

LAMINATIONS IN HOLOMORPHIC DYNAMICS

MIKHAIL LYUBICH AND YAIR MINSKY

CONTENTS

1. A missing line in the dictionary.	1
2. Laminations: general concepts	5
3. Natural extension and its regular part.	6
4. The Type Problem and affine structure on the leaves.	11
5. Post-critically finite maps.	17
6. Hyperbolic 3-laminations.	21
7. Universal orbifold laminations	25
8. Convex-cocompactness, non-recurrence and conical points	32
9. Quasi-isometries and rigidity	38
10. Further program	41
11. Appendix 1: Circle and polynomial-like maps	42
12. Appendix 2: Background material	43
References	47

1. A MISSING LINE IN THE DICTIONARY.

There is an intriguing dictionary between two branches of conformal dynamics: the theory of Kleinian groups and dynamics of rational maps. This dictionary was introduced by Sullivan, and led him in the early 80's to the no wandering domains theorem, deformation theory and geometric measure theory for holomorphic maps. Thurston's rigidity and realization theory, developed at the same time, was also motivated by this analogy. More recently, McMullen has made important contributions to the renormalization theory motivated by the analogy with 3-manifolds which fiber over the circle [33, 31].

However, the translation from one language to another, as in usual life, is not automatic. There are concepts and methods in each of these fields which only barely allow translation to the other one. And even when it is possible, the results achieved are often complementary (see Sullivan's table in [46] of the results on the structural stability and hyperbolicity problems).

In this paper we explore a construction which attempts to provide an element of the dictionary that has so far been missing: an explicit object that plays for a rational map the role played by the hyperbolic 3-orbifold quotient of a Kleinian group. To build this object we replace the notion of manifold by "lamination", which is a topological object whose local structure is the product of Euclidean space by a (possibly complicated) transverse space.

Another goal of this work is to study the space of backward orbits of a rational function. Since Fatou and Julia, inverse branches of iterated rational functions have played a crucial role in the theory. Unfortunately, the space of such branches, with its natural topology, is wild (should be compared with the Hénon attractor), and may deserve to be called a "turbulation". By imposing

a finer topology and completing, we turn this space into an affine lamination, in the hope that this will tame it.

Laminations were introduced into conformal dynamics by Sullivan, whose *Riemann surface laminations* play a role similar to that of Riemann surfaces for Kleinian groups (see [47, 48] or §3.5 and Appendix 1 of this paper). These are objects which locally look like a product of a complex disk times a Cantor set. Sullivan associated such a holomorphic object to any C^2 -smooth expanding circle map. The construction involves “conformal extension” of a non-analytic one-dimensional map (see Appendix 1).

In this paper we go one dimension up and make a “hyperbolic three dimensional extension” of a non-Möbius map. This object is called a *hyperbolic orbifold 3-lamination* and can be constructed in the following way.

Step 1: the natural extension. Consider the full natural extension $\hat{f} : \mathcal{N}_f \rightarrow \mathcal{N}_f$ of a rational map f (points of \mathcal{N}_f are backward orbits $\hat{z} = (\cdots \mapsto z_{-1} \mapsto z_0)$ of f).

Step 2: the regular leaf space. Restrict \hat{f} to the “regular part” \mathcal{R}_f of \mathcal{N}_f where the inverse iterates branch only finitely many times. This space is a union of leaves which are non-compact Riemann surfaces, simply connected except for Herman rings, that is, hyperbolic or parabolic planes. It can be viewed as a Riemann surface with uncountably many sheets where all inverse iterates f^{-n} live simultaneously.

Step 3: affine orbifold lamination. Consider the subset $\mathcal{A}_f^{\text{aff}}$ of \mathcal{R}_f consisting of parabolic leaves. The parabolic leaves possess a canonical affine structure preserved by the map. However this structure is not necessarily continuous in the transverse direction. To make it continuous, we refine the topology on $\mathcal{A}_f^{\text{aff}}$, obtaining a space \mathcal{A}_f^{ℓ} with a laminar structure¹. We then complete \mathcal{A}_f^{ℓ} to obtain a final object \mathcal{A}_f some of whose new leaves may be 2-orbifolds.

This step is technically the hardest.

Step 4: three-dimensional extension. Each affine leaf is naturally the boundary of a three-dimensional hyperbolic space (in the half-space model). The union of these spaces forms a hyperbolic orbifold 3-lamination \mathcal{H}_f with \hat{f} acting properly discontinuously, and by isometries on the leaves.

Step 5: quotient. Finally taking the quotient \mathcal{H}_f/\hat{f} of this lamination by \hat{f} we obtain the desired hyperbolic orbifold 3-lamination.

We also define the convex core of the lamination \mathcal{H}_f/\hat{f} and prove that it is compact if and only if f is critically non-recurrent without parabolic points. Using this criterion, we prove a rigidity theorem 9.1 for critically non-recurrent maps which extends Thurston’s rigidity theorem for post-critically finite rational maps (see Douady-Hubbard [17]). Our three-dimensional proof gives an explicit connection between Thurston’s and Mostow’s rigidity theorems.

The structure of the paper is as follows:

§2. Basic notions of laminations and orbifold laminations.

§3. The natural extension \mathcal{N}_f , and its regular part \mathcal{R}_f . The space \mathcal{R}_f consists of backward orbits which have neighborhoods whose pullbacks hit the critical points only finitely many times. This space can be decomposed into leaves that admit a natural conformal structure. We show that (with the exception of Herman rings) the leaves of \mathcal{R}_f are simply connected non-compact Riemann surfaces, i.e. either hyperbolic or parabolic planes.

We discuss criteria for when \mathcal{R}_f is all of \mathcal{N}_f except for a finite set, and when \mathcal{R}_f is open in \mathcal{N}_f . This discussion crucially depends on a theorem by R. Mañé on the behaviour of non-recurrent critical points [26].

¹A different approach to this part was independently suggested by Meiyu Su who imposed a laminar topology associated to the transversal measure structure [43].




FIGURE 1. A caricature of the hyperbolic 3-lamination for $z^2 + \epsilon$. Sullivan’s solenoidal Riemann surface lamination is shown as the mapping torus of the doubling map on the solenoid (final gluing indicated by the double arrow), and the product of this with the interval is suggested. See discussion in §6.

§4. Here we discuss the affine part \mathcal{A}_f^n of \mathcal{R}_f , which leads us to the type problem for the leaves. This problem seems to be intimately related to the geometry of the Julia set. Parabolicity of leaves reflects “some” (but not necessarily uniform) expansion – see Lemma 4.4. We give several simple criteria for parabolicity and apply them to some special cases. In particular, all leaves of the real Feigenbaum quadratic are parabolic. This follows from an expansion property of f with respect to a hyperbolic metric (compare McMullen [30]). The only examples known to us of hyperbolic leaves are the invariant lifts of Siegel disks and Herman rings.

We also give an explicit formula for the affine coordinate on a parabolic leaf. It generalizes the classical formulas for the linearizing Königs and Leau-Fatou coordinates near repelling and parabolic points. From this point of view the affine structures on the leaves are just the linearizing coordinates along the backward orbits of f .

§5. Here we carry out Step 3 of the construction for the post-critically finite case, the construction of an affine orbifold lamination \mathcal{A}_f . We refine the topology on \mathcal{A}_f^n to separate leaves which branch inconsistently over the sphere, and enlarge \mathcal{A}_f^n to \mathcal{A}_f by making several copies of the post-critical periodic leaves, and replacing the original affine structure on some of them by an orbifold affine structure. This is the price we pay for having the affine structure transversally continuous, while keeping the lamination complete (in an appropriate sense).

§6. We define the notion of an *orbifold affine extension* $\hat{f} : \mathcal{A} \rightarrow \mathcal{A}$ of a rational map f , and show that it is naturally the boundary at infinity for an orbifold hyperbolic 3D extension $\hat{f} : \mathcal{H} \rightarrow \mathcal{H}$.

We prove that the action of \hat{f} on \mathcal{H} is properly discontinuous, so that the quotient \mathcal{H}/\hat{f} inherits the structure of a hyperbolic orbifold 3-lamination.

Then we introduce and discuss the notion of the convex core \mathcal{C}_f in \mathcal{H}_f/\hat{f} , which will play a key role in the rigidity argument.

We also describe the topological structure of the 3-lamination associated to quadratics $p_\epsilon : z \mapsto z^2 + \epsilon$ with ϵ inside of the main cardioid of the Mandelbrot set. We show that it is homeomorphic to $\mathcal{S} \times (0, 1)$ where \mathcal{S} is the Sullivan lamination. So, like in the case of quasi-Fuchsian groups, the 3-lamination connects the 2-laminations associated to the attracting basins of p_ϵ .

At the end of this section we discuss the “scenery flow” introduced by A. Fisher as an analogue of the geodesic flow on 3-manifolds. The phase space of this flow, constructed in [5] for rational maps satisfying axiom A, is loosely speaking the set of “pictures”, that is all possible rescalings of the infinitesimal germs of the Julia set. This scenery flow is topologically equivalent to the “vertical geodesic flow” on the 3-lamination over the lifted Julia set.

This vertical geodesic flow is an extra piece of structure which makes a difference between 3-laminations of rational maps and 3-manifolds of Kleinian groups. An equivalent way of viewing this structure is by saying that there exists a preferred \hat{f} -invariant cross-section, “ ∞ ”, at the boundary of the lamination \mathcal{H}_f .

§7. In this section we give a general construction of the affine and hyperbolic orbifold laminations associated to a rational map. The main hurdle is, as in the post-critically finite case, the fact that a sequence of disks in \mathcal{R}_f whose projection to the sphere is branched can limit onto a disk on which the projection is univalent. In the general case sorting out the different branching types is more involved since the set of points where this happens is no longer finite. Thus many copies of a leaf, possibly a continuum, must be added. One can keep track of this, and define an appropriate topology, using the affine structures themselves and their limiting behavior.

The self-organizing idea for this construction is to observe that the natural projection $\pi : \mathcal{N}_f \rightarrow \bar{\mathbb{C}}$ gives a meromorphic function on each leaf of \mathcal{A}_f^n , and this family of functions has a natural topology which induces a topology for \mathcal{A}_f^n . In fact, the space of non-constant meromorphic functions on \mathbb{C} with the right action of the affine group serves as a “universal” lamination on which every rational function acts. For any fixed f the structure of \mathcal{A}_f and \mathcal{H}_f can be extracted from the attractor of f in this universal space.

In conclusion we prove (using Ahlfors’ five islands theorem) that every lamination \mathcal{H}_f is minimal, except for the Chebyshev and Lattès maps. In these special cases, the lamination becomes minimal after removing the invariant isolated leaf. This is the characteristic property of these remarkable maps from the lamination point of view.

§8. In this section we prove that f is convex co-compact (that is, its convex core \mathcal{C}_f is compact) if and only if it is critically non-recurrent and does not have parabolic periodic points. Note that thus convex co-compactness differs from hyperbolicity, while these two notions are equivalent for Kleinian groups (one more illustration of the loose nature of the dictionary).

We also define the conical limit set and give in these terms a criterion of convex cocompactness. We then study ergodic properties of the conical limit set by means of the blow-up technique (on the lamination level) and Ahlfors’ harmonic extension method. Along the lines we obtain the lamination insight on the existence of invariant line fields for the Lattès examples: it comes from the existence of the isolated leaves.

§9. This section contains the three-dimensional proof of rigidity for convex co-compact maps.

We start by lifting the topological equivalence between the maps to a quasi-isometry \hat{h} between their 3-laminations (using the convex co-compactness). It follows that \hat{h} is quasi-conformal on the leaves of the affine extension. This reduces the problem to the existence of invariant line fields on the Julia set of \mathcal{A}_f , which was analyzed in the previous section.

§10. Conjectures and further program.

§11. In the first appendix we outline Sullivan’s construction of the Riemann surface lamination associated to an expanding map of the circle. We also give a globalization construction for the natural extension of polynomial-like maps via the inductive limit procedure.

§12. Appendix 2 fills in some necessary background, all of which is well-known to those who work in either dynamics or geometry, but not always to both. It also fixes some terminology and notation.

Acknowledgements. Since the first drafts of this manuscript appeared in spring 1994, we have received valuable feedback from many people, and we take this opportunity to thank all of them. John Milnor read parts of the manuscript and made many useful suggestions and comments. He also gave a simpler proof of Lemma 3.4 on non-compactness of the leaves. Alberto Baider simplified the proof of Lemma 3.3 on simple connectivity of the leaves. We thank Dennis Sullivan for several enjoyable discussions and stimulating questions, and Curt McMullen for some useful comments on renormalization. We also enjoyed talks with Alby Fisher about scenery flows, as well as discussions with Meiyu Su who is independently working on the structure of the natural extension of a rational map from a measure-theoretical point of view. After the preprint was issued (IMS at Stony Brook, 1994/20) we have found out that Jeremy Kahn has independently considered the idea of a 3D lamination and its possible application to the measure problem.

2. LAMINATIONS: GENERAL CONCEPTS

In this paper, a *lamination* will be a Hausdorff topological space \mathcal{X} equipped with a covering $\{U_i\}$ and coordinate charts $\phi_i : U_i \rightarrow T_i \times D_i$, where D_i is homeomorphic to a domain in \mathbb{R}^n and T_i is a topological space. The transition maps $\phi_{ij} = \phi_i \circ \phi_j^{-1} : \phi_j(U_i \cap U_j) \rightarrow \phi_i(U_i \cap U_j)$ are required to be homeomorphisms that take leaves to leaves (see Sullivan [48] and Candel [11]).

Subsets of the form $\phi_i^{-1}(\{t\} \times D)$ are called *local leaves*. The requirement on the transition maps implies that the local leaves piece together to form *global leaves*, which are n -manifolds immersed injectively in \mathcal{X} .

As usual we may restrict the class of transition maps to obtain finer structures on \mathcal{X} . If D_i are taken to lie in \mathbb{C} and ϕ_{ij} are conformal maps, we call \mathcal{X} a *Riemann surface lamination* and note that the global leaves have the structure of Riemann surfaces. If ϕ_{ij} are further restricted to be complex affine maps $z \mapsto az + b$, then we call \mathcal{X} a (*complex*) *affine lamination*, and the global leaves have a (complex) affine structure. If the leaves of an affine lamination are isomorphic to the complex plane, we also call it a \mathbb{C} -lamination. One can similarly consider real affine laminations, but as they will not play a role in this paper we shall assume from now on that “affine” means “complex affine”.

If D are taken to lie in \mathbb{H}^n and ϕ_{ij} are hyperbolic isometries, then \mathcal{X} is an *n -dimensional hyperbolic lamination*, or *hyperbolic n -lamination*. In the case when all leaves of the lamination are hyperbolic spaces, let us call it an \mathbb{H}^n -lamination.

When the laminated space \mathcal{X} is a (smooth/analytic) manifold, the lamination is usually called a *foliation*. It is called *smooth/analytic* if there is a smooth/analytic atlas of laminar local charts.

We shall need the notion of distance between affine structures on a Riemann surface. Let S be a Riemann surface supplied with two affine structures \mathcal{A}_1 and \mathcal{A}_2 . Let ϕ_1 and ϕ_2 be any two local charts of the structures \mathcal{A}_1 and \mathcal{A}_2 respectively, $\psi = \phi_1 \circ \phi_2^{-1} : U_1 \rightarrow U_2$ be the transition function. Then we may define

$$\text{dist}(\mathcal{A}_1, \mathcal{A}_2) = \sup_{\phi_1, \phi_2} \text{Dis}(\phi_1 \circ \phi_2^{-1}),$$

where Dis stands for the distortion (see Appendix 2, §12).

We will encounter situations where a Riemann surface lamination \mathcal{R} can be refined to give an affine lamination. Suppose that the global leaves of \mathcal{R} admit affine structure – that is, each global

leaf L admits a collection of conformal coordinate charts with affine transition maps. We say that these affine structures *vary continuously* in \mathcal{R} if, for any product box $U = T \times D$ the induced family of affine structures on D vary continuously with T , in the sense of the above notion of distance.

In other words, continuity of affine structure means that for each coordinate chart $\phi : U \rightarrow T \times D$ there is a choice of coordinate $\psi_t : \phi^{-1}(\{t\} \times D) \rightarrow \mathbb{C}$ for each $t \in T$, so that ψ_t is a restriction of an affine coordinate chart on a global leaf, and so that the family $\psi_t \circ \phi^{-1}(t, \cdot) : D \rightarrow \mathbb{C}$ varies continuously with t . The following is easy to check:

Lemma 2.1. *A continuous family of affine structures on the global leaves of a Riemann surface lamination \mathcal{R} induces an affine lamination structure on \mathcal{R} compatible with the original structure.*

Similarly the Riemann surface lamination can be viewed as a topological lamination with transversally continuous conformal structure on the leaves.

2.1. Orbifold laminations. In analogy with Thurston's notion of orbifolds (see Thurston [51], Scott [42] and also Satake's similar notion of V-manifolds, [41]), we may define an orbifold lamination to be a space for which every point has a neighborhood that is either homeomorphic to a standard product box neighborhood in a lamination, or to a quotient of such a box by a finite leaf-preserving group (called an *orbifold box*).

If the covering box has an affine (or conformal, or hyperbolic) structure which is preserved by the finite group, then we say that the orbifold box inherits an orbifold affine (or conformal, or hyperbolic) structure.

For example, let $T \times D$ be a product box, with D a two dimensional disk, let $\sigma : T \rightarrow T$ be a finite-order map and let $\rho : D \rightarrow D$ be a finite order rotation of D . Then the map $\sigma \times \rho$ generates a finite cyclic group action on $T \times D$ and the quotient is an orbifold box. Cycles of σ of order not divisible by the order of ρ (fixed points, for example) give rise to quotient leaves with orbifold points.

See also [49, 17] for the use of (regular 2-dimensional) orbifolds in the context of post-critically finite maps.

Example 2.2. This example illustrates how orbifold boxes will arise in §5. Let K be a Cantor set and $K' = K \setminus \{a\}$ for some $a \in K$, and let $\pi : \hat{D} \rightarrow D$ be a doubly branched map. Let B denote $(K' \times \hat{D}) \cup (\{a\} \times D)$, topologized so that a sequence (b_i, z_i) in $K' \times \hat{D}$ converges to (a, z) in $\{a\} \times D$ if and only if $b_i \rightarrow a$ and $\pi(z_i) \rightarrow z$.

We can then express B as an orbifold box, by letting T be the double of K , with both copies of a identified, and $\sigma : T \rightarrow T$ the map that interchanges copies. Let $\rho : \hat{D} \rightarrow \hat{D}$ be the involution that interchanges pairs of preimages of points in D . Then $(T \times \hat{D})/(\sigma \times \rho)$ is exactly B .

3. NATURAL EXTENSION AND ITS REGULAR PART.

3.1. Natural extension. Let $f : \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a rational endomorphism of the Riemann sphere. Let us consider the space of its backward orbits:

$$\mathcal{N} = \mathcal{N}_f = \{\hat{z} = (z_0, z_{-1}, \dots) : z_0 \in \bar{\mathbb{C}}, f : z_{-(n+1)} \mapsto z_{-n}\},$$

with topology induced by the product topology in $\bar{\mathbb{C}} \times \bar{\mathbb{C}} \times \dots$. This is a compact space projected down to $\bar{\mathbb{C}}$ by $\pi : \hat{z} \mapsto z_0$. The endomorphism f naturally lifts to a homeomorphism $\hat{f} : \mathcal{N} \rightarrow \mathcal{N}$ as $\hat{f}(\hat{z}) = (fz_0, z_0, z_{-1}, \dots)$. (The inverse map forgets the first coordinate of the backward orbit). Moreover, $\pi \circ \hat{f} = f \circ \pi$. In dynamics the map \hat{f} is usually called *the natural extension* of f . In algebra this object is also called *the projective (or inverse) limit* of

$$\bar{\mathbb{C}} \xleftarrow{f} \bar{\mathbb{C}} \xleftarrow{f} \bar{\mathbb{C}} \xleftarrow{f} \dots$$

One can also think of a point $\hat{z} \in \mathcal{N}$ as a *full orbit* $\{z_n\}_{n=-\infty}^{\infty}$, where $f : z_n \mapsto z_{n+1}$. (But don't confuse them with *grand orbits* generated by the equivalence relation $z \sim \zeta$ if there exist natural m and n such that $f^m z = f^n \zeta$). Along with the projection $\pi \equiv \pi_0$ let us also consider projections $\pi_n : \mathcal{N}_f \rightarrow \bar{\mathbb{C}}$ such that $\pi_n(\hat{z}) = z_n$. Clearly $\pi_n = f^{n-m} \circ \pi_m$ for $n \geq m$.

Given a (forward) invariant set $X \subset \bar{\mathbb{C}}$, let $\hat{X} \subset \mathcal{N}_f$ denote its *invariant lift* to \mathcal{N}_f , that is, the set of orbits $\{z_n\} \subset X$. This is nothing but the natural extension of $f|_X$. Note that it differs from $\pi^{-1}X$, unless X is completely invariant (that is, $f^{-1}X = X$).

Let $\hat{z} = (z_0, z_{-1}, \dots) \in \mathcal{N}_f$, D be a topological disk containing z_0 , and N be a natural number. Consider the pullback D_0, D_{-1}, \dots of D along \hat{z} . That is, D_{-n} is the component of $f^{-n}(D)$ containing z_{-n} . Let us define the following "boxes":

$$\begin{aligned} B(D, \hat{z}, N) &= \pi_{-N}^{-1}(D_{-N}) \\ &= \{\hat{\zeta} = (\zeta_0, \zeta_{-1}, \dots) \in \mathcal{N}_f : \zeta_{-N} \in D_{-N}\}, \end{aligned} \quad (3.1)$$

which form a basis of the topology in \mathcal{N}_f . For $N = 0$ we will shorten the notation as $B(D, \hat{z}) \equiv B(D, \hat{z}, 0)$.

3.2. The regular leaf space. Let us say that a point $\hat{z} = (z_0, z_{-1}, \dots) \in \mathcal{N}$ is *regular* if there is neighborhood U of z_0 in $\bar{\mathbb{C}}$ whose pullback U_{-n} along the backward orbit (z_0, z_{-1}, \dots) is eventually univalent. Let $\mathcal{R} = \mathcal{R}_f$ denote the set of regular points of the natural extension. This set is clearly completely invariant. Moreover, if z_0 is outside the ω -limit set $\omega(C)$ of the critical points, then $\hat{z} \in \mathcal{R}_f$ (see Appendix 2).

The path connected components of \mathcal{R} will be called the *leaves* and denoted by $L(\hat{z})$ for $\hat{z} \in \mathcal{R}$.

Lemma 3.1. *The leaves $L(\hat{z})$ possess an intrinsic topology and analytic structure such that the projection $\pi : L(\hat{z}) \rightarrow \bar{\mathbb{C}}$ is analytic. The branched points are the backward orbits passing through critical points. Moreover $\hat{f} : L(\hat{z}) \rightarrow L(\hat{f}\hat{z})$ is a biholomorphic isomorphism.*

Proof. Let $\hat{z} = (z_0, z_{-1}, \dots) \in \mathcal{R}$. Then there is a neighborhood $U \ni z_0$ whose pull-back U_{-n} along the orbit z_{-n} is eventually univalent. Let us take $\hat{U} = \{\hat{\zeta} = (\zeta_0, \zeta_{-1}, \dots) : \zeta_{-n} \in U_{-n}\}$ as a base neighborhood of \hat{z} (also called a leafwise neighborhood).

Let $f : U_{-(n+1)} \rightarrow U_{-n}$ be univalent for $n \geq N$. Then the map $\pi_{-N} : \hat{\zeta} \mapsto \zeta_{-N}$ is a homeomorphism between \hat{U} and U_{-N} . Let it be our local chart. The transition functions are just appropriate iterates of f , so that this provides us with a complex structure.

The last two statements are obvious. □

We may characterize the leaves in dynamical terms via the following observation.

Lemma 3.2. *Two points \hat{z} and $\hat{\zeta}$ in \mathcal{R}_f belong to the same leaf iff the following holds. There is a sequence of paths (γ_{-n}) in $\bar{\mathbb{C}}$ such that γ_{-n} connects z_{-n} to ζ_{-n} , and $f(\gamma_{-n}) = \gamma_{-n+1}$. Furthermore, for n sufficiently large there are neighborhoods U_{-n} of γ_{-n} such that there is a branch g of f^{-1} defined on U_{-n} and $f(U_{-n}) = U_{-n-1}$. In particular ζ_{-n} can be obtained from z_{-n} by analytic continuation of f^{-1} along γ_{-n+1} .*

Proof. Assume that \hat{z} and $\hat{\zeta}$ are on the same leaf and let $\hat{\gamma}$ be a path connecting them. We may represent any such path as a sequence of paths (γ_{-n}) in $\bar{\mathbb{C}}$ such that γ_{-n} connects z_{-n} to ζ_{-n} , and $f(\gamma_{-n}) = \gamma_{-n+1}$. Since each point in \mathcal{R}_f has a neighborhood whose projections are eventually univalent, we take a finite covering of $\hat{\gamma}$ and consider its projections by π_{-n} for n sufficiently large. These are the neighborhoods U_{-n} .

Conversely, given the sequence γ_{-n} satisfying the conditions, it is immediate that the path $\hat{\gamma}$ in \mathcal{N}_f that they define in fact lies in \mathcal{R}_f . □

By *local leaves* in a box $B(D, \hat{z}, N)$ we will mean the components of intersection of the global leaves with this box.

Unfortunately these boxes in general don't have a product structure, so that \mathcal{R}_f is not always a Riemann surface lamination. For this reason \mathcal{R}_f will be called a (*conformal*) *leaf space*, that is, a space which is decomposed into the union of leaves supplied with smooth (conformal) structure. Actually the leaves behave so wildly (keep in mind the Henon map) that one might rather call the space a "*turbulation*".

However, if the orbits of the critical points don't meet D_{-N} then $B(D, \hat{z}, N) \approx T \times D_{-N}$, where T may be identified with the fiber $\pi^{-1}(z_{-N})$. The local leaves in this box correspond to slices $\{t\} \times D_{-N}$.

3.3. Topology of the leaves. Our main tool in this section will be the Shrinking Lemma (see Appendix 2), which states roughly that, in a uniform sense, backward iterates of a region on which the branching of f is bounded have (spherical) diameters that shrink to 0. This holds except if the iterates remain in a *rotation domain* – a Siegel disk or Herman ring – for all time.

Let us first consider some exceptional cases: If a component W of the Fatou domain F_f is a rotation domain, then its invariant lift \hat{W} , consisting of all orbits which remain in W for all time, is a full leaf of \mathcal{R}_f , and $\pi : \hat{W} \rightarrow W$ is a conformal equivalence. The second part is obvious since $f|_W : W \rightarrow W$ is a 1-1 conformal map. It only remains to check that \hat{W} is not properly contained in a leaf. That is, we must check that any point on $\partial\hat{W}$ in \mathcal{N}_f does not lie in \mathcal{R}_f . Such a point \hat{w} is an orbit which stays in ∂W for all time, and in particular is on the Julia set. If w_0 had a neighborhood D_0 which pulled back along \hat{w} eventually univalently, then by the Shrinking Lemma (after possibly trimming D_0 to a slightly smaller disk), $\text{diam}(D_{-n}) \rightarrow 0$. However $D_0 \cap W$ is being pulled back by the univalent map $f|_W$, and so the diameter of $D_{-n} \cap W$ cannot shrink.

We shall adopt the convention of using rotation domain, Siegel disk or Herman ring, to refer also to the leaves of \mathcal{R}_f which are invariant lifts of these domains.

Except in the case of rotation domains, the structure of a leaf reflects the behavior of f at *small scales* – this is another consequence of the Shrinking Lemma. The following two lemmas show that, barring the obvious exception, all leaves are topologically trivial.

Lemma 3.3. *All leaves of \mathcal{R}_f which are not Herman rings are simply connected.*

Proof. The invariant lift of a Siegel disk is, by the above discussion, a disk, so we may from now on consider a leaf L which is not either kind of rotation domain. That is, for $\hat{z} \in L$ there is some n for which z_{-n} is not in a rotation domain.

Let $\hat{\gamma} : S^1 \rightarrow L$ be a simple closed smooth curve on L , which does not pass through the branched points of π . We need to show that $\hat{\gamma}$ bounds a disk on the leaf L . Let us consider the corresponding sequence of smooth curves on the Riemann sphere: $\gamma_{-n} = \pi_{-n} \circ \hat{\gamma} : S^1 \rightarrow \bar{\mathbb{C}}$. Deforming $\hat{\gamma}$ slightly, we can get γ_0 to have only finitely many points of self intersection, all of which are double points. Clearly, the γ_{-n} have no more points of self intersection than γ_0 , since if $\gamma_{-n}(a)$ is a simple point for some $a \in S^1$, so is $\gamma_{-(n+1)}(a)$.

Let us now consider a point of self intersection, $\gamma_0(a) = \gamma_0(b)$, where $a, b \in S^1$, and $a \neq b$. Since $\hat{\gamma}(a) \neq \hat{\gamma}(b)$, there is an n_0 such that $\gamma_{-n}(a) \neq \gamma_{-n}(b)$ for $n \geq n_0$, so that γ_{-n} has strictly fewer points of self intersection than γ_0 . It follows that eventually all the curves γ_{-n} are simple.

Furthermore, by the Shrinking lemma, $\text{diam} \gamma_{-n} \rightarrow 0$ as $n \rightarrow \infty$. Let D_{-n} be the component of $\mathbb{C} \setminus \gamma_{-n}$ of small diameter. Then it contains at most one critical point of f for n sufficiently big. If $D_{-(n+1)}$ actually contained a critical point, the curve $\gamma_{-(n+1)}$ (obtained by analytic continuation of f^{-1} along the simple curve γ_{-n}) would not be closed. Hence the D_{-n} eventually don't contain the critical points.

It follows that the maps $f : D_{-(n+1)} \rightarrow D_{-n}$ are univalent for n sufficiently big: $n \geq N$. Hence the set \hat{D} of backward orbits $\{(z_{-n}) : z_{-n} \in D_{-n} \text{ for } n \geq N\}$ represents a topological disc in L bounded by $\hat{\gamma}$ (with a homeomorphic projection $\pi_{-N} : \hat{D} \rightarrow D_{-N}$). \square

The following lemma excludes elliptic leaves (that is, conformal spheres).

Lemma 3.4. *If $\deg f > 1$, there are no compact leaves in the lamination \mathcal{R} .*

Proof. Assume that a leaf L is compact. Then the projection $\pi : L \rightarrow \bar{\mathbb{C}}$ is a finite-sheeted branched covering. However, we can also express π as $f^n \circ \pi \circ \hat{f}^{-n}$, so $\deg \pi \geq (\deg f)^n$ for any n . This is a contradiction. \square

3.4. Criteria for regularity. Let us consider some cases in which we can say which part of \mathcal{N}_f is regular.

Axiom A case. (See Appendix 2 for definitions) We will call these functions ‘‘Axiom A’’ instead of the more common ‘‘hyperbolic’’ in order to avoid sentences like ‘‘in the hyperbolic case all leaves are parabolic’’.

If f satisfies axiom A then $\mathcal{R}_f = \mathcal{N}_f \setminus \{\text{finite set of points}\}$, namely the attracting cycles of \hat{f} . Note that the backward orbits like $(\alpha, \dots, \alpha, \beta, \dots)$, where α is an attracting fixed point and $\beta \neq \alpha$ is another preimage, are included into \mathcal{R}_f , since $\beta \notin \omega(C)$.

Critically non-recurrent case. We will use the notation $\alpha(\hat{z}) \subset \bar{\mathbb{C}}$ for the limit set of the backward orbit $\hat{z} = (z_{-n})_{n>0}$.

Lemma 3.5. *Let $\hat{z} = (z_{-n}) \in \mathcal{N}_f$ be a backward orbit satisfying the property that for some N , z_{-N} does not belong to an attracting or parabolic cycle, nor to the ω -limit set of a recurrent critical point. Then $\hat{z} \in \mathcal{R}_f$.*

Proof. Let C_1 be the set of critical points such that for $c \in C_1$, $z_{-n} \in \omega(c)$, $n = 0, 1, \dots$, and C_2 be the complementary set of critical points. Without loss of generality we can assume that already z_0 does not belong to an attracting or parabolic cycle, nor to the closure $\text{cl}(\text{orb}(c))$ for any $c \in C_2$.

By the assumption, C_1 consists of non-recurrent points. Hence there is an $\epsilon > 0$ such that $\text{dist}(z_{-n}, C_1) \geq \epsilon$, $n = 0, 1, \dots$. For $\delta > 0$ let $U_0 = D(z_0, \delta)$, and U_{-n} be the pull-back of U_0 along z_{-n} . By Mañé’s Theorem ([26] and §12.1), there is a $\delta > 0$ such that $\text{diam } U_{-n} < \epsilon$. Hence U_{-n} does not hit the critical points of C_1 .

Moreover, if δ is sufficiently small, the orbits of the critical points $c \in C_2$ clearly don’t meet U_0 . Hence U_{-n} don’t hit these critical points either, so that the pull-back $\{U_{-n}\}$ is univalent. \square

When we refer to an attracting/parabolic etc. cycle in \mathcal{N}_f , we mean the invariant lift of the corresponding cycle in $\bar{\mathbb{C}}$. Let us recall from Appendix 2 that C_r denotes the set of recurrent critical points in the Julia set.

Lemma 3.6. *The closure of the set $\mathcal{N}_f \setminus \mathcal{R}_f$ of irregular points in \mathcal{N}_f coincides with the invariant lift $\hat{\omega}(C_r)$ together with attracting and parabolic cycles.*

Proof. If $\hat{z} \notin \hat{\omega}(C_r)$, nor is an attracting or parabolic periodic point, then it follows from lemma 3.5 that $B(D, \hat{z}) \subset \mathcal{R}_f$ for sufficiently small neighborhood $D \ni z$. Thus $\hat{z} \in \text{int } \mathcal{R}_f$.

Vice versa, let $\hat{z} \in \hat{\omega}(C_r)$. Let D be a neighborhood of z_0 , $N > 0$ be any integer, and $B_0 = \text{cl } B(D, \hat{z}, N)$ be a closed neighborhood of \hat{z} . We should show that B_0 contains an irregular point.

Since $D_{-N} \cap \omega(C_r) \neq \emptyset$, there is a critical point $c \in C_r$ such that $f^{n_1}c \in D_{-N}$ for some $n_1 > 0$. Let $\hat{z}^{(1)}$ be any backward orbit with $z_{-(N+n_1)}^{(1)} = c$, and

$$B_1 \equiv \text{cl } B(D, \hat{z}^{(1)}, N + n_1) \subset B_0.$$

Then all leaves of B_1 over D are at least double branched.

Let us now consider a neighborhood base $D \equiv D^1 \supset D^2 \supset \dots$ of z , and let $D_{-(N+n_1)}^2 \ni c$ be the pullback of D^2 along the orbit $\{f^k c\}_{k=0}^{N+n_1}$. Since c is recurrent, there is an n_2 such that $f^{n_2} c \in D_{-(N+n_1)}^2$. Take any backward orbit $\hat{z}^{(2)}$ with $z_{-(N+n_1+n_2)}^{(2)} = c$, and consider the closed box $B_2 = \text{cl } B(D^2, \hat{z}^{(2)}, N+n_1+n_2) \subset B_1$. All leaves of B_2 over D^2 are at least triple branched.

Proceeding in this way, we will construct a nest $B_0 \supset B_1 \supset B_2 \supset \dots$ of closed boxes, such that all leaves of B_n are at least n times branched over D^n . Hence the intersection of these boxes consist of irregular points. \square

Let us call a map f *critically non-recurrent* if all its critical points on the Julia set are non-recurrent. The following fact was proved by Carleson, Jones and Yoccoz [13] (in different language).

Corollary 3.7. *A map f is critically non-recurrent if and only if*

$$\mathcal{R}_f = \mathcal{N}_f \setminus \{\text{attracting and parabolic cycles}\}.$$

Let us call a map f *persistently recurrent* if any backward orbit U_0, U_{-1}, \dots of a neighborhood U_0 along $\omega(C_r)$ hits a critical point. In other words, all points of $\hat{\omega}(C_r)$ are irregular. Lemma 3.6 also yields the following criterion of openness of the regular leaf space.

Corollary 3.8. *The regular leaf space \mathcal{R}_f is open in \mathcal{N}_f if and only if f is either critically non-recurrent or persistently recurrent. In the latter case*

$$\mathcal{R}_f = \mathcal{N}_f \setminus (\hat{\omega}(C_r) \cup \text{parabolic and attracting cycles})$$

3.5. The Julia and Fatou sets. Let us consider the pull-backs $\mathcal{J}_f^r \equiv \mathcal{J}^r = \pi^{-1}J \cap \mathcal{R}_f$ and $\mathcal{F}_f^r \equiv \mathcal{F}^r = \pi^{-1}F$ of the Julia set J and the Fatou set F to the space \mathcal{R}_f .

Note first that \mathcal{F}^r is obtained from the pullback of F to \mathcal{N}_f just by removing the attracting cycles. Also, if we remove from \mathcal{F}^r the invariant lifts of Siegel disks and Herman rings, then we obtain a Riemann surface lamination. Indeed, if U is compactly contained in the Fatou set, and a backward trajectory U_0, U_{-1}, \dots eventually does not meet either attracting cycles, Siegel disks or Herman rings, then there is an N such that $f^{-k}U_{-N}$ does not meet the critical points for $k \geq 0$. In particular, the boxes $B(U_0, \hat{z}, N)$ have a product structure if $\hat{z} \in \mathcal{F}^r \setminus \{\text{Siegel disks and Herman rings}\}$ and N is large.

Note further that \hat{f} acts properly discontinuously on \mathcal{F}^r with Siegel disks and Herman rings removed. Indeed for any $\hat{z} \in \mathcal{F}^r$ which is not in a rotation domain, either z_{-n} lie in an attracting or parabolic basin and pull back toward its boundary, or eventually end up in preimages of a periodic domain. Thus there is a neighborhood V of z_{-N} for some N such that all further pullbacks of V accumulate onto J . It follows that $\hat{f}^{-n}B(V, z_{-N})$ eventually escapes every compact subset of \mathcal{F}^r .

Thus, \mathcal{F}^r/\hat{f} is a Hausdorff topological space, and in fact a Riemann surface lamination, since it inherits its local structure from \mathcal{R} .

So to each basin of the Fatou set we can associate a Riemann surface lamination. These play the role of the Riemann surfaces associated to a Kleinian group.

In [48, 47] Sullivan considered the natural extension of the attracting basin of infinity for a polynomial, and obtained a ‘‘solenoidal Riemann surface lamination’’, called \mathcal{S} (see Appendix 1). A similar object appears as a subset of \mathcal{F}^r/\hat{f} in general. Let us consider the topological structure of these laminations in somewhat greater detail.

Attracting domains. Consider a cycle of basins $U_1 \xrightarrow{f} \dots \xrightarrow{f} U_m \xrightarrow{f} U_1$ for an attracting (or super-attracting) cycle, and let \mathcal{G} denote the subset of \mathcal{F}^r consisting of orbits \hat{z} that are attracted (in forward time) to this cycle. This sublamination divides naturally into two pieces: let \mathcal{G}_1 contain

orbits which stay in $\cup U_i$ for all time, and let \mathcal{G}_2 consist of orbits which, before some time, lie outside the U_i .

Suppose that all of the domains are simply connected. We claim that \mathcal{G}_1/\hat{f} is Sullivan's solenoidal Riemann surface lamination (of appropriate degree), and \mathcal{G}_2/\hat{f} is a finite union of copies of (plane domain) \times (Cantor set), which accumulates onto the solenoidal part. The full quotient \mathcal{G}/\hat{f} is, in particular, compact.

We can study \mathcal{G}_1/\hat{f} by considering just the return map f^m to U_1 , and the quotient of the set of orbits of this map that stay in U_1 for all time. On a neighborhood of ∂U_1 , f^m is topologically conjugate to $z \mapsto z^d$ acting on a neighborhood of the boundary in the unit disk \mathbb{D} , and every orbit in \mathcal{G}_1 accumulates in backward time onto ∂U_1 (note that we omit the orbit which remains on the attracting periodic cycle, since it does not lie in \mathcal{R}_f). It follows that the quotient \mathcal{G}_1/\hat{f} is homeomorphic to the quotient of the Fatou domain of 0 (or ∞) for the lamination of $z \mapsto z^d$, namely Sullivan's solenoidal Riemann surface lamination.

Now consider an orbit \hat{z} which escapes $\cup U_i$ in backward time. Let \tilde{z} denote the full orbit $(\dots, z_{-1}, z_0, z_1, \dots)$. There is a finite list V_1, \dots, V_p of preimages of U_1 such that no V_i contains a post-critical point, and every full orbit \tilde{z} with $\hat{z} \in \mathcal{G}_2$ passes through a unique V_i . Let q denote the smallest integer for which $z_q \in \cup V_i$. Since \mathcal{G}_2/\hat{f} is just the space of these full orbits modulo shift, we can identify it with $(\cup V_i) \times \Sigma$, where Σ is a Cantor set, so that the $\cup V_i$ component is z_q and the Σ component specifies the preimages of V_i which contain the preimages z_{q-n} , $n = 1, 2, \dots$.

It remains to see that the closure of \mathcal{G}_2/\hat{f} is in \mathcal{G}_1/\hat{f} . Let A denote a fundamental annulus in U_1 . This is a compact annulus, surrounding the fixed point of f^m , through which every full orbit of \mathcal{G}_1 passes exactly once (or twice if on the boundary). Now if we consider \hat{z} in \mathcal{G}_2 , such that $z_q \in V_i$, we see that z_{q+N} passes through A where $N > 0$ gets larger as z_q approaches ∂V_i . Thus \hat{z} is very close to some orbit $\hat{w} \in \mathcal{G}_1$ which agrees with \hat{z} for all moments n where $z_n \in \cup U_i$. It follows that \mathcal{G}_2/\hat{f} accumulates on \mathcal{G}_1/\hat{f} , and in fact that all of \mathcal{G}_1/\hat{f} is obtained this way.

If the domains U_i in the cycle are not simply connected the topological structure of the quotient is more complicated and we shall not describe it here. However let us sketch an argument showing that it is compact. Let D be a small closed disk around the attracting fixed point for f^m in U_1 , so that D maps univalently to $f^m(D) \subset D$. For any orbit \hat{z} attracted to the cycle there is a first moment $q \in \mathbb{Z}$ when z_q lies in D .

Let $A = D \setminus \text{int}(f^m(D))$; this is the same fundamental annulus described above. Let \mathcal{D}_1 denote all orbits $\hat{z} \in \mathcal{G}$ for which $z_0 \in A$. Let \mathcal{D}_2 denote all orbits $\hat{z} \in \mathcal{G}$ for which $z_0 \in f^m(D)$ and $z_{-n} \notin D$ for $n > 0$. Then modulo the action of \hat{f} every orbit is uniquely represented in $\mathcal{D}_1 \cup \mathcal{D}_2$, except for some identifications on the boundaries. Since both \mathcal{D}_1 and \mathcal{D}_2 are compact, it follows that \mathcal{G}/\hat{f} is compact.

Leau (parabolic) domains. For a cycle of domains with a parabolic periodic point, the quotient of the corresponding lamination is not compact. One should think of these as obtained from the solenoidal Riemann surface laminations by a "pinching", but we will not try to elaborate on this case in this paper.

4. THE TYPE PROBLEM AND AFFINE STRUCTURE ON THE LEAVES.

By Lemmas 3.3 and 3.4 every leaf of \mathcal{R}_f is either a parabolic (affine) or hyperbolic plane, except possibly for (invariant lifts of) Herman rings, which are hyperbolic annuli. Siegel disks are the only example we know of hyperbolic planes.

Type Problem. *Are there any other cases of hyperbolic leaves except Siegel disks and Herman rings?*

4.1. **Criteria for parabolicity of leaves.** Let us look at the type problem in some special cases.

Repelling fixed point. Let α be a repelling fixed point for f with multiplier λ , and $\hat{\alpha} = (\alpha, \alpha, \dots)$ be its invariant lift to \mathcal{N}_f . Let us consider the invariant leaf $L(\hat{\alpha}) = \{\hat{z} : z_{-n} \rightarrow \alpha\}$ through $\hat{\alpha}$. This leaf is parabolic since the quotient of $L \setminus \{\hat{\alpha}\}$ by the action of \hat{f} is a torus. Similar reasoning applies to the case of a repelling periodic point.

Parabolic fixed point. Let now α be a parabolic fixed point with combinatorial rotation number p/q . Then f^q has $s = ql$ invariant repelling petals P_i . Let us consider the set $L_i = L_i(\hat{\alpha})$ consisting of backward orbits \hat{z} such that the suborbit z_{-qn} , $n = 0, 1, \dots$, eventually lands in P_i . (Observe that $\hat{\alpha}$ itself does not belong to these leaves.) The map \hat{f} permutes the leaves L_i organizing them into cycles of order q . These leaves are parabolic since their quotients by the \hat{f}^q -action are ‘‘Ecalte-Voronin cylinders’’ with infinite modulus (that is, conformally equivalent to \mathbb{C}^*). The case of parabolic periodic points is treated similarly.

General conditions. Let us now give a couple of general conditions for a leaf to be parabolic. Let $D(z, \epsilon)$ denote the spherical disk of radius ϵ centered at z , and $\hat{D}(\hat{z}, \epsilon)$ denote the component of $L(\hat{z}) \cap \pi^{-1}D(z, \epsilon)$ containing \hat{z} .

Lemma 4.1. *Let a backward orbit $\hat{z} = \{z_0, z_{-1}, \dots\} \in \mathcal{R}_f \setminus$ (rotation sets) satisfy the following property. There is an $\epsilon > 0$ and a subsequence $\{n(k)\}$ such that the disk $D(z_{-n(k)}, \epsilon)$ can be univalently pulled back along the rest of the orbit, $\{z_{-m}\}_{m \geq n(k)}$. Then the leaf $L(\hat{z})$ is parabolic.*

Remark. In terms of the natural extension the assumption of the lemma means that the $\hat{D}_{-k} \equiv \hat{D}(\hat{f}^{-n(k)}\hat{z}, \epsilon)$ univalently project down to the sphere.

Proof. Assume without loss of generality that $n(0) = 0$. By the Shrinking Lemma, $\text{diam}(\pi_{-m}\hat{D}_0) = \delta(m) \rightarrow 0$ as $m \rightarrow \infty$. Hence for sufficiently large k the annulus $\hat{D}_{-k} \setminus \hat{f}^{-n(k)}\hat{D}_0$ is univalently mapped to an annulus on the sphere containing a round annulus with outradius ϵ and inradius $\delta(n(k))$. Its modulus can therefore be estimated via

$$\text{mod}(\hat{D}_{-k} \setminus \hat{f}^{-n(k)}\hat{D}_0) \geq \frac{1}{2\pi} \log(\sin c\epsilon/\delta(n(k))) \rightarrow \infty,$$

where the constant c accounts for distortion between spherical and Euclidean metrics. This is equal to the modulus of its univalent image, $\hat{A}_k = \hat{f}^{n(k)}(\hat{D}_{-k} \setminus \hat{D}_0)$, which is an annulus in $L(\hat{z})$ surrounding \hat{D}_0 . Since $\text{mod}(\hat{A}_k) \rightarrow \infty$, the leaf $L(\hat{z})$ must be parabolic. \square

Recall that C denotes the set of critical points of f . The following is an immediate consequence of Lemma 4.1.

Corollary 4.2. *If a backward orbit $\hat{z} = \{z_0, z_{-1}, \dots\} \in \mathcal{R}$ does not converge to $\omega(C)$, then the leaf $L(\hat{z})$ is parabolic.*

Note that the set C can be replaced here by the set C_r of recurrent critical points.

Lemma 4.3. *Let $\hat{z} \in \mathcal{R}_f$. Assume that for some sequence $n(k)$ there exist annuli $\hat{A}_{-n(k)} \subset L_{-n(k)} = L(\hat{f}^{-n(k)}\hat{z})$ enclosing $\hat{f}^{-n(k)}\hat{z}$ and a branched point of the projection $\pi : L_{-n(k)} \rightarrow \bar{\mathbb{C}}$, whose moduli stay away from 0. Then the leaf $L(\hat{z})$ is parabolic.*

Proof. Let $B_{-n} \subset L_{-n}$ be the set of branched points for the projection $\pi : L_{-n} \rightarrow \bar{\mathbb{C}}$. Since every branched point is represented by a backward orbit finitely many times passing through a critical point, $\hat{f}^{-1}B_{-n} \supset B_{-(n+1)}$, and moreover for any $\hat{c} \in B_0$ there is an n such that $\hat{f}^{-n}\hat{c}$ is not a branched point any more. Let $P_n = \hat{f}^n B_{-n} \subset L_0$. Then $P_0 \supset P_1 \supset P_2 \supset \dots$, and $\cap P_n = \emptyset$. As P_0 is discrete, the sets P_n escape to ∞ . Thus, if the leaf L_0 were hyperbolic, then the modulus of any annulus R_n enclosing \hat{z} and a point of P_n would tend to 0 as $n \rightarrow \infty$, which would contradict our assumption. \square

Lemma 4.4. *Consider a backward orbit $\hat{z} = \{z_0, z_{-1}, \dots\} \in \mathcal{R}$ which does not hit the set $\omega(C)$. Assume that $\|Df^{-n}(z)\| \rightarrow 0$ as $n \rightarrow \infty$ where f^{-n} is the branch of the inverse map which sends z to z_{-n} , and $\|\cdot\|$ means the hyperbolic metric in $\mathbb{C} \setminus \omega(C)$. Then the leaf $L(\hat{z})$ is parabolic.*

Proof. Let $L_{-n} = L(\hat{f}^{-n}\hat{z})$. Then the projection

$$\pi : L_{-n} \setminus \pi^{-1}\omega(C) \rightarrow \bar{\mathbb{C}} \setminus \omega(C)$$

is a covering map, and hence a local hyperbolic isometry (with respect to the corresponding hyperbolic metrics).

Assume now that the leaf L_0 is hyperbolic. Then all L_{-n} are also hyperbolic. Since the inclusion $i : L_{-n} \setminus \pi^{-1}\omega(C) \rightarrow L_{-n}$ is a hyperbolic contraction, the projection π is expanding from the hyperbolic metric of L_{-n} to the hyperbolic metric of $\bar{\mathbb{C}} \setminus \omega(C)$.

Note finally that $\hat{f}^{-n} : L_0 \rightarrow L_{-n}$ is a hyperbolic isometry. Hence $\|Df^{-n}(z)\| \geq \|D\pi(\hat{z})\|^{-1} > 0$ where the last norm is measured from the hyperbolic metric on L_0 to the hyperbolic metric of $\bar{\mathbb{C}} \setminus \omega(C)$. Contradiction. \square

Remark: We don't know whether the above contracting property along the backward orbits is always satisfied (unlike the expansion property along the forward orbits: see McMullen [30], Theorem 3.4). See also Lemma 4.6 below.

Axiom A case. Let f satisfy Axiom A. Let us consider a backward orbit $\hat{z} = \{z_{-n}\} \in \mathcal{R}_f$. Then this backward orbit converges to the Julia set, and hence stays bounded distance away from $\omega(C)$. By Corollary 4.2, all leaves of \mathcal{R}_f are parabolic.

Critically non-recurrent case.

Proposition 4.5. *Assume that all critical points on the Julia set are non-recurrent. Then*

$$\mathcal{A}_f^{\text{p}} = \mathcal{R}_f = \mathcal{N}_f \setminus \{\text{attracting and parabolic cycles}\},$$

so that all regular leaves are parabolic.

Proof. The second equality $\mathcal{R}_f = \mathcal{N}_f \setminus \{\text{attracting and parabolic cycles}\}$ was proved above (Corollary 3.7).

In order to prove the first one, let us consider the following ordering on the set of critical points in $J(f)$: $c_1 \succ c_2$ if $\text{cl}(\text{orb}(c_1)) \ni c_2$. Given a $\hat{z} \in \mathcal{R}_f$, let \tilde{C} denote the set of critical points belonging to $\alpha(\hat{z})$.

Assume first that $\tilde{C} \neq \emptyset$. Then let us take a critical point $a \in \tilde{C}$ which is a maximal element of this ordering. Let $\epsilon > 0$ be such that z_{-n} stay distance at least ϵ from all critical points $c \notin \tilde{C}$. By Mañé's theorem (Appendix 2), there is a $\delta > 0$ such that for all n all components of $f^{-n}D(a, \delta)$ have diameter at most ϵ . Hence if $z_{-k} \in D(a, \delta)$, and we pull $D(a, \delta)$ back along $\{z_{-(k+n)}\}_n$, then we don't hit the critical points $c \notin \tilde{C}$. Clearly we will not hit the critical points c of \tilde{C} either (provided δ is small enough), since their forward orbits don't accumulate on a . Hence this pull back is univalent.

Select now a sequence $k(l)$ such that $z_{-k(l)} \rightarrow a$, and apply Lemma 4.1. (Note that the lemma applies since there can be no rotation domains in the non-recurrent case).

If $\tilde{C} = \emptyset$ then take any point $a \in \alpha(\hat{z})$, and repeat the above argument. \square

Remark. By a minor modification of the above argument one can check that *the leaf $L(\hat{z})$ is parabolic, provided \hat{z} is not an attracting or parabolic cycle, and $\alpha(\hat{z})$ is not contained in $\omega(C_r)$, where C_r is the set of recurrent critical points.*

Invariant measures with positive characteristic exponent. Let μ be an invariant measure of f , and suppose that for μ -a.e. z the characteristic exponent

$$\chi(z) = \lim \frac{1}{n} \log |Df^n(z)|$$

exists and is positive ($|\cdot|$ means the spherical norm).

Let $\hat{\mu}$ be the lift of μ to the natural extension. The Pesin local unstable manifolds for $\hat{\mu}$ are the sets $\hat{D}(\hat{z}, \epsilon(\hat{z})) \subset L(\hat{z})$ which univalently project down to the sphere and whose backward orbits shrink exponentially. Moreover, $\epsilon(z) > 0$ $\hat{\mu}$ -a.e.

Let $X_\epsilon = \{\hat{z} : \epsilon(\hat{z}) > \epsilon\}$. It follows from the Poincaré recurrence theorem that for $\hat{\mu}$ -a.e. \hat{z} there is an $\epsilon > 0$ such that the backward orbit $\hat{f}^{-n}\hat{z}$ infinitely many times visits X_ϵ . By Lemma 4.1 the leaves $L(\hat{z})$ are parabolic for $\hat{\mu}$ -almost all \hat{z} (compare [4], [53]).

Infinitely renormalizable quadratics. We refer to the papers of Douady and Hubbard [16] and McMullen [30], [31] for the background in holomorphic renormalization theory. Here we will briefly recall the basic concepts.

Let U' and U be two topological disks such that $\text{cl } U' \subset U$. A double branched covering map $f : U' \rightarrow U$ is called *quadratic-like*. We assume that its critical point is located at the origin 0. The set $K(f) = \{z : f^n z \in U', n = 0, 1, \dots\}$ is called the *filled Julia set*; its boundary is called the Julia set $J(f)$. The Julia set is connected iff the critical point 0 is non-escaping, that is, $0 \in K(f)$.

Any quadratic polynomial can be viewed as a quadratic-like map with U being a round disk of sufficiently large radius, and U' being its pullback. By the Straightening Theorem of Douady and Hubbard any quadratic-like map $f : U' \rightarrow U$ is quasi-conformally conjugate to some quadratic polynomial $z \mapsto z^2 + c$. Moreover, if $\text{mod}(U' \setminus U) \geq \epsilon > 0$ then there is a conjugacy with dilatation bounded by $K(\epsilon)$.

We can specify a distinguished fixed point of f as follows: Take an arc $\gamma \subset U \setminus K$ with endpoints a and $f(a)$. Choosing appropriate pullbacks of this arc by f we obtain a curve $\Gamma \supset \gamma$ such that $f(\Gamma \setminus \gamma) = \Gamma$. It turns out that if $J(f)$ is connected then this curve lands at a specific fixed point of f , usually denoted by β . This point is repelling for any quadratic polynomial $z \mapsto z^2 + c$ except $c = 1/4$.

A quadratic-like map f is called *renormalizable* under the following circumstances:

- Some iterate $g = f^p$ ($p > 1$) restricted to an appropriate topological disk $U' \ni 0$ is quadratic-like with connected Julia set;
- The sets $f^k K(f)$, $k = 1, \dots, p-1$, do not touch $K(f)$ except perhaps for the β -fixed point of g .

Under these circumstances the map g is called a *renormalization* of f . If there is a sequence of renormalizations $g_n : U'_n \rightarrow U_n$ with increasing periods p_n , the map f is called *infinitely renormalizable*. If this sequence can be selected in such a way that the ratios p_{n+1}/p_n are bounded, then one says that f is of *bounded type*.

We say that f is an infinitely renormalizable map with *a priori bounds* if there is a sequence of renormalizations as above and an $\epsilon > 0$ such that

$$\text{mod}(U_n \setminus U'_n) \geq \epsilon, \quad n = 0, 1, \dots$$

Let us say that a map is *Feigenbaum-like* if it is infinitely renormalizable of bounded type with a priori bounds. Any infinitely renormalizable real quadratic of bounded type is *Feigenbaum-like*: complex a priori bounds were established by Sullivan (see [47], [35]).

For a Feigenbaum-like map the set $\omega_f(0)$ is a Cantor set of bounded geometry, and $f|_{\omega_f(0)}$ is an invertible minimal dynamical system (conjugate to a translation on a group). In particular, f is persistently recurrent and hence, by Corollary 3.8, $\mathcal{R}_f = \mathcal{N}_f \setminus \hat{\omega}_f(0)$.

Lemma 4.6. *Let f be a Feigenbaum-like quadratic polynomial. Then all leaves of the lamination \mathcal{R}_f are parabolic.*

Proof. Let $\hat{z} = \{z_0, z_{-1}, \dots\} \in \mathcal{R}$. Then this orbit eventually stays out of the set $\omega(0)$, so we can assume that $z_0 \in \mathbb{C} \setminus \omega(0)$.

Let $g : U' \rightarrow U$ be a renormalized map, $\text{mod}(U \setminus U') > \epsilon > 0$. Let β be its distinguished fixed point, as above. Since f is of bounded type, g is $K(\epsilon)$ -quasi-conformally conjugate to a polynomial $z \mapsto z^2 + c$ with $|c - 1/4| > \delta > 0$ (with δ depending on the type and a-priori bound). Hence

$$\|Dg(\beta)\| \geq \lambda > 1. \tag{4.1}$$

Because of Corollary 4.2, we can assume that z_{-n} converge to $\omega(0)$. Let $\omega_g(0)$ be the closure of the postcritical set of g . Then there is a backward orbit ζ_{-l} of g converging to $\omega_g(0)$ which is a part of the backward orbit \hat{z} . Let us take the second element ζ_{-1} of this backward orbit. Clearly $\zeta_{-1} \in U' \setminus U''$, where U'' is the g -pullback of U' .

The set of Feigenbaum-like maps is compact in the Caratheodory topology (see McMullen [30]). Hence there is a path γ in $U \setminus \omega_g(0)$ joining ζ_{-1} and β of bounded hyperbolic length, such that the analytic continuation of g^{-1} which fixes β carries ζ_{-1} to ζ_{-2} . It follows from this and (4.1) that $\|Dg(\zeta_{-2})\| \geq \theta > 1$ for some θ depending on ϵ , λ and the hyperbolic length of γ only.

Hence $\|Df^{-n}(z)\| \rightarrow 0$ as $n \rightarrow \infty$ (where $f^{-n}z = z_{-n}$), and Lemma 4.4 yields the desired result. \square

Remark: The above way to get hyperbolic contraction is, modulo the details, due to Curt McMullen (compare [31], Proposition 5.9). It is actually possible to weaken the assumptions of the lemma: McMullen has an argument showing that his notion of “robustness” suffices to give the desired contraction.

4.2. Affine structures on the leaves and linearization. Being unable to resolve the type problem in full generality, let us define a new leaf space $\mathcal{A}_f^\mathbb{A}$ by throwing away from \mathcal{R}_f all hyperbolic leaves. All leaves in $\mathcal{A}_f^\mathbb{A}$ are conformally equivalent to the complex plane \mathbb{C} and hence possess a unique affine structure compatible with their conformal structure.

We can express this affine structure as a limit of rescalings of backward branches of f :

Lemma 4.7. *Let f be a rational map with ∞ a critical point. Given $\hat{z} \in \mathcal{A}_f^\mathbb{A}$, there exists a sequence of similarities $A_n(w) = \alpha_n w + \beta_n$ such that the maps*

$$\phi_n = A_n \circ \pi \circ \hat{f}^{-n} : L(\hat{z}) \rightarrow \mathbb{C}$$

converge (uniformly on compact sets) to a conformal isomorphism $\phi : L(z) \rightarrow \mathbb{C}$.

Remark: The condition that ∞ is critical can always be arranged by conjugation with an appropriate Möbius transformation. For polynomials it is automatic.

Proof. Take a disk neighborhood $\hat{U} = (U_0, U_{-1}, \dots)$ of \hat{z} in $L(\hat{z})$ with compact closure. Since the leaf $L(\hat{z})$ is parabolic, for any $M > 0$ \hat{U} is contained in a disk $\hat{V} = \hat{V}(M)$ with modulus $\text{mod}(\hat{V} \setminus \hat{U}) = M$.

Let $l = l(M)$ be such that $\pi_{-n} = \pi \circ \hat{f}^{-n}$ is univalent on \hat{V} for $n \geq l$ (possible by definition of \mathcal{R}_f). Thus for $n > l$, $V_{-n} = \pi_{-n}(\hat{V})$ contains no critical points and in particular lies in \mathbb{C} . Choose the similarity A_n such that $A_n(z_{-n}) = 0$ and $A'_n(z_{-n}) = (\pi'_{-n}(\hat{z}))^{-1}$. Therefore $\phi_n = A_n \circ \pi_{-n}$ have been normalized by $\phi_n(\hat{z}) = 0$, $\phi'_n(\hat{z}) = 1$. (In order for these derivatives to make sense on L we should fix some local coordinate chart).

For $n > l$ and $k > 0$, we can write $\phi_{n+k} = \phi_n \circ G_{n,k}$, where $G_{n,k} = A_{n+k} \circ f_n^{-k} \circ A_n^{-1}$, with f_n^{-k} denoting the branch of f^{-k} taking V_{-n} to V_{-n-k} . Note that $G_{n,k}$ is defined and univalent on $A_n(V_{-n}) = \phi_n(\hat{V})$, and is normalized so that $G_{n,k}(0) = 0$, $G'_{n,k}(0) = 1$. By the Koebe 1/4 theorem

$\phi_n(\hat{U})$ contains a disk of definite radius $\delta > 0$. Since ϕ_n is univalent the modulus of $\phi_n(\hat{V}) \setminus \phi_n(\hat{U})$ is M , so by the Koebe distortion theorem (see appendix 2) the nonlinearity of $G_{n,k}$ on the δ -disk around 0 is small, and goes to 0 as $M \rightarrow \infty$, independently of k .

Letting M , and therefore l and n , go to ∞ , it follows that $G_{n,k} \rightarrow \text{id}$ uniformly on a $\delta/2$ disk around 0 as $n \rightarrow \infty$. Thus $\{\phi_n\}$ form a Cauchy sequence, and so converge uniformly on a neighborhood of \hat{z} . Since $\{\phi_n|_{\hat{U}}\}$ is a normal family, they must converge on all of \hat{U} .

Applying this argument to a sequence of disks \hat{U}_m exhausting $L(\hat{z})$, we conclude that ϕ_n converge uniformly on compact sets to a global map $\phi : L(\hat{z}) \rightarrow \mathbb{C}$, which is univalent. Since $L(\hat{z})$ is parabolic its image must be all of \mathbb{C} , so ϕ is an isomorphism. \square

In the particular case where π is already univalent on a leafwise neighborhood of $\hat{z} \in \mathcal{R}_f$ (i.e. no z_{-n} is a critical point for $n > 0$), we can identify this neighborhood with a neighborhood of z_0 and obtain this local formula for the affine chart:

$$\phi_z(\hat{\zeta}) = \lim_{n \rightarrow \infty} (f^n)'(z_{-n})(\zeta_{-n} - z_{-n}) \quad (4.2)$$

(if f is appropriately normalized, e.g., if ∞ is critical). In the case of a leaf corresponding to a repelling fixed point this exactly corresponds to the classical formulas for the linearizing coordinate. Note however that *uniform* expansion is not necessary for this formula to hold.

Namely, if α is a repelling fixed point, then the affine map $\phi : L(\hat{\alpha}) \rightarrow \mathbb{C}$ is given by the classical Königs linearizing function:

$$\phi(\hat{\zeta}) = \lim \lambda^{-n}(\zeta_{-n} - \alpha).$$

Note that $\phi(\hat{f}^{-1}\hat{\zeta}) = \lambda^{-1}\phi(\hat{\zeta})$, $\zeta \in L(\hat{\alpha})$, so that ϕ conjugates \hat{f}^{-1} on the leaf to the linear map $z \mapsto \lambda^{-1}z$.

Let now α be a parabolic fixed point with combinatorial rotation number p/q . An explicit affine map from the associated leaves $L_i(\hat{\alpha})$ to \mathbb{C} is given by the Leau-Fatou linearizing function:

$$\phi(\hat{\zeta}) = \lim (h(\frac{1}{(\zeta_{-nq} - \alpha)^s}) - n),$$

where s is the number of petals at α , and h is an appropriate local chart at a sectorial region at ∞ (compare Milnor [36], §7). This function conjugates \hat{f}^{-q} on the leaf to a translation $z \mapsto z + a$. This corresponds to a variation on the construction in Lemma 4.7, where the rescaling map A_n is precomposed with a fixed local chart in $\bar{\mathbb{C}}$, in this case $w \mapsto h(1/(w - \alpha)^s)$.

In general, affine structure on the leaves of \mathcal{A}_f^n can be viewed as a *simultaneous linearization* of the dynamics along the backward orbits. Indeed, \hat{f} becomes an affine map between the leaves. In the affine local charts (4.2) these maps become just multiplications by the derivative at the base point:

$$\phi_{f\hat{z}}(\hat{f}\hat{\zeta}) = f'(z) \cdot \phi_{\hat{z}}(\hat{\zeta}), \quad (4.3)$$

provided no z_{-n} is critical for $n > 0$.

4.3. Density of leaves. Let us say that a leaf space X is *minimal* if all leaves are dense in X .

Lemma 4.8. *Any parabolic leaf L is dense in \mathcal{N}_f . Thus the leaf space \mathcal{A}_f^n is minimal. Moreover, $\mathcal{J}^r \cap L$ is dense in the pullback $\pi^{-1}J$ of the Julia set to \mathcal{N}_f .*

Proof. Since π_{-n} is a non-constant analytic map on the parabolic leaf L , it can miss at most two points in $\bar{\mathbb{C}}$. Now consider any $\hat{z} \in \mathcal{N}_f$, and large $n > 0$. Since $\pi_{-n}(L)$ is dense, there is some $\hat{w} \in L$ with w_{-n} as close as we like to z_{-n} . If it is sufficiently close then the spherical $\text{dist}(w_{-j}, z_{-j})$ will be small for all $0 \leq j \leq n$. Thus L is dense. If $\hat{z} \in \pi^{-1}J$ then clearly \hat{w} can be selected from \mathcal{J}^r . \square

We remark that it seems plausible that L is dense even if it is hyperbolic, provided that it is not a rotation domain.

Let \mathcal{J}_f^n denote $\mathcal{J}_f^r \cap \mathcal{A}_f^n$, the Julia set in the affine leaf space.

Corollary 4.9. *The Julia set $\mathcal{J}_f^n \subset \mathcal{A}_f^n$ is compact if and only if f is critically non-recurrent.*

Proof. If f is critically non-recurrent then by Proposition 4.5 $\mathcal{J}_f^n = \mathcal{J}_f^r = \pi^{-1}(J)$, which is a closed subset of \mathcal{N}_f , and thus compact.

Otherwise, by Lemma 3.6, f has irregular points on $\pi^{-1}J$. On the other hand, by Lemma 4.8, $\mathcal{J} = \mathcal{J}_f^n$ is dense in $\pi^{-1}J$. Hence \mathcal{J} is not closed in \mathcal{N}_f , thus not compact. \square

4.4. Local leaves on a global leaf. Let $\hat{\alpha}$ be a repelling periodic point, and $L = L(\hat{\alpha})$ the leaf of $\hat{\alpha}$. Let $D \subset \bar{\mathbb{C}}$ be a topological disk which does not contain $\alpha = \pi(\hat{\alpha})$, and let Δ be a topological disk compactly contained in D . Let \hat{D}_i be the connected components of $\pi^{-1}(D) \cap L$ which univalently project down onto D , and let $\hat{\Delta}_i \subset \hat{D}_i$ be the corresponding components of $\pi^{-1}(\Delta) \cap L$.

The following lemma is a “natural extension” of the Shrinking Lemma (see Appendix), and will be applied in §5.

Lemma 4.10. *The size of the $\hat{\Delta}_i$ shrinks relative to their distance to $\hat{\alpha}$:*

$$\frac{\text{diam}_L \hat{\Delta}_i}{\text{dist}_L(\hat{\Delta}_i, \hat{\alpha})} \rightarrow 0, \quad \text{as } i \rightarrow \infty, \quad (4.4)$$

where diam_L and dist_L are measured in any uniformizing chart $\phi : L \rightarrow \mathbb{C}$.

Remarks. 1. Clearly the ratio in (4.4) does not depend on the choice of uniformizing map ϕ .

2. The result is still valid if we take all components \hat{D}_i with some uniform bound on their branching over D .

Proof. Clearly we can assume that $\hat{\alpha}$ is fixed. Note then that the $\hat{\Delta}_i$ escape to infinity in L (that is, eventually don't intersect any given leaf-compact subset in L), since they have disjoint collars $\hat{D}_i \setminus \hat{\Delta}_i$ of definite modulus.

Let $\hat{U} \subset L$ be a leafwise neighborhood of $\hat{\alpha}$ such that $\text{cl} \hat{U} \subset \hat{f}(\hat{U})$, and $\hat{f}\hat{U}$ is univalently projected down onto $U \subset \bar{\mathbb{C}}$. Let n_i be the first positive integer such that $\hat{f}^{-n_i}(\hat{\Delta}_i) \subset \hat{U}$. As Δ_i escape to infinity in L , n_i goes to ∞ .

The disks $\pi \hat{f}^{-n_i}(\hat{\Delta}_i)$ are univalent pullbacks of Δ which are not contained in a rotation domain, so by the Shrinking lemma their (spherical) diameters go to 0 as $i \rightarrow \infty$. As $\pi|_{\hat{U}}$ has bounded distortion (from the affine structure on \hat{U} to the spherical structure on U), we have

$$\frac{\text{diam}_L(\hat{f}^{-n_i} \hat{\Delta}_i)}{\text{dist}_L(\hat{f}^{-n_i} \hat{\Delta}_i, \hat{\alpha})} \rightarrow 0 \quad \text{as } i \rightarrow \infty.$$

But since \hat{f} preserves the affine structure on L , the ratio in the last equation is equal to the ratio in (4.4). \square

5. POST-CRITICALLY FINITE MAPS.

The affine leaf space \mathcal{A}_f^n which we have constructed so far is not, in general, a lamination. The missing ingredients are both topological – the lack of a local product structure – and geometric – non-continuity of the affine structures in the transverse direction, even where there is a product structure. As we shall see, these two problems are related.

In this section we will give an explicit rearrangement of \mathcal{A}_f^n – a change of topology and the addition of new leaves – in the special case of post-critically finite rational maps. This should serve

as a motivating example, an indication of the kind of structure that arises, and a demonstration of how orbifold leaves appear in a natural way.

In §7, we will give a completely general construction of an affine orbifold lamination for any rational map, emerging naturally from the affine group action on a space of meromorphic functions. Thus one could read that section without first reading this one, but the reader may find that the explicit examples given here help to illuminate the more abstract approach.

We will first construct the orbifold lamination topologically, and then discuss continuity of affine structures.

5.1. Topological orbifold lamination. To fix ideas, assume for the moment that f is a postcritically finite quadratic polynomial, and moreover that the critical point is actually pre-fixed: there is a fixed point α such that $f^l c = \alpha$ for some $l > 1$. (We will discuss the general postcritically finite case in §5.3). It is standard that α is a repelling fixed point (see discussion at the end of §5.3).

As usual, let $\hat{\alpha} = \{\alpha, \alpha, \dots\}$ denote the invariant lift of α to \mathcal{N}_f , and let $L \equiv L(\hat{\alpha})$ denote the \hat{f} -invariant leaf of $\hat{\alpha}$ in \mathcal{N}_f .

Recall that \mathcal{A}_f^n and \mathcal{R}_f are both equal to $\mathcal{N}_f \setminus \hat{\alpha}$ (Lemma 3.7). Our orbifold lamination \mathcal{A}_f will consist of \mathcal{A}_f^n , with the leaf L replaced by two copies named L^r and L^s . The topology τ_ℓ , and orbifold structure, are described as follows.

Let $q : \mathcal{A}_f \rightarrow \mathcal{R}_f$ be the map that re-identifies L^r and L^s . Let us consider the pull-back topology $q^{-1}\tau_n$, where τ_n is the natural topology of \mathcal{R} as a subset of \mathcal{N} . Note that \hat{f} is naturally lifted to a homeomorphism \tilde{f} of \mathcal{A}_f with this topology. However the pull-back topology is not Hausdorff since it does not separate the leaves L^r and L^s . The actual topology τ_ℓ will be the minimal strengthening of the pullback topology $q^{-1}\tau_n$, which separates these leaves, keeps \tilde{f} as a homeomorphism, and gives \mathcal{A}_f the structure of an orbifold lamination.

Let $\hat{z} = (z_0, z_{-1}, \dots) \in \mathcal{N}_f$, D be a topological disk containing z_0 and at most one postcritical point $f^k c$, $1 \leq k \leq l$, and let N be a natural number. Let $B(D, \hat{z}, N)$ and $B(D, z_0) = \pi^{-1}D \equiv B(D, \hat{z}, 0)$ be the τ_n box neighborhoods defined as in (3.1), and recall that D_0, D_{-1}, \dots are the pullbacks of D along \hat{z} .

If D_{-N} does not intersect the postcritical set then $B(D, \hat{z}, N)$ has a natural product structure $T \times D_{-N}$. Moreover, the projection $\pi : B(D, \hat{z}, N) \rightarrow D$ is either univalent or two-to-one branched covering on all leaves. The latter occurs when D contains a postcritical point $f^k c$, and then this point is the projection of the branched point on any leaf.

This situation always occurs if $\hat{z} \neq \hat{\alpha}$, and N is sufficiently high. It is more complicated for $\hat{z} = \hat{\alpha}$. In this case some of the leaves are univalent and some are branched, so that $B(D, \hat{\alpha}, N)$ does not have a natural box structure.

Let us call a backward orbit \hat{z} with $z_0 = \alpha$ *singular* if it contains c (i.e. it is a branch point of π), and *regular* if it does not contain c , and is not equal to $\hat{\alpha}$.

Given a topological disk $D \ni \alpha$ (not containing other postcritical points), let $B^r(D, \hat{\alpha}, N)$ consist of the union of local leaves in $B(D, \hat{\alpha}, N)$ containing regular orbits, and $B^s(D, \hat{\alpha}, N)$ be the union of local leaves containing singular orbits. These are disjoint open sets in \mathcal{N}_f with a natural product structure. Moreover, together with the local leaf $\hat{D} = \hat{D}(\hat{\alpha})$ containing the fixed point $\hat{\alpha}$, they make up all of $B(D, \hat{\alpha}, N)$. We set $B^\mu(D) \equiv B^\mu(D, \hat{\alpha}, 0)$, where μ stands for r or s .

Given a set $X \subset \mathcal{N}_f$, let $\tilde{X} \subset \mathcal{A}_f$ denote $q^{-1}X$. Let also \tilde{D}^μ denote the component of $q^{-1}(\hat{D})$ lying in the corresponding leaf L^μ . The similar meaning is given to a point $\tilde{z}^\mu \in \tilde{D}^\mu$ corresponding to $\hat{z} \in \hat{D}$.

Let

$$Q^\mu(D, \hat{\alpha}, N) = \tilde{B}^\mu(D, \hat{\alpha}, N) \cup \tilde{D}^\mu \quad \text{and} \quad Q^\mu(D) \equiv Q^\mu(D, \hat{\alpha}, 0). \quad (5.1)$$

These sets are going to be neighborhood bases for points $\tilde{\alpha}^\mu$.

Let us now define the topology τ_ℓ as the minimal strengthening of the pull-back topology $q^{-1}\tau_n$ for which the sets $\tilde{f}^n(Q^\mu(D))$ are open for all $n \geq 0$.

Lemma 5.1. *With the new topology, \mathcal{A}_f is an orbifold lamination with one singular point, $\tilde{\alpha}^s$. The projection $q : \mathcal{A}_f \rightarrow \mathcal{N}_f$ is continuous, and \tilde{f} acts homeomorphically.*

Proof. Given a $\hat{z} \in \hat{D}$ and a topological disk $\Delta \subset D$ containing z_0 , let

$$Q^\mu(\Delta, \hat{z}, N) = Q^\mu(D) \cap \tilde{B}(\Delta, \hat{z}, N). \quad (5.2)$$

When $\hat{z} = \hat{\alpha}$, we go back to the sets $Q^\mu(D, \hat{\alpha}, N)$ introduced above.

Let $\tilde{\mathcal{B}}$ be the family of all sets $\tilde{B}(\Delta, \hat{z}, N)$ for $\hat{z} \in \mathcal{R}_f$ and any disk $\Delta \ni z_0$. Let \mathcal{Q} be the family of sets $Q^\mu(\Delta, \hat{z}, N)$, where $\hat{z} \in \tilde{D}^\mu$ and $\Delta \subset D$ contains z_0 . Let $\mathcal{T} = \tilde{\mathcal{B}} \cup \bigcup_{n \geq 0} \tilde{f}^n \mathcal{Q}$. We claim that \mathcal{T} is a neighborhood basis for the topology τ_ℓ .

All elements of \mathcal{T} are open in τ_ℓ , by definition. We need to check that, for any U, V in \mathcal{T} and $x \in U \cap V$ there is some $W \in \mathcal{T}$ such that $x \in W \subset U \cap V$.

Clearly the sets $\tilde{B}(\Delta, \hat{z}, N)$ form a basis for the pullback topology $q^{-1}\tau_n$. Also, restricting Δ or increasing N without changing other parameters clearly makes a set from \mathcal{T} smaller. Taking additionally into account (5.2), we conclude that it is enough to check the case when $U = \tilde{f}^m Q^\mu(D)$ and $V = \tilde{f}^n Q^\nu(D)$ (Here μ and ν are independently either r or s). By pulling back, we may assume that $m = 0$.

Assume that x does not belong to the local leaf \tilde{D}^μ of $\tilde{\alpha}^\mu$ in $Q^\mu(D)$. Note that $\Omega = Q^\mu(D) \setminus \tilde{D}^\mu$ is open in the pullback topology. Hence $\tilde{f}^{-n}\Omega$ contains a basic $X \ni \tilde{f}^{-n}x$ of family $\tilde{\mathcal{B}}$. By (5.2) $Q^\nu(D) \cap X \in \mathcal{Q}$. Thus $\tilde{f}^n(Q^\nu(D) \cap X) \in \mathcal{T}$ is a desired set W .

Assume now that $x \in \tilde{D}^\mu$. We can select the basis of disks D in such a way that $\hat{f}\tilde{D}^\mu$ overflows \hat{D}^μ . Then $\tilde{f}^{-n}x \in \tilde{D}^\mu$, and hence $\nu = \mu$. Moreover, the component X of $Q^\mu(D) \cap \tilde{f}^{-n}Q^\mu(D)$ containing $\tilde{f}^{-n}x$ is just $Q^\mu(D, \hat{\alpha}, n)$, a set of family \mathcal{Q} . Now the desired statement follows.

It is clear that \mathcal{A}_f is Hausdorff – the doubled points have separating neighborhoods, by the construction. Note that, away from the postcritical points, the topology has been changed only in the fiber direction, where a dense set of fibers has been doubled.

Let us now check that $\tilde{f} : \mathcal{A}_f \rightarrow \mathcal{A}_f$ is a homeomorphism. Obviously \tilde{f}^{-1} is continuous. To verify that \tilde{f} is continuous, it is enough to check that $\tilde{f}^{-1}Q^\mu(D)$ are open. Let $f^{-1}D = D_0 \cup D_1$ where $D_0 \ni \alpha$ while $D_1 \ni f^{l-1}c$. Then

$$\tilde{f}^{-1}Q^\mu(D) = Q^\mu(D_0) \cup B(D_1, f^{l-1}c).$$

Finally let us check that all sets of the basis \mathcal{T} are orbifold boxes. Indeed, all sets $B(D, \hat{z}, N) \approx D \times T$ are regular lamination boxes. Hence the sets $\tilde{B}(D, \hat{z}, N) \approx D \times \tilde{T}$ are also regular boxes with \tilde{T} obtained from T by doubling points corresponding to the leaf L .

The sets $Q^r(D, \hat{\alpha}, N)$ are also regular boxes $D \times T$ with the transversal T consisting of all backward orbits $\alpha, \dots, \alpha, \dots$ (at least N α 's) which never pass through c .

Let us now consider the set K of singular backward orbits $\alpha, \dots, \alpha, \dots$ (at least N α 's) together with the point $a = \bar{\alpha}$ (in the natural extension topology). Then $Q^s(D, \hat{\alpha}, N)$ is homeomorphic to the orbifold box with transversal (K, a) described in Example 2.2.

Finally if $\hat{z} \neq \hat{\alpha}$ and $D \not\ni \alpha$ then the sets $Q^\mu(D, \hat{z}, N)$ are regular boxes with the same transversal K .

Thus \mathcal{A}_f is indeed an orbifold lamination. \square

5.2. Orbifold affine structure. By Corollary 3.7 all leaves of the lamination \mathcal{A}_f are parabolic. Let us supply all leaves except L^s with their unique affine structure. As to the leaf L^s , let us consider a branched double covering $p : \Lambda^s \rightarrow L^s$ with a single branched point over $\tilde{\alpha}^s$. Then Λ^s is

a parabolic plane which hence has a unique affine structure. Pushing this structure down to L^s we obtain an orbifold affine structure on L^s with one singular point at $\tilde{\alpha}^s$.

There is no ambiguity in the above construction as the double covering p is uniquely defined up to pre- and post-compositions with affine maps. So after appropriate selection of the affine coordinates z and ζ on L^s and Λ^s correspondingly, p just becomes the quadratic map $z = \zeta^2$. But z is a linearizing coordinate on the leaf L^s (see §4.2). Thus the *orbifold affine coordinate* ζ on L^s can be viewed as the square root of the linearizing coordinate.

Let S_N denote the family of affine structures on the leaves of \mathcal{A}_f^n , and let S_L denote the family of orbifold affine structures on the leaves of \mathcal{A}_f . Let us also consider the pullback affine structures $q^{-1}S_N$ on the leaves of \mathcal{A}_f . They coincide with S_L on all leaves except the singular leaf L^s .

Lemma 5.2. *The orbifold affine structures S_L on the leaves of \mathcal{A}_f make it an affine orbifold lamination.*

Proof. We need to check that the affine structure depends continuously on the leaf. We will use the box basis of \mathcal{A}_f described above and the explicit formula for the affine coordinates of §4.2.

Take an $x \in \mathcal{A}_f$ with $qx = \hat{z} = (z_0, z_{-1}, \dots) \in \mathcal{N}_f$. Let us first assume that \hat{z} does not lie on the invariant leaf $L = L(\hat{\alpha})$. Then there is a subsequence $z_{-n(k)}$ staying distance at least an $\epsilon > 0$ from the postcritical set.

Take now a neighborhood $D \ni z_0$ containing at most one point of the postcritical set. Let D_{-k} denote the pullback of D along \hat{z} . Let us consider boxes $B_n = B(D, \hat{z}, n)$. Take a big k and let $\hat{\zeta} \in B_{n(k)}$. By Lemma 4.7 the affine structure on the leaf $\hat{D}(\hat{\zeta})$ of such a box is given by rescaling $\pi_{-m} = \pi \circ \hat{f}^{-m}$ and passing to limit. But for $m > n(k)$, $\pi_{-m} = f^{-(m-n(k))} \circ \pi_{-n(k)}$ for an appropriate branch of the inverse function. Since $\text{diam}(D_{-n})$ is small for n sufficiently large, and $f^{-(m-n(k))}$ allows analytic extension in the ϵ neighborhood of $z_{-n(k)}$, by the Koebe Distortion Theorem it is almost linear on $D_{-n(k)}$ uniformly in $\hat{\zeta}$. It follows that the variation of the affine structure on the leaves of $B_{n(k)}$ is small, provided k is sufficiently large.

Hence the variation of the pullback structure $q^{-1}S_N$ on the leaves of the box $\tilde{B}_{n(k)} \equiv \tilde{B}(D, \hat{z}, n(k))$ is also small for big k . Thus the variation of the structure S_L on all regular leaves of $\tilde{B}_{n(k)}$ is small as well.

Let now $y \in \tilde{B}_{n(k)} \cap \tilde{L}^s$, and $\tilde{D}^s(y)$ be the singular local leaf of y in $\tilde{B}_{n(k)}$. Let $\phi : (\tilde{L}^s, \tilde{\alpha}) \rightarrow (\mathbb{C}, 0)$ be a regular uniformization of \tilde{L}^s . Then according to the discussion preceding this lemma, an orbifold affine chart on $\tilde{D}^s(y)$ is given by $\sqrt{\phi}$. But the image $\phi(\tilde{D}^s(y))$ escapes to ∞ in \mathbb{C} when $y \rightarrow x$. Moreover, by Lemma 4.10 its size relative to the distance to the origin is vanishing. Hence the non-linearity of the square root map on this set goes to zero. Thus the affine structure S_L on $\tilde{D}^s(y)$ is close to the pullback structure $q^{-1}S_N$ on this local leaf. Consequently, it is close to the affine structure on the leaf $\tilde{D}(x)$ when y is close enough to x . We are done with the case when $\hat{z} \notin L$.

Let now $\hat{z} \in L$, so that $x \in \tilde{L}^\mu$ for $\mu = r$ or $\mu = s$. We wish to check that the affine structures S_L on leaves $\tilde{\Delta}(y)$ of a box $f^n Q^\mu(\Delta, \hat{z}, N)$ defined by (5.2) approach the affine structure on $\tilde{\Delta}(x)$. By pulling back and enlarging the box, we see that it is enough to check this for $x = \tilde{\alpha}^\mu$ and boxes $Q^\mu(D)$ defined in (5.1).

Let first $x = \tilde{\alpha}^r$. Let us consider a regular orbit

$$\hat{\zeta} = (\alpha, \dots, \alpha, \zeta_{-(N+1)}, \dots) \in B^r(D, \hat{\alpha}, N).$$

Then the inverse branches of $f^{-(n-N)} : (D_{-N}, \alpha) \rightarrow (D_{-n}, \zeta_{-n})$ along $\hat{\zeta}$ allow a uniform $\epsilon > 0$ -enlargement, and hence have small non-linearity for big N . It follows that the affine structure S_N on the local leaf $\hat{D}(\hat{\zeta})$ is close to the regular affine structure on $\hat{D}^r(\hat{\alpha})$ (given by the linearizing

coordinate near α). Now we can pass to the orbifold structures on $Q^r(D)$ in the same way as in the above case $\hat{z} \notin L$.

Finally let $x = \tilde{\alpha}^s$. Let us now consider a singular orbit

$$\hat{\zeta} = (\alpha, \dots, \alpha, \zeta_{-(N+1)}, \dots, \zeta_{-n} = c, \dots) \in B^s(D, \hat{\alpha}, N),$$

where $n = N + l$. Then for sufficiently large N the rescaled branch $f^{-(n-1)} : D \rightarrow D_{-(n-1)} \ni fc$ is close to the linearizing coordinate near α . The next inverse branch $f^{-1} : D_{-(n-1)} \rightarrow D_{-n}$ is almost the square root map (since D_{-n} is small), while all further inverse iterates are almost linear on D_{-n} . It follows that the affine coordinate on the local leaf $\hat{D}^s(\hat{\zeta})$ is close to the square root of the linearizing coordinate, which is exactly the orbifold affine coordinate on $\hat{D}^s(\hat{\alpha})$.

Now we again can pass from the box $B^s(D)$ to $Q^s(D)$ in the same way as above. □

5.3. General post-critically finite construction. Let f be an arbitrary post-critically finite map. If a critical point lands in a cycle then the cycle is either repelling or super-attracting (contains a periodic critical point) – see e.g. [24, Thms. 1.4, 1.6, Prop. 1.11]. In the latter case this cycle is omitted from \mathcal{R}_f . Thus we need only consider repelling cycles. It also follows that there is a uniform bound on the branching index of π at all points in \mathcal{R}_f , since a backward orbit in \mathcal{R}_f can only hit the critical set a bounded number of times. Given a postcritical repelling periodic point α , let us consider all occurring branching indices $1 = d_1(\alpha) < \dots < d_{l(\alpha)}(\alpha)$ of the leaves over α .

A general construction of the orbifold lamination for a post-critically finite map has the following differences as compared with the previous particular case:

- Make $l(\alpha)$ copies of the post-critical periodic leaf $L(\hat{\alpha})$.
- Supply these copies with orbifold structures of degrees $d_i(\alpha)$.
- Organize the leaves of the lamination over $\hat{\alpha}$ into the boxes according to their branching indices and then compactify them by adding the corresponding orbifold leaves. These boxes will be open in the new topology.

5.4. Structure of the Chebyshev and Lattès laminations. Let us consider the quadratic Chebyshev polynomial $p : z \mapsto 2z^2 - 1$, $J(p) = [-1, 1]$. Let $T(z) = z^2$, and $\phi(z) = 1/2(z + 1/z)$. Then $\phi \circ T = p \circ \phi$, so that p is conformally equivalent to T on the quotient space of \mathbb{C}^* by the involution $\sigma : z \mapsto 1/z$.

Then the natural extension \mathcal{N}_p is the quotient of the natural extension \mathcal{N}_T modulo the involution $\hat{\sigma} : (z_0, z_{-1}, \dots) \mapsto (\sigma z_0, \sigma z_{-1}, \dots)$. The only invariant leaf of this involution is the invariant leaf $L = L(\hat{1})$ of \hat{p} . Since $\mathcal{A}_T^n = \mathcal{R}_T$ is a regular affine lamination, we obtain a natural orbifold affine lamination structure on \mathcal{A}_p^n (with one singular leaf L). The orbifold lamination \mathcal{A}_p constructed above is obtained from this one by adding an isolated copy of L (with regular affine structure).

The situation for the higher degree Chebyshev polynomials is completely analogous.

Similarly, the regular leaf associated with the post-critical fixed point of a Lattès map is isolated. Proposition 12.1 shows that these are the only postcritically finite maps with isolated leaves (see Proposition 7.6 for a more general statement). After removing this leaf, the lamination becomes the quotient of the “torus solenoid” (that is, the natural extension of the torus endomorphism) modulo an involution.

6. HYPERBOLIC 3-LAMINATIONS.

6.1. Affine extensions in the abstract. In this section, let us forget the specific construction of Section 5 and take an “axiomatic” approach to what we call affine extensions. The general construction of Section 7 will yield objects of this type.

Let $f : \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a rational map. An *affine extension of f* is an affine (orbifold) 2-lamination \mathcal{A} with simply connected leaves, together with a homeomorphism $\hat{f} : \mathcal{A} \rightarrow \mathcal{A}$, acting by conformal automorphisms on leaves, and a projection $\pi : \mathcal{A} \rightarrow \bar{\mathbb{C}}$, such that

1. $f \circ \pi = \pi \circ \hat{f}$
2. π is continuous, and restricted to any leaf is non-constant and complex-analytic.

Condition (1) immediately implies that π factors through a map $p : \mathcal{A} \rightarrow \mathcal{N}_f$, given by

$$p(\mathbf{z}) = (\pi(\mathbf{z}), \pi\hat{f}^{-1}(\mathbf{z}), \pi\hat{f}^{-2}(\mathbf{z}), \dots).$$

Let $\pi_k = \pi \circ \hat{f}^k$, as usual.

In fact p is continuous by (2), and we immediately see that $p(\mathcal{A})$ is contained in \mathcal{R}_f : on any leaf L , π factors through f^n for any $n > 0$ and so the pullbacks $\pi \circ \hat{f}^{-n}(U)$ for any disk $U \subset L$ with compact closure are eventually unbranched. Thus p restricted to each leaf is a complex analytic map to a leaf of \mathcal{R}_f , and we further conclude that the leaf must be parabolic. Hence $p : \mathcal{A} \rightarrow \mathcal{A}_f^n$.

The construction in §5 yields just such an object, and as in that case the map p need not be injective: it re-identifies the leaves which we separated in our construction.

6.2. Extending to three dimensions. Even before we consider the action on \mathcal{A} we can associate to it a naturally defined \mathbb{H}^3 -lamination \mathcal{H} , by attaching a copy of hyperbolic 3-space, realized as its upper half-space model, to (a finite cover of) every leaf. Since transition maps for affine charts on the leaves of \mathcal{A} are affine, they extend naturally to isometries on the hyperbolic 3-spaces.

In particular given two affine charts $\phi, \phi' : \mathbb{C} \rightarrow L$, the corresponding transition map from \mathbb{H}^3 to \mathbb{H}^3 multiplies heights by the norm of the derivative of $\phi^{-1} \circ \phi'$. Thus we can consider a copy of \mathbb{H}^3 for each chart ϕ , and define the leaf H_L attached to L as the identification of all these copies via the transition maps. However, we prefer to make the following definition, which will be easier to work with:

Consider the group Aff of complex-affine maps $A : \mathbb{C} \rightarrow \mathbb{C}$ (henceforth just “affine”). We can identify the complex plane \mathbb{C} and the hyperbolic space $\mathbb{H}^3 \equiv \mathbb{C} \times (0, \infty)$ (with a preferred point at ∞) as homogeneous spaces for Aff , namely $\mathbb{C} \cong \text{Aff}/\mathbb{C}_*$ and $\mathbb{H}^3 \cong \text{Aff}/S^1$. In other words, consider the projections $p_1 : \text{Aff} \rightarrow \mathbb{C}$ and $p_2 : \text{Aff} \rightarrow \mathbb{H}^3$ given by

$$p_1 : g \mapsto g(0) \in \mathbb{C} \tag{6.1}$$

and

$$p_2 : g \mapsto g(0, 1) = (g(0), |g'|) \in \mathbb{H}^3. \tag{6.2}$$

Fibres of p_1 are orbits of the *right action* of the subgroup $\mathbb{C}_* = \text{Fix}(0) \subset \text{Aff}$, that is $\{z \mapsto \alpha z : \alpha \neq 0\}$. Fibres of p_2 are orbits of the right action of S^1 , that is the group $\{z \mapsto \lambda z : |\lambda| = 1\}$. The left-action of Aff on itself projects to complex-affine maps on \mathbb{C} , and to hyperbolic isometries on \mathbb{H}^3 .

Now suppose first that \mathcal{A} has no orbifold leaves, and for a leaf L consider the set $\{\phi : \mathbb{C} \rightarrow L\}$ of affine isomorphisms (“charts”) from \mathbb{C} to L , which admits a fixed-point-free right-action by Aff . We may identify L with $\{\phi\}/\mathbb{C}_*$ by taking ϕ to $\phi(0)$. The space $\{\phi\}/S^1$ is naturally identified with \mathbb{H}^3 as above, and we call this the hyperbolic leaf H_L associated to L .

Thus, the total space \mathcal{H} may be defined as $\{\phi : \mathbb{C} \rightarrow \mathcal{A}\}/S^1$, where the maps ϕ vary over all charts for leaves of \mathcal{A} . This clearly inherits the structure of a hyperbolic 3-lamination. We will usually write $[\phi]$ for an equivalence class of charts modulo rotation in S^1 .

The same construction works for the orbifold leaves, with the charts replaced by finite coverings. Thus an orbifold affine 2-lamination extends to a hyperbolic 3-orbifold lamination.

One should think of a chart $\phi : \mathbb{C} \rightarrow L$ as determining a point and a choice of *scale* for the leaf L . Changes of scale correspond to vertical motion in the upper half-space model. Indeed, let $e^{\mathbb{R}}$ denote the subgroup of \mathbb{C}_* acting by scaling without rotation. The \mathbb{R} -action induced on \mathcal{A} by

the right-multiplication $r : [\phi] \mapsto [\phi \circ e^r]$ is simply the vertical geodesic flow in each leaf, where r measures arclength and increasing r corresponds to increasing heights in each leaf (as is evident from (6.2)).

Finally, we remark that this extension of an orbifold affine lamination to an orbifold \mathbb{H}^3 -lamination is unique, in the sense that if \mathcal{H}' is another orbifold \mathbb{H}^3 -lamination with a projection $\mathcal{H}' \rightarrow \mathcal{A}$ such that on each leaf, fibres of points are geodesics with a common endpoint at infinity, then \mathcal{H} and \mathcal{H}' are related by an isomorphism fixing \mathcal{A} .

6.3. Proper discontinuity of actions. The action \hat{f} on \mathcal{A} (even without assuming that it projects to a rational map) extends naturally to an action, which we also call \hat{f} , on \mathcal{H} by hyperbolic isometries, namely

$$\hat{f} : [\phi] \mapsto [\hat{f} \circ \phi]. \quad (6.3)$$

It is useful to note that we now have two commuting actions on \mathcal{A} : a \mathbb{Z} -action generated by \hat{f} on the left, and an \mathbb{R} -action, the vertical geodesic flow, generated by $e^{\mathbb{R}}$ on the right. These actions have a certain *coherence*: forward iterates of \hat{f} tend to increase heights, as a result of the general expansive properties of the rational map f . Let us make this precise with the following statement:

Lemma 6.1. *Let $(\mathcal{A}, \hat{f}, \pi)$ be an affine extension of a rational map f , and let \mathcal{H} be the hyperbolic 3-lamination associated to \mathcal{A} . For any two points $p, q \in \mathcal{H}$ there are neighborhoods U_p, U_q for which the following holds: if n_i, r_i are sequences such that*

$$(\hat{f}^{n_i} \circ U_p) \cap (U_q \circ e^{r_i}) \neq \emptyset$$

then $n_i \rightarrow +\infty$ if and only if $r_i \rightarrow +\infty$, and $n_i \rightarrow -\infty$ if and only if $r_i \rightarrow -\infty$.

In other words, whenever high forward/backward iterate of $\mathbf{z} \in U_p$ is comparable with $\zeta \in U_q$ in the sense that these points lie on the same vertical geodesic, the former point is much higher/lower than the latter.

Proof. Represent p by a chart $\phi : \mathbb{C} \rightarrow \mathcal{A}$ and q by a chart $\psi : \mathbb{C} \rightarrow \mathcal{A}$. Recall that $\pi \circ \phi$ is analytic, and hence its image misses at most two points in $\bar{\mathbb{C}}$. Thus there exists an open set W which meets the Julia set J_f , and a disk $D \subset \mathbb{C}$ around 0 such that $W \subset \pi \circ \phi(D)$. Let U_p be small enough that for any $[\phi'] \in U_p$, $\pi(\phi'(D))$ contains W . This is possible since π is continuous in \mathcal{A} . In addition choose U_p small enough that there is some upper bound on the degree of $\pi \circ \phi'$ in D (here by “degree” we mean the maximal degree over any point in the image).

Now we can see that $\pi \circ \hat{f}^n \circ \phi'(D)$ will tend to blow up as $n \rightarrow \infty$, and down (in diameter) as $n \rightarrow -\infty$. Indeed, there is some n_0 such that $f^n(W)$ contains all of J_f for $n > n_0$, and as $n \rightarrow \infty$ the degree of f^n on W increases without bound – hence the same is true for $\pi \circ \hat{f}^n \circ \phi'$ on D , for any $[\phi'] \in U_p$.

To see what happens to $\pi \circ \hat{f}^{-n} \circ \phi'(D) = \pi_{-n}(\phi'(D))$ as $n \rightarrow \infty$, we may invoke the Shrinking Lemma given in Appendix 2, once we observe two things: (1) the degree of f^n on $\pi_{-n}\phi'(D)$ is bounded by the degree of π on $\phi'(D)$, and hence uniformly over U_p . (2) Since every leaf of \mathcal{A} is affine, $\pi_{-n}\phi'(D)$ is eventually outside the closure of the rotation domains, so that $\text{diam}(\pi_{-n}\phi'(D)) \rightarrow 0$ as $n \rightarrow \infty$, uniformly for all $[\phi'] \in U_p$.

Now let us choose U_q such that for $[\psi'] \in U_q$ the degree of $\pi \circ \psi'$ on D is uniformly bounded. Suppose that we have $\hat{f}^{n_i} \circ \phi'_i \equiv \psi'_i \circ e^{r_i}$ in \mathcal{A} for $[\phi'_i] \in U_p$ and $[\psi'_i] \in U_q$. Then if $n_i \rightarrow \infty$ then $r_i \rightarrow \infty$ as well, so that the degree of $\pi \circ \psi'_i$ on $e^{r_i}(D)$ can go to infinity. Conversely, suppose that $n_i \rightarrow -\infty$. Then the above Shrinking Lemma argument implies that $\text{diam} \pi \circ \psi'_i(e^{r_i}D) \rightarrow 0$, and so r_i must go to $-\infty$. The other two implications are similar. \square

Recall that a group action on a space X is *proper* if, for any two points $p, q \in X$, there exist neighborhoods $U \ni p$ and $V \ni q$ so that the subset of group elements g such that $gU \cap V \neq \emptyset$ has compact closure in the group. If the group has the discrete topology this set must be finite, and we say the action is *properly discontinuous*.

A consequence of the Shrinking Lemma, as used in Lemma 6.1, is the following central fact:

Proposition 6.2. *Let $(\mathcal{A}, \hat{f}, \pi)$ be an affine extension of a rational map f , and let \mathcal{H} be the hyperbolic 3-lamination associated to \mathcal{A} . The induced action of \hat{f} on \mathcal{H} is properly discontinuous. Similarly, the vertical geodesic flow on \mathcal{H} is a proper \mathbb{R} -action.*

Proof. Given p and q in \mathcal{H} , choose U_p and U_q as in lemma 6.1. Then in particular $(\hat{f}^n \circ U_p) \cap U_q = \emptyset$ for all but finitely many n . Geometrically, we say that forward iterates of U_p cannot continue to intersect U_q , since their heights are going to infinity whenever comparison is possible.

Similarly, $U_p \cap (U_q \circ e^r) = \emptyset$ for all but a bounded set of r . □

At this moment we can conclude that the quotient \mathcal{H}/\hat{f} is a Hausdorff space which inherits the structure of the hyperbolic orbifold 3-lamination. On the other hand, the quotient via the flow action recovers the affine lamination \mathcal{A} . This duality between \mathcal{H}/\hat{f} and \mathcal{A} will be useful in what follows (see Proposition 8.5).

6.4. Convex Hulls. In analogy with the situation in Kleinian groups, we denote by $\mathcal{C}(\mathcal{J})$ the *convex hull* in \mathcal{H} of the lift $\mathcal{J} = \pi^{-1}(J)$ of the Julia set to \mathcal{A} . This is simply the union of the convex hulls of $(\mathcal{J} \cap L) \cup \{\infty\}$ in H_L for every leaf H_L of \mathcal{H} bounded by a leaf L of \mathcal{A} . The quotient $\mathcal{C}(\mathcal{J})/\hat{f}$ can be called the *convex core* of \mathcal{H}_f/\hat{f} .

Using lemma 12.3 we can obtain the following. Let \mathcal{C}_δ denote the leafwise δ -neighborhood of $\mathcal{C}(\mathcal{J})$.

Corollary 6.3. *For $\delta > 0$, \mathcal{C}_δ inherits the structure of a 3-lamination with boundary. In fact \mathcal{C}_δ is homeomorphic to $\mathcal{H} \cup \mathcal{F}/\hat{f}$, where $\partial\mathcal{C}_\delta$ is taken to the “boundary at infinity” \mathcal{F}/\hat{f} .*

Except when the Julia set of f is smooth, the above holds for $\delta = 0$, and we note also that $\partial\mathcal{C}$ inherits a metric from \mathcal{H}_f which makes it into a hyperbolic 2-lamination.

Proof. Since \mathcal{J}_f is the pullback of J_f by π and π varies continuously in the transverse direction, for any product box $T \times D$, if T is sufficiently small then the intersections of \mathcal{J}_f with the local leaves $\{t\} \times D$ are close to each other in the Hausdorff topology in D . The same applies to the finite covers of orbifold boxes, and hence for any (large) closed disk on a leaf we can take a small transversal neighborhood so that the Julia sets vary only slightly in the Hausdorff topology.

It follows, applying lemma 12.3, that any point $x \in \mathcal{C}_\delta$ has an (orbifold) box neighborhood in \mathcal{H} which intersects \mathcal{C}_δ in a set of the form $T \times C$, up to bilipschitz homeomorphism. Here C is the intersection of \mathcal{C}_δ with a leafwise neighborhood of x .

The homeomorphism from \mathcal{C}_δ to $\mathcal{H} \cup \mathcal{F}/\hat{f}$ comes directly from the leafwise homeomorphism discussed in section 12.2. The case where J_f is smooth corresponds exactly to the case where $\mathcal{J} \cap L$ is contained in a straight line in L (again, the Shrinking Lemma), and this is the only case where the discussion fails for $\delta = 0$. □

The $z^2 + \epsilon$ case. The simplest possible quadratic polynomial is $f(z) = z^2 + \epsilon$ where ϵ is small (more precisely, let ϵ lie in the main cardioid of the Mandelbrot set, so that f has one attracting fixed point). In this case J is a quasi-circle, the Fatou domain lifts to two components in \mathcal{R}_f , each of which has quotient homeomorphic to Sullivan’s solenoidal Riemann surface lamination \mathcal{S} (see Appendix 1), and in fact \mathcal{R}_f is already an affine lamination (Proposition 4.5).

In each leaf, \mathcal{J} is a quasi-line separating the plane into two components, where one component projects to the outside and one to the inside of J in \mathbb{C} . We claim that the convex core \mathcal{C} is (for $\epsilon \neq 0$)

simply a product $\mathcal{S} \times [0, 1]$. This makes concrete the analogy between $z^2 + \epsilon$ and a quasi-Fuchsian group.

To prove this, or the equivalent fact that $\mathcal{H} \cup \mathcal{F}/\hat{f} \cong \mathcal{S} \times [0, 1]$, make the following leafwise construction. Let H be a leaf of \mathcal{H} bounded by L . Foliate each component D of $\mathcal{F} \cap L$ by Poincaré geodesics coming from infinity in L (vertical geodesics in the upper half-plane uniformization). Above each such ray r lies a “curtain” in H , bounded by the vertical line above the point $r \cap \mathcal{J}$. Let l be the union of two rays on opposite sides meeting at \mathcal{J} . The curtain above l is, in the induced metric, isometric to \mathbb{H}^2 , and the vertical line v above $l \cap \mathcal{J}$ is a geodesic. Use the orthogonal projection to v in this surface to define a product structure. This varies continuously with the lines l in L , and varies continuously in the transverse direction of the lamination. Thus it gives a product structure for the entire lamination, which is also preserved by \hat{f} , since it is clearly affinely invariant.

Note in particular that the convex core is compact. In section 8 we will discuss this phenomenon more generally.

6.5. The scenery flow. In Bedford-Fisher-Urbanski [5] a construction called the “scenery flow” is discussed, which is related to the constructions of this paper, in the case of an axiom A rational map f .

The scenery flow is, roughly, the set of all “pictures” of the Julia set at small scales. That is, one considers all (complex) affine rescalings (and rotations) of J in \mathbb{C} and takes limits in the Hausdorff topology. The resulting collection of subsets of \mathbb{C} is indexed by backward orbits of f , using the linearization formula (4.2): for each backward orbit \hat{z} we consider the Hausdorff limit $J(\hat{z})$ of the sets $J_n = A_n(J)$ where $A(z) = (f^n)'(z_{-n})(z - z_{-n})$. The natural action of f on such a set is $\hat{f}(J(\hat{z})) = J(\hat{f}(\hat{z})) = f'(z_0) \cdot J(\hat{z})$. The flow on the set of pictures is defined by $J(\hat{z}) \mapsto e^t J(\hat{z})$.

Translating this into our terminology, given $\hat{z} \in \mathcal{J}$ a point in \mathcal{H} lying above \hat{z} can be written as $[\phi]$ where $\phi: \mathbb{C} \rightarrow L(\hat{z})$ is a chart such that $\phi(0) = \hat{z}$. The corresponding picture $J(\hat{z})$ is given by $\phi^{-1}(\mathcal{J} \cap L(\hat{z}))$, and the scaling flow is exactly the vertical geodesic flow (this interpretation of rescaling as geodesic flow was part of the original motivation for [5]). Thus the scenery flow is taken to the “curtain” above the lift of the Julia set.

7. UNIVERSAL ORBIFOLD LAMINATIONS

In this section we introduce the machinery for a general construction of an affine orbifold 2-lamination and accompanying hyperbolic orbifold 3-lamination, for any rational map.

In the original construction we were faced with the following issue: A small disk D in the affine part could be approached (in \mathcal{N}_f) by a sequence of disks D_i in such a way that the projections π on D_i do not converge in any sensible way to the projection on D . For example there could be branching on the D_i whereas D projects univalently. We resolved this in the post-critically finite case by creating new leaves and redefining the topology so as to sort out the different branching possibilities.

From the point of view of this section, the projection maps themselves from the affine leaves to the sphere will be the basic objects, so that the topology will automatically include convergence of the maps. Thus we will consider a space of meromorphic functions, with an associated action by affine transformations of the domain which gives rise to a leaf structure (that is, we can think of a leaf as a set of choices of basepoint for a meromorphic function, with precomposition with affine maps giving the change of basepoint). This will be our “universal” orbifold foliation, and any rational map will act naturally on it and give rise to an invariant lamination in which our original affine space $\mathcal{A}_f^{\text{aff}}$ will be a subset. The new topology \mathcal{A}_f^{ℓ} is induced from this space, and a final closure step will yield the added leaves.

7.1. Leaves of the affine action in the Universal space. Let $\tilde{\mathcal{U}}$ denote the space of meromorphic functions on \mathbb{C} , with the topology of uniform convergence on compact sets, and let \mathcal{U} denote the open subset of non-constant functions. Since $\tilde{\mathcal{U}}$ is a complex vector space, \mathcal{U} can be viewed as an infinite-dimensional complex analytic manifold (analyticity amounts to analytic dependence on Taylor coefficients).

The space \mathcal{U} admits two natural commuting analytic actions: a *left-action* $\psi \mapsto f \circ \psi$ by the semigroup of rational maps $f : \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$, and a *right-action* $\psi \mapsto \psi \circ A$ by the group Aff of complex-affine maps $A : \mathbb{C} \rightarrow \mathbb{C}$.

Let us first consider the structure of the individual orbits $\psi \circ \text{Aff}$ of the right-action of Aff on \mathcal{U} , and later show that they fit together into a foliation. On each orbit we place the *leafwise topology*, in which open neighborhoods are sets of the form $u \circ V$ where $u \in \mathcal{U}$ and V is an open set in Aff . Note that this may be a stronger topology (more open sets) than the induced topology from \mathcal{U} , since a leaf may accumulate on itself in \mathcal{U} .

The map $\text{Aff} \rightarrow \psi \circ \text{Aff}$ is locally non-singular – that is, the derivative map $D_\psi : T_{id}(\text{Aff}) \rightarrow T_\psi(\mathcal{U})$ is non-singular, as one may check by explicit computation. Note that the tangent space $T_\psi(\mathcal{U})$ can be identified with the space $\tilde{\mathcal{U}}$. It follows that, for h sufficiently close to but not equal to the identity, $\psi \circ h \neq \psi$. Thus the isotropy subgroup $\Gamma_\psi = \{\delta \in \text{Aff} : \psi \circ \delta = \psi\}$ is discrete in Aff . We may therefore make the identification

$$\psi \circ \text{Aff} \cong \Gamma_\psi \backslash \text{Aff}.$$

which is a homeomorphism if $\psi \circ \text{Aff}$ is taken with the leafwise topology. (The quotient is on the left since Aff acts on the right, so that $\psi \circ g = \psi \circ h \iff g = \delta \circ h$ for $\delta \in \Gamma_\psi$).

Note also that Γ_ψ must in fact consist of isometries of \mathbb{C} since a non-constant meromorphic function cannot be invariant under a dilation.

Now as in Section 6, we may think of \mathbb{C} and \mathbb{H}^3 as the quotients $\mathbb{C} \cong \text{Aff} / \mathbb{C}_*$ and $\mathbb{H}^3 \cong \text{Aff} / S^1$, with associated left-action of Aff .

Since right and left actions commute, we may form the quotients

$$L_{aff}(\psi) \equiv \psi \circ \text{Aff} / \mathbb{C}_* \cong \Gamma_\psi \backslash \text{Aff} / \mathbb{C}_* \cong \Gamma_\psi \backslash \mathbb{C},$$

which is a Euclidean 2-orbifold, and

$$L_{hyp}(\psi) \equiv \psi \circ \text{Aff} / S^1 \cong \Gamma_\psi \backslash \text{Aff} / S^1 \cong \Gamma_\psi \backslash \mathbb{H}^3,$$

which is a hyperbolic 3-orbifold. Note that the singularities, always arising from rotations in Γ_ψ , are cone axes.

The natural projection $L_{hyp}(\psi) \rightarrow L_{aff}(\psi)$ is a one-dimensional fiber bundle whose leaves are the orbits of the *vertical flow* $V_r : \phi \rightarrow \phi \circ e^r$ (this flow is well-defined since \mathbb{C}_* is commutative).

As an example, consider $\psi(z) = z^m$, so that Γ_ψ is the cyclic group generated by $\xi : z \mapsto e^{2\pi i/m} z$. The leaf $\psi \circ \text{Aff} / \mathbb{C}_*$ is then the orbifold $\langle \xi \rangle \backslash \mathbb{C}$ with one order m cone point. More interesting examples are Chebyshev polynomials associated with trigonometric functions and Lattès maps associated with elliptic functions.

A local (orbifold) affine chart on a leaf $L_{aff}(\phi)$ near ϕ is given by translations $t \mapsto \phi(z + t)$, where $t \in \mathbb{C}$ is small. In these coordinates the map $f : L_{aff}(\phi) \rightarrow L_{aff}(f \circ \phi)$ becomes the identity. Thus f is affine on the leaves. Hence it is automatically a covering.

Similar statements are valid for the hyperbolic leaves of $L_{hyp}(\phi)$, with local charts $(t, e^r) \mapsto \psi(e^r z + t)$.

7.2. Foliation structure. With this point of view on individual leaves, let us consider how they fit together into the total space \mathcal{U} , and its quotients.

Lemma 7.1. *The Aff action supplies the space \mathcal{U} with an analytic foliation with two complex dimensional leaves.*

Proof. This is a generality about any non-singular analytic Lie group action. However, rather than using deep Implicit Function Theorems (see [22]), we can check the statement directly.

Let $\phi \in \mathcal{U}$. Without loss of generality we can assume that $\phi'(0) \neq 0$. Let

$$T = \{\theta \in \mathcal{U} : \theta(0) = \phi(0); \theta'(0) = \phi'(0)\}.$$

We will show that T is a local transversal to the action of Aff . Indeed take a $\psi \in \mathcal{U}$ near ϕ and a $\gamma \in \text{Aff}$ near id , $\gamma(z) = az + b$. The condition that $\psi \circ \gamma \in T$ amounts to the following system of two equations for a and b :

$$\psi(b) = \phi(0), \quad a\psi'(b) = \phi'(0). \tag{7.1}$$

If ψ is close to ϕ then the first equation has a unique root b near 0 by the Hurwitz theorem. Thus the second equation has a unique root a near 1.

It follows that the Aff action has a local product structure near ϕ given by the map $T \times \text{Aff} \rightarrow \mathcal{U}$, $(\theta, \gamma) \mapsto \theta \circ \gamma$ near (ϕ, id) . This structure is analytic since this map is so. The inverse map is also analytic as the solutions of (7.1) analytically depend on the Taylor coefficients of ψ (by the Implicit Function Theorem). \square

Now we may form the quotients $\mathcal{U}^a = \mathcal{U}/\mathbb{C}_*$ and $\mathcal{U}^h = \mathcal{U}/S^1$, and we claim that they are orbifold 2- and 3-foliations, respectively.

This follows from the following general fact. Suppose \mathcal{L} is a lamination with finite-dimensional smooth leaves and a Lie group G acts on \mathcal{L} (say from the right) preserving leaves. We call the action smooth if its leafwise derivative exists, and is continuous in \mathcal{L} (in the transverse direction as well).

Lemma 7.2. *Let \mathcal{L} be a lamination with finite-dimensional leaves, admitting a nonsingular proper smooth action by a Lie group G . Then \mathcal{L}/G is an orbifold lamination, where the leaves have dimension equal to the codimension of G -orbits in the leaves of \mathcal{L} . If \mathcal{L} is actually an analytic foliation and the action of G is analytic then \mathcal{L}/G is an analytic orbifold foliation as well.*

Proof. Note first that the properness of the action ensures that the quotient \mathcal{L}/G is Hausdorff.

Let now $p \in B \subset \mathcal{U}$ where B is a product box $B = T \times V$, with V a leafwise neighborhood. Write $p = (t, v) \in B$. Then because the G -action is smooth and non-singular, we can find a continuous family of transversals K_s to the G -orbits in each local leaf $\{s\} \times V$. The union K is a transversal to the G action, which itself has a product box structure.

The subgroup G_p fixing p is discrete by the non-singularity assumption, and is finite by the properness assumption. We now get a ‘‘first return’’ action of G_p on a small enough neighborhood of p in the transversal K , and the quotient of this neighborhood by this action is our orbifold box in the quotient \mathcal{L}/G . To see this, note that if K' is a sufficiently small neighborhood of p in K then for any $q \in K'$ and $g \in G_p$, qg is in the original neighborhood B , and hence can be uniquely pushed to K along its G -orbit in B . Thus each element of G_p induces a map $K' \rightarrow K$ fixing p and altogether we obtain a finite group action on the union of images of K' .

Finally, it is obvious that if the lamination \mathcal{L} and the action of G have some transversal regularity (e.g., analytic), then the quotient lamination inherits it. \square

Let us summarize the above discussion:

Corollary 7.3. *The quotient $\mathcal{U}^h = \mathcal{U}/S^1$ is a hyperbolic orbifold 3-foliation. The quotient $\mathcal{U}^a = \mathcal{U}/\mathbb{C}_*$ is an affine orbifold 2-foliation, and the projection $\mathcal{U}^h \rightarrow \mathcal{U}^a$ is a fiber bundle. On the leaves of \mathcal{U}^h it is identified with the vertical projection in each half-space to the bounding plane.*

Remark: The projection $\mathcal{U} \rightarrow \mathcal{U}^h$ is similar to a Seifert fibration. A function admitting rotational symmetries around 0 gives rise to a singular fiber: its S^1 orbit is finitely covered by the S^1

action, whereas for nearby functions without the symmetry the orbit is an injective image of S^1 . However, note that singular fibers are not isolated, as they are in Seifert-fibred three-manifolds.

Proof. To apply Lemma 7.2 we need to check that the actions of \mathbb{C}_* and S^1 are proper. For S^1 this is clear since it is compact. For \mathbb{C}_* we just have to consider the vertical flow $\psi \mapsto \psi \circ e^r$, $r \in \mathbb{R}$ (as in Section 6). But it is easy to see that as r goes to $\pm\infty$ $\psi \circ e^r$ diverges in \mathcal{U} – it becomes a constant in one direction, and blows up at every point in the other. In fact for a small enough neighborhoods $U \ni \psi$ and $V \ni \phi$ there is a fixed R so that for $|r| > R$ the rescaling $u \circ e^r$ is outside V for any $u \in U$. This proves that \mathbb{C}_* acts properly.

Note finally the local affine charts are transversally analytic, so that we obtain an affine foliation. Indeed taking an analytic transversal K to the foliation \mathcal{U}^a , the map $(\psi, t) \mapsto \psi(z+t)$, where $\psi \in K$, $t \in \mathbb{C}$ is small, provides us with an orbifold affine box. Similarly, the hyperbolic structure on \mathcal{U}^h is transversally analytic. \square

We remark that it is easy to see that \mathcal{U} is metrizable, and in fact one can give it a complete metric which is invariant under the right \mathbb{C}_* -action. However, we shall not need this explicitly.

7.3. Characteristic laminations. Now given a rational map $f : \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$, we extract from our universal space \mathcal{U} the characteristic orbifold laminations for f . First consider the “global attractor”

$$\mathcal{K}_f \equiv \mathcal{K} = \bigcap_{n \geq 0} f^n(\mathcal{U})$$

which is the maximal invariant subset for which $f : \mathcal{K} \rightarrow \mathcal{K}$ is surjective. Note also that \mathcal{K} is naturally a sublamination since it is leafwise saturated.

Let us show that \mathcal{K} is closed in \mathcal{U} . It is enough to check that for any rational map g , $g(\mathcal{U})$ is closed. Let $g \circ \phi_n \rightarrow \psi$. Then $\{\phi_n\}$ is a normal family. Indeed, given any point $a \in \mathbb{C}$, consider two neighborhoods $U \ni V \ni a$. Then eventually for all n , $\phi_n(V) \subset g^{-1} \circ \psi(U)$. Take U so small that the complement of $g^{-1} \circ \psi(U)$ has non-empty interior. By Montel’s Theorem, the family $\{\phi_n\}$ is normal on V . As normality is a local property, $\{\phi_n\}$ is normal. Let $\phi : \mathbb{C} \rightarrow \bar{\mathbb{C}}$ be any limit function. Then $\psi = g \circ \phi$, and we are done.

It is not necessarily true that $f|_{\mathcal{K}}$ is injective (see remark below). Thus we take the natural extension, or inverse limit, of the system $\mathcal{K} \xleftarrow{f} \mathcal{K} \xleftarrow{f} \dots$. Call this new system $\hat{f} : \hat{\mathcal{K}}_f \rightarrow \hat{\mathcal{K}}_f$.

Elements in $\hat{\mathcal{K}} \equiv \hat{\mathcal{K}}_f$ are simply sequences $\hat{\psi} = \{\psi_n \in \mathcal{U}\}_{n \leq 0}$ such that $\psi_{n+1} = f \circ \psi_n$.

Note that $\hat{\mathcal{K}}$ is still naturally a leaf space: \mathcal{K} is invariant under the right action of Aff , which then extends to $\hat{\mathcal{K}}$ via $\{\psi_n\} \circ A = \{\psi_n \circ A\}$, so that the leaves project down to (orbifold) cover leaves in \mathcal{K} . We want to check that $\hat{\mathcal{K}}$ is in fact a lamination, i.e. that there is a local product structure.

Lemma 7.4. *$\hat{\mathcal{K}}_f$ is a lamination whose leaves are the right Aff -orbits. The projection from $\hat{\mathcal{K}}_f$ to \mathcal{K}_f is an orbifold covering on leaves. Similarly, $\hat{\mathcal{K}}_f^a \equiv \hat{\mathcal{K}}_f/\mathbb{C}_*$ and $\hat{\mathcal{K}}_f^h \equiv \hat{\mathcal{K}}_f/S^1$ are orbifold affine 2- and hyperbolic 3-laminations, respectively.*

Proof. Fix $\hat{\psi} = \{\psi_n\}$ in $\hat{\mathcal{K}}$. We will describe the structure of a neighborhood of $\hat{\psi}$ as follows. Let D be some disk in \mathbb{C} on which ψ_0 is univalent, and let $D'' \Subset D' \Subset D$ be nested open disks. Let U be an open neighborhood of ψ_0 in \mathcal{K} for which any $u \in U$ is univalent in D' and such that $\psi_0(D'') \subset u(D')$. Note that we may assume U is a product neighborhood of the form $T \times V$ where V is a neighborhood of the identity in Aff .

The preimage \hat{U} of U in $\hat{\mathcal{K}}$ consists of sequences $\hat{u} = \{u_n\}$ for which $u_0 \in U$. For any such \hat{u} , notice that, since $u_0 = f^n \circ u_{-n}$ for any $n \geq 0$, u_{-n} is univalent on D' and f^n is univalent on $u_{-n}(D')$. Thus \hat{u} determines an infinite sequence of univalent pullbacks of $u_0(D')$, and hence of $W_0 \equiv \psi_0(D'')$. Conversely \hat{u} is determined by this sequence and u_0 , using analytic continuation.

Let Σ denote the set of all possible infinite pullback sequences W_0, W_{-1}, \dots where f is univalent at each step. With the natural topology, this is a closed subset of the set of all possible pullback sequences for W_0 , and hence a closed subset of a Cantor set. We thus have an injection of \hat{U} into $\Sigma \times U$. Let us show it is also a topological embedding. Suppose for a sequence \hat{u}^i that the images $((W_{-n}^i, u_0^i)$ in $\Sigma \times U$ converge. Then u_0^i converge as meromorphic functions, and for each $n > 0$ eventually W_{-n}^i are constant. The functional equation $u_0^i = f^n \circ u_{-n}^i$ then implies that u_{-n}^i converge locally on D'' as $i \rightarrow \infty$, and in fact form a normal family so that they converge globally. Thus \hat{u}^i converge. The other direction is easy.

The subset of $\Sigma \times U$ obtained is saturated in the leaf direction, since any composition $u_0 \circ A$ lying in U (with $A \in \text{Aff}$) can be pulled back along the same sequences as u_0 . Hence there is some subset $Q \subset \Sigma \times T$ so that we may identify \hat{U} with $Q \times V$. This is the desired product box.

The fact that $\hat{\mathcal{K}}/\mathbb{C}_*$ and $\hat{\mathcal{K}}/S^1$ are orbifold laminations now follows by another application of Lemma 7.2. \square

Remark. We expect that in most cases the natural extension step is unnecessary; that is, f is already injective on \mathcal{K} . Counterexamples are maps with symmetry: for example, the leaf of $\psi_0(z) = e^z$ in \mathcal{U}/\mathbb{C}_* is a cylinder $\mathbb{C}/2\pi i$, and if $f(z) = z^d$ then f is a d -fold cover from this leaf to itself. It follows that the lift of this leaf to $\hat{\mathcal{K}}_f/\mathbb{C}_*$ is a solenoidal Riemann surface lamination (in fact it is just the original \mathcal{A}_f^n in this case). We conjecture that non-injectivity only happens when f or a power of f has a Möbius symmetry.

7.4. Completion. There is an equivariant inclusion of \mathcal{A}_f^n into our new object $\hat{\mathcal{K}}^a$, as follows. Let \hat{z} be a point on a leaf L of \mathcal{A}_f^n , and let $\phi : \mathbb{C} \rightarrow L$ be an isomorphism such that $\phi(0) = \hat{z}$. Then for $n \leq 0$, $\pi_n \circ \phi$ is an element of \mathcal{K} (compare §6). The choice of ϕ was determined only up to precomposition by \mathbb{C}_* , so that \hat{z} determines a well-defined sequence in \mathcal{K}^a , which gives an element $\iota(\hat{z}) \in \hat{\mathcal{K}}^a$.

The map ι takes leaves to leaves, since another element of L can be written as $\phi(A(0))$ with $A \in \text{Aff}$. ι is injective, since at least one of the coordinates z_n must differ for different points on \mathcal{A}_f^n .

On the other hand, $\hat{\mathcal{K}}^a$ is also an affine orbifold extension of f , in the sense of Section 6.1, and hence there is also a continuous, equivariant projection $p : \hat{\mathcal{K}}^a \rightarrow \mathcal{A}_f^n$. That is, for any $[\hat{\psi}] = \{[\psi_n]\} \in \hat{\mathcal{K}}^a$, let $p([\hat{\psi}])$ be the backward orbit $\{\psi_n(0)\}$.

It is immediate from the definitions that $p \circ \iota$ is the identity, but we note that the opposite is false, since in fact p is not injective and hence ι is not surjective. Indeed, let $g : \mathbb{C} \rightarrow \mathbb{C}$ be any non-affine entire function and L a leaf of \mathcal{A}_f^n with chart $\phi : \mathbb{C} \rightarrow L$. Then the sequence $\{\pi_n \circ \phi \circ g\}$ is on a leaf of $\hat{\mathcal{K}}/\mathbb{C}_*$ which projects to L but is different from $\iota(L)$.

Note that the topology on \mathcal{A}_f^n induced from $\hat{\mathcal{K}}^a$ is in general stronger than its own topology, induced from \mathcal{N}_f (so that the inclusion ι is discontinuous). This is in fact the main point of the construction. Let \mathcal{A}_f^ℓ denote $\iota(\mathcal{A}_f^n)$, with the topology induced from $\hat{\mathcal{K}}^a$. We also think of \mathcal{A}_f^n and \mathcal{A}_f^ℓ as being the same underlying space, with different topologies, which we call “natural” and “laminar”.

Our final step is to take the closure, in $\hat{\mathcal{K}}^a$, of \mathcal{A}_f^ℓ , obtaining automatically an affine orbifold extension of f (in the sense of §6) which we call \mathcal{A}_f . We think of \mathcal{A}_f as a *completion* of \mathcal{A}_f^ℓ .

Going through the same construction replacing \mathbb{C}_* -action with S^1 -action, we obtain the hyperbolic 3D-extension $\mathcal{H}_f \subset \hat{\mathcal{K}}/S^1 \equiv \hat{\mathcal{K}}^h$, with the hyperbolic action of \hat{f} .

Remark. The laminated space \mathcal{A}_f inherits from the universal space \mathcal{U} the quality of a metrizable separable space. Moreover, it has a natural uniform structure coming from the linear structure

of \mathcal{U} , and complete with respect to it. However, \mathcal{A}_f may presumably inherit from \mathcal{U} also the bad fortune of not being locally compact.

7.5. Induced topology. Let us give a dynamical description of the new laminar topology \mathcal{A}_f^ℓ on the leaf space \mathcal{A}_f^n .

By a local leaf $L_{loc}(\hat{z}, V)$ over a domain $V \subset \bar{\mathbb{C}}$ containing $\pi(\hat{z})$ we mean the connected component of $L(\hat{z}) \cap \pi^{-1}V$ containing \hat{z} .

Proposition 7.5. *A sequence $\hat{z}^n \in \mathcal{A}_f^\ell$ converges to $\hat{\zeta} \in \mathcal{A}_f^\ell$ if and only if the following hold:*

- (i) $\hat{z}^n \rightarrow \hat{\zeta}$ in the natural topology,
- (ii) For any N and a neighborhood V of ζ_{-N} , if the local leaf $L_{loc}(\hat{f}^{-N}\hat{\zeta}, V)$ is univalent over V , then for sufficiently large n , $L_{loc}(\hat{f}^{-N}\hat{z}^n, V)$ is univalent over V as well.

We remark that convergence to a point in $\mathcal{A}_f \setminus \mathcal{A}_f^\ell$ is more subtle to characterize in general. Proposition 8.2 does this in the post-critically non-recurrent case.

Proof. Assume first that conditions (i) and (ii) are satisfied.

Represent $\hat{\zeta}$ as a sequence $\hat{\psi} = \{\psi_j\}$ in $\hat{\mathcal{K}}_f$, and each \hat{z}^n as $\hat{\phi}^n = \{\phi_j^n\}$, in particular noting $\psi_j(0) = \zeta_j$ and $\phi_j^n(0) = z_j^n$.

The statement that $L_{loc}(\hat{f}^{-N}\hat{\zeta}, V)$ is univalent over V is equivalent to saying that ψ_{-N} is univalent in the component W of $\psi_{-N}^{-1}(V)$ containing 0 (henceforth we say “ ψ_{-N} is locally univalent over V ”), and similarly for $\hat{f}^{-N}\hat{z}^n$ and ϕ_{-N}^n .

Whenever, for some n, N, V , both ψ_{-N} and ϕ_{-N}^n are locally univalent over V , there is a unique univalent map $h^n : W \rightarrow \mathbb{C}$ satisfying $\psi_{-N} = \phi_{-N}^n \circ h^n$ on W . Note that, applying f a finite number of times, we have

$$\psi_{-j} = \phi_{-j}^n \circ h^n \tag{7.2}$$

on W for any $j \leq N$. Thus if we increase V or change N (but preserve the local univalence), we obtain h^n equal to the original on the original domain, or in other words h^n is locally independent of N and V . Choose the normalization of each $\hat{\phi}^n \pmod{\mathbb{C}_*}$ so that $(h^n)'(0) = 1$.

Because $\hat{\zeta} \in \mathcal{A}_f^\ell$, for any disk D_r around 0 there is some $N(r)$ for which ψ_{-N} is univalent on D_r whenever $N > N(r)$. Let $V = \psi_{-N}(D_r)$. For sufficiently large $n(r)$, by (i) we have that $\phi_{-N}^n(0) = z_{-N}^n \in V$, and by (ii) that ϕ_{-N}^n is locally univalent over V . Thus we have h^n defined as above on D_r if $n > n(r)$.

If we let x^n be the preimage of z_{-N}^n in D_r by ψ_{-N} (note that x^n is independent of N if $N > N(r)$), we note that $h^n(x^n) = 0$, and by (i), $x^n \rightarrow 0$ as $n \rightarrow \infty$.

Thus, the sequence of functions h^n now has these properties: $(h^n)'(0) = 1$, $h^n(x^n) = 0$ where $\lim_{n \rightarrow \infty} x^n = 0$, and h^n is eventually defined on any compact set in \mathbb{C} . It is an application of the Koebe distortion lemma now to show that h^n converges to the identity on compact sets, and indeed that the image of h^n eventually contains any compact set in \mathbb{C} so that $(h^n)^{-1}$ converges to the identity on compact sets as well.

Applying (7.2) for any j , we conclude that $\phi_j^n \rightarrow \psi_j$ on compact sets for all j . Thus $\hat{z}^n \rightarrow \hat{\zeta}$ in \mathcal{A}_f^ℓ .

Conversely, let $\hat{z}^n \in \mathcal{A}_f^\ell$ converge to $\hat{\zeta} \in \mathcal{A}_f^\ell$. Assertion (i) is obvious; it is just the statement that $p : \hat{\mathcal{K}}_f \rightarrow \mathcal{A}_f^n$ is continuous, which we have already observed.

For (ii), let V be a neighborhood of ζ_{-N} such that $L_{loc}(\hat{f}^{-N}\hat{\zeta}, V)$ is univalent over V , and let $\psi_{-N} : \mathbb{C} \rightarrow \bar{\mathbb{C}}$ be as above. If W is the component of $\psi_{-N}^{-1}(V)$ containing 0 we then have ψ_{-N} univalent on W .

Because a slight enlargement V' of V (so that $V \Subset V'$) pulls back along the rest of $\hat{\zeta}$ with bounded branching (by definition of \mathcal{A}_f^n), it follows that W has compact closure in \mathbb{C} . Let $W \Subset W' \Subset W''$ be a pair of enlargements of W , also with compact closure. By definition of convergence in \mathcal{A}_f^ℓ , there are representatives $\hat{\varphi}^n = \{\varphi_j^n\}$ of \hat{z}^n in $\hat{\mathcal{K}}_f$ such that φ_{-N}^n converges on W'' to ψ_{-N} . It follows that for large enough n , φ_{-N}^n is univalent on W' and $\varphi_{-N}^n(W')$ contains V , and thus (ii) holds. \square

7.6. Uniqueness. Let us now consider, for an abstract affine orbifold extension \mathcal{A} of f in the sense of Section 6.1, what properties force it to be equal to our universal construction $\hat{\mathcal{K}}^a$.

There is a natural map $I : \mathcal{A} \rightarrow \hat{\mathcal{K}}^a$, defined similarly to ι : for any $\mathbf{z} \in \mathcal{A}$, let $\phi : \mathbb{C} \rightarrow L(\mathbf{z})$ be (the inverse of) any affine chart for the leaf of \mathbf{z} that takes 0 to \mathbf{z} . Then the sequence $\{[\pi_n \circ \phi]\}$ gives a well-defined element of $\hat{\mathcal{K}}^a$, where π_n are the projections of \mathcal{A} to $\bar{\mathbb{C}}$. The difference between I and ι is that I is automatically continuous, because of the transverse continuity of the affine structures in \mathcal{A} .

We now observe that $I(\mathcal{A})$ is equal to $\hat{\mathcal{K}}^a$ if the following conditions hold:

1. The map I is an embedding,
2. \mathcal{A}_f^ℓ is dense in $I(\mathcal{A})$, and
3. $I(\mathcal{A})$ is closed

In particular, the first condition reduces to checking that I is both injective and proper: i.e. that an element of \mathcal{A} is determined uniquely by the sequence of functions $\pi_n \circ \phi$, and that convergence in \mathcal{A} follows from convergence of the sequence of functions.

For the construction of Section 5 of orbifold laminations for post-critically finite maps, these properties evidently hold, and therefore the general construction produces the same object.

7.7. Minimality. Let us show that the laminations we constructed are minimal. Note that this does not follow from Lemma 4.8 since topology of \mathcal{A}_f is stronger than that of \mathcal{N}_f .

Proposition 7.6. *The laminations \mathcal{A}_f and \mathcal{H}_f are minimal except for the Chebyshev and Lattès examples. In those cases the lamination becomes minimal after removing the isolated invariant leaf associated with a post-critical fixed point.*

Hence every open set K of either lamination contains a global cross-section for it (except the isolated leaves in the above special cases).

Proof. Clearly it is enough to consider \mathcal{A}_f . Since \mathcal{A}_f^ℓ is dense in \mathcal{A}_f , it suffices to demonstrate density of leaves in \mathcal{A}_f^ℓ .

Let us first show that any invariant leaf L is dense. Take a point $\hat{z} = \{z_0, z_{-1}, \dots\}$ in \mathcal{A}_f^ℓ , and a finitely branched pullback of neighborhoods $\{U_0, U_{-1}, \dots\}$ along it. In the case when f is Lattès or Chebyshev assume that \hat{z} is not a postcritical fixed point. Then Proposition 12.1 and the expansion property of f on the Julia set easily yield existence of a limit point $a \in \bar{\mathbb{C}}$ for \hat{z} and $\epsilon > 0$ such that one of the local leaves L_{loc} over $D(a, \epsilon)$ is not branched.

For sufficiently large N , U_{-N} pulls back univalently along the rest of \hat{z} , and by the Shrinking Lemma, there is a sequence $N_i \rightarrow \infty$ such that $U_{-N_i} \subset D(a, \epsilon)$.

Thus L_{loc} is univalent over U_{-N_i} . Let \hat{b}^i be the point on this local leaf which projects to z_{-N_i} . Then by Proposition 7.5, the sequence $\hat{f}^{N_i} \hat{b}^i \in L$ converges to \hat{z} in the \mathcal{A}_f topology as $i \rightarrow \infty$, which proves density of L .

Replacing f by its iterate, we conclude that every periodic leaf is dense in \mathcal{A}_f .

Let us now show that every leaf $L(\hat{z}) \subset \mathcal{A}_f^\ell$ accumulates on some periodic leaf. To this end take five periodic points α_k and associated periodic leaves $L_k \equiv L(\alpha_k)$. Select five disjoint topological discs $D_k \ni \alpha_k$. By Ahlfors' Five Islands Theorem (see [52], Theorem VI.8), for any n , each $\hat{f}^{-n}L(\hat{z})$

has a univalent local leaf over one of the domains D_k . Take a k for which this happens for infinitely many n 's. Then by the same argument as above $L(\hat{z})$ accumulates on the periodic leaf L_k . \square

8. CONVEX-COCOMPACTNESS, NON-RECURRENCE AND CONICAL POINTS

Define the Julia set \mathcal{J}_f in \mathcal{A}_f to be the pullback of J_f by $\pi : \mathcal{A}_f \rightarrow \bar{\mathbb{C}}$. Let \mathcal{J}_f^ℓ denote $\mathcal{J}_f \cap \mathcal{A}_f^\ell$. Note that \mathcal{J}_f^ℓ and \mathcal{J}_f^n have the same underlying set and different topologies, and that \mathcal{J}_f is the closure of \mathcal{J}_f^ℓ .

We say that f is *convex cocompact* if the quotient $\mathcal{C}(\mathcal{J}_f)/\hat{f}$ of the convex hull is compact. In this section we prove several criteria for convex cocompactness. The main criterion is the following:

Theorem 8.1. *A rational map f is convex cocompact if and only if it is postcritically non-recurrent and has no parabolic points.*

Remark. This criterion is closely related to the ‘‘John domain criterion’’ given by Carleson, Jones and Yoccoz for polynomials [13]. See also McMullen [34] for the connection between convex-cocompactness and the John condition in the setting of Kleinian groups.

8.1. Convergence and compactness. For a critically non-recurrent map f without parabolics, we can give a dynamical criterion for convergence in \mathcal{A}_f (note that Proposition 7.5 only applied to convergence within \mathcal{A}_f^ℓ . This criterion includes the possibility that a bounded amount of branching persists in the limit and yields a point outside \mathcal{A}_f^ℓ). Let $p : \mathcal{A}_f \rightarrow \mathcal{A}_f^\ell$ denote the natural projection.

Proposition 8.2. *Let f be critically non-recurrent without parabolics. A sequence of points $\hat{z}^n \in \mathcal{A}_f^\ell$ converges to $\zeta \in \mathcal{A}_f$, with $p(\zeta) = \hat{\zeta}$, if and only if*

- (i) $\hat{z}_n \rightarrow \hat{\zeta}$ in the natural topology and
- (ii) For any N and a neighborhood V of ζ_{-N} , if the local leaf $L_{loc}(\hat{f}^{-N}\hat{\zeta}, V)$ is univalent over V , then the following holds:

There is a finite set of points $\{c_k\} \subset V$ such that for any neighborhood Ω of $\{c_k\}$ there exists $M = M(\Omega)$ so that, if $n > M$, the local leaf $L_{loc}(\hat{z}^n, V \setminus \Omega)$ covers $V \setminus \Omega$ without branching, and for any $n, m > M$ the coverings are topologically equivalent.

Moreover, the projection $L(\zeta) \rightarrow L(\hat{\zeta})$ is a finitely branched covering with uniformly bounded degree.

Proof. By Mañé’s Theorem, there is a neighborhood W of J_f , and $\epsilon_0 > 0$ and K_0 with the following property: for any backward trajectory $\hat{z} = \{z_0, z_{-1}, \dots\} \in \mathcal{N}_f$ with $z_0 \in W$, the pullback of the disk $D(z_0, \epsilon_0)$ along \hat{z} branches at most K_0 times. (Compare the proofs of Lemma 3.5 and Proposition 4.5).

Furthermore, for any \hat{z} which is not an attracting cycle, there is an $N_0(\hat{z})$ such that $z_{-n} \in W$ for $n > N_0$.

Assuming that (i) and (ii) hold, represent $\hat{\zeta}$ using a sequence $\{\psi_{-N}\} \in \hat{\mathcal{K}}_f$. For any disk $D \subset \mathbb{C}$, for large enough $N > N_0$ the map ψ_{-N} is univalent in D and has image in W , and in fact in $D(\zeta_{-N}, \epsilon_0)$. For sufficiently close \hat{z}^n to $\hat{\zeta}$, z_{-N}^n is also in W , and hence the pullback of $D(z_{-N}^n, \epsilon_0)$ along the rest of \hat{z}^n has uniformly bounded branching.

Condition (ii) now gives a branched cover of D which is conformally equivalent to the coverings of $L_{loc}(\hat{f}^{-N}\hat{z}^n, \psi_{-N}(D)) \rightarrow \psi_{-N}(D)$ away from a small neighborhood of the critical points, for large enough n . This branching is uniformly bounded no matter how large D is taken, so we obtain a polynomial $h : \mathbb{C} \rightarrow \mathbb{C}$. The sequence $\{\psi_{-N} \circ h\} \in \hat{\mathcal{K}}_f$ will represent the limit ζ of the \hat{z}^n in \mathcal{A}_f , by an argument similar to that in Proposition 7.5, where condition (ii) keeps the branching consistent.

More precisely, let $D' = h^{-1}D$ and assume that D' is large enough that the (finite) set C of critical points of h is separated from $\partial D'$ by an annulus of modulus $M > 0$. For each n represent

\hat{z}^n by a sequence of functions $\phi_{-N}^n : \mathbb{C} \rightarrow \bar{\mathbb{C}}$, normalized so its 1-jet agrees with $\psi_{-N} \circ h$ at a fixed non-critical point $w \in D'$. Let Y be a neighborhood of C so that $D' \setminus Y$ contains an annulus of modulus M around each puncture. Then condition (ii) gives, for large enough n , a univalent map $u_n : D' \setminus Y \rightarrow \mathbb{C}$ such that $\psi_{-N} \circ h = \phi_{-N}^n \circ u_n$, and $u_n(w) = w, u_n'(w) = 1$. Note that u_n , once defined on $D' \setminus Y$, remains the same there as we enlarge D' , shrink Y and increase N , and that $\psi_{-N} \circ h = \phi_{-N}^n \circ u_n$ wherever it is defined. Again using Koebe distortion (this time on a multiply connected domain), we have $u_n \rightarrow \text{id}$ on compact subsets of $\mathbb{C} \setminus C$. It follows that for every N , $\phi_{-N}^n \rightarrow \psi_{-N} \circ h$, so that $\hat{z}^n \rightarrow \zeta$ as $n \rightarrow \infty$.

Moreover, $p : L(\zeta) \rightarrow L(\hat{\zeta})$ is a finitely branched covering with bounded degree since h is.

Conversely, suppose that the sequence \hat{z}^n converges in \mathcal{A}_f . Since by the same discussion the branching over each disk $D(z, \epsilon_0)$, $z \in J_f$, is eventually uniformly bounded, there must be some subsequence of the \hat{z}^n for which the branching converges in the sense of (ii), and so the limit is equal to the limit defined in the previous paragraph. It follows that the same holds for any subsequence, so that in fact (ii) holds for the whole sequence. \square

Corollary 8.3. *Let f be critically non-recurrent without parabolics. A set $K \subset \mathcal{A}_f^\ell$ is pre-compact in \mathcal{A}_f if and only if its closure in the natural extension \mathcal{N}_f does not contain attracting cycles.*

Proof. If a sequence $\{\hat{z}^n\} \subset \mathcal{A}_f^\ell$ does not accumulate on attracting cycles then Mane's Theorem easily yields existence of a subsequence satisfying (i) and (ii) of the previous proposition. \square

8.2. Proof of Theorem 8.1. Let us split the proof into two steps represented by the following two criteria.

Lemma 8.4. *The Julia set \mathcal{J}_f is compact if and only if f is critically non-recurrent and has no parabolic points.*

Proof. If \mathcal{J}_f is compact then \mathcal{J}_f^n is compact in \mathcal{A}_f^n , since $\mathcal{J}_f^n = p(\mathcal{J}_f)$ where $p : \mathcal{A}_f \rightarrow \mathcal{A}_f^n$ is the natural continuous projection. Hence by Corollary 4.9 f is critically non-recurrent without parabolic points.

Vice versa, if f is critically non-recurrent without parabolic points then compactness of \mathcal{J}_f follows from Corollary 8.3. \square

Proposition 8.5. *A rational map f is convex cocompact if and only if the Julia set \mathcal{J}_f is compact.*

Proof. Let $\mathcal{V} = \mathcal{V}(\mathcal{J})$ denote the ‘‘curtain’’ over $\mathcal{J} = \mathcal{J}_f$ in \mathcal{H} . That is, the union of vertical geodesics over points of \mathcal{J} . We will first show that \mathcal{V}/\hat{f} is compact if and only if \mathcal{J} is compact.

Observing that \mathcal{J} is just the quotient $\mathcal{V}/e^{\mathbb{R}}$ by the vertical geodesic flow, we may view this equivalence in slightly generalized terms:

Let X be a Hausdorff space admitting commuting actions by two closed non trivial subgroups G and H of \mathbb{R} . Let G act on the left and H on the right, for clarity. Suppose G and H both act properly, and that they are *coherent* in the sense of Lemma 6.1: any $x, y \in X$ are contained in neighborhoods U_x, U_y for which $g_i U_x \cap U_y h_i \neq \emptyset$ only if g_i, h_i both remain bounded, both go to $+\infty$ or to $-\infty$. We claim that $G \backslash X$ is compact if and only if X/H is compact.

Suppose without loss of generality that $G \backslash X$ is compact, and let $K \subset X$ be a compact fundamental domain, i.e. $GK = X$. Let $x \in K$ and consider the *positive return time* $g_+(x)$ for the orbit xH to return to KH under G . That is, let g_+ be the smallest positive element of G such that $g_+ xH \cap KH \neq \emptyset$. We claim this is bounded for $x \in K$. Choose $h < 0$ in H sufficiently far from 0 that $Kh \cap K = \emptyset$ (by the proper action of H), and that $gK \cap Kh \neq \emptyset$ only for $g < 0$ (this is possible by coherence, after covering K with a finite number of neighborhoods U_p).

Thus the point xh is not in K so there is some $g_+ > 0$ such that $g_+ xh \in K$. Thus $g_+ Kh \cap K \neq \emptyset$ for each g_+ , so that fixing h we have an upper bound for g_+ independent of x , by the proper action of G .

Reversing the signs in the argument we also obtain a bounded *negative* return time for every $x \in K$. We conclude that, in the action of G on X/H , every point has a bounded negative and positive return time to the projection KH of K . Since X/H is covered by G -translates of KH , it follows that there is a bounded subset I of G such that IKH covers X/H . Thus X/H is compact.

In our situation the groups are \mathbb{Z} and \mathbb{R} , and we conclude that \mathcal{V}/\hat{f} is compact if and only if $\mathcal{V}/e^{\mathbb{R}} = \mathcal{J}$ is compact (note: it would be more consistent to denote the first quotient $\hat{f}\backslash\mathcal{V}$). It remains to check that compactness of \mathcal{V}/\hat{f} is equivalent to compactness of the convex core quotient. Since the curtain is closed in the convex core, one implication is clear. Conversely, if we know that \mathcal{V}/\hat{f} is compact, we need only to observe that the convex core lies in a bounded neighborhood of the curtain. That is, let $p \in \mathcal{C}$ be some point, represented as (z, t) in a half-space model of the leaf $L_{hyp}(p)$ of \mathcal{H} . If z_0 is the nearest point to z in the local Julia set $\mathcal{J} \cap L$, then $t > |z - z_0|$ because otherwise p lies in a hemisphere over z disjoint from \mathcal{J} , and therefore outside the convex hull. It follows that the hyperbolic distance from p to (z_0, t) , which lies in \mathcal{V} , is less than 1.

It is easy to check that a leafwise 1-neighborhood of a compact subset of a hyperbolic 3-lamination is itself compact, so this concludes the proof. \square

8.3. Conical points. Given a point $\mathbf{z} \in \mathcal{J}$, let $\gamma_{\mathbf{z}}$ be the vertical geodesic in \mathcal{H}_f terminating at \mathbf{z} . By analogy with Kleinian groups, let us say that $\mathbf{z} \in \mathcal{A}$ is a *conical point* if the projection of the geodesic $\gamma_{\mathbf{z}}$ to the quotient lamination \mathcal{H}/\hat{f} does not escape to infinity (which means that there is a sequence of points $p_n \in \gamma_{\mathbf{z}}$ tending to \mathbf{z} whose projection to \mathcal{H}/\hat{f} converges). Note that in this definition the vertical geodesic can be replaced by any geodesic terminating at \mathbf{z} since all of them are asymptotic in the hyperbolic metric.

Equivalently, $\mathbf{z} \in \mathcal{A}_f$ is conical iff its forward orbit $\{\hat{f}^n \mathbf{z}\}_{n=0}^{\infty}$ is non-escaping in \mathcal{A}_f , that is, the ω -limit set $\omega(\mathbf{z})$ is non-empty. Indeed, given two proper commuting group actions G and H on a space X , the G -orbit of a point $x \in X$ is non-escaping in the quotient by H if and only if its H -orbit is non-escaping in the quotient by G (since either is equivalent to non-escaping of the double orbit GxH in X). In our situation we have a \mathbb{Z} -action by \hat{f} on \mathcal{H}_f , and the \mathbb{R} -action of the vertical geodesic flow (as in Section 6). The directionality of our statement (forward \hat{f} -orbits accumulate in $\mathcal{H}_f/e^{\mathbb{R}} = \mathcal{A}_f$ if and only if backward \mathbb{R} -orbits accumulate in \mathcal{H}_f/\hat{f}) comes directly from the coherence of the actions, lemma 6.1.

Let $\Lambda = \Lambda_f$ denote the set of conical points.

We further note that the property of being conical depends only on the projection to $\bar{\mathbb{C}}$. Let us say that a set $X \subset \mathcal{A}$ is *fiber saturated* if $X = \pi^{-1}(\pi(X))$. The reason is that the fibers play the role of local stable manifolds for \hat{f} (the proof below makes precise the sense of this statement).

Proposition 8.6. *The set of conical points is fiber saturated.*

Proof. Let us show that the ω -limit sets of \mathbf{z} and \mathbf{w} are equal, up to finite branched cover. Represent \mathbf{z} and \mathbf{w} in $\hat{\mathcal{K}}_f^{\mathbf{a}}$ by sequences of meromorphic functions $\{\phi_n\}$ and $\{\psi_n\}$ such that $\phi_0(0) = \psi_0(0) = \pi(\mathbf{z})$. To compute the ω -limit sets it suffices to consider just the first coordinate functions, $\phi = \phi_0$ and $\psi = \psi_0$. Suppose first that π is non-singular at \mathbf{z} and \mathbf{w} , so that we may assume $\psi'(0) = \phi'(0) = 1$.

Now suppose that h is a limit point of $f^n \circ \phi$ in $\mathcal{U}^{\mathbf{a}}$. This means that for some sequence n_i , and $\lambda_i \in \mathbb{C}^*$, $f^{n_i} \circ \phi \circ \lambda_i$ converges to h on compact subsets of \mathbb{C} . By Lemma 6.1, we know that $|\lambda_i| \rightarrow 0$. Now fixing a disk $D \subset \mathbb{C}$ around 0, we see that for i sufficiently large, ϕ and ψ are both invertible in $\lambda_i D$, and by the Koebe distortion theorem, the combined map $(\psi \circ \lambda_i)^{-1}(\phi \circ \lambda_i)$ converges to the identity on D . It follows that $f^{n_i} \circ \psi \circ \lambda_i$ also converges to h .

If, on the other hand, ψ and/or ϕ have branched points at 0, say with degrees k and m respectively, let $d = \text{lcm}(k, m)$ and write $\tilde{\psi} = \psi \circ b_{d/k}$ and $\tilde{\phi} = \phi \circ b_{d/m}$, where $b_j(z) = z^j$. Now $\tilde{\psi}$ and $\tilde{\phi}$

both have degree d at 0, and for small $|\lambda_i|$ we still make sense of $(\tilde{\psi} \circ \lambda_i)^{-1}(\tilde{\phi} \circ \lambda_i)$ as a univalent map. Hence the Koebe distortion argument goes through and we may conclude that $\{f^n \circ \tilde{\psi}\}$ and $\{f^n \circ \tilde{\phi}\}$ have the same ω -limit points in \mathcal{U}^a . \square

Let $\Delta = \pi\Lambda \subset J_f \subset \bar{\mathbb{C}}$. By the above proposition, it is justified to call the points of this set conical as well. Let us show that it is trapped in between two well-studied sets.

First, let Δ_1 denote the set of points $z \in J_f$ such that there is an $r > 0$ and a sequence $n_i \rightarrow \infty$ (depending on z) such that the multi-valued inverse branch $f^{-n_i} : D(f^{n_i}z, r) \rightarrow U_i \ni z$ has a bounded degree (compare [23]).

The second set, Δ_0 is the union of all expanding subsets of the Julia set (a compact invariant set $X \subset \bar{\mathbb{C}}$ is called *expanding* if $f : X \rightarrow X$ is surjective and some iterate $f^n|_X$ has spherical derivative strictly greater than 1).

Proposition 8.7. $\Delta_0 \subset \Delta \subset \Delta_1$.

Proof. Let us start with the right-hand inclusion. Let $z = \pi(\mathbf{z})$ for $\mathbf{z} \in \Lambda_f$. Then there exists a sequence $n_i \rightarrow \infty$ such that $\hat{f}^{n_i}\mathbf{z} \rightarrow \zeta \in \mathcal{A}_f$. Translation of this to the language of meromorphic functions provides us with a desired family of inverse branches with bounded degree.

For the left-hand inclusion, take a point z in an expanding set $X \subset \bar{\mathbb{C}}$. First notice that by Lemma 4.1 any backward orbit \hat{z} in the invariant lift $\hat{X} \subset \hat{\mathcal{N}}_f$ belongs to a parabolic leaf. Then, take any convergent subsequence $f^{n_i}\hat{z} \rightarrow \hat{\zeta} \in \hat{X}$ in the natural topology and apply Proposition 8.2 to see that it is convergent to the same point in the laminar topology as well (the local leaves in condition (ii) of this proposition can be selected univalent). \square

Proposition 8.8. *If f is convex cocompact then all points of the Julia set \mathcal{J}_f are conical.*

Conversely, if the lamination \mathcal{A}_f is locally compact and all points of the Julia set \mathcal{J}_f are conical then f is convex cocompact.

Proof. Assume f is convex cocompact, that is, the convex core \mathcal{C}_f/\hat{f} is compact. Since $\gamma_{\mathbf{z}} \subset \mathcal{C}_f$ for any $\mathbf{z} \in \mathcal{J}_f$, the conical property of \mathbf{z} follows.

For the converse, suppose the lamination \mathcal{A}_f is locally compact. Then there is a compact set K with non-empty interior. By Proposition 7.6, K meets every leaf of the lamination. Since the set $K \cap L_{hyp}(p)$ is closed in the intrinsic leaf topology, for any $p \in \mathcal{H}$, there is a length minimizing geodesic Γ_p joining p to K . Let $\text{dist}(p, K)$ denote the hyperbolic length of this geodesic. It can be also defined as the infimum of lengths of all curves joining p and K .

Given a set $X \subset \mathcal{H}$, let $N(X, r) = \{p \in \mathcal{H} : \text{dist}(p, X) < r\}$ denote the leafwise R -neighborhood of X . Then any compact set $Q \subset \mathcal{H}$ is covered by some $N(K, R)$. Indeed, for every $q \in Q$ there is a curve γ joining q with a point in the interior of K . If the length of this curve is r then all sufficiently nearby points can be joined with $\text{int } K$ by a curve of length less than $r + \epsilon$. Now compactness of Q yields the statement.

Note also that the space \mathcal{G} of one-sided geodesics beginning in K is parametrized by the unit tangent bundle over K and hence is compact.

Assuming that the convex core \mathcal{C}_f/\hat{f} is not compact let us construct in it an escaping geodesic. Consider a sequence of points $q_n \in \mathcal{C}_f/\hat{f}$ escaping to ∞ and the corresponding minimizing geodesics $\Gamma_n \equiv \Gamma_{q_n}$. By compactness of \mathcal{G} , there is a limit geodesic Γ beginning at K . Let us show that this geodesic escapes to ∞ .

Indeed, otherwise there is a compact set $Q \subset \mathcal{C}_f/\hat{f}$ which Γ does not escape. Let us consider the leafwise 1-neighborhood $N(Q, 1)$ of Q . Its closure is compact and hence is contained in some leafwise neighborhood $N(K, R)$ of K .

Since the Γ_n accumulate on Γ , for some n there are two points $a, b \in \Gamma_n \cap N(Q, 1)$ such that the distance between them along Γ_n is greater than R . On the other hand, there is a curve from b to K of length less than R which contradicts the minimality of Γ_n . \square

Let us say that a set $X \subset \mathcal{A}$ is *locally fiber saturated* if for any point $p \in X$ there is a box neighborhood $U \ni p$ such that if $q \in U \cap X$ then the whole fiber $\pi^{-1}(\pi q) \cap U$ belongs to X . We can then say that such a set X is measurable and has “zero”, “positive” or “full” measure if the corresponding property is satisfied leafwise, that is for its intersection with every leaf. Note that these notions are well defined on the affine leaves though the Lebesgue measure is not. Note also that they don’t require any transversal measure.

Given a measurable locally fiber saturated \hat{f} -invariant set $X \subset \mathcal{A}$, we say that $\hat{f}|X$ is *ergodic* if every measurable locally fiber saturated \hat{f} -invariant subset $Y \subset X$ has either zero or full measure.

An *invariant line field* on \mathcal{A} is a measurable real one-dimensional distribution in the tangent bundle $T\mathcal{A}$ over a set of positive measure, which is transversally continuous in measure and invariant under \hat{f} . We say that the line field is *constant* if it is constant in the affine chart on any leaf. Note, if we are considering an orbifold leaf then this must take place in a finite cover – this allows the case of an orbifold point of order two, and a line field with a simple pole singularity. This is exactly what occurs for the deformable Lattès example.

Given a measurable set X and a set of positive Lebesgue measure Y on an affine leaf L , let $\text{dens}(X|Y) = \text{meas}(X \cap Y)/\text{meas}(Y)$ (note that this is a well defined quantity). Let us formulate some general ergodic properties of the conical set:

Proposition 8.9. • *The set Λ_f of conical points has either zero or full Lebesgue measure.*

- *In the latter case f is ergodic, except for the Lattès examples.*
- *Any invariant line field on Λ_f is constant, except for the isolated leaves of Lattès examples.*

Proof 1. This proof demonstrates how the blow-up method works in the lamination context.

Take any invariant locally fiber saturated set $X \subset \Lambda_f$ of positive measure. Then $X \cap L$ has positive measure for any leaf $L \subset \mathcal{A}$. Take a leaf L and a density point \mathbf{z} of X in L . Since \mathbf{z} is conical, there is a convergent sequence $\hat{f}^{n(k)}\mathbf{z} \rightarrow \zeta$. Take an arbitrary round disc $D \subset L(\mathbf{z})$ and a box neighborhood $D \times T$ of \mathbf{z} . Let $\zeta = (\zeta, \tau)$, $\mathbf{z}_{n(k)} = (z_k, t_k) \in D \times T$.

Let us consider round discs $\Delta_k = \hat{f}^{-n(k)}(D \times t_k)$ on the leaf $L(\mathbf{z})$. By the Shrinking Lemma, they shrink to \mathbf{z} and hence $\text{dens}(X|\Delta_k) \rightarrow 1$. Since \hat{f} is leafwise affine, $\text{dens}(X|(D \times t_k)) \rightarrow 1$ as $k \rightarrow \infty$. Since X is fiber saturated, $\text{dens}(X|(D \times \tau)) = 1$. Since the disc D is arbitrarily big, X has full measure on the leaf $L(\zeta)$.

Since the leaf $L(\zeta)$ is dense in \mathcal{A}_f (except the isolated leaves in the Lattès examples) and X is locally fiber saturated, it has full measure on every leaf. This proves the first two statements, except for the Lattès examples.

Finally, the first statement holds for the Lattès examples since $\Lambda_f = \mathcal{A}_f$ by Theorem 8.1 and Proposition 8.8. The second statement fails for the trivial reason that the isolated leaf is an invariant locally fiber saturated subset of \mathcal{A}_f . However, the previous argument shows that this leaf and its complement are the only subsets like this.

If now X supports an invariant line field μ , take \mathbf{z} to be a Lebesgue continuity point for this field on the leaf $L(\mathbf{z})$, so that μ is almost constant on $\Delta_k \cap X \setminus Y$ where $\text{dens}(Y|\Delta_k) \rightarrow 0$ as $k \rightarrow \infty$. It follows that $\mu|(D \times t_k)$ accumulates in measure on constant line fields. Since μ is transversally continuous in measure, $\mu|(D \times \tau)$ is constant almost everywhere, and hence almost everywhere on the leaf $L(\zeta)$. As this leaf is dense in \mathcal{A} , except for the isolated leaves of Lattès examples, the last statement follows as well. \square

Proof 2. This proof (for first two statements only) exploits Ahlfors’ harmonic extension method. Namely, let $X \subset \Lambda$ be a locally fiber saturated set of positive measure. Then we construct a

harmonic function on \mathcal{H}/\hat{f} by solving a Dirichlet problem on each leaf. That is, given an affine leaf $L \subset \mathcal{A}$ and the attached hyperbolic leaf $H_L \subset \mathcal{H}$, construct harmonic $h : H_L \rightarrow \mathbb{R}_+$ whose boundary values are 0 on \mathcal{F} and 1 on $\Lambda \cap L$. It is perhaps best to think of $h(x)$ for $x \in H_L$ as the area of $\Lambda \cap L$ as measured in the “visual metric” at x . That is, we map $H_L \cup L$ by Möbius transformation to the unit ball taking x to 0, and measure the area of the image of $\mathcal{J} \cap L$ on the unit sphere. This is the same as integrating the Poisson kernel against the characteristic function of $\Lambda \cap L$.

One must check that h is continuous (in the transverse direction). But if we fix a box neighborhood $T \times D$ for a large D in L (in the orbifold case this should be the finite cover of a box neighborhood), then for points near x (in the transverse direction) the visual measure induced on the leaves near L changes continuously on D (and if we choose D large enough, is very small on the complement of D in each leaf). The intersections of Λ with nearby leaves is a continuous family of analytic branched covers. It follows that area measure on Λ varies continuously in the transverse direction, and therefore so does its integral with respect to the visual measure.

Consider now a density point $\mathbf{z} \in X$ and the geodesic $\gamma_{\mathbf{z}} \subset L(\mathbf{z})$ terminating at this point. Then $h(p) \rightarrow 1$ as $p \rightarrow \mathbf{z}$ along $\gamma_{\mathbf{z}}$, since the visual area of X as seen from p is going to 1.

Observing also that h is invariant by \hat{f} , we obtain a continuous leafwise harmonic function g on the quotient \mathcal{H}/\hat{f} . Since \mathbf{z} is conical, the projection of $\gamma_{\mathbf{z}}$ to \mathcal{H}/\hat{f} has a limit point q . By continuity, $g(q) = 1$. By the Maximum Principle, g is identically equal to 1 on the whole leaf $L(q)$. By Proposition 7.6, this leaf is dense, except for the Lattès examples, and thus g is identically equal to 1 on the whole lamination. It follows that X has full measure. \square

Remark. Given Proposition 8.7, the results of the above proposition are not really new (compare [23], [7, Lemma 10], [30, Theorem 3.9]). However, the laminations give a new insight on them, and strengthen the connection to the corresponding results for Kleinian groups.

Corollary 8.10. *If f is not Lattès, then there are no invariant line fields on Λ_f which come from the sphere \mathbb{C} .*

Proof. It is easy to see that one can always find two leaves $L(\hat{z}_1)$ and $L(\hat{z}_2)$, with $\pi(\hat{z}_1) = \pi(\hat{z}_2) \equiv z$ such that $L(\hat{z}_1)$ is branched at \hat{z}_1 while $L(\hat{z}_2)$ is regular at \hat{z}_2 . Then the push-forward of a constant line field from $L(\hat{z}_1)$ has a singularity, while the push-forward from $L(\hat{z}_2)$ does not.

The only case when this does not lead to a contradiction is when one of the above leaves is isolated, so that the invariant line field is not necessarily constant on it. But this may happen only for the Lattès examples. \square

8.4. Elliptic structure of the Lattès examples. Let us show in conclusion how the invariant line field imposes the “elliptic structure” of the Lattès examples. We have seen that the invariant line field may exist only if there is an isolated leaf L^r . But then there should exist a non-isolated orbifold leaf L^s with an orbifold-constant line field.

Considering the projection $\pi : L^s \rightarrow \bar{\mathbb{C}}$ we see that the line field on $\bar{\mathbb{C}}$ is locally (a.e.) the image of the constant line field under a branched cover. It follows that the branching of π can be at most degree 2, and that the line field on $\bar{\mathbb{C}}$ can only have isolated index $-1/2$ (pole) singularities. By the index theorem on line fields, there must be exactly four of these. Thus $\bar{\mathbb{C}}$ has the structure of an orbifold with four order-2 singular points, the (2,2,2,2) orbifold (this is exactly Thurston’s orbifold for this map).

Let $X \subset \bar{\mathbb{C}}$ denote the above set of four singular points. It is clearly forward invariant under f . The property that the leaf L^r is isolated means that all backward orbits \hat{z} with $z_0 \in X$ eventually escaping X hit a critical point. In other words, $\pi : L^s \rightarrow \bar{\mathbb{C}}$ is double branched at all points of $L^s \cap \pi^{-1}X$, except the singular periodic point. Thus this map is an orbifold cover. (See e.g. Thurston [51] or Scott [42] for a discussion of orbifolds and orbifold covers).

Let $q : \tilde{L}^s \rightarrow L^s$ be the double covering associated to the orbifold structure of L^s , $\tilde{L}^s \approx \mathbb{C}$. It follows that $\pi \circ q : \tilde{L}^s \rightarrow (\bar{\mathbb{C}}, X)$ is an orbifold universal cover. The group of deck translations for such a cover is generated by a lattice of translations and the involution $z \mapsto -z$.

Let m be a period of the leaf L^s . Note that $\hat{f}^m : L^s \rightarrow L^s$ lifts (in two ways, because of choice of sign) to a multiplication map $g : z \mapsto nz$ on \tilde{L} . The constant n must be real, since g preserves the line field. On the other hand g commutes with $\pi \circ q$, so it preserves the lattice. Hence n is an integer. In other words the original map f is the projection of an integral torus endomorphism, i.e. a deformable Lattès example.

9. QUASI-ISOMETRIES AND RIGIDITY

9.1. Rigidity. In this section we will use the convex-cocompactness of the quotient 3-lamination to prove rigidity of critically non-recurrent maps without parabolic points, which extends Thurston's rigidity theorem (see [17]).

Theorem 9.1. *Let f and g be two critically non-recurrent rational maps without parabolic periodic points.*

1. *If f and g are topologically conjugate then they are quasi-conformally conjugate.*
2. *If the conjugacy is equivariantly homotopic to conformal on the Fatou sets, then f and g are Möbius conjugate, except for the Lattès examples.*

In particular, the second case holds automatically when the Julia sets of f and g coincide with the whole sphere.

Remarks. Thurston's proof of rigidity for post-critically finite maps used a contraction principle on a Teichmüller space, which is another aspect of the connection between rational maps and Kleinian groups (see [32, 33]).

Our proof uses another familiar scheme from both dynamics and hyperbolic geometry, which is roughly as follows. In step one, a topological conjugacy is promoted to a quasi-conformal conjugacy, using some geometric information. In step two, the quasi-conformal conjugacy is found to be conformal by an ergodic reasoning, because it induces an invariant line field on the Julia set.

In the convex cocompact case, the topological conjugacy is almost immediately quasi-conformal, because it gives rise to a homeomorphism on compact sets (the convex cores), which is automatically a quasi-isometry of the 3-laminations. This is directly analogous to the proof of Mostow's rigidity theorem in the case where the Fatou domain is empty, and to Marden's isomorphism theorem otherwise.

The second step, absence of invariant line fields, follows from the properties of the conical limit set given in the previous sections.

Proof. Let \mathcal{A}_f and \mathcal{A}_g be the affine orbifold laminations constructed from the natural extensions of f and g , and let \mathcal{H}_f and \mathcal{H}_g be the hyperbolic orbifold 3-laminations built over \mathcal{A}_f and \mathcal{A}_g .

Let $\Phi : \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be the homeomorphism conjugating f to g . Let $\hat{\Phi} : \mathcal{N}_f \rightarrow \mathcal{N}_g$ denote the natural extension of Φ , which conjugates the action of \hat{f} to that of \hat{g} . This map admits a continuous extension to a homeomorphism, which we also call $\hat{\Phi}$, from \mathcal{A}_f to \mathcal{A}_g , again conjugating \hat{f} to \hat{g} , and preserving orbifold affine structure. Indeed, Proposition 8.2 describes convergence in \mathcal{A}_f in dynamical terms which are respected by topological conjugacy. (Note: this is not obvious and maybe not true for critically recurrent maps.)

We may assume that Φ is quasi-conformal on the Fatou set $F(f)$, possibly after applying an equivariant homotopy. Let us give a sketch of this well-known procedure. Let \bar{a} be an attracting cycle. If it is not superattracting, we may choose a fundamental annulus around one of its points a . On this annulus we may homotope Φ , fixing it on the post-critical points, to some C^1 diffeomorphism which conjugates f to g on the boundary. This homotopy can then be transported by the action

of f and g to the rest of the attraction basin B of \bar{a} . By a Poincaré length argument the tracks of the homotopy have vanishing Euclidean length near ∂B , so that it can be extended as the identity to ∂B . Finally, Mané's Theorem implies that the diameters of the Fatou components tend to 0, so that the homotopy can be extended as the identity to the rest of the sphere.

If \bar{a} is superattracting, the Böttcher coordinate provides us with an invariant circle foliation in a punctured neighborhood of a . Moreover, this foliation is affine (that is, there is a canonical affine structure on the leaves), as the Böttcher coordinate is unique up to scaling and rotation. Select now a fundamental annulus, with the affine circle foliation inside and marked post-critical points. There is a homotopy of Φ in the fundamental annulus to some diffeomorphism, which respects this extra structure, and conjugates f and g on the boundary. By means of dynamics this homotopy can be spread around the whole basin B . By the same reason as above it can be extended to the rest of the sphere as the identity.

In the post-critically finite case the action of a power of f on the immediate basin of a (that is, the component of D containing a) is conjugate to $z \mapsto z^d$, and similarly for g (see [24], Theorem 1.6). Then Φ can be homotoped in the fundamental annulus to a diffeomorphism which is linear in the Böttcher coordinates. Hence it is conformal on the basin, and we are in case (2) of the theorem.

We next extend $\hat{\Phi}$ to a conjugacy of the 3-laminations, using the following elementary fact:

Lemma 9.2. *For any homeomorphism $\phi : \mathbb{C} \rightarrow \mathbb{C}$ there is a homeomorphism $e(\phi) : \mathbb{H}^3 \cup \mathbb{C} \rightarrow \mathbb{H}^3 \cup \mathbb{C}$ which restricts to ϕ on \mathbb{C} , such that the following are satisfied:*

1. *The extension is affinely natural: If α, β are (complex) affine maps of \mathbb{C} then $e(\alpha)$ and $e(\beta)$ are the unique possible similarities of \mathbb{H}^3 , and*

$$e(\alpha \circ \phi \circ \beta) = e(\alpha) \circ e(\phi) \circ e(\beta).$$

2. *$e(\phi)$ depends continuously on ϕ , in the compact-open topology on maps of \mathbb{C} and \mathbb{H}^3 .*
3. *$e(\phi)^{-1}$ depends continuously on ϕ or, equivalently, on ϕ^{-1} .*

Proof. The definition of $e(\phi)$ is the following:

$$e(\phi)(z, t) = \left(\phi(z), \max_{|w|=t} |\phi(z+w) - \phi(z)| \right).$$

Note in particular that the vertical line over each $z \in \mathbb{C}$ is mapped homeomorphically to the vertical line over $\phi(z)$, since the max is monotonic in t as a result of the assumption that ϕ is a homeomorphism. Hence the map is a homeomorphism. The other properties follow easily. Note that part (3) is not completely automatic since $e(\phi)^{-1}$ is not in general equal to $e(\phi^{-1})$. \square

As a corollary, we can extend $\hat{\Phi}$ leafwise to a map $\hat{E} : \mathcal{H}_f \rightarrow \mathcal{H}_g$, which is a homeomorphism on every leaf. The extension is well-defined because it is affinely natural. Note that, on the orbifold leaves, we must apply the lemma to the appropriate branched cover of the leaf. Since the map back to the orbifold leaf is quotient by rotations, the affine naturality of the extension implies that the extension is well-defined downstairs.

Continuity of \hat{E} follows from part (2) of lemma 9.2, applied to a local trivialization, i.e. a product-box (or orbifold-box) neighborhood in \mathcal{H}_f and in \mathcal{H}_g . Continuity of \hat{E}^{-1} follows from the same argument, using part (3) of lemma 9.2. Thus \hat{E} is a homeomorphism.

Again the affine naturality of the extension and the fact that \hat{f} and \hat{g} act by affine isomorphisms on the leaves imply that \hat{E} conjugates \hat{f} to \hat{g} . We conclude that it projects to a homeomorphism

$$E : \mathcal{H}_f / \hat{f} \rightarrow \mathcal{H}_g / \hat{g}.$$

We next show that the E can be deformed to a quasi-isometry:

Lemma 9.3. *There exist $K, \delta > 0$ and a map $\hat{E}' : \mathcal{H}_f \rightarrow \mathcal{H}_g$, which agrees with \hat{E} on \mathcal{A}_f and is a (K, δ) -quasi-isometry on each leaf.*

Proof. Note that to show a map $h : \mathbb{H}^3 \rightarrow \mathbb{H}^3$ is a quasi-isometry it suffices to show that there exist ϵ_1, ϵ_2 such that for all balls B of radius ϵ_1 , $\text{diam}(h(B)) \leq \epsilon_2$, and similarly for h^{-1} . Let us call this property quasi-Lipschitz, so that quasi-isometry is equivalent to quasi-Lipschitz in both directions.

Consider first the case that f (and therefore g) has no Fatou domain. In this case the convex cores are the entire quotients, and by Theorem 8.1 \mathcal{H}_f/\hat{f} and \mathcal{H}_g/\hat{g} are both compact. If we fix $\epsilon_1 > 0$ then the function $x \mapsto \text{diam}(\hat{E}(B(x, \epsilon_1)))$ is continuous in $x \in \mathcal{H}_f$ – as one can see by considering a local trivialization of the lamination. (Here $B(x, \epsilon_1)$ is a leafwise hyperbolic ball of radius ϵ_1 , and diam refers to diameter measured inside a leaf.) By compactness, then, it has a finite upper bound. Since we can do the same for \hat{E}^{-1} , we are done in this case.

In the case where the convex core \mathcal{C}_f is not all of \mathcal{H}_f , we first adjust the map so that it takes a small neighborhood of \mathcal{C}_f to \mathcal{C}_g .

Let $\mathcal{C}_f(\epsilon)$ denote the closed ϵ -neighborhood of \mathcal{C}_f , by which we mean the union of leafwise ϵ -neighborhoods. Note that $\mathcal{C}_f(\epsilon)/\hat{f}$ is still compact. Recall the product structure on $\mathcal{H} \setminus \mathcal{C}_f(\epsilon)$, discussed in Appendix 2 for the leafwise case, but extended to the global lamination by virtue of the discussion in §6.4 and lemma 12.3 on continuous variation of convex hulls. This product structure (in particular projection along the gradient lines) gives a C^1 identification between $\partial\mathcal{C}_f(\epsilon)$ and \mathcal{F}_f , and moreover we obtain a homeomorphism $P_f : \mathcal{H}_f \cup \mathcal{F}_f \rightarrow \mathcal{C}_f(\epsilon)$ which is the identity on \mathcal{C}_f , and equal to Π_ϵ on \mathcal{F}_f . On each leaf P_f is the map $h_{\epsilon, J}^{-1}$ discussed in the proof of lemma 12.3. Because the construction is natural, P_f commutes with \hat{f} .

Letting P_g denote the corresponding construction for g , we then have (fixing $\epsilon > 0$) a map

$$\hat{E}' = P_g \circ \hat{E} \circ P_f^{-1} : \mathcal{C}_f(\epsilon) \rightarrow \mathcal{C}_g(\epsilon)$$

which is a homeomorphism that restricts to a C^1 diffeomorphism on $\partial\mathcal{C}_f(\epsilon)$, and conjugates \hat{f} to \hat{g} . We can extend this to a map, also called \hat{E}' , on all of \mathcal{H}_f , using the product structure; that is, sending gradient lines to gradient lines at unit speed.

This map is the desired quasi-isometry. On $\mathcal{C}_f(\epsilon)$ it is quasi-Lipschitz as before, by the same compactness argument on the quotient; and similarly for $(\hat{E}')^{-1}$ on $\mathcal{C}_g(\epsilon)$. In the exterior, proposition 12.2 determines the metric up to bilipschitz homeomorphism in terms of the metric on the boundary of $\mathcal{C}_f(\epsilon)$ (or $\mathcal{C}_g(\epsilon)$). It follows that it is bilipschitz on the exterior, since \hat{E}' is a C^1 diffeomorphism on the boundary. (We are also using the fact that $\partial\mathcal{C}_f(\epsilon)/\hat{f}$ is compact to bound the derivatives of the map on the boundary).

Since \hat{E}' is a quasi-isometry it extends continuously to a quasiconformal homeomorphism on the boundary at infinity, namely \mathcal{A}_f . It remains to check that the boundary values of \hat{E}' agree with the original ones of \hat{E} , namely $\hat{\Phi}$. In the Fatou domain this is automatic from the construction. For any point in \mathcal{J}_f , we note that it lies in the closure of \mathcal{C}_f . For any point $x \in \mathcal{C}_f$, the maps \hat{E} and \hat{E}' differ by an application of P_g , so their leafwise distance is (again by compactness of the quotient) uniformly bounded. It follows that the two maps have identical boundary values on \mathcal{J}_f . \square

We can now complete the proof of theorem 9.1. Lemma 9.3 implies that $\hat{\Phi}$ extends to a quasi-isometry of the 3-laminations – that is, a map which is a quasi-isometry on every leaf, with uniform constants. and therefore (lemma 12.4) $\hat{\Phi}$ is in fact a quasiconformal map on every leaf, with uniform constant. Since $\hat{\Phi}$ is just the lift of the original conjugacy Φ , we conclude that Φ itself is quasiconformal.

This concludes step one of the proof (that topological conjugacy implies quasi-conformal), which is case (1) of the theorem. To finish the proof we need to show that a quasi-conformal conjugacy

which is conformal on the Fatou set is Möbius, except for the Lattès examples. But this is equivalent to the absence of invariant line fields on the Julia set which follows from Proposition 8.8 and Corollary 8.10.

10. FURTHER PROGRAM

Let us outline some possible directions for further development, problems and conjectures.

1. *Regular leaf space.* Study the regular leaf space \mathcal{R}_f in more detail. What is the behaviour of the leaves of \mathcal{R}_f near irregular points? In particular, look at the Feigenbaum case. What happens to \mathcal{R}_f at a parabolic bifurcation? Other than rotation domains, are there any leaves which are not dense? (Lemma 4.8 shows that all parabolic leaves are dense.) Can it happen that a leaf other than a rotation domain does not intersect the Julia set?

2. *Type Problem* (see §4). Are there hyperbolic leaves in \mathcal{R}_f except for Siegel disks and Herman rings? It seems that the right place to look for hyperbolic leaves are maps with non-locally connected Julia set (Cremer points or infinitely renormalizable polynomials of highly unbounded type, see [37]). Prove that all leaves of a “fake Feigenbaum” quadratic (that is, a rational map which is topologically equivalent to the Feigenbaum quadratic) are parabolic. Conjecturally there are no fake Feigenbaum maps (a special case of the rigidity problem), but this would be the first step of trying to apply the laminations to this problem. More generally does the topological type of the map determine the conformal types of the leaves?

3. *Uniqueness problem.* In general, can one reconstruct f from its 3-lamination? How does the lamination detect the difference between polynomial and polynomial-like maps?

4. *Geometric finiteness.* There are many definitions of geometrically finite Kleinian groups, all equivalent for dimensions 2 and 3 (see Maskit [29], Bowditch [8]). The definition in terms of finite-sided fundamental domain (see Ahlfors [1]) seems to fail altogether in the lamination context; it is also not equivalent to the others for hyperbolic manifolds in higher dimensions [8]. The definition in terms of conical and parabolic points (Beardon-Maskit [3]) can be translated into the lamination setting. We expect it to pick out critically non-recurrent maps with or without parabolic points. Thurston’s definition in terms of finite volume of a neighborhood of the convex core, or compact thick part of the convex core (similar also to Marden’s definition in [27]) seems harder to transport to laminations. Is there a good replacement for the notions of volume and injectivity radius which would make this translation work?

5. *Deformation theory.* Describe the space of \mathbb{H}^3 laminations, or affine 2-laminations, or just those arising from rational maps. A fundamental difficulty here is that there is no common “universal cover”, as there is for hyperbolic manifolds.

6. *Topology of \mathcal{H}_f/\hat{f} .* What is the topological structure of \mathcal{H}_f and \mathcal{H}_f/\hat{f} ? Does \mathcal{H}_f/\hat{f} always have two ends for quadratic f ?

Particular cases are the Axiom A polynomials (take $z \mapsto z^2 - 1$ first) and the Feigenbaum quadratic. Is there an internal structure to \mathcal{H}_f that mirrors the sequence of bifurcations going from $z \mapsto z^2$ to f (degree 2 case)?

Let us consider the following model. Let $f_c : z \mapsto z^2 + c$, $c \in [c_0, 0]$, where c_0 is the Feigenbaum point, or any point preceding it. Let K_c and J_c denote the filled Julia set and the Julia set for f_c . Consider their lifts \mathcal{K}_c and \mathcal{J}_c to \mathcal{R}_f . Consider the set $\mathcal{M} = \{(c, \hat{z}) : c_0 \leq c \leq 0, \hat{z} \in \mathcal{K}_c\}$.

There is a natural projection from J_c onto J_{c_0} , since J_{c_0} is obtained from J_c by some “pinchings” (compare Douady [15]). This induces a projection $r_c : \mathcal{J}_c \rightarrow \mathcal{J}_{c_0}$. Let us consider the quotient \mathcal{M}/\sim where the equivalence relation \sim identifies (c, \hat{z}) , $\hat{z} \in \mathcal{J}_c$ with $(c_0, r_c \hat{z})$. The map f induces a self-map \tilde{f} of \mathcal{M}/\sim .

Is $\tilde{f} : \mathcal{M}/\sim \rightarrow \mathcal{M}/\sim$ topologically equivalent to $\hat{f}_{c_0} : \mathcal{H}_{f_{c_0}} \cup \mathcal{A}_{f_{c_0}} \rightarrow \mathcal{H}_{f_{c_0}} \cup \mathcal{A}_{f_{c_0}}$?

7. *Geometry of \mathcal{H}_f/\hat{f} .* Give a quasi-isometric model for \mathcal{H}_f/\hat{f} . Does topology of this lamination determine its geometry? (It is certainly a quite strong version of the Rigidity Problem).

Can one place “pleated solenoids” inside \mathcal{H}_f/\hat{f} , and use them in analogy with pleated surfaces in hyperbolic 3-manifolds? (In the Feigenbaum case, one can consider the pullback of the little Julia set $J(R^n f)$ to \mathcal{A}_f (where R denotes the renormalization operator), take the boundary of its convex hull in \mathcal{H}_f , and spread it around by iterates of \hat{f}).

8. *Spectral Theory.* We define the three dimensional Poincaré series of \hat{f} by taking a transversal K of \mathcal{H}_f , averaging $\exp(-\rho(\hat{f}^{-n}x, K))$ along a natural transversal measure on K (where ρ stands for the leafwise hyperbolic distance), and summing up over n (see Su [43] for a discussion of the transversal measure). Is it true that the corresponding critical exponent coincides with the Hausdorff dimension of the conical limit set? A natural further project is to develop a spectral theory on the lamination \mathcal{H}_f/\hat{f} , and to study measure and dimension of the Julia sets from this point of view (compare Sullivan [44, 45], Canary [10], Bishop-Jones [6], Denker-Urbanski [18]). The Ahlfors-type argument used in §8 of this paper is a first step in this direction.

9. *Added leaves of \mathcal{A}_f .* Can it happen that \mathcal{A}_f is not locally compact? This problem requires understanding of the added leaves of \mathcal{A}_f . What one can say about the entire function corresponding to the leaf projection $p : L_{aff}(\mathbf{z}) \rightarrow L_{aff}(p(\mathbf{z}))$? Can it have asymptotic values? (In the critically non-recurrent case it is polynomial.)

10. *Action of rational functions in the Universal space.* It would be interesting to have a general idea of this action. What is the structure of the characteristic attractor K_f ? Is a generic $f : \mathcal{U} \rightarrow \mathcal{U}$ injective? More precisely, let us consider a functional equation $f \circ \phi = f \circ \psi$ where $\phi, \psi \in \mathcal{U}$ are meromorphic. Is it true that any solution of this equation has a form $\phi = \gamma \circ \psi$ where γ is a *symmetry* of f (that is, a Möbius transformation such that $f \circ \gamma = f$), or $\phi = \psi \circ \delta$ where δ is a rotation? See Fatou [20] and Ritt [40] for further discussion of this problem (we are grateful to A. Eremenko for providing these references).

11. APPENDIX 1: CIRCLE AND POLYNOMIAL-LIKE MAPS

11.1. **Sullivan’s laminations for circle maps.** Let $f : S^1 \rightarrow S^1$ be a C^2 expanding map of the circle of degree $d > 1$. The expanding property means that there exist constants $C > 0$ and $\lambda > 1$ such that $|Df^n(x)| \geq C\lambda^n$, $n = 0, 1, \dots$. Sullivan’s construction goes as follows (see Sullivan [48, 47], and de Melo-van Strien [35]):

Step (i). Consider the natural extension $\hat{f} : \mathcal{N}_f \rightarrow \mathcal{N}_f$. Topologically \mathcal{N}_f is the standard solenoid over the circle. Dynamically \hat{f} is a hyperbolic (in the sense of Anosov and Smale) map with one-dimensional unstable leaves.

Step (ii). Supply the leaves with the affine structure by means of the explicit formula (4.2) (existence of the limit follows from the standard distortion estimates for hyperbolic maps). The map \hat{f} preserves this structure.

Step (iii). Attach hyperbolic planes to the leaves and extend \hat{f} to the corresponding hyperbolic 2-lamination \mathcal{H}_f^2 acting isometrically on the leaves.

Step (iv). Take the quotient \mathcal{H}^2/\hat{f} . This is Sullivan’s Riemann surface lamination associated to f . Topologically it is a solenoidal fibration over the circle.

The main difference between this construction and the one outlined in the Introduction is related to the critical points on the Julia set. These tend to distort the affine structures and complicate the transversal behavior of the leaves. Also, as we have seen, even in the Axiom A case the topological structure of the 3-lamination is not at all obvious.

Sullivan constructed 2-laminations to build up the deformation space of expanding circle maps. We try to study rigidity phenomenon by means of 3-laminations. This is a usual philosophical difference between dimensions two and three.

11.2. Polynomial-like maps: globalization of the leaves. Polynomial-like maps are not globally defined, and certainly cannot be in general extended to the whole sphere. However, such a globalization can be carried out on the natural extension level. Lemma 11.1 shows that it leads to the same object, provided the map was a priori globally defined.

Let U and V be two open sets of \mathbb{C} such that $\text{cl}U \subset V$, and $f : U \rightarrow V$ be an analytic branched covering. Keep in mind Douady-Hubbard polynomial-like [16] maps, generalized polynomial-like maps [25], or a rational function R restricted on the sphere minus an invariant neighborhood of attracting cycles.

For such a map we can consider the space \mathcal{N}_f of backward orbits, and lift f^{-1} to this space as the map which forgets the first coordinate: $\hat{g} \equiv \hat{f}^{-1} : \mathcal{N}_f \rightarrow \mathcal{N}_f$. This map is injective but not surjective: its image consists of the orbits which start with a $z_0 \in U$.

To make it invertible, let us consider the *inductive (direct) limit* of

$$\mathcal{N} \xrightarrow{\hat{g}} \mathcal{N} \xrightarrow{\hat{g}} \mathcal{N} \xrightarrow{\hat{g}} \dots,$$

which is defined in the following way. Take infinitely many copies \mathcal{N}^m of the same space \mathcal{N} . Let us embed \mathcal{N}^m into \mathcal{N}^{m+1} by means of the map

$$i_m \equiv \hat{g} : \mathcal{N}^m = \mathcal{N} \rightarrow \mathcal{N} = \mathcal{N}^{m+1}.$$

In other words, we identify a point $\hat{z} \in \mathcal{N}^m$ with the point $i_m \hat{z} \in \mathcal{N}^{m+1}$. Thus we obtain an increasing sequence of the spaces

$$\mathcal{N}^0 \hookrightarrow \mathcal{N}^1 \hookrightarrow \mathcal{N}^2 \hookrightarrow \dots \tag{11.1}$$

Let $\mathcal{D} \equiv \mathcal{D}_f = \cup \mathcal{N}^m$. To define a topology on \mathcal{D} , let us call a set $W \subset \mathcal{D}$ open if $W = \cup W_i$ where W_i is an open set in \mathcal{N}^i .

The map $\hat{g} : \mathcal{N}^k \rightarrow \mathcal{N}^k$ respects the embeddings $i_m : \mathcal{N}^m \hookrightarrow \mathcal{N}^{m+1}$, and hence induces the self-map of \mathcal{D} , which we will denote by the same letter. Moreover, \hat{g} homeomorphically maps \mathcal{N}^m onto $i_{m-1} \mathcal{N}^{m-1}$, $m > 0$, so that it is invertible on \mathcal{D} . We will keep the notation \hat{f} for \hat{g}^{-1} .

Lemma 11.1. *Assume that a branched covering $f : U \rightarrow V$ is the restriction of a rational endomorphism $R : \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ such that $\mathbb{C} \setminus V$ is contained in the basin of attraction of a finite attracting set A . Then $\hat{f} : \mathcal{D}_f \rightarrow \mathcal{D}_f$ is naturally conjugate to $\hat{R} : \mathcal{N}_R \setminus \hat{A} \rightarrow \mathcal{N}_R \setminus \hat{A}$.*

Proof. Let us consider the following commutative diagram:

$$\begin{array}{ccccccc} \mathcal{N}^0 & \hookrightarrow & \mathcal{N}^1 & \hookrightarrow & \mathcal{N}^2 & \hookrightarrow & \dots \\ & & \downarrow \text{id} & & \downarrow \hat{R} & & \downarrow \hat{R}^2 \\ \mathcal{N} & \hookrightarrow & \hat{R}\mathcal{N} & \hookrightarrow & \hat{R}^2\mathcal{N} & \hookrightarrow & \dots \end{array}$$

where $\mathcal{N} \equiv \mathcal{N}_f$, the upper line is the sequence (11.1) for \hat{f} , while the lower one is the sequence of natural inclusions. It induces a homeomorphism between \mathcal{D}_f and $\cup \hat{R}^n \mathcal{N} = \mathcal{N}_R \setminus \hat{A}$, which is the desired conjugacy. \square

12. APPENDIX 2: BACKGROUND MATERIAL

12.1. Dynamics. We assume the following background in holomorphic dynamics:

- Classification of periodic points as *attracting*, *repelling*, *parabolic*, *Siegel* and *Cremer*, and the local dynamics near these points;
- Notions of the Julia set $J(f)$ and the Fatou set $F(f)$;
- Classification of components of the Fatou set as *attracting basins*, *parabolic basins*, *Siegel disks* and *Herman rings*; Siegel disks and Herman rings will be also called the *rotation sets*.

- The notion of an *Axiom A* or *hyperbolic* rational function. There are two equivalent definitions of this property:
 - All critical points are in basins of attracting cycles;
 - The map is uniformly expanding on the Julia set, that is, there exist constants $A > 0$ and $\lambda > 1$ such that for any $z \in J(f)$,

$$\|Df^n(z)\| \geq A\lambda^n, \quad n = 0, 1, 2, \dots,$$

where $\|\cdot\|$ denotes the spherical metric.

All this material can be found in any book or survey in holomorphic dynamics – e.g. [12, 24, 36].

As usual, $\omega(z) \equiv \omega_f(z)$ denotes the ω -limit set of a point z . A point z is called *recurrent* if $z \in \omega(z)$. Given a set Z , let

$$\text{orb}(Z) = \bigcup_{z \in Z} \text{orb } z, \quad \omega(Z) = \bigcup_{z \in Z} \omega(z).$$

Let C denote the set of critical points of f , and C_r the set of recurrent critical points.

The critical values of f^n are the points of $f^k C$, $1 \leq k \leq n$. So if a simply connected neighborhood U does not meet $\text{orb } C$ then all inverse branches of f^{-n} are well defined univalent functions in U .

The *non-linearity*, or *distortion* of a conformal map $\psi : U \rightarrow \mathbb{C}$ is defined as

$$\text{Dis}(\psi) = \sup_{z, \zeta \in U} \log \left| \frac{\psi'(z)}{\psi'(\zeta)} \right|.$$

Koebe Distortion Theorem. *Let $\psi : B(a, r) \rightarrow \mathbb{C}$ be a conformal map, $k < 1$. Then the distortion of ψ in $B(a, kr)$ is bounded by a constant $C(k)$ independent of ψ . Moreover $C(k) = O(k)$ as $k \rightarrow 0$.*

Let $U \subset \bar{\mathbb{C}}$ be any domain. Let us select a base point $z \in U$, and count its n -fold preimages: z_i^n . Let U_i^{-n} denote a component of $f^{-n}U$ containing z_i^n . This specifies a “multi-valued branch” f_i^{-n} of the inverse map. (The reader can think of these branches as functions living on appropriate Riemann surfaces, or as equivalence relations, or just as a convenient way of describing the situation). Singular points for an inverse branch are critical values for the direct map. There is a natural way of composing and restricting the inverse branches (with an appropriate adjustment of the base points, which may change only the way of counting).

The following lemma is a variation of a well-known fact (compare [24], Proposition 1.10). As it plays a crucial role for this paper, we will include the proof.

Shrinking Lemma. *Let f be a rational map of degree $d > 1$. Let $U \subset \mathbb{C}$ be a domain which is not contained in any rotation set of f , and let k be a natural number. Let us consider a family $\{f_i^{-n}\}$ of all inverse branches in U with at most k singular points (counting with multiplicities). Then for any domain W compactly contained in U , $\text{diam}(f_i^{-n}|_W) \rightarrow 0$ as $n \rightarrow \infty$ independently of i (where diam denotes spherical diameter).*

Proof. We first consider the case that U , and every pullback U_i^{-n} , are disks. Let $z \in U$ be a point outside any rotation domain of f .

Let $\Phi_{n,i} : D \rightarrow U_i^{-n}$ be a Riemann mapping taking 0 to a preimage of z , where D is the unit disk. Then $\pi_{n,i} = f^n \circ \Phi_{n,i}$ is a proper branched covering from D to U , with at most k critical points counted with multiplicity. (One can think of the disk D here as the Riemann surface over U for the corresponding branch of the inverse function.)

Let $\alpha_1, \dots, \alpha_k$ be a periodic cycle of f of length at least 3, not meeting some neighborhood of z . Then no preimage of this neighborhood meets the cycle either. By normality of the family $\{\pi_{n,i}\}$, there must be some disk D' compactly contained in D such that $\Phi_{n,i}(D')$ omits $\{\alpha_j\}$ for all n, i , and such that $\pi_{n,i}(D') \ni z$. Thus $\{\Phi_{n,i}\}$ is a normal family on D' .

Because of the bound k on the number of critical points of $\pi_{n,i}$, there is some δ such that the disk $B = B(z, \delta)$ is contained in $\pi_{n,i}(D')$ for all n, i (one can show this for example by noting that $\pi_{n,i}^{-1}(U \setminus B(z, \delta))$ contains an annulus whose modulus is bounded below depending only on k and δ , and goes to ∞ as $\delta \rightarrow 0$). We now claim that the diameters $\text{diam}(B_i^{-n})$ go to 0 uniformly.

If not, we can extract a convergent subsequence $\Phi_{n_k, i_k}|_{D'}$, and conclude that for the limit point $z_\infty = \lim \Phi_{n_k, i_k}(0)$ there is a neighborhood B_∞ whose images under arbitrarily high iterates are in U . This implies in particular that B_∞ (and therefore B) is disjoint from the Julia set (as any neighborhood intersecting the Julia set covers it under some iterate of f). By a smaller choice of δ we may assume it is compactly contained in the Fatou set. Thus, either forward iterates of B_∞ under f limit to an attracting/parabolic periodic cycle, or B_∞ is contained in a rotation domain. The former is impossible since $f^{n_k}(B_\infty)$ limits onto all of B . The latter is ruled out by the choice of z .

It now follows that $\text{diam}(W_i^{-n}) \rightarrow 0$ for any W compactly contained in U , since $\Phi'_{n,i}$ must converge to 0 uniformly on compact sets.

To treat the general case, take a finite covering of W by disks D compactly contained in U , none of which are contained in a rotation domain. We must consider the possibility that some of the pullbacks D_i^{-n} are not disks. For any $\epsilon > 0$ there exists $N = N(D, \epsilon) > 0$ such that, if D_i^{-n} is a disk and $n \geq N$, then $\text{diam} D_i^{-n} \leq \epsilon$. For if not, we could find a subfamily of pullbacks, all disks, whose diameters fail to shrink to 0. The previous argument applies, so this is impossible.

Thus, let ϵ be less than half the distance between any two critical values of f . Then the preimage of any disk of diameter less than ϵ is a disjoint union of disks. It follows that, if some D_i^{-n} is not a disk then some image D_i^{-m} of it, with $0 \leq m \leq N$, is also not a disk. That is, the transition from disk to non-disk occurs in the first N levels. Thus, if we remove from consideration the finite number of non-disks D_j^{-n} with $n \leq N$, and all their preimages, we are left with a family in which all preimages are disks. For this subfamily, we have uniform shrinking by the previous arguments.

For each of the finitely many non-disks D_j^{-n} ($n \leq N$), we can now repeat the argument, covering W_j^{-n} with disks not contained in rotation domains, and so on. However now the bound on the number of singular points is $k - 2$, since in the transition from disk to non-disk at least two singular points must be used. We can therefore obtain a uniform rate of shrinking for this family, by induction on k . This concludes the proof. \square

A key result on critically non-recurrent rational maps is the following theorem of Mañé [26] closely related to the Shrinking Lemma.

Mañé's Theorem [26]. *Let $f : \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$ be a rational map. If a point $x \in J(f)$ is neither a parabolic periodic point, nor belongs to the ω -limit set of a recurrent critical point then, for all $\epsilon > 0$, there exists a neighborhood U of x such that for all $n \geq 0$ every connected component of $f^{-n}(U)$ has diameter $\leq \epsilon$.*

Chebyshev and Lattès examples. Let us finally dwell on the remarkable examples of rational functions whose dynamics often present some special features.

The *Chebyshev polynomial* p_d of degree d can be defined by means of the functional equation $p_d(\cos z) = \cos(dz)$. In other words, consider the dilation map $T_d : z \mapsto dz$ on the cylinder $C = \mathbb{C}/2\pi\mathbb{Z}$. Then p_d is the quotient of this map via the involution $z \mapsto -z$.

The Julia set of p_d coincides with the interval $[-1, 1]$. The endpoint 1 is always fixed, while -1 is either fixed (for odd degrees) or pre-fixed (for even degrees). Any critical point is mapped by p_d to one of the endpoints.

Similarly, the Lattès examples come from the functional equations $f_d(P(z)) = P(dz)$, where $P : \mathbb{C} \rightarrow \bar{\mathbb{C}}$ is a Weierstrass P -function, $\deg f_d = |d|^2$ where d is not necessarily integer. They can be viewed as quotients of torus endomorphisms. That is, let $\mathbb{T} = \mathbb{C}/\Lambda$ be a torus, where Λ is a

lattice. Then identifying z with $-z$ sends \mathbb{T} to $\bar{\mathbb{C}}$ via a two-fold branched cover. If $T_d(\Lambda) \subset \Lambda$ then the dialtion T_d induces a torus endomorphism, which further projects to a rational map of $\bar{\mathbb{C}}$ of degree $|d|^2$. (This occurs for all integer d 's on any torus, but also for some special tori and special non-real values of d : take, e.g., the standard lattice $\Lambda = \mathbb{Z}^2$ and $d = 1 + i$).

The Julia set of the Lattès examples is the whole sphere. Like in the Chebyshev case, every critical point of a Lattès map is pre-fixed.

The following dynamical characterization of these examples is well-known:

Proposition 12.1. *Assume that a rational map f has a periodic point $a \in J(f)$ such that every backward trajectory $a = a_0, a_{-1}, \dots$ which passes through a only finitely many times hits a critical point. Then f is either Chebyshev or Lattès.*

We will see in this paper how this property manifests itself in the lamination structure.

For integer values of d the Lattès maps are quasi-conformally deformable, since Λ may be varied (or, since the constant line field on the torus is dilation invariant). Conjectually they are the only examples which admit quasi-conformal deformations on the Julia set. We will see a lamination reasoning behind this conjecture.

12.2. Geometry.

Hyperbolic geometry and convex hulls. We assume familiarity with the hyperbolic space \mathbb{H}^3 and its boundary at infinity the Riemann sphere. (See e.g. Beardon [2], Thurston [51]). Most natural for us will be the upper half space model $\mathbb{C} \times \mathbb{R}_+$.

We recall some fundamental facts about hyperbolic convex hulls. Most of these facts appear in Epstein-Marden [19], or can be obtained from that paper with a small amount of effort.

The convex hull $C = C(E) \subset \mathbb{H}^3$ of a closed set E on the Riemann sphere $\bar{\mathbb{C}}$ is defined as the smallest convex set in \mathbb{H}^3 whose closure in $\mathbb{H}^3 \cup \bar{\mathbb{C}}$ contains E . Equivalently, C is the intersection of all closed half-spaces in \mathbb{H}^3 containing E at infinity. Provided E is not contained in a round circle, $C \cup E$ is homeomorphic to a closed 3-ball, and ∂C is a subsurface of \mathbb{H}^3 , which is isometric to a complete hyperbolic surface, using the metric of shortest paths in ∂C .

The geometry of the complement $\mathbb{H}^3 - C$ is well-understood. We begin with the projection $\Pi : \mathbb{H}^3 \rightarrow C$ assigning to $x \in \mathbb{H}^3$ the point in C nearest to x , which is unique by the convexity of C . This projection also extends continuously to $\bar{\mathbb{C}} - E$.

Let $d : \mathbb{H}^3 \rightarrow [0, \infty)$ be the distance function $d(x) = d_{\mathbb{H}^3}(x, C)$. This is a C^1 function in $\mathbb{H}^3 - C$, and its gradient is the unit vector tangent to the geodesic through x and $\Pi(x)$, and pointing away from $\Pi(x)$ (lemma 1.3.6 in [19]). In fact these geodesics are the integral lines of this gradient field, and they foliate $\mathbb{H}^3 \setminus C$. The gradient vector field itself is Lipschitz, with a uniform constant outside a neighborhood $C_\epsilon = d^{-1}([0, \epsilon])$, for any fixed $\epsilon > 0$ (see §2.11 in [19]).

The level surfaces $S_\epsilon = d^{-1}(\epsilon)$ are, therefore, C^1 submanifolds for $\epsilon > 0$, and are all homeomorphic via the gradient flow. Since each gradient line terminates at infinity, the level surfaces can be identified with $\bar{\mathbb{C}} \setminus E$, which we may label S_∞ . Thus we have a natural product structure identifying $\mathbb{H}^3 \cup \bar{\mathbb{C}} \setminus (C \cup E)$ with $(0, \infty] \times S_\epsilon$ for $\epsilon \in (0, \infty]$.

The identification between S_ϵ and S_∞ is a quasiconformal map, and in fact the following is a consequence of Theorem 2.3.1 in [19]:

Proposition 12.2. *Let σ denote the Poincaré metric on $S_\infty = \bar{\mathbb{C}} - E$. Let ρ denote the metric on $(0, \infty) \times S_\infty$ given infinitesimally as*

$$d\rho^2 = dr^2 + (\cosh^2 r)d\sigma^2$$

where $r \in (0, \infty)$ is the first coordinate. The identification of $(\epsilon, \infty) \times S_\infty$ with $\mathbb{H}^3 - C_\epsilon(E)$ is bilipschitz with constant L depending only on $\epsilon > 0$.

The dependence of $C(E)$ (or $C_\delta(E)$) on E is *continuous*, with respect to the Hausdorff topology on closed subsets of the ball $\mathbb{H}^3 \cup \bar{\mathbb{C}}$. This is easy in our setting; a proof for a more general context appears in Bowditch [9]. In fact more is true: on compact sets in \mathbb{H}^3 , a small variation of E produces a locally homeomorphic deformation of C_δ :

Lemma 12.3. *Let there be given a closed $E_0 \subset \bar{\mathbb{C}}$, a hyperbolic R -ball $B(x, R)$ around a point $x \in \mathbb{H}^3$, and $\delta > 0$. For each $\epsilon > 0$ there is a neighborhood U of E_0 in the Hausdorff topology on closed subsets of $\bar{\mathbb{C}}$ such that, for any $E \in U$, there is a $(1 + \epsilon)$ -bilipschitz map $\Psi_E : B(x, R) \rightarrow \mathbb{H}^3$ fixing x , such that $\Psi_E^{-1}(C_\delta(E)) = C_\delta(E_0) \cap B(x, R)$.*

Remarks. (1) In particular, note that (δ -neighborhoods of) convex hulls of sufficiently nearby sets are, locally, homeomorphic, *even if the sets themselves are not homeomorphic*. (2) We take C_δ rather than C itself here in order to avoid the exceptional case when E_0 lies on a round circle. Then the convex hull fails to have interior, and is not homeomorphic to convex hulls of nearby sets. In all other situations the lemma holds for $C_0 = C$.

Proof. We give only a sketch, and refer the reader to [19] for a thorough treatment of the techniques.

Using the product structure on $\mathbb{H}^3 - C_\delta(E)$ discussed above, there is a homeomorphism $h_{\delta, E} : C_\delta(E) \rightarrow \mathbb{H}^3 \cup S_\infty(E)$, which expands segments to gradient lines, and is the identity on $C(E)$. Now note that, for a fixed ball $B(x, R)$ and E sufficiently close to E_0 , the image $h_{\delta, E_0}(B(x, R) \cap C_\delta(E_0))$ misses E . Therefore the map $h_{\delta, E}^{-1} \circ h_{\delta, E_0}$ is defined on $B(x, R) \cap C_\delta(E_0)$. Extend to the rest of $B(x, R)$, again using the product structure. \square

Quasi-isometries and QC maps. We call a map $h : \mathbb{H}^3 \rightarrow \mathbb{H}^3$ a (K, δ) quasi-isometry if the following holds for all $p, q \in \mathbb{H}^3$:

$$\frac{1}{K}d(p, q) - \delta \leq d(h(p), h(q)) \leq Kd(p, q) + \delta.$$

The connection (in one direction) of quasi-isometries to quasi-conformal maps is given by the following lemma. For a proof, see Thurston [50] or (in the more general context of hyperbolic spaces in the sense of Gromov) [14, 21].

Lemma 12.4. *Given (K, δ) there exists L so that any (K, δ) -quasi-isometry $h : \mathbb{H}^3 \rightarrow \mathbb{H}^3$ extends continuously to an L -quasiconformal homeomorphism $\tilde{h} : \bar{\mathbb{C}} \rightarrow \bar{\mathbb{C}}$.*

REFERENCES

- [1] L. Ahlfors, *Fundamental polyhedrons and limit point sets for Kleinian groups*, Proc. Nat. Acad. Sci. **55** (1966), 251–254.
- [2] A. F. Beardon, *The geometry of discrete groups*, Springer-Verlag, 1983.
- [3] A. F. Beardon and B. Maskit, *Limit sets of Kleinian groups and finite sided fundamental polyhedra*, Acta Math. **132** (1974), 1–12.
- [4] E. Bedford, M. Lyubich, and J. Smillie, *Polynomial diffeomorphisms of \mathbf{C}^2 : The measure of maximal entropy and laminar currents*, Inventiones Math. **112** (1993), 77–125.
- [5] T. Bedford, A. M. Fisher, and M. Urbanski, *The scenery flow for hyperbolic Julia sets*, (in preparation).
- [6] C. Bishop and P. Jones, *Hausdorff dimension and Kleinian groups*, Stony Brook IMS Preprint #1994/5.
- [7] A. M. Blokh and M. Yu. Lyubich, *Attractors of the maps of the interval*, in “Dynamical Systems and Ergodic Theory”, Banach Center Publ., **23** (1989), 427–442.
- [8] B. Bowditch, *Geometrical finiteness for hyperbolic groups*, J. Func. Anal. **113** (1993), 245–317.
- [9] ———, *Some results on the geometry of convex hulls in manifolds of pinched negative curvature*, Comment. Math. Helv. **69** (1994), 49–81.
- [10] R. D. Canary, *On the laplacian and the geometry of hyperbolic 3-manifolds*, J. Differential Geom. **36** (1992), 349–367.
- [11] A. Candel, *Uniformization theorem for surface laminations*, Ann. scient. Éc. Norm. Sup., 4^e série t. **26** (1993), 489–516.
- [12] L. Carleson and T. W. Gamelin, *Complex dynamics*, Springer-Verlag, 1993.

- [13] L. Carleson, P. W. Jones, and J.-C. Yoccoz, *Julia and John*, Bol. Soc. Brasil. Mat. (N.S.) **25** (1994), 1–30.
- [14] M. Coornaert, T. Delzant, and A. Papadopoulos, *Géométrie et théorie de groupes: les groupes hyperboliques de Gromov*, Springer-Verlag, 1990.
- [15] A. Douady, *Descriptions of compact sets in \mathbf{C}* , Topological Methods in Modern Mathematics (L. Goldberg and A. Phillips, eds.), Publish or Perish, 1993, pp. 429–466.
- [16] A. Douady and J. Hubbard, *On the dynamics of polynomial-like mappings*, Ann. Sci. Éc. Norm. Sup. **18** (1985), 287–344.
- [17] ———, *A proof of Thurston's topological characterization of rational maps.*, Acta Math. **171** (1993), 263–297.
- [18] M. Denker and M. Urbanski. *On Sullivan's conformal measure for rational maps of the Riemann sphere*. Non-linearity, **4** (1991), 365–384.
- [19] D. B. A. Epstein and A. Marden, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces*, Analytical and Geometric Aspects of Hyperbolic Space, Cambridge University Press, 1987, London Math. Soc. Lecture Notes Series no. 111, pp. 113–254.
- [20] P. Fatou, *Sur un groupe de substitutions algébriques*, C. R. de l'Acad. Sci. **173** (1921) 694–696.
- [21] E. Ghys and P. de la Harpe, *Sur les groupes hyperboliques d'après Mikhael Gromov*, Birkhäuser, 1990.
- [22] R. S. Hamilton, *The Inverse Function Theorem of Nash and Moser*, Bull. Amer. Math. Soc. **7** (1982), #1, 65–222.
- [23] M. Lyubich, *On typical behavior of the trajectories of a rational mapping of the sphere*, Soviet Math. Dokl. **27** (1983), 22–25.
- [24] ———, *The dynamics of rational transforms: the topological picture*, Russian Math. Surveys **41:4** (1986), 43–117.
- [25] ———, *On the Lebesgue measure of the Julia set of a quadratic polynomial*, Stony Brook IMS Preprint #1991/10.
- [26] R. Mañé, *On a theorem of Fatou*, Bol. Soc. Brasil. Mat., **24** (1993), 1–12.
- [27] A. Marden, *The geometry of finitely generated Kleinian groups*, Ann. of Math. **99** (1974), 383–462.
- [28] A. Marden and B. Maskit, *On the isomorphism theorem for Kleinian groups*, Invent. Math. **51** (1979), 9–14.
- [29] B. Maskit, *Kleinian groups*, Springer-Verlag, 1988.
- [30] C. McMullen, *Complex dynamics and renormalization*, Annals of Math. Studies no. 135, Princeton Univ. Press, 1994.
- [31] ———, *Renormalization and 3-manifolds which fiber over the circle*, Annals of Math. Studies no. 142, Princeton Univ. Press, 1996.
- [32] ———, *Iteration on Teichmüller space*, Invent. Math. **99** (1990), 425–454.
- [33] ———, *Rational maps and Kleinian groups*, Proceedings of the International Congress of Mathematicians 1990, Springer-Verlag, 1991, pp. 889–899.
- [34] ———, *Kleinian groups and John domains*, Topology **37** (1998), 485–496.
- [35] W. de Melo and S. van Strien, *One-dimensional dynamics*, Springer-Verlag, 1993.
- [36] J. Milnor, *Dynamics in One Complex Variable: Introductory Lectures*, SUNY Stony Brook IMS Preprint # 1990/5.
- [37] ———, *Local connectivity of Julia sets: Expository lectures*, Stony Brook IMS Preprint #1992/11.
- [38] G. D. Mostow, *Quasiconformal mappings in n -space and the rigidity of hyperbolic space forms*, Publ. I.H.E.S. **34** (1968), 53–104.
- [39] ———, *Strong rigidity of locally symmetric spaces*, Ann. of Math. Stud. no. 78, Princeton University Press, 1973.
- [40] J.F. Ritt, *Equivalent rational substitutions*, Trans. of the Amer. Math. Soc. **26** (1924), 221–229.
- [41] I. Satake, *On a generalization of the notion of a manifold*, Proc. Nat. Acad. Sci. U.S.A **42** (1956), 359–363.
- [42] P. Scott, *The geometries of 3-manifolds*, Bull. London Math. Soc. **15** (1983), 401–487.
- [43] M. Su, *Measured solenoidal Riemann surfaces and holomorphic dynamics*, J. Differential Geom. **47** (1997) 170–195.
- [44] D. Sullivan, *The density at infinity of a discrete group of hyperbolic motions*, Publ. I.H.E.S. **50** (1979), 419–450.
- [45] ———, *Entropy, hausdorff measures old and new, and limit sets of geometrically finite Kleinian groups*, Acta. Math. **153** (1984), 259–277.
- [46] ———, *Quasiconformal homeomorphisms and dynamics II: Structural stability implies hyperbolicity for Kleinian groups*, Acta Math. **155** (1985), 243–260.
- [47] ———, *Bounds, quadratic differentials, and renormalization conjectures*, Mathematics into the Twenty-first Century, Volume 2, Amer. Math. Soc., 1991.
- [48] ———, *Linking the universalities of Milnor-Thurston, Feigenbaum and Ahlfors-Bers*, Topological Methods in Modern Mathematics (L. Goldberg and A. Phillips, eds.), Publish or Perish, 1993.
- [49] W. Thurston, *On the combinatorics of iterated rational maps*, preprint 1982.
- [50] ———, *The geometry and topology of 3-manifolds*, Princeton University Lecture Notes, 1982.

- [51] ———, *Three-Dimensional Geometry and Topology*, Draft copy, University of Minnesota Geometry Center preprint, 1990.
- [52] M. Tsuji. Potential theory in modern function theory. Chelsea Publ. Co.
- [53] H. Wu, *Complex stable manifolds of holomorphic diffeomorphisms*. Indiana Univ. Math. J., 4 (1993), 1349-1358.

MATHEMATICS DEPARTMENT AND IMS, SUNY STONY BROOK