appropriate branch of g^{-1}. Using the commutativity of the diagram, we choose to lift $\theta_1(f(s))$ to $\theta_2(s)$ under g^{-1}. With the branch of the inverse chosen, we can now lift all of I_1 to a unique path I_2

with endpoints $\theta_2(s)$ and $g^{-1}(\theta_2(f(s)))$ (using the Path-Lifting Property). In fact, we lift the entire isotopy θ_1^t via g^{-1} to a new isotopy. This defines a homeomorphism $\theta_3 := g^{-1} \circ \theta_2 \circ f$ and an isotopy via families of homeomorphisms θ_t^2 where $t \in [2,3]$ with $\theta_2^2 = \theta_2$ and $\theta_3^2 = \theta_3$. This completes the above diagram's bottom portion.

We now build the next levels of the diagram inductively in a similar manner. In particular, suppose $\theta_{i-1} \sim \theta_i$ for $i \geq 2$ are such that the following diagram commutes.

$$\begin{array}{ccc} K_f & \xrightarrow{\theta_i} & K_g \\ \Big\downarrow f & & g \Big\downarrow \\ K_f & \xrightarrow{\theta_{i-1}} & K_g \end{array}$$

Let $s \in K_f$. Given an isotopy θ_{i-1}^t with $t \in [i-1, i]$ such that $\theta_{i-1}^{i-1}(s) = \theta_{i-1}(s)$ and $\theta_{i-1}^{i}(s) = \theta_i(s)$, let f(s) run through θ_{i-1}^t for each t. This yields a path, I_{i-1}, with endpoints $\theta_{i-1}(f(s))$ and $\theta_i(f(s))$. We can now define θ_{i+1} by taking the g-preimage of $\theta_{i-1}(s)$ equal to $\theta_i(s)$ (using the assumption that the above diagram commutes). Again, using the Path-Lifting Property, we can now uniquely lift I_{i-1} under g^{-1} to a path with endpoints $\theta_i(s)$ and $g^{-1}(\theta_i(f(s)))$ where the branch of the inverse of g is uniquely determined. We define $\theta_{i+1} \coloneqq g^{-1}(\theta_i(f(s)))$.

We claim that the sequence $\{\theta_i\}_{i=1}^{\infty}$ converges uniformly to a map θ_∞. To this end, we show that the sequence θ_i is uniformly Cauchy and therefore converges to a continuous map θ_∞. Let $\lambda > 1$ be the uniform expanding factor for g. Let $M = \sup_{x \in S^2}(\mathrm{diam}\{\gamma_1(x) \mid \gamma_1(x)$ is a path from $\theta_1(x)$ to $\theta_2(x)\}$. Then, for any i, if γ_i is a path from $\theta_i(x)$ to $\theta_{i+1}(x)$ obtained by pulling some $\gamma_1(x)$ back under g $i-1$ times, then $\mathrm{diam}(\gamma_i) \leq \frac{M}{\lambda^{i-1}}$ (since g is uniformly expanding by a factor of λ). Fix $\varepsilon > 0$. Choose N such that $\frac{M}{\lambda^{N-2}(\lambda-1)} < \varepsilon$. This is possible since $\lambda > 1$ and M is fixed. Now, let $n, m > N$ and $x \in S^2$. We will show that $|\theta_m(x) - \theta_n(x)| < \varepsilon$. Without loss of generality, we will assume that $n > m$. Then,

$$
\begin{aligned}
|\theta_m(x) - \theta_n(x)| &\leq \sum_{i=m}^{n-1} |\theta_i(x) - \theta_{i+1}(x)| \\
&\leq \sum_{i=m}^{n-1} \frac{M}{\lambda^{i-1}} \\
&< \sum_{i=m}^{\infty} \frac{M}{\lambda^{i-1}} \\
&= \frac{M}{\lambda^{m-2}(\lambda - 1)} \\
&< \frac{M}{\lambda^{N-2}(\lambda - 1)} \\
&< \varepsilon.
\end{aligned}
$$

Therefore, we see that the sequence of θ_i converges uniformly, and thus θ_∞ is continuous. We can, in a similar manner as above, construct maps θ_i^{-1} which converge uniformly to a continuous map θ_∞^{-1}. This is completed by using the inverse maps θ_1^{-1} and θ_2^{-1} and the following commuting diagram.

$$
\begin{array}{ccc}
K_f & \xleftarrow{\theta_2^{-1}} & K_g \\
\Big\downarrow f & & g \Big\downarrow \\
K_f & \xleftarrow{\theta_1^{-1}} & K_g
\end{array}
$$

The fact that θ_∞ is a bijection follows from the fact that θ_1 and θ_2 themselves are bijections. At each stage in the construction we constructed the next θ_i by uniquely determining the lift of each point under g^{-1} so that θ_i and θ_i^{-1} are in fact inverses of one another. Hence, θ_∞ and θ_∞^{-1} are inverses of one another, making θ_∞ a homeomorphism. Therefore, we see that f and g are topologically conjugate on K_f and K_g via the map θ_∞.

Now, we enlarge our spaces K_f and K_g by taking out smaller disks than we did before around the grand orbit of the critical points of g (and the appropriate θ_1 and θ_2-preimages of these disks) from S^2. We run the exact same arguments as above to extend θ_∞ on this larger space, noting that θ_∞ is defined the exact same way on K_f

and K_g in this new space. We continue inductively by running the arguments on a sequence of spaces whose union is the complement of the grand orbit of the critical points of g on which we have the map θ_∞ defined. We can then continuously extend this map to all of S^2 to obtain the map $\theta_\infty : S^2 \to S^2$, which conjugates f and g on S^2. The following commuting diagram depicts this relationship.

$$\begin{array}{ccc} S^2 & \xrightarrow{\theta_\infty} & S^2 \\ \downarrow f & & g \downarrow \\ S^2 & \xrightarrow{\theta_\infty} & S^2 \end{array}$$

□

3. Equal Laminations up to Rotation

THEOREM 3.13. *Consider the maps f, g, and θ_∞ from above. Then, $\theta_\infty(J_f) = J_g$.*

PROOF. We will show that the basin of infinity, $B_\infty(f)$, of f maps onto the basin of infinity, $B_\infty(g)$ of g under θ_∞. The result will follow since $J_f = \partial(B_\infty(f))$ and $J_g = \partial(B_\infty(g))$ and boundaries are sent to boundaries under continuous mappings.

Let $x \in B_\infty(f)$ and suppose by way of contradiction that $\theta_\infty(x) \notin B_\infty(g)$. Thus, $\theta_\infty(x) \in F_g$, the filled-in Julia set of g. Now, let $V \subset B_\infty(g)$ be an open set with $\infty \in V$. Let U be the component of $\theta_\infty^{-1}(V)$ containing ∞. Then, for some n, $f^n(x) \in U$. Thus, $\theta_\infty(f^n(x)) \in \theta_\infty(U) \subset V \subset B_\infty(g)$. On the other hand, since F_g is invariant under g, $g^n(\theta_\infty(x)) \in F_g$. Now, by using the commutativity of f and g via θ_∞, we see that $\theta_\infty(f^n(x)) = g^n(\theta_\infty(x))$. This is a contradiction since $F_g = S^2 \setminus B_\infty(g)$. Thus, $\theta_\infty(B_\infty(f)) \subset B_\infty(g)$. Since θ_∞ is a homeomorphism, similar arguments show that $B_\infty(g) \subset \theta_\infty(B_\infty(f))$. The result follows. □

Before proceeding, we will state a few helpful and well-known lemmas found in [9]. Throughout these lemmas, the map we are considering is a monic polynomial, P, with connected Julia set, J_P. The following is Lemma 18.9 of [9].

LEMMA 3.14. *If a fixed ray $R_t = P(R_t)$ lands at z_0, then z_0 is either a repelling or a parabolic fixed point.*

The next lemma is Lemma 18.12 of [9].

Lemma 3.15. *If one periodic ray lands at z_0, then only finitely many rays land at z_0, and these rays are all periodic of the same period (which may be larger than the period of z_0).*

Theorem 3.16. *Suppose f and g are as in the above. Then, the laminations associated with f and g are equal up to rotation by a fixed point.*

Proof. Let $\mathcal{L}_f$ and $\mathcal{L}_g$ be the laminations corresponding to f and g, respectively. Choose the fixed point $z_0 = f(z_0) \in J_f$ such that the external ray with angle 0, R_0, lands at z_0. By Lemma 3.15, the only other rays that can land at z_0 are finitely many fixed rays. In $\mathcal{L}_f$, R_0 corresponds to the fixed point $0 \in \mathbb{S}$. Now, since θ_∞ conjugates f and g, $\theta_\infty(R_0)$ is a fixed ray under g. Thus, $\theta_\infty(R_0)$ corresponds to a σ_d-fixed point $x_g \in \mathbb{S}$. Note that we view x_g as a part of $\mathcal{L}_g$.

Now we will rotate $\mathcal{L}_g$ so that the fixed point x_g becomes 0. Call this rotated lamination $\widehat{\mathcal{L}_g}$. This process induces a map $\widehat{\theta}_\infty : \mathbb{S} \to \mathbb{S}$. Our first goal is to show that $\widehat{\theta}_\infty$ is the identity map on $\mathbb{S}$. By the construction, we set $\widehat{\theta}_\infty(0) = 0$. Next, we show that we can construct $\widehat{\theta}_\infty$ to be the identity on all σ_d^n-preimages of 0. This follows from the fact that our maps are order-preserving. That is, if $x_1 \in \mathbb{S}$ is the first preimage of 0 after 0 in the positive circular order, then in the image, $\widehat{\theta}_\infty(x_1)$ is the first preimage of 0 located after 0 in the positive circular order. We do this for all first preimages of 0. Then, we use the same reasoning to inductively conclude that $\widehat{\theta}_\infty$ is the identity on $\sigma_d^{-n}(0)$ for all n. This identity map on a dense subset of $\mathbb{S}$ uniquely extends continuously to the identity map on all of $\mathbb{S}$.

Now our goal is to show that $\mathcal{L}_f = \widehat{\mathcal{L}_g}$. Indeed, let $\ell = \overline{ab} \in \mathcal{L}_f$. Then, we can approximate a by a sequence $\{a_i\}_{i=1}^\infty$ of preimages of 0. Similarly, we can approximate b by a sequence $\{b_i\}_{i=1}^\infty$ of preimages of 0. Each a_i and b_i corresponds to external rays R_{a_i} and R_{b_i}, respectively. The R_{a_i} rays converge to the external ray R_a corresponding to a. The R_{b_i} rays converge to the external ray R_b corresponding to b. Note that

$R_a \neq R_b$ since $\overline{ab}$ is a leaf of $\mathcal{L}_f$. Furthermore, these rays land at a common point x_{ab}. Mapping these rays forward under θ_∞, we see that the sequence of $\theta_\infty(R_{a_i})$s converges to $\theta_\infty(R_a)$ and the sequence of $\theta_\infty(R_{b_i})$s converges to $\theta_\infty(R_b)$. We see that then $\theta_\infty(R_a)$ and $\theta_\infty(R_b)$ must land at the point $\theta_\infty(x_{ab})$. This holds true since the Julia set of g is locally connected. Now, $\theta_\infty(R_a)$ and $\theta_\infty(R_b)$ correspond to points $a_g, b_g \in \mathbb{S}$ in $\mathcal{L}_g$. There must be a leaf connecting $\ell_g = \overline{a_g b_g} \in \mathcal{L}_g$ since $\theta_\infty(R_a)$ and $\theta_\infty(R_b)$ land at a common point. After rotating ℓ_g by the same rotation as we rotated $\mathcal{L}_g$ to $\widehat{\mathcal{L}_g}$, we see that $\ell \in \widehat{\mathcal{L}_g}$. This follows since the induced map $\widehat{\theta}_\infty$ on $\mathbb{S}$ is the identity. □

This concludes our discussion on the realization of hyperbolic laminations as complex polynomials. To recap, given a hyperbolic lamination $\mathcal{L}$, there exists a complex polynomial P such that its corresponding lamination $\mathcal{L}_P$ equals $\mathcal{L}_r$ where either $\mathcal{L}_r = \mathcal{L}$ or $\mathcal{L}_r$ is a rotation of $\mathcal{L}$ by a fixed point of σ_d.

CHAPTER 4
THE STRUCTURE OF PERIODIC CRITICAL SETS OF 3-INVARIANT LAMINATIONS

Now that we know that hyperbolic laminations are realizable as complex polynomials, a natural question to ask is, "What properties do hyperbolic laminations have?" For such laminations, the restrictions on the critical sets give us reasons to study periodic critical sets. Given a hyperbolic lamination $\mathcal{L}$, since each critical set G must eventually map to a critical set (which are finite in number) under σ_d, there must exist a periodic critical set in $\mathcal{L}$. In this section, we focus our efforts on understanding periodic critical sets for 3-invariant laminations. The reason for this focus is because the larger the degree of an invariant lamination is, the more complicated the types and structures of its critical sets gets.

We begin by observing the possible types of full sibling collections for leaves in a 3-invariant lamination $\mathcal{L}$. Suppose $\ell = \overline{ab} \in \mathcal{L}$ is a nondegenerate leaf. The *length*, $||\ell||$, of ℓ is given by the minimum of the lengths of the arcs (a, b) and (b, a). The goal is to understand the sizes and layout of a full sibling collection of leaves which map to ℓ. Consider the σ_3-preimages a_1, a_2, a_3 of a and the σ_3-preimages b_1, b_2, b_3 of b. Since σ_3 is locally expanding by a factor of 3 and the maximum distance between two points on $\mathbb{S}$ is $\frac{1}{2}$, for each a_i, there is a b_j such that $|b_j - a_i| \leq \frac{1}{6}$. For simplicity, we will assume that $a_1 < b_1 < a_2 < b_2 < a_3 < b_3$ and $|b_i - a_i| \leq \frac{1}{6}$ for all i. Now, there are a few ways in which we can connect the preimages of a to the preimages of b to form a pairwise disjoint collection of leaves mapping to ℓ. Each of these preimages falls under one of three categories as presented in the next definition.

1. Leaf Length

Definition 4.1. Let $\mathcal{L}$ be a 3-invariant lamination. Let $\ell \in \mathcal{L}$ be a nondegenerate leaf. Then, ℓ is one of the following:

(1) If $0 < ||\ell|| < \frac{1}{6}$, we say ℓ is *short*.
(2) If $\frac{1}{6} \leq ||\ell|| < \frac{1}{3}$, we say ℓ is *medium*.
(3) If $\frac{1}{3} < ||\ell|| \leq \frac{1}{2}$, we say ℓ is *long*.

For completeness, we say that a leaf of length $\frac{1}{3}$ is a *critical leaf*.

With this classification of leaf lengths, we can now see the three possible ways to configure the preimages of a given leaf. Let $\mathcal{L}$ be a 3-invariant lamination and $\ell \in \mathcal{L}$ a nondegenerate leaf. If ℓ_1, ℓ_2, and ℓ_3 are such that $\sigma_d(\ell_i) = \ell$ for $i = 1, 2, 3$, then one of the following holds:

(1) ℓ_1, ℓ_2, and ℓ_3 are all short (we denote this case by SSS)
(2) ℓ_1, ℓ_2, and ℓ_3 are all medium (we denote this case by MMM)
(3) each of ℓ_1, ℓ_2, and ℓ_3 are of a different type, namely, one is long, one is medium, and one is short (we denote this case by LMS).

Under σ_3, leaves grow and shrink in length according to the *leaf length function*, which is given by the following formula:

$$||\sigma_3(\ell)|| = \begin{cases} 3||\ell|| & 0 \leq ||\ell|| \leq \frac{1}{6} \\ |3||\ell|| - 1| & \frac{1}{6} < ||\ell|| \leq \frac{1}{2} \end{cases}$$

Figure 4.1 is the graph of the leaf length function and is helpful for knowing when a leaf grows or shrinks. The dotted diagonal is the identity map. The points of intersection of the leaf length function and the diagonal indicate the fixed leaf lengths. Leaves less than $\frac{1}{4}$ in length must grow under σ_3 while leaves longer than $\frac{1}{4}$ must shrink.

Definition 4.2. Let $\mathcal{L}$ be a d-invariant lamination. If $\ell \in \mathcal{L}$ is a leaf closest to critical in length among all leaves in $\mathcal{L}$, then we call ℓ a *main major* or a *major of*

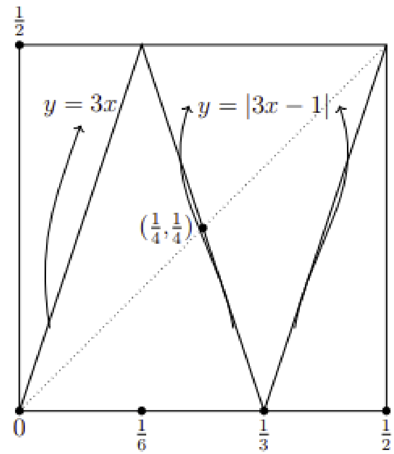

FIGURE 4.1. The leaf length function's graph for degree 3. The x-axis corresponds to a leaf's length while the y-axis corresponds to the length of the leaf's image under σ_3. Image from [12].

$\mathcal{L}$. If U is a gap of $\mathcal{L}$ and $M \in \mathcal{L}$ is a leaf on the boundary of U which is closest to critical in length among all leaves on the boundary of U, then we call M a *major of* U. The image of M, $\sigma_d(M)$ is called a *minor*. Given a major M of a gap U, if $\sigma_d^i(M)$ is the iterate of M which is closest to critical in length among all of its forward orbit, then we say that $\sigma_d^i(M)$ is the *primary major* of M.

DEFINITION 4.3. Let $\mathcal{L}$ be a d-invariant lamination and U a periodic gap of period n. If ℓ is a leaf (possibly degenerate) on the boundary of U such that $\sigma_d^n(\ell) = \ell$, then we say that ℓ is a *refixed leaf*.

2. Short Strips

DEFINITION 4.4 (Short Strips and Long Strips). Let $\mathcal{L}$ be a d-invariant lamination. Suppose $\ell \in \mathcal{L}$ is a leaf which is not critical. Consider all the siblings $\ell_i \in \mathcal{L}$ of ℓ. Consider a component, C, of $\overline{\mathbb{D}} \setminus \bigcup \ell_i$. If C has a nondegenerate arc of $\mathbb{S}$ of length less than $\frac{1}{2d}$ on its boundary, then C is called a *short strip of* ℓ. If C has exactly one sibling of ℓ on its boundary and $||\ell|| < \frac{1}{2d}$, we call C a *degenerate short strip of* ℓ. If C has an arc of $\mathbb{S}$ of length greater than $\frac{1}{2d}$ on its boundary, then C is called a *long strip of* ℓ. If C has exactly one sibling of ℓ on its boundary and $||\ell|| > \frac{1}{2d}$, we call

C a *degenerate long strip of* ℓ. We will use the terms short/long strips to indicate nondegenerate short/long strips unless otherwise noted.

DEFINITION 4.5. Let $\mathcal{L}$ be a lamination with no critical leaves. Consider all critical sets of $\mathcal{L}$. For a critical set U, let M be a major on the boundary of U. We say that a U is a *short gap* if the siblings of M (of length greater than $\frac{1}{2d}$) on the boundary of U form a short strip. Otherwise, U is a *long gap*.

DEFINITION 4.6. Let $\mathcal{L}$ be a d-invariant lamination. A *collapsing polygon*, P, is a collection of leaves in $\mathcal{L}$ which forms a finite polygon in $\overline{\mathbb{D}}$ and $\sigma_d(P)$ is a non-degenerate leaf in $\mathcal{L}$. Note that P is not necessarily a gap of $\mathcal{L}$, for there may be a leaf in the interior of P.

DEFINITION 4.7. Let $\ell = \overline{ab}$ be a leaf of a lamination $\mathcal{L}$. We define the *open leaf* corresponding to ℓ as $\ell^\circ = \ell \setminus \{a, b\}$. Denote the critical leaves in $\mathcal{L}$ (if there are any) by $\overline{c}_i(\mathcal{L}) = \overline{c}_i$.

The following lemma is a result from [4] and tells us when it is necessary for a given d-invariant lamination to contain a collapsing polygon.

LEMMA 4.8. *Let $\mathcal{L}$ be a d-invariant lamination and $\ell = \overline{ab} \in \mathcal{L}$ be a leaf. If C is a component of $(\sigma_d)^{-1}(\ell) \setminus \cup_i \overline{c}_i^\circ$ and G is the convex hull of $\partial(C)$, then G is a leaf or a collapsing polygon of $\mathcal{L}$.*

COROLLARY 4.9. *Suppose $\mathcal{L}$ is a lamination containing no critical leaves. If two distinct leaves, $\ell_1, \ell_2 \in \mathcal{L}$ such that $\sigma_d(\ell_1) = \sigma_d(\ell_2)$, share a common endpoint, then ℓ_1 and ℓ_2 are contained in a collapsing polygon.*

LEMMA 4.10. *Short strips and long strips are well-defined, i.e. if S is a short strip, then it is not a long strip. Furthermore, if L is a long strip, then it is not a short strip.*

Proof. Let S be a short strip of ℓ. We prove that all of the circle arcs on the boundary of S are *short* (of length less than $\frac{1}{2d}$). Suppose $\sigma_d(\ell) = \overline{ab}$ where, without loss of generality, the length of the circle arc (a, b) is equal to $\varepsilon < \frac{1}{2}$.

Let (a_1, b_1) be an arc of length less than $\frac{1}{2d}$ which is on the boundary of S. Order the preimages of a and b in the positive circular order so that $a_1 < b_1 < a_2 < b_2 < ... < a_d < b_d < a_1$ and for each $1 \leq i \leq d$, the circular arc (a_i, b_i) has length equal to $\frac{\varepsilon}{d}$. Note that for all i, the circular arc (b_i, a_{i+1}) has length $\frac{1-\varepsilon}{d}$ (since the length of (b, a) is $a - \varepsilon$). Since (a_1, b_1) is a circle arc on the boundary of S and we have a full collection of preimages of a and b, there is a leaf ℓ_1 on the boundary of S having b_1 as an endpoint. The other endpoint of ℓ_1 is a_{k_1} for some $k_1 \neq 1$. By Corollary 4.9, no two leaves on the boundary of S meet at an endpoint. Therefore, there is not another leaf on the boundary of S coming out of a_{k_1}. Therefore, the short circular arc (a_{k_1}, b_{k_1}) is on the boundary of S. Indeed, this circular arc is a maximal arc because there is necessarily a sibling of ℓ_1 with b_{k_1} as one of its endpoints. In a similar fashion, let ℓ_2 be the leaf on the boundary of S which has b_{k_1} as an endpoint. The other endpoint of ℓ_2 is a_{k_2} for some $k_2 \neq k_1$. By the same reasoning, (a_{k_2}, b_{k_2}) is a maximal short circle arc on the boundary of S. Continuing by induction, we see that all circle arcs on the boundary of S are short. Thus, there are no arcs of $\mathbb{S}$ of length greater than $\frac{1}{2d}$ on the boundary of S, and so S is not a long strip.

It immediately follows that, if L is a long strip of ℓ, then there are no arcs of $\mathbb{S}$ of length less than $\frac{1}{2d}$ on its boundary. This is because if there was one such arc, then each arc on the boundary of L would also be less than $\frac{1}{2d}$, a contradiction that L has at least one arc of $\mathbb{S}$ of length greater than $\frac{1}{2d}$. □

A key result in understanding how a leaf maps in relation to its own short strip is given by Theorem 2.9 of [6], the Central Strip Lemma, which can be used for d-invariant laminations. We will refer to this lemma as the Short Strip Lemma, as the name Central Strip Lemma is derived from its original form in the degree $d = 2$ case where a short strip corresponds to a strip containing the origin (and is hence a *central strip*). We now present this lemma.

LEMMA 4.11. *Let C be a short strip of leaf ℓ and its siblings with 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 as $\frac{\eta}{d^{j-k}}$ to a critical chord in $\overline{\mathbb{D}} \setminus C$.

Throughout the remainder of this chapter, we utilize the Short Strip Lemma in order to know when a leaf can enter its own short strip as a short leaf and when it must map into its own short strip as a non-short leaf.

3. Periodic Critical Sets

Consider a short critical gap U. U is either a critical polygon or a Fatou gap of degree greater than 1. If U is a critical polygon, then it is a finite gap which may map to a gap or a leaf. In the case that U is a Fatou gap, then U is (pre)-periodic. We are interested in the case when U is a periodic Fatou gap.

LEMMA 4.12. *Suppose U is a short critical gap of a d-invariant lamination, $\mathcal{L}$, which does not contain any critical leaves or diameters. Then, $\sigma_d(U)$ has a unique minor.*

PROOF. Suppose, by way of contradiction, there are two minors m_1 and m_2. By definition of minors, $|m_1| = |m_2|$. Since U is a short gap, $\sigma_d(U)$ is under both m_1 and m_2, implying m_1 is under m_2 and m_2 is under m_1. The only way for this to occur is if the two minors coincide, a contradiction. $\square$

Remark 4.13. Now, we will order the critical short gaps of $\mathcal{L}$ by the length of their minors. Consider the lengths of the minors of all critical short gaps of $\mathcal{L}$. Label these lengths as ε_i so that $\varepsilon_1 < \varepsilon_2 < ... < \varepsilon_l$. For some i, there may be multiple minors with length ε_i. For each i, we call each of the corresponding minors a *level i minor* and denote them by m_i^k for $1 \leq k \leq j$ where j is the number of distinct minors of length ε_i. If m_i^k is a level i minor of a critical short gap, then we label that gap $U_i^{k,n}$ and call it a *level i gap*. Here, the index n is used to distinguish critical short gaps which have the same minor. Note that $\sigma_d(U_1^{i,n}) \neq \sigma_d(U_1^{k,m})$ for $i \neq k$. This is because $U_1^{i,n}$ and $U_1^{k,m}$ must map under m_1^i and m_1^k, respectively, and if they mapped to the same gap, then m_1^i would coincide with m_1^j, a contradiction.

Our goal is to observe how leaves on the boundary of periodic Fatou gaps of degree greater than 1 must map.

Theorem 4.14. *Suppose U_1 is a periodic level 1 critical gap. Then, a major, M_1, on the boundary of U_1 (or one of its siblings on the boundary of U) is a refixed leaf. Moreover, if ℓ is a leaf on the boundary of any gap in the orbit of U_1, then ℓ must eventually map to the refixed leaf (M_1 or one of its siblings on the boundary of U).*

Proof. Suppose U_1 is a periodic level 1 critical gap of period n. Consider a major M_1 on the boundary of U_1 as well as its siblings which are also on the boundary of U_1. As we map these leaves forward, they can never map into the short strip of M_1 short. Otherwise, M_1 would have had to map to a leaf which was closer to critical in length than M_1, a contradiction to the choice of M_1. Therefore, M_1 and its siblings on the boundary of U_1 must each map to either M_1 or one of its siblings under σ_d^n. This implies that there exists a refixed leaf among M_1 and the siblings of M_1 on the boundary of U_1. Moreover, each leaf on the boundary of U_1 must grow until one of its iterates is not short and thus eventually maps to M_1 or its sibling on the boundary of U_1, implying there is one orbit of leaves on the boundary of U_1. $\square$

We now discuss the properties of level 2 critical sets in the case when $d = 3$.

Theorem 4.15. *Let U_2 be a periodic level 2 short critical gap of a 3-invariant lamination $\mathcal{L}$ such that $\mathcal{L}$ does not contain any diameters or critical leaves. Then, there is a major on the boundary of U_2 which is either periodic or eventually maps to a non-degenerate refixed leaf. Moreover, there are at most two orbits of leaves on the boundary of U_2.*

Proof. Fix a 3-invariant lamination $\mathcal{L}$ with no diameters or critical leaves. Let U_1 be a degree 2 level 1 gap and U_2 be a degree 2 level 2 gap of $\mathcal{L}$. First, suppose that U_1 and U_2 are periodic and are in the same orbit. Let n be the period of U_1 and let M_1 and M_1' be a major and its sibling, respectively, of U_1. When U_1 returns to itself after n iterates, M_1 and M_1' can only map to either M_1 or M_1'. This is because M_1 is a leaf closest to critical in all of $\mathcal{L}$, and so it cannot enter its own short strip as a short leaf. Therefore, either M_1 or M_1' is a refixed leaf. Without loss of generality, let M_1 be the refixed leaf. Now consider a major M_2 of U_2. At the moment U_2 maps to U_1, M_2 can only map to either M_1 or M_1'. M_2 cannot map to a short leaf on the boundary of U_1 since the short strip of M_1 is narrower than the short strip of M_2. First, suppose M_2 maps to M_1. Then, M_2 itself is a refixed leaf. If, instead, M_2 maps to M_1', then M_2 is not refixed. Rather, there is a short leaf on the boundary of U_2 which is refixed. Now, since short leaves on the boundary of U_1 must grow to a length greater than $\frac{1}{6}$ under some iterate of σ_3^n, they must all eventually map to either M_1 or M_1'. Similarly, all short leaves on the boundary of U_2 must eventually map to M_2 or M_2'. Therefore, combining these results, we see that there is only one orbit of leaves on the boundaries of U_1 and U_2. The following theorem gives a further restriction on the leaves on the boundaries of U_1 and U_2.

Now consider the critical periodic gap U_2 with period n as before; however, suppose U_1 and U_2 are not in the same periodic orbit. In fact, U_1 may not even be periodic. First, suppose M_2 is the primary major in its own orbit. Then, by similar reasoning as before, when U_2 returns to itself under σ_3^n, M_2 must map either to itself or its sibling, M_2', on the boundary of U_2 (this is since M_2 is the closest to critical in length in its forward orbit by assumption). Thus, either M_2 or M_2' is a refixed leaf. Also as before, each short leaf on the boundary of U_2 must grow under σ_3^n so that it maps to either M_2 or M_2'. Therefore, there is only one orbit of leaves on the boundary of U_2.

Now suppose the major M_2 of U_2 is not the primary major in the orbit of M_2. Let $\sigma_3^i(M_2) = M_2^i$ be the primary major in the orbit of M_2 and let $U_2^i = \sigma_3^i(U_2)$. U_2^i can be situated in two ways. It is either contained in a short strip of M_2^i or in a long strip of M_2^i. In either case there is a leaf, N^i, distinct from M_2^i on the boundary of U_2^i which is not short. There are at most two possible choices for N^i, but we choose N^i to be the leaf closer to critical in length. We call the other leaf D.

First consider the case when U_2^i is in a short strip of M_2^i. It is necessary that U_2^i is contained in a long strip of N^i. Similar to before, each short leaf on the boundary of U_2^i must grow under σ_3^n until it is greater than $\frac{1}{6}$ in length. Thus, each of these short leaves map to either M_2^i or N^i. Furthermore, when U_2^i returns to itself under σ_3^n, M_2^i must also return to itself and is thus refixed. This is because M_2^i cannot map short into its own short strip (otherwise we would contradict the fact that M_2^i is the primary major in the orbit of M_2). By the same reason, M_2^i cannot map to N^i (since N^i is closer to critical than M_2^i). Therefore, either M_2 or its sibling M_2' is a refixed leaf on the boundary of U_2. Without loss of generality, we will say that M_2 is refixed. Now observe how N^i can map. Eventually, its short iterates will begin to grow in length under σ_3^n so that N^i eventually maps either to itself or M_2^i. If N^i maps M_2^i, then we have that the orbit of U_2 contains exactly one orbit of leaves (since each leaf eventually maps to M_2^i and thus to M_2). If N^i eventually maps back to itself, then, there are exactly two orbits of leaves on the boundary of U_2 (namely those leaves which map to M_2 and those which map to N, where N is the σ_3^{-i} preimage of N^i on the boundary of U_2 which is periodic).

Now consider the case when U_2^i is contained in a long strip of M_2^i. U_2^i is contained in either a long strip of N^i or a short strip of N^i.

First suppose that U_2^i is contained in a long strip of N^i. We first prove that the primary majors of both N^i and D are both further to critical than M_2^i if they do not map to a non-short sibling of M_2^i. Indeed, suppose without loss of generality that N^i eventually gets closer to critical than M_2^i. For this to happen, N^i would eventually have to map to a non-short leaf inside a short strip of M_2^i. At this moment, the gap U_2^i would map inside this short strip which is narrower than the short strip of M_2^i. This contradicts the fact that M_2^i is the primary major in its own orbit. This is because M_2^i can't map to the non-short leaves since they would be closer to critical than M_2^i, and M_2^i cannot map to a short leaf at this moment because a primary major cannot enter its own short strip as a short leaf. Therefore, we have shown that M_2^i must be closer to critical than any leaves in the orbits of N^i and D. In particular, M_2^i must be closer to critical than N^i. Therefore, in this case, we have that U_2^i is in the long strips of both M_2^i and N^i and N^i is further to critical than M_2^i. Now, each short leaf on the boundary of U_2^i must grow under σ_3^n so that it eventually maps to one of the non-short leaves on the boundary of U_2^i, namely M_2^i, N^i, or D. We claim that N^i (or D) must map to M_2 or M_2' since, if not, the primary major of N^i (or D) must be farther from critical length than M_2 and, hence cannot map to a short leaf of U_2. Hence, either M_2^i maps to M_2 (or M_2'), and that leaf is refixed or M_2^i eventually maps to N^i or D and M_2 (or M_2') is periodic.

Lastly, suppose U_2^i is contained in a short strip of N^i. Under σ_3^n, all short leaves on the boundary of U_2^i must eventually grow to a length of at least $\frac{1}{6}$ and hence all leaves on the boundary of U_2^i eventually map to either to M_2^i or N^i. Therefore, N_i must also eventually map to itself or to M_2^i. If N^i maps to itself, then N^i is a refixed leaf. M_2^i either eventually maps to itself or to N_i as well. If M_2^i maps to itself, then M_2^i is periodic. So, U_2 has a periodic major on its boundary. Thus, there can be up to two orbits of leaves on the boundary of U_2. If M_2^i maps to N^i, then the leaves on the boundary of U_2^i form one orbit of leaves.

Now suppose N^i eventually maps to M_2^i. If M_2^i eventually maps to N^i, then M_2^i (and hence a major of M_2) is periodic, and there is one orbit of leaves on the boundary of M_2. Similarly, if M_2^i eventually maps to itself, then there is a periodic major on the boundary of U_2, and there is only orbit of leaves on its boundary. □

COROLLARY 4.16. *If the refixed leaf on the boundary of U_2 is degenerate, then a major of U_2 is periodic.*

PROOF. This follows immediately from Theorem 4.15 since M_2 is either periodic or eventually maps to a nondegenerate refixed leaf on the boundary of U_2. □

THEOREM 4.17. *Let U be a periodic Fatou gap of a 3-invariant lamination of period n whose leaves on its boundary eventually map to the same periodic leaf $M \in U$, i.e. the leaves on the boundary of U form one orbit. Then, no two leaves on the boundary of U intersect.*

PROOF. Suppose by way of contradiction that ℓ and ℓ' are on the boundary of U and they intersect at a point. Then, $\sigma_3^k(\ell) = M$ for some k and $\sigma_3^{k'}(\ell') = M$ for some k'. Assume without loss of generality that $k < k'$. Notice that M and $\sigma_3^k(\ell')$ also intersect at a point. Under iterates of σ_3^n, $\sigma_3^k(\ell')$ will never map to M as it will always only intersect M at a point. This contradicts the assumption that there is only one orbit leaves which all map to M eventually. □

4. Fat Gaps

We now turn our attention to *fat gaps* of 3-invariant laminations. The location of the origin will play a key role in the characterization of such gaps.

LEMMA 4.18. *Suppose ℓ and ℓ' is a pair of medium and long siblings of a 3-invariant lamination $\mathcal{L}$. Then, the origin is not contained in the interior of the short strip between ℓ and ℓ'.*

Proof. To simplify the argument, we will assume that the short sibling of ℓ and ℓ' is the leaf $\overline{\frac{-1}{12}\frac{1}{12}}$. We do this since $\frac{1}{6}$ is the longest a short leaf can be. The two sibling points of $\frac{-1}{12}$ are $\frac{1}{4}$ and $\frac{7}{12}$. The two sibling points of $\frac{1}{12}$ are $\frac{5}{12}$ and $\frac{3}{4}$. Therefore, it is clear that the long sibling is a diameter which separates the medium and short siblings. Now, if we suppose that the short sibling was shorter than $\frac{1}{6}$, then the long sibling would be located further away from the short sibling than in the first case. Thus, the origin is located in the region of the disk between the long and short siblings. The conclusion of the lemma follows. ☐

Definition 4.19. Let $\mathcal{L}$ be a d-invariant lamination. We say that a gap U of $\mathcal{L}$ is a *fat gap* if its image $\sigma_d(U)$ is not under any leaf of $\mathcal{L}$.

Lemma 4.20. *If U is a fat gap which contains the origin, then it is invariant.*

Proof. Since $\sigma_d(U)$ is not under any leaves of $\mathcal{L}$, it also contains the origin. Thus, since gaps map onto gaps, we have that $U = \sigma_d(U)$ as desired. ☐

Corollary 4.21. *If U is a fat critical set, then it is invariant.*

Proof. Since U is critical, each leaf on its boundary has a sibling which is also on the boundary of U. Since U is a fat gap, it is not contained in a short strip of any leaf. Therefore, there is a short/long sibling pair of leaves on the boundary of U. By Lemma 4.18 U contains the origin. Therefore, U is invariant by Lemma 4.20. ☐

Lemma 4.22. *If U is a fat gap, then $\sigma_d(U)$ contains the origin.*

Proof. This follows directly from the definition of fat gaps. ☐

Theorem 4.23. *Suppose U_2 is a periodic critical level 2 gap. Then, no forward σ_d-image of U is a fat gap.*

Proof. Suppose, by way of contradiction, That $\sigma_3^i(U_2) = U_2^i$ is a fat gap. There are two cases to consider: 1) U_2^i is contained under a leaf of either U_2 or a level 1 critical set U_1 or 2) U_2^i is contained in the component of $\overline{\mathbb{D}} \setminus (U_1 \cup U_2)$ between U_1 and U_2.

In either case, $\sigma_3(U_2^i)$ will contain the origin by Lemma 4.22. Furthermore, $\sigma_3(U_2^i)$ must be located in a short strip with two sibling leaves (one medium and one long) on its boundary. Otherwise, $\sigma_3(U_2^i)$ would be invariant by Lemma 4.20. It is not the case that U_2^i is contained in a short strip with three medium siblings on its boundary. Therefore, U_2^i is a critical set. This is a contradiction to the fact that all criticality has been accounted for in U_1 and U_2. □

Definition 4.24. Given a d-invariant lamination $\mathcal{L}$, the unique gap of $\mathcal{L}$ containing the origin is called the *central gap*.

Theorem 4.25. *Let $\mathcal{L}$ be a d-invariant lamination. Then, the boundary of the central gap O of $\mathcal{L}$ cannot consist only of short leaves.*

Proof. By way of contradiction, suppose that all leaves on the boundary of O are short. First, observe that $\sigma_d(O) \neq O$. This is because each leaf on the boundary of O must eventually grow to a length larger than $\frac{1}{2d}$ under σ_d, contradicting that each leaf is short. Since O is not forward invariant, O must map under a leaf on the boundary of O. This is again a contradiction to the fact that each leaf on the boundary of O must grow. Therefore, O must contain leaves of length greater than $\frac{1}{2d}$. □

We now classify all fat critical sets of a given 3-invariant lamination.

Theorem 4.26. *Let $\mathcal{L}$ be a 3-invariant lamination. If G is a fat critical set in $\mathcal{L}$, then G is a quadratic invariant gap, i.e. an invariant gap which maps to itself with degree 2 under σ_3.*

Proof. Suppose $\mathcal{L}$ contains a fat critical set G. Then, $\sigma_3(G) = O$ where O is the central gap of $\mathcal{L}$. This proof hinges on the number and type of non-short leaves on the boundary of O. By Lemma 4.25, there is at least one non-short leaf on the boundary of O. Therefore, we consider the following three cases: 1) there is exactly one long leaf and no medium leaves on the boundary of O, 2) there are two long leaves on the boundary of O, and 3) there is a medium leaf on the boundary of O. We will show that case 1 is only possible when G is a quadratic invariant gap. We will show that cases 2 and 3 are not possible.

Case 1) Let L be the long leaf on the boundary of O. We first claim that O must be critical. By Lemma 4.18, the origin is located in the region of $\overline{\mathbb{D}}$ between L and its short sibling L'. Suppose by way of contradiction that O is not critical. Then, there is a short leaf, ℓ' on the boundary of O which separates L and L', i.e. L' is under ℓ'. Note that the sibling portrait for ℓ' cannot be SSS. For if that was the case, then the siblings of ℓ' would cross L and L's medium sibling. Therefore, ℓ' has a long sibling ℓ. Again, by Lemma 4.18, the origin is located between ℓ and ℓ'. Moreover, since L is under ℓ', ℓ is located outside the short strip of L. Therefore, ℓ separates L and ℓ'. This is a contradiction since L and ℓ' are both on the boundary of the gap O. Therefore, O is a critical set.

Now we claim that O is invariant. $\mathcal{L}$ contains a second critical set H (distinct from O). H is located in the short strip between L and L's medium sibling, since this is where the remaining criticality is located. Therefore, H is not a fat critical set (since it maps under $\sigma_3(L)$). Therefore, the given fat critical set G must coincide with the critical set O. Since O is a fat set, its image is itself. Thus, O is invariant. Therefore, O is a quadratic invariant gap.

Case 2) Let L_1 and L_2 be long leaves on the boundary of O. Then, L_1 and L_2 have medium siblings M_1 and M_2, respectively, which are located in the complement of O. Therefore, there is a critical set located in the short strip between L_1 and M_1, and there is a critical set located in the short strip between L_2 and M_2. Since both critical sets are inside short strips, neither of them are fat gaps. Therefore, there are no fat critical gaps in a lamination whose central gap contains two long leaves on its boundary.

Case 3) Let M be a medium leaf on the boundary of O. The sibling portrait for M is MMM. This is because if the sibling portrait was LMS, then by Lemma 4.18, M would not be on the boundary of O. Therefore, all criticality in $\mathcal{L}$ is located in the cubic short strip of M. Thus, the image of any critical set in $\mathcal{L}$ is located under $\sigma_3(M)$. Therefore, there are no fat critical sets in this case.

Thus, if G is a fat critical set, then G is a quadratic invariant gap. $\square$

CHAPTER 5

LIMIT LAMINATIONS

Current research in the field of polynomial dynamics is concerned with understanding the space of degree d polynomials. The end goal of this research is to be able to model the connectedness locus of the space of all degree d polynomials via understanding an appropriate space of d-invariant laminations. This chapter is motivated by work completed in [5]. In that paper, it is described that an appropriate space to use in studying d-invariant laminations is that of the space of degree d *limit laminations*. In this chapter, we describe the notion of limit laminations and prove a few properties for such laminations for degree d as well as a few properties for the degree 3 case.

As we have seen, we are able to associate to each complex polynomial with locally connected Julia set a q-lamination. However, we would also like to associate an invariant lamination to complex polynomials with non-locally connected Julia sets. To this end, consider a degree d polynomial P with non-locally connected Julia set. Then, perturb P so that we have a sequence of polynomials P_i, each with a locally connected Julia set, such that $P_i \to P$. Since the Julia set of each P_i is locally connected, we have the associated q-lamination $\mathcal{L}_i$ for each P_i. If the the $\mathcal{L}_i$s converge in an appropriate metric, then the resulting limit is a d-invariant lamination, $\mathcal{L}$, but is not necessarily a q-lamination. Note that the resulting *limit lamination* $\mathcal{L}$ is not necessarily unique; it depends on how we perturb the polynomial P. The work in [5] resolves this issue in the $d = 2$ case.

In what follows, our goal is to study two things with regards to limit laminations: 1) how a limit lamination and its close by q-laminations are related (with respect to periodic and rotational gaps) and 2) how 3-invariant limit laminations can look like with respect to the presence of invariant objects. We begin with some formal definitions and notation.

Definition 5.1. Let $\mathbb{L}_d^q$ be the family of all d-invariant q-laminations. Let $\overline{\mathbb{L}_d^q}$ be the closure of $\mathbb{L}_d^q$ in the compact space of all subcontinua of $\overline{\mathbb{D}}$ with the Hausdorff metric. We call elements of $\overline{\mathbb{L}_d^q}$ *limit laminations*.

To elaborate on the metric on the space of laminations, given a lamination $\mathcal{L}$, consider each leaf $\ell \in \mathcal{L}$ as a point in the set of all subcontinua of $\overline{\mathbb{D}}$ (denoted $C(\overline{\mathbb{D}})$). Thus, $\mathcal{L}$ corresponds to a compact subset $K \subset C(\overline{\mathbb{D}})$. Suppose we have $\mathcal{L}_1$ and $\mathcal{L}_2$ which correspond to the compact subsets $K_1, K_2 \subset C(\overline{\mathbb{D}})$, respectively. Then we can measure the distance between K_1 and K_2 via the Hausdorff metric. Thus, if K_1 and K_2 are ε-close, then we say that $\mathcal{L}_1$ and $\mathcal{L}_2$ are ε-close.

Any intuitive way of measure the "closeness" of two laminations $\mathcal{L}_1$ and $\mathcal{L}_2$ is by seeing if for every $\ell_1 \in \mathcal{L}_1$, there is a leaf $\ell_2 \in \mathcal{L}_2$ which is close to ℓ_1 and vice versa.

The reason for this perhaps cumbersome way of defining a metric on the space of laminations (instead of just using the Hausdorff metric on $\bigcup \mathcal{L}$) is to bypass an issue that arises in the degree 3 case. In particular, there are two 3-invariant laminations, known as the *vertical lamination* and *horizontal lamination*. The vertical lamination consists of no gaps and each leaf is vertical in $\overline{\mathbb{D}}$. The horizontal lamination also has no gaps and each leaf is horizontal. Clearly these two laminations are very different and present different dynamics. Thus, with the metric we choose, we want these two laminations to have a nonzero distance between them. However, if we simply use the Hausdorff on the union of leaves of the laminations, then we see that their distance from each other is 0 (since the union of their leaves both give us $\overline{\mathbb{D}}$).

The reason to use limit laminations rather than q-laminations is two fold. First, it allows us to assign limit laminations to polynomials with non-locally connected Julia sets. The second is that we want to use the Hausdorff topology, as described above, on the set of limit laminations to correspond to the topology on the set of polynomials. This later fact fails if we use q-laminations. However, by requiring that the limit laminations we use are limits of non-constant q-laminations, we will achieve this goal. For example, our space of limit lamination will not contain q-laminations with cubic Fatou gaps nor q-laminations with two distinct cycles of degree two Fatou gaps. For such polynomials we will instead assign appropriate limit laminations to them.

1. Rigidity Properties of Limit Laminations

In [5], one property that is well explored is that of *rigidity*. We briefly describe this concept as well as some important results regarding rigidity which are proved in that paper. The main focus in the section is on how periodic and rotational gaps of a limit lamination affect the structure of close by q-laminations.

DEFINITION 5.2. A leaf (or gap) G of a limit lamination $\mathcal{L} \in \overline{\mathbb{L}_d^q}$ is *rigid* if any q-laminations close to $\mathcal{L}$ has G as its leaf (or gap).

The following is Lemma 2.5 of [5].

LEMMA 5.3. *Let $\mathcal{L} \in \overline{\mathbb{L}_d^q}$ and let $\ell = \overline{ab}$ be a periodic leaf of $\mathcal{L}$ that is not an edge of a gap of $\mathcal{L}$. Then ℓ is rigid.*

Given a gap G of a limit lamination $\mathcal{L} \in \overline{\mathbb{L}_d^q}$, we would like to find a corresponding gap $G(\widehat{\mathcal{L}})$ in a given close by q-lamination $\widehat{\mathcal{L}}$. This is because, if done appropriately, we can conclude that $G(\widehat{\mathcal{L}})$ has certain properties. We simply define $G(\widehat{\mathcal{L}})$ as the gap of $\widehat{\mathcal{L}}$ such that the area of $G(\widehat{\mathcal{L}}) \cap G$ is greater than half the area of G and $G(\widehat{\mathcal{L}})$ is close to G in the Hausdorff metric (if such a gap exists). Otherwise, set $G(\widehat{\mathcal{L}}) = \emptyset$. We can now state part of Lemma 2.5 and Lemma 2.10 of [5].

LEMMA 5.4. *Let $\mathcal{L} \in \overline{\mathbb{L}_d^q}$, and let G be a gap of $\mathcal{L}$. Then for any lamination $\widehat{\mathcal{L}} \in \mathbb{L}_d^q$ close to $\mathcal{L}$, the gap $G(\widehat{\mathcal{L}})$ is non-empty and such that $G(\widehat{\mathcal{L}}_i) \to G$ as $\widehat{\mathcal{L}}_i \to \mathcal{L}$.*

LEMMA 5.5. *Suppose that G is a periodic Fatou gap of a lamination $\mathcal{L} \in \overline{\mathbb{L}_d^q}$. If no image of G has critical edges, then G is rigid.*

2. Stability of Rotational Sets

We now wish to make a claim regarding limit laminations with a rotational gap. This first requires definitions of lifts, order-preserving, and rotational sets. From there, we can assign a rotation number to rotation sets. Then, we are able to see how limit laminations containing rotational gaps must relate to nearby q-laminations. Note that these definitions are adapted from [3].

Definition 5.6. Let $e : \mathbb{R} \to \mathbb{S}$ be the map $e(x) = e^{2\pi i}$. Consider a map $f : \mathbb{S} \to \mathbb{S}$. A *lift* of f is defined to be a function $\widehat{f} : \mathbb{R} \to \mathbb{R}$ such that $(e \circ \widehat{f})(x+m) = (f \circ e)(x)$ for all $x \in [0,1]$ and $m \in \mathbb{Z}$.

Definition 5.7. A map $g : \mathbb{R} \to \mathbb{R}$ is *degree 1* if $g(x+1) = g(x)+1$ for all $x \in \mathbb{R}$. A map $f : \mathbb{S} \to \mathbb{S}$ is *degree 1* if there exists a degree 1 lift $\widehat{f} : \mathbb{R} \to \mathbb{R}$ of f.

Definition 5.8. Let $f : \mathbb{S} \to \mathbb{S}$ be a map. Then f is *order-preserving* if f is monotone and of degree 1. If $A \subset \mathbb{S}$ is closed, then f is *order-preserving* on A if $f|_A$ can be extended to an order-preserving map $F : \mathbb{S} \to \mathbb{S}$ such that $F|_A = f|_A$.

Definition 5.9. Let $f : \mathbb{S} \to \mathbb{S}$ be a map and A be an invariant subset of $\mathbb{S}$ under f. Then we say that A is *rotational* if A is closed and f is order-preserving on A.

Definition 5.10. Given a degree 1 monotone map $f : \mathbb{S} \to \mathbb{S}$, let the *rotation number of* f be defined as $\rho(f) = \lim_{n\to\infty} \frac{\widehat{f}^n(\widehat{x})}{n} \pmod 1$ where $x \in \mathbb{S}$, $\widehat{x}$ is any preimage of x under e, and $\widehat{f}$ is any lift of f.

Definition 5.11. Given a rotational set A and map $f : \mathbb{S} \to \mathbb{S}$, we can now assign a unique *rotation number* $\rho(A)$, to A as follows: $\rho(A) = \lim_{n\to\infty} \frac{\widehat{F}^n(\widehat{a})}{n} \pmod 1$ where $\widehat{a} \in A$ and $\widehat{F}$ is any lift of an order-preserving extension F of $f|_A$.

It is known that the rotation number of A is well-defined, i.e. the rotation number is the same regardless of our choice of $\widehat{a}$ and $\widehat{F}$. The notion of a map being order-preserving on a set A has an equivalent form which is perhaps more useful when discussing laminations.

The following lemma is Theorem 2.2.1 of [8], the so-called Hole Theorem.

Lemma 5.12. *Let A be a closed subset of $\mathbb{S}$. Then, σ_d is order-preserving on A if and only if $\mathbb{S} \setminus A$ contains $d-1$ pairwise disjoint open intervals, each of length $\frac{1}{d}$.*

The $d-1$ pairwise disjoint open intervals from Lemma 5.12 are called *holes*.

Lemma 5.13. *Suppose G is a periodic gap of a d-invariant lamination, $\mathcal{L}$, with period n. Then G is rotational under σ_d if and only if σ_d^n is of degree one on G.*

Proof. First, suppose G is rotational under σ_d. Then, G maps back to itself order-preserving. Thus, by Lemma 5.12, $G \cap \mathbb{S}$ is contained in arcs of the circles such that the sum of their arc lengths is less than or equal to $\frac{1}{d}$. In other words, all criticality in $\mathcal{L}$ is located in the holes of G. Therefore, σ_d is of degree one on G. The other direction is a symmetric argument. $\square$

This next result on limits of sibling collections is Lemma 3.19 of [4].

Lemma 5.14. *Take sequences of d sibling leaves $\overline{a_i^j b_i^j}, 1 \leq j \leq d, i = 1, 2, \ldots$ such that $\overline{a_i^1 b_i^1} \to \overline{a^1 b^1} = \ell^1, \overline{a_i^2 b_i^2} \to \overline{a^2 b^2} = \ell^2, \ldots, \overline{a_i^d b_i^d} \to \overline{a^d b^d} = \ell^d$ and $\sigma_d(\ell^1)$ is not degenerate. Then $\ell^j, 1 \leq j \leq d$ are siblings with non-degenerate image.*

Before we state and prove the following Theorem 5.16, we first describe a construction and state a result by [8] regarding a parameterization for σ_3-rotational sets. Given a σ_3-rotational set G, G has two disjoint holes of length $\frac{1}{3}$. Starting from 0 and moving in the counterclockwise direction, let a be the first endpoint of the first hole. Then, let b be the second endpoint of the second hole. The collection of these parameter pairs (a, b) is then compactified to give a parameter space, Δ_3 for space of σ_3-rotational sets. Consider any such parameter pair $(a, b) \in \Delta_3$. Then, we define $\rho_3(a, b) = \rho(f)$, where f is the degree 1 monotone extension of σ_3 restricted to the complement of the holes defined by (a, b) and $\rho(f)$ is the rotation number of f. With this construction in mind, we can now state the following result, which is part of Theorem 3.2.1 of [8].

LEMMA 5.15. *Let $\rho_3 : \Delta_3 \to [0,1]$ be defined as above. Then,*

(1) ρ_3 is continuous.

(2) ρ_3 is monotone.

With these results in mind, we can now understand how gaps in nearby q-laminations behave with respect to periodicity, degree, and rotation.

THEOREM 5.16. *Let $\mathcal{L} \in \overline{\mathbb{L}_d^q}$ and let G be a periodic gap of $\mathcal{L}$ with period n. Then, for all q-laminations, $\widehat{\mathcal{L}}$ sufficiently close to $\mathcal{L}$, $G(\widehat{\mathcal{L}})$ is periodic and $\sigma_d^n(G(\widehat{\mathcal{L}})) = G(\widehat{\mathcal{L}})$. Furthermore, the degree of the first return map on G equals the degree of the first return map on $G(\widehat{\mathcal{L}})$. Moreover, if $d = 3$ and if the degree of the first return map is 1, then the rotation number of $G(\widehat{\mathcal{L}})$ is close to the rotation number of G.*

PROOF. Suppose $G(\widehat{\mathcal{L}})$ is so close to G such that $\sigma_d^i(G)$ is close to $\sigma_d^i(G(\widehat{\mathcal{L}}))$ for all $1 \le i \le n$. Then, if by way of contradiction, we suppose $\sigma_d^n(G(\widehat{\mathcal{L}})) \ne G(\widehat{\mathcal{L}})$, then we see that $\sigma_d^n(G(\widehat{\mathcal{L}}))$ is not sufficiently close to $\sigma_d^n(G)$ as assumed. Therefore, $G(\widehat{\mathcal{L}})$ is periodic and $\sigma_d^n(G(\widehat{\mathcal{L}})) = G(\widehat{\mathcal{L}})$.

Now, suppose the degree σ_d^n on G is k. Suppose, by way of contradiction, that the degree of σ_d^n on $G(\widehat{\mathcal{L}})$ is $m > k$. Thus, for any non-critical leaf, $\widehat{\ell}$ on the boundary of $G(\widehat{\mathcal{L}})$, there are m siblings of $\widehat{\ell}$ on the boundary of $G(\widehat{\mathcal{L}})$. Suppose ℓ is a non-critical leaf on the boundary of G to which $\widehat{\ell}$ is close. By Lemma 5.14, each of these m siblings is sufficiently close to m distinct siblings of ℓ which are all on the boundary of G. Thus, the degree of σ_d^n on G is m, a contradiction. If $m < k$, then not all of the siblings of ℓ on the boundary of G are approximated by siblings of $\widehat{\ell}$ on the boundary of $G(\widehat{\mathcal{L}})$, a contradiction. Therefore, the degrees of the first return map on G and $G(\widehat{\mathcal{L}})$ are equal.

The final statement follows from Lemma 5.13 and the continuity of the map ρ_3 in Lemma 5.15. ☐

REMARK 5.17. Theorem 5.16 introduces the notion of *stability* for rotational gaps. From this theorem, we see that a rotational gap G of a degree 3 limit lamination

$\mathcal{L}$ is *stable* in the sense that nearby q-laminations must have rotational gaps which approximate G. Furthermore, the rotation numbers of these approximating gaps must converge to the rotation number of G. As we can see, stability is a much weaker property than rigidity. Rigidity properties exist when approximating q-laminations have the exact same object (leaf/gap) as in the limit lamination they are approximating. Stability properties exist when approximating q-laminations have the same type of object (e.g. rotational gaps) but not necessarily the same exact object as in the limit lamination. In the case of stability of rotational gaps, we can lose key features when we move from a limit lamination to an approximating q-lamination. For example, if a limit lamination $\mathcal{L}$ has a rotational gap G with an irrational rotation number, then given a close by q-lamination $\widehat{\mathcal{L}}$ with a rotational gap $G(\widehat{\mathcal{L}})$ which approximates G, the rotation number of $G(\widehat{\mathcal{L}})$ may be rational.

3. Invariant Sets

We now transition into a discussion on the existence of invariant objects (leaves/gaps) of a given 3-invariant lamination. We begin with 3-invariant laminations in general and provide an assumption which guarantees the existence of either an invariant leaf or invariant gap. We then narrow our focus to that of 3-invariant limit laminations and observe what happens when we lift these assumptions.

Definition 5.18. Suppose $\ell = \overline{ab}$ is a non-invariant leaf in a 3-invariant lamination $\mathcal{L}$. We say that ℓ *moves forward* if the endpoints of $\sigma_3(\ell)$ are located in the closure of the component of $\mathbb{S} \setminus \{a, b\}$ not containing 0.

Definition 5.19. We say that a leaf $\ell = \overline{ab}$ *separates* a leaf $\ell' \neq \ell$ from a point $x \in \mathbb{S}$ if the endpoints of ℓ' are in the closure of the component of $\mathbb{S} \setminus \{a, b\}$ not containing x.

Theorem 5.20. *Suppose $\mathcal{L}$ is a 3-invariant lamination which does not contain any critical leaves with a fixed endpoint. Then, $\mathcal{L}$ contains an invariant gap or an invariant leaf.*

PROOF. First assume that 1) both fixed points, 0 and $\frac{1}{2}$, are not on the boundary of a gap and 2) no leaves come out of either 0 or $\frac{1}{2}$. Under these assumptions, we see that 0 is a limit of a sequence of nondegenerate leaves, ℓ_i, such that for all i sufficiently large, 0 is located under ℓ_i. An analogous fact is true for $\frac{1}{2}$, and for the upcoming arguments we deal only with 0, but all arguments will hold for $\frac{1}{2}$. Now, since the fixed points are repelling, for those $\ell_i s$ sufficiently small in length, we have that ℓ_i is located under $\sigma_3(\ell_i)$. In particular, $\sigma_3(\ell_i)$ is contained in the component of $\overline{\mathbb{D}} \setminus \ell_i$ which does not contain 0, i.e. these ℓ_is move forward.

Let B be the set of all leaves in $\mathcal{L}$ which move forward. Let $A \subset B$ be defined as follows: $\ell \in A$ if and only if $\ell \in B$ and for all leaves $\ell' \in \mathcal{L}$ which separate ℓ from 0, $\ell' \in B$. We can now define a natural linear order, $<_A$, on A. We will say that, for any two distinct leaves $\ell, \ell' \in A$, $\ell <_A \ell'$ if and only if ℓ separates ℓ' from 0.

Let $\widehat{\ell} = \sup A$ under $<_A$. We first show that $\widehat{\ell}$ is a nondegenerate leaf. Suppose not, i.e. $\widehat{\ell}$ is a point. Then, the only way for $\widehat{\ell}$ to move forward is for it to be a fixed point. But then, there is either an invariant gap with $\widehat{\ell}$ on its boundary, a contradiction, or there is a sequence of leaves in B over $\widehat{\ell}$ which converge to $\widehat{\ell}$. The latter case is impossible since leaves arbitrarily close to a fixed point must expand under σ_3 and thus do not move forward, a contradiction since those close leaves are in B.

Now we show that $\widehat{\ell} \in A$ if it is not invariant. First, suppose there exists a gap G in the component of $\overline{\mathbb{D}} \setminus \widehat{\ell}$ containing 0 with $\widehat{\ell}$ on its boundary. That is to say that $\widehat{\ell}$ is not a limit of leaves in A, and hence $\widehat{\ell} \in A$. Now suppose $\widehat{\ell}$ is a limit of leaves in the component of $\overline{\mathbb{D}} \setminus \widehat{\ell}$ containing 0. Then, there are arbitrarily close leaves from A converging to $\widehat{\ell}$ from inside the component of $\overline{\mathbb{D}} \setminus \widehat{\ell}$ containing 0. Thus, if we suppose $\widehat{\ell} \notin A$, then either $\widehat{\ell}$ is an invariant leaf (and we are done) or $\sigma_d(\widehat{\ell})$ is in the component of $\overline{\mathbb{D}} \setminus \widehat{\ell}$ containing 0. But then those leaves from A sufficiently close to $\widehat{\ell}$ also map into that component since, by continuity, the their images are close to $\sigma_3(\widehat{\ell})$. This is a contradiction since those leaves from A must move forward.

Now, assuming that an invariant gap or leaf has not been located up to this point, we will argue to existence of such an object from this construction. First, assume that $\widehat{\ell}$ is the limit of a sequence of leaves ℓ_i located in the component of $\overline{\mathbb{D}} \setminus \widehat{\ell}$ which does not contain 0. Then, if $\widehat{\ell}$ is not invariant, $\widehat{\ell}$ moves forward since it is in A. Thus, sufficiently close ℓ_is will have images close to $\sigma_3(\widehat{\ell})$ and hence those ℓ_is move forward as well. Those, those ℓ_is are clearly in A. Furthermore, by construction, each ℓ_i is such that $\widehat{\ell} <_A \ell_i$, a contradiction to $\widehat{\ell} = \sup A$.

Now assume that $\widehat{\ell}$ is on the boundary of a gap G located in the component of $\overline{\mathbb{D}} \setminus \widehat{\ell}$ not containing 0. If G is invariant, we are done. If not, then there exists a leaf ℓ' on the boundary of G such that $\sigma_3(G)$ is located in the component of $\overline{\mathbb{D}} \setminus \ell'$ not containing 0. This follows since $\widehat{\ell}$ must move forward. Either ℓ' is a fixed leaf which flips (and we are done) or ℓ' maps into the component of $\overline{\mathbb{D}} \setminus \ell'$ not containing 0 and hence moves forward. Moreover, $\ell' \in A$ since $\widehat{\ell}$ is the only leaf on the boundary of G which separates ℓ' from 0, and we know that $\widehat{\ell} \in A$. Thus, by the established linear order, $\widehat{\ell} <_A \ell'$, which contradicts that $\widehat{\ell} = \sup A$.

Now, we will assume without loss of generality that 0 is on the boundary of a gap G and no leaves have 0 as an endpoint. Since 0 is fixed and on the boundary of exactly one gap, namely G, G is invariant.

Now assume, without loss of generality, that there is leaf ℓ with 0 as an endpoint. If $\ell = \overline{0\frac{1}{2}}$ we are done since ℓ is a fixed leaf. Suppose $\ell = \overline{0a_0}$ where $a_0 \neq 0, \frac{1}{2}$. Then, the forward images of ℓ form a sequence of leaves such that $\sigma_3^i(\ell) = \overline{0a_i}$. Note that this is a finite sequence because, otherwise, the a_is would have to converge to a fixed point which contradicts the fixed points being repelling. Supposing that no leaf $\overline{0a_i}$ is the fixed leaf $\overline{0\frac{1}{2}}$ or a critical leaf (necessarily with a fixed endpoint at 0), such a collection of leaves cannot exist. This follows from the fact that the a_is must map forward order preserving.

Therefore in any case, there is an invariant leaf or gap in $\mathcal{L}$. □

We now turn to the case when $\mathcal{L}$ contains a critical leaf, c, with a fixed endpoint. This case differs from the previous cases because there are choices to be made about the number and type of pullbacks of the critical leaf. We will see that we can force the existence of an invariant gap with c on its boundary so long as there are only a minimal number of pullbacks of c in $\mathcal{L}$. Before proceeding we present the notion of *cone* and part of an important lemma (Lemma 2.17 from [5]) regarding the motion of leaves which all share a common periodic endpoint in a limit lamination.

Definition 5.21. A family C of leaves $\overline{ab}$ sharing the same endpoint a is said to be a *cone* of leaves. We call the point a the vertex of C. We denote the collection of endpoints of leaves in C by C'.

Lemma 5.22. *Let $\mathcal{L} \in \overline{\mathbb{L}_d^q}$. Let C be an infinite cone of $\mathcal{L}$ with periodic vertex v of period n. Let $\overline{va_1}, ..., \overline{va_k}$ be all leaves in C with $v = a_0 < a_1 < ... < a_k < v = a_{k+1}$ and $\sigma_d^n(a_i) = a_i$ for each i. If, for some i, $C' \cap (a_i, a_{i+1}) \neq \emptyset$, then there are the following cases.*

(1) *The map σ_d^n moves all points of $C' \cap (a_i, a_{i+1})$ in the positive direction except for those which are mapped to v.*
(2) *The map σ_d^n moves all points of $C' \cap (a_i, a_{i+1})$ in the negative direction except for those which are mapped to v.*

Remark 5.23. Given a 3-invariant limit lamination $\mathcal{L}$ which contains a critical leaf with a fixed endpoint (we consider the critical leaf $c = \overline{0\frac{1}{3}}$ without loss of generality), we use Lemma 5.22 to rule out some cases regarding which pullbacks of the critical leaf can exist in $\mathcal{L}$. In particular, if the leaf $\overline{0\frac{1}{2}}$ is not in $\mathcal{L}$, Lemma 5.22 prevents a *collapsing motion* for pullbacks of the critical leaf. For example, if $\mathcal{L}$ contains the leaf $\overline{0\frac{1}{9}}$, then $\mathcal{L}$ necessarily contains an infinite cone C with a fixed vertex at 0. C is generated by continuing to pull back successive eventual preimages of the critical leaf with endpoints located in the interval $[0, \frac{1}{3}]$. If we then suppose $\mathcal{L}$ contains another preimage, ℓ, of c (distinct from $\overline{0\frac{1}{9}}$) with 0 as an endpoint, then clearly $\ell \in C$. However, by Lemma 5.22, this is impossible because ℓ and $\overline{0\frac{1}{9}}$ collapse to the critical leaf c. By collapsing in this way, we contradict that all endpoints of leaves in C (except the vertex 0) must move in the same direction (since $\overline{0\frac{1}{2}} \neq C$).

THEOREM 5.24. *Let $\mathcal{L} \in \overline{\mathbb{L}_3^q}$ be a 3-invariant limit lamination containing a critical leaf c with a fixed endpoint a. If there are at least 2 preimages of c having a as an endpoint, then $\mathcal{L}$ contains an invariant gap or invariant leaf.*

PROOF. To simplify the argument, we consider the case when $a = 0$ and $c = \overline{0\frac{1}{3}}$. The other three possible choices of critical leaves with a fixed endpoint have similar arguments. Notice that if $\overline{0\frac{1}{2}} \in \mathcal{L}$ then we are done since it is a fixed leaf. By the considerations in Remark 5.23, $\overline{0\frac{1}{9}}$ cannot be one of the preimages of c coming out of 0. Therefore, the only possible choices for the pullback leaves of c coming of 0 are $\overline{0\frac{4}{9}}$ and $\overline{0\frac{7}{9}}$. There must be a sibling leaf to $\overline{0\frac{4}{9}}$ coming out of $\frac{1}{9}$ in order to complete a full sibling collection. The only choice for such a leaf is $\overline{\frac{1}{9}\frac{1}{3}}$. Successively pulling $\overline{\frac{1}{9}\frac{1}{3}}$ back under σ_3 yields a countable concatenation of leaves, D, converging back to 0. Under σ_d, the leaves in D click forward to the next leaf in D (*next* in the sense of positive circular order). Thus, $D \cup \{0\}$ is invariant. We claim that the leaves in D, together with the point 0, form the boundary of an invariant gap G. Indeed, we cannot insert any additional leaves in the interior of the convex hull $\mathrm{CH}(D)$ of D. This is because if we insert a leaf $\ell = \overline{pq}$ which connects two nonadjacent endpoints in D, then since p and q both click forward to the next endpoint in D, $\sigma_3(\ell)$ crosses ℓ. Therefore, $D \cup \{0\}$ forms the boundary leaves of a gap G of $\mathcal{L}$. In particular, we see that G is an invariant caterpillar gap, as desired. $\square$

REFERENCES

[1] M. Aziz. *Canonical Laminations for Fixed Point Portraits*. Masters Dissertation UAB 2023.

[2] B. Bielefeld, Y. Fisher, and J. Hubbard. *The Classification of Critically Preperiodic Polynomials as Dynamical Systems*, Volume 5, Number **4** of Journal of the American Mathematical Society, 721-762, 1992.

[3] A. Blokh, J. Malaugh, J. Mayer, L. Oversteegen, and D. Parris. *Rotational Subsets of the Circle Under z^d*. Topology and its Applications, **153**, 1540-1570, 2006.

[4] A. Blokh, D. Mimbs, L. Oversteegen, and K. Valkenburg. *Laminations in the Language of Leaves*. Transactions of the American Mathematical Society, **365**, 5367-5391, 2013.

[5] A. Blokh, L. Oversteegen, V. Timorin, and R. Ptacek. *The combinatorial Mandelbrot set as the quotient of the space of geolaminations*. arXiv:1503.00351, 2015.

[6] D. Cosper, J. Houghton, J. Mayer, and J. Olson. *Central Strips of Sibling Leaves in Laminations of the Unit Disk*. Top. Proc. **48**, 69-100, 2016.

[7] J. Kiwi. *Real laminations and the topological dynamics of complex polynomials*. Adv. Math. 184, no. **2**, 207-267, 2004.

[8] J. Malaugh. *Rotational Sets of the Circle under z^d*. Ph.D. Dissertation UAB, 2003.

[9] J. Milnor. *Dynamics in One Complex Variable*. Third edition, Annals of Mathematics Studies, Number **160**, Princeton University Press, Princeton, New Jersey, 2006.

[10] W. Thurston. *Polynomial Dynamics from Combinatorics to Topology*. p. 1-109 in *Complex Dynamics: Families and Friends*, ed. Dierk Schleicher, A K Peters: Wellesley, MA, 2009.

[11] W. Thurston. *The combinatorics of iterated rational maps* (1985), with appendix by D. Schleicher, *Laminations, Julia sets, and the Mandelbrot set*. Complex Dynamics: Families and Friends, edited by D. Schleicher and A. Peters, 1-137, 2009.

[12] S. Vejandla. *Cubic Symmetric Laminations*. Ph.D. Dissertation UAB, 2021.