Big Mapping Class Groups and Complex Dynamics

Danny Calegari

University of Chicago, Chicago, Ill 60637 USA
E-mail address: `dannyc@math.uchicago.edu`

Preliminary version – July 31, 2019

Contents

Preface

Part 1

Theory

CHAPTER 1

The Shift

This section consists of the analysis of a very simple but fundamental family of complex dynamical systems that we call shifts. Although the domain of definition of a shift is a Riemann surface of infinite type (actually, a disk minus a Cantor set), every shift admits a canonical extension to a dynamical system in $\mathbb{D}$.

1.1. Annuli

Since the Riemann surfaces we consider are of infinite type, it's important to get some kind of control on their holomorpic geometry. One elementary method involves the moduli of annuli.

Let $0 \leq r < R \leq \infty$ and let $A(r, R)$ denote the annulus in $\mathbb{C}^*$ consisting of the set of z with $r < |z| < R$. An annulus is *round* if it is of the form $A(r, R)$. The *modulus* of $A(r, R)$ is

$$\mathcal{M}(A(r,R)) := \begin{cases} \log(R/r) \text{ if } 0 < r < R < \infty \\ \infty \text{ otherwise} \end{cases}$$

If A is any Riemann surface homeomorphic to an annulus, then it can be uniformized as $A(r, R)$ for some r, R and the modulus of A, denoted $\mathcal{M}(A)$, is by definition equal to the modulus of $A(r, R)$.

An annulus has infinite modulus if and only if at least one end is holomorphically equivalent to a puncture. Two annuli with finite modulus are holomorphically equivalent if and only if the moduli are the same.

The following Lemma is well-known:

LEMMA 1.1.1 (Reverse triangle inequality). *Let A be an annulus with $\mathcal{M}(A)$ finite, and suppose A_1, A_2 are disjoint annuli contained in A so that each inclusion $A_i \to A$ is a homotopy equivalence. Then $\mathcal{M}(A) \geq \mathcal{M}(A_1) + \mathcal{M}(A_2)$ with equality if and only if $\alpha := A - (A_1 - A_2)$ is a circle, and the uniformization of A takes α to a round circle.*

PROOF. This can be proved by extremal length. □

DEFINITION 1.1.2. Let U be a planar connected Riemann surface, and let $\alpha \subset U$ be a properly embedded arc. A holomorphic embedding $U \to A$ is *essential on* α if α is properly embedded in A and intersects the core of A transversely in one point. Define

$$\mathcal{M}(U, \alpha) := \inf_A \mathcal{M}(A)$$

where the infimum is taken over all essential holomorphic emdeddings $U \to A$.

DEFINITION 1.1.3 (Modulus of separation). Let Ω be a connected Riemann surface and let $X, Y \subset \Omega$ be disjoint compact contractible subsets of Ω. Let $[\alpha]$ be a homotopy class of proper paths in $\Omega - (X \cup Y)$ from X to Y. The *modulus of separation* of X and Y *along* $[\alpha]$, denoted $\mathcal{M}_\Omega(X, Y, [\alpha])$, is the supremum

$$\mathcal{M}_\Omega(X, Y, [\alpha]) := \sup_{(U,\alpha)} \mathcal{M}(U, \alpha)$$

over all *embedded* arcs α in the class of $[\alpha]$ properly contained in planar open sets $U \subset \Omega - (X \cup Y)$. We omit the subscript Ω if it is understood.

The modulus of separation is monotone with respect to inclusion:

LEMMA 1.1.4. *If $X \cup Y \cup \alpha \subset \Omega' \subset \Omega$ then*

$$\mathcal{M}_\Omega(X, Y, [\alpha]) \geq \mathcal{M}_{\Omega'}(X, Y, [\alpha])$$

PROOF. Every planar U in Ω' is also in Ω. □

COROLLARY 1.1.5. *If Ω is planar then*

$$\mathcal{M}_\Omega(X, Y, [\alpha]) = \mathcal{M}(\Omega - (X \cup Y), \alpha)$$

The following two lemmas give us techniques to estimate $\mathcal{M}(X, Y)$.

LEMMA 1.1.6 (disjoint annuli). *Let X and Y be disjoint compact contractible subsets of Ω, and suppose there is an embedded arc α from X to Y and a family of disjoint annuli A_i so that α intersects each A_i transversely in a single point. Then*

$$\mathcal{M}(X, Y, [\alpha]) \geq \sum_i \mathcal{M}(A_i)$$

PROOF. The union of the A_i with a neighborhood of α is a planar surface U in which α sits properly. If $(U, \alpha) \to A$ is a holomorphic embedding in an annulus, essential on α then the A_i are all disjoint essential annuli in A. By Lemma 1.1.1 we have $\mathcal{M}(A) \geq \sum \mathcal{M}(A_i)$ and therefore $\mathcal{M}(U, \alpha) \geq \sum \mathcal{M}(A_i)$. □

LEMMA 1.1.7 (two points finite). *Let X and Y be disjoint compact contractible subsets of Ω. If $\tilde{\Omega}$ is the universal cover of Ω, and $\tilde{\alpha}$ is a lift of α running between lifts $\tilde{X}$, $\tilde{Y}$ of X, Y then*

$$\mathcal{M}_\Omega(X, Y, [\alpha]) \leq \mathcal{M}_{\tilde{\Omega}}(\tilde{X}, \tilde{Y}, [\tilde{\alpha}])$$

In particular, if X and Y both have at least two points, $\mathcal{M}_\Omega(X, Y, [\alpha])$ is finite for any embedded α from X to Y.

PROOF. Let U be open and planar in $\Omega - (X \cup Y)$ containing α, and let $\tilde{U}$ be the cover of U in Ω containing $\tilde{\alpha}$. Then $\tilde{U}$ is planar, because $\tilde{\Omega}$ is either a plane or a sphere. Now, $\tilde{\Omega} - (\tilde{X} \cup \tilde{Y})$ is already planar, so we can embed it holomorphically in $\hat{\mathbb{C}}$, and then $\hat{\mathbb{C}} - (X \cup Y)$ is an annulus with finite modulus (because X and Y both have at least two points). □

REMARK 1.1.8. Every finite type Riemann surface has a finite dimensional moduli space, and Fenchel–Nielsen coordinates on its universal cover (i.e. Teichmüller space) are the analog of moduli of annuli as we've defined them. For arbitrary Ω and contractible subsets X, Y some proper homotopy classes $[\alpha]$ do not contain any embedded representative, and in fact there may be no representative that's contained in a planar subsurface of Ω.

Nevertheless, sometimes some representative is contained in a subsurface Σ of finite type, and one can define moduli of separation in terms of the uniformization of such subsurfaces.

However, in every case, the noncompactness of these finite dimensional moduli spaces arises from the existence of embedded annuli of bigger and bigger moduli, and therefore it can be detected by the more naive tools we've introduced here.

1.2. Pants, Slots, Shifts

We begin with the definitions of pants and slots.

DEFINITION 1.2.1 (Pants). A *k-pants* P is a (marked) disk with k holes.The boundary of the disk is called the *waist*, and the other k boundary components are called the *cuffs*. We denote the waist by $\partial_0 P$ or just ∂_0 if P is understood, and the jth cuff by $\partial_j P$ or ∂_j. A k-pants is *analytic* if it has the structure of a Riemann surface with real analytic boundary.

A 2-pants is colloquially known as a *pair of pants*. One doesn't want to think too hard about this analogy when $k > 2$, but if one must, let's call them *monkey pants*.

DEFINITION 1.2.2 (Slot). For $k \geq 1$ a *k-slot* is the data of an analytic k-pants P, and for $1 \leq j \leq k$ real analytic orientation-reversing diffeomorphisms $f_j : \partial_0 \to \partial_j$ called *insertions*.

A slot is like a kind of operad. We can build a Riemann surface from copies of P by gluing new copies along the cuffs of the old by the insertions f_j. We denote the set of f_j by the letter f where this does not cause ambiguity. Let's formalize this.

CONSTRUCTION 1.2.3 (Slot to shift). Given a k-slot (P, f) we can build a Riemann surface Ω as follows. Let Γ be the free semigroup on k fixed generators, and let's fix an identification between elements of Γ and finite strings γ on the alphabet $\{f_1, \cdots, f_k\}$. We take the product $P \times \Gamma$ and glue each $P \times \gamma$ to the k surfaces $P \times \gamma f_j$ by the maps f_j. Formally, Γ acts by right multiplication on the product $P \times \Gamma$ and we let Ω be the quotient under the identification $(x, \gamma f_j) \sim (f_j(x), \gamma)$ for $x \in \partial_0$ and $\gamma \in \Gamma$.

Since P is analytic and the f_j are real-analytic diffeomorphisms, Ω admits the natural structure of a Riemann surface. Furthermore, Ω is homeomorphic to a disk minus a Cantor set. Let $\mathcal{E} := \{1, \cdots, k\}^{\mathbb{N}}$. This is a Cantor set, which may be canonically identified with the space of ends of Γ, or equivalently with the space of ends of Ω.

Finally, the *left* action of Γ on itself gives an action on $P \times \Gamma$ which descends to an action on Ω. By abuse of notation we denote the generators by $f_j : \Omega \to \Omega$ and observe that they are holomorphic injections with disjoint image.

We call the holomorphic dynamical system (Ω, f) the *shift* obtained from the slot (P, f), and conversely we call (P, f) the *fundamental domain* for (Ω, f).

Which holomorphic dynamical systems (Ω, f_j) arise as shifts obtained from slots? The answer is given in the next lemma.

LEMMA 1.2.4 (Fundamental domain is slot). *Let Ω' be a connected Riemann surface with boundary ∂_0. Let $f_j : \Omega' \to \Omega'$ be holomorphic injections with disjoint image. Let P denote the closure of $\Omega' - \cup_j f_j(\Omega')$. Suppose P is a k-pants. Then the restriction (P, f) is a slot, and if $\Omega \subset \Omega'$ is the shift obtained from P, then either*

(1) $\Omega' = \Omega$; *or*
(2) Ω' *is a disk identified in a natural way with the (topological) end-completion of* Ω. *In other words,* $\Omega' = \Omega \cup \mathcal{E}$.

Proof. The f_j generate an action of a semigroup Γ. We claim first of all that this action is free. Equivalently, the translates of P are interior disjoint. By induction, for each $\gamma \in \Gamma$, the translates $f_j\gamma(\Omega)$ are all disjoint and contained in $\gamma(\Omega)$ and the complement of their union in $\gamma(\Omega)$ is $\gamma(P)$. This proves the claim.

Next we define $\Omega := \cup_\gamma \gamma(P) \subset \Omega'$. We claim second of all that $\Omega' - \Omega$ is totally disconnected.

Suppose not. Let Y be a small round ball in $\Omega' - \Omega$ and let X be a small round ball in P near ∂_0. Let α be an embedded arc from X to Y. By Lemma 1.1.7 we have $\mathcal{M}(X, Y, [\alpha]) < \infty$.

On the other hand, every end of Ω corresponds to an infinite sequence $\gamma_i \in \Gamma$. Any arc α from X to Y must exit some end of Ω. If A is an annulus in P separating ∂_0 from every $f_j\partial_0$ then any arc in the homotopy class of α from X to Y must intersect all the translates $\gamma_i(A)$. We can choose α embedded that intersects the core of every $\gamma_i(A)$ transversely in exactly one point. Then by Lemma 1.1.6

$$\mathcal{M}(X, Y, [\alpha]) \geq \sum_i \mathcal{M}(\gamma_i(A)) = \sum_i \mathcal{M}(A) = \infty$$

This contradiction proves the claim.

It follows that $\Omega' - \Omega$ is totally disconnected, and may therefore be identified with a subset of $\mathcal{E}$. But the action of Γ on $\mathcal{E}$ is minimal, so $\Omega' - \Omega$ is either empty or equal to $\mathcal{E}$. □

1.3. Skinning and Realization

Let (Ω, f) be the k-shift obtained from a k-slot (P, f). For any proper extension (Ω', f) with $\partial\Omega' = \partial\Omega$, Lemma 1.2.4 says that Ω' is a disk, homeomorphic to the end completion of Ω.

Now, the end completion of Ω always exists as a topological space, but it is not a priori clear that it can be given the structure of a Riemann surface compatibly with the action of the f_j. However, it turns out that such an Ω' always exists and is unique.

Let $\mathbb{D}$ denote the unit disk in $\mathbb{C}$.

Definition 1.3.1 (Realization). A *realization* of a k-shift (Ω, f) is a holomorphic embedding $\psi : \Omega \to \mathbb{D}$ with complement a Cantor set Λ so that the f_j extend to holomorphic injections from $\mathbb{D}$ to $\mathbb{D}$.

By abuse of notation we will denote the extensions of f_j to $\mathbb{D}$ also by f_j.

A realization of a k-shift is a dynamical system $(\mathbb{D}, f)$, and the Cantor set Λ is the limit set (equivalently: the *attractor*, or the *Julia set*) of the dynamical system.

Theorem 1.3.2 (Realization Exists). *Let* (Ω, f) *be a* k*-shift. Then there is a realization* $\psi : \Omega \to \mathbb{D}$, *unique up to composition with a Möbius transformation of the disk.*

Proof. Existence is proved by skinning. Uniqueness is proved the same way.

Let (P,f) be the fundamental k-slot for (Ω,f). Choose once and for all three points on $\partial_0 P$ in cyclic order. Let $\mathcal{T}$ denote the set of holomorphic embeddings $P\to\mathbb{D}$ sending ∂_0 to S^1 and the three marked points to the third roots of unity. Define a (Teichmüller) distance on $\mathcal{T}$ as follows. Given two maps $\psi_1 : P\to\mathbb{D}$, $\psi_2 : P\to\mathbb{D}$ then

$$d_{\mathcal{T}}(\psi_1,\psi_2)=\frac{1}{2}\inf_\phi \log(\text{dilatation of }\phi)$$

where the infimum is taken over all quasiconformal maps ϕ from $\mathbb{D}$ to $\mathbb{D}$ such that $\phi\psi_1$ and ψ_2 agree on S^1.

Define the *skinning map* $\sigma:\mathcal{T}\to\mathcal{T}$ as follows. Given ψ we define $\sigma(\psi)$ by cutting out the disks of $\mathbb{D}-\psi(P)$, gluing back k new copies of $\mathbb{D}$ by the insertions f_j, and uniformizing the result. By construction the map $\sigma(\psi):P\to\mathbb{D}$ extends holomorphically to $P\cup_j f_j(P)$, and we can push forward f_j to holomorphic maps from $\psi(P)$ to $\mathbb{D}$. By abuse of notation we denote these pushforwards also by f_j. Notice further that for any word τ of length at most n in the f_j the map τ pushes forward to $\tau:\sigma^n(\psi)(P)\to\mathbb{D}$.

We prove the following two claims:

(1) σ is strictly distance decreasing; and
(2) every orbit of σ is bounded.

Together these claims imply that σ is uniformly distance decreasing on every orbit, and therefore has a unique fixed point Ψ, and for any other ψ the orbit $\sigma^n(\psi)$ must converge to Ψ at a geometric rate.

The first claim is standard. If $\phi:\mathbb{D}\to\mathbb{D}$ is extremal for ψ_1, ψ_2 then we may obtain $\sigma(\phi):\mathbb{D}\to\mathbb{D}$ by extending the identity map on P by ϕ on each of k disks glued to the cuffs by the f_j. Then $\sigma(\phi)$ and ϕ have the same dilatation, but the dilatation of $\sigma(\phi)$ may be reduced slightly by relaxation.[1]

We now prove the second claim. The f_j restrict to maps from the waist to the cuffs of P. Pushing these maps forward by ψ gives maps f_j from S^1 to the images of the cuffs by ψ. We construct $F_j:\mathbb{D}\to\mathbb{D}$ with the following properties:

(1) F_j agrees with f_j on S^1;
(2) the maps F_j are quasiconformal;
(3) there is a uniform K for which every element of the semigroup generated by the F_js is K-quasiconformal.

Of course there is no difficulty in finding F_j satisfying the first two bullets; all the subtlety is in the third bullet, that we can uniformly control the quasiconformality of any finite composition.

Let D_j be the regions in $\mathbb{D}$ bounded by $\psi(\text{cuffs})$, and inductively for each word τ in Γ let $D_{\tau j}$ be the subdisks of D_τ equal to the images of D_j under τ as defined so far. Now, for all words τ of length 2 let D_τ be a round disk, and let F_j restricted to D_τ be linear, i.e. of the form $z\to\alpha z+\beta$ for some $|\alpha|<1$ and β satisfying the topological requirements. Evidently any composition in F_j is the composition of a word of length at most 2 with a linear map, and therefore these are all uniformly quasiconformal.

[1] isometries of universal Teichmüller space

Let L be the limit set of the F_j in $\mathbb{D}$ and let Δ be the complement. Then $\psi : P \to \mathbb{D}$ extends uniquely to a map $\psi : \Omega \to \Delta$ conjugating f_j to F_j, and by construction ψ as above is K-quasiconformal.

When we apply the skinning map to replace ψ by $\sigma(\psi)$ we can construct new maps $F_j^1 : \mathbb{D} \to \mathbb{D}$ as above by defining F_j^1 to be equal to f_j on P, and to be equal to the conjugate of F_j by f on $\mathbb{D} - \sigma(\psi)(P)$. The maps F_j^1 have a new limit set L^1 with complement Δ^1 and get a new $\psi^1 : \Omega \to \Delta^1$ also K-quasiconformal, conjugating f_j to F_j^1.

Continue by induction to define F_j^n, L^n, Δ^n. Because each F_j^n is actually linear on a neighborhood of L^n, these limit sets all have Hausdorff dimension strictly less than 2 and therefore Lesbesgue measure 0. Consequently the maps $\psi^n(\psi^m)^{-1}$ extend uniquely to maps from $\mathbb{D}$ to $\mathbb{D}$ which are $2K$-quasiconformal. Furthermore, these maps take $\sigma^m(\psi)$ to $\sigma^n(\psi)$. In particular, the $\sigma^n(\psi)$ all lie in a compact subset of $\mathcal{T}$ and therefore σ is a strict contraction on the orbit of ψ. This proves the claim and the theorem. $\square$

1.4. Multipliers

1.4.1. Multipliers. Let (Ω, f) be a k-shift. Fix some basepoint $p \in \partial$. By Theorem 1.3.2 there is a canonical extension to (D, f) where D is a disk. For each j the map $f_j : D \to D$ has a unique fixed point p_j, and we denote $\lambda_j := f_j'(p_j)$. The quotient of $D - p_j$ by f_j is an annulus, whose universal cover can be uniformized as $\mathbb{C}^*$. This determines for each j a unique holomorphic embedding $e_j : \Omega \to D \to \hat{\mathbb{C}}$ taking p_j to 0 and p to 1, and conjugating f_j to the map $z \to \lambda_j z$. We call $\lambda := \{\lambda_j\}$ the (set of) *multipliers* of the k-shift. Note that $0 < |\lambda_j| < 1$.

1.4.2. Dynamical cocycles.

1.4.3. Multiplier inequalities. The Schwarz Lemma says if $f : \mathbb{D} \to \mathbb{D}$ is any holomorphic map, then at any interior fixed point p we have $|f'(p)| \le 1$ with equality if and only if f is an isomorphism. If f is a holomorphic injection, the Koebe $1/4$ Theorem (due to Bieberbach!) says that if we change coordinates so that $p = 0$ then the image $f(\mathbb{D})$ contains the disk of Euclidean radius at least $|f'(0)|/4$.

If $f_j : \mathbb{D} \to \mathbb{D}$ is a *collection* of holomorphic maps with disjoint images and (unique) fixed points p_j then each $|f_j'(p_j)| < 1$ but otherwise these numbers can be arbitrary. On the other hand, the functions $|f_j'|$ are subject to several norm inequalities that we call *multiplier inequalities* which can be thought of as a generalization of the Schwarz Lemma.

First of all, by disjointness of images we have $\int_{\mathbb{D}} \sum_j |f_j'|^2 d\text{area} < 1$. More subtly, if r is the Hausdorff dimension of the limit set Λ then $\int_\Lambda \sum_j |f_j'|^r d\mathcal{H}_r \le 1$, where $\mathcal{H}_r$ denotes r-dimensional Hausdorff measure.

1.5. Multishifts and Splitting

1.5.1. Multishifts.

Definition 1.5.1 (Multislot). A *multislot* is a real analytic surface P which is a finite disjoint union of pants (called a *multipants*) and a collection of real analytic orientation-reversing diffeomorphisms f from the waists of P to the cuffs of P (called *insertions*) so that every waist is the domain of at least one insertion, and every cuff is the range of exactly one insertion.

The insertions f generate a free semigroupoid Γ, where two insertions f_1 f_2 are composable if the range of f_2 is in the pants containing the domain of f_1. Parallel to Construction 1.2.3 we can build a Riemann surface Ω by taking the product $P \times \Gamma$ and gluing $P \times \gamma$ to $P \times \gamma f_j$ by $(x, \gamma f_j) \sim (f_j(x), \gamma)$ when this makes sense. The Riemann surface Ω is homeomorphic to a finite union of disks minus Cantor sets, one for each component of P. The left action of Γ on itself gives an action on Ω, which we denote (Ω, f) and call the result the *multishift* obtained from (P, f).

The Realization Theorem 1.3.2 holds for multishifts, and with the same proof via skinning.

1.5.2. Branched surfaces. It is convenient to describe Multislots in the language of *branched surfaces*. A branched surface is a certain kind of 2-complex assembled by gluing surfaces (*sheets*) together along graphs (the *branch locus*) together with a choice of co-orientation along the branch locus so that it makes sense to distinguish between the sheets that are incoming and the sheets that are outgoing on either side. For an introduction to the theory of branched surfaces, especially as used in 3-manifold topology, see e.g. [**20**] or [**11**].

The branched surfaces that arise in this theory are very special in the context of the general theory.

Definition 1.5.2 (Shift-like branched surface). A compact branched surface B whose sheets are Riemann surfaces with real analytic branch locus is *shift-like* if it satisfies the following properties:

(1) the branch locus is an embedded collection of cirles;
(2) each sheet is a planar surface with exactly one boundary component where the co-orientation points inwards, and at least two boundary components where the co-orientation points outwards;
(3) at each branch circle there is exactly one outgoing sheet and at least two incoming sheets.

Given a multislot (P, f) we glue waists to cuffs by the maps f. The result is a shift-like branched surface B. Conversely given a shift-like branched surface B we can cut it open along the branch locus into its constituent sheets P. These sheets are pants, and the co-orientation determines which is the waist. Remembering how these sheets were glued together in B determines maps f giving (P, f) the structure of a multislot. Thus there is a precise correspondence between multislots and shift-like branched surfaces.

Lemma 1.5.3. *A branched surface B corresponds to a slot if and only if it has exactly one branch circle.*

Proof. There is a bijection between branch circles of B and components of P. □

1.5.3. Splitting. Since branched surfaces and multislots are equivalent, the advantage of working with one over the other is more about psychology than mathematics.

Construction 1.5.4 (Elementary Splitting). Let B be a shift-like branched surface, and let γ be a component of the branch locus. Let S be the sheet on the outgoing side of γ and let $R \subset S$ be a subsurface with $\partial R \cap$ branch locus $= \gamma$; in other words, the surface R is not allowed to have other boundary components in the branch locus.

The operation of *elementary splitting* B along R is a new branched surface B' obtained as follows. Suppose there are $k > 1$ incoming sheets S_i along γ. Cut R out of B. Take k disjoint copies of R and glue S_i to the ith copy of R along γ. Then glue every other boundary component of every copy of R to wherever it was originally glued in B.

For each subset of the S_i with cardinality 1 we get a copy of γ with exactly one incoming and one outgoing sheet, so we can erase this copy from the branch locus. On the other hand, any component of ∂R that was not in the branch locus in B becomes part of the branch locus in B'. The branched surface B' is shift-like.

We denote this operation by $B \rightsquigarrow_R B'$.

If there is a sequence of elementary splittings

$$B_0 \rightsquigarrow_{R_0} B_1 \rightsquigarrow_{R_1} B_2 \rightsquigarrow_{R_2} \cdots \rightsquigarrow_{R_{n-1}} B_n$$

then we write $B_0 \rightsquigarrow B_n$ and say B_n is obtained from B_0 by *splitting.*

Each sheet S of B is a pants, and each subsurface R of S is a subpants, which is to say a sphere with at least two boundary components. If R has more than 3 boundary components we can write it as a union of thrice-punctured pants R_i so that the result of splitting along R is the composition of splitting along R_i.

1.5.4. Splitting and Embedding.

Lemma 1.5.5 (Splittings give embeddings). *Suppose B and B' are branched surfaces associated to (Ω, f) and (Ω', f') where $B \rightsquigarrow B'$. Then there is an embedding $\Omega' \to \Omega$ so that the restriction of f to Ω' is f'.*

Proof. By induction we can assume $B \rightsquigarrow_R B'$. Let (P, f) be the multipants associated to B, and let (Ω, f) be the associated multishift. Then $R \subset P \to \Omega - f(\Omega)$ and we can set $\Omega' = \Omega - R$. □

1.5.5. Building Multishifts. Let f be a collection of holomorphic embeddings $f_j : \mathbb{D} \to \mathbb{D}$ such that the closure of $f_j(\mathbb{D})$ is contained in $\mathbb{D}$. Each f_j is strictly contracting for the hyperbolic metric, so there is some $\epsilon > 0$ so that each f_j is $(1-\epsilon)$-Lipschitz. Let $\mathcal{E}$ denote the set of right infinite words in the f_j; we may topologize this as a Cantor set. Because of the strict contraction there is a continuous map $\phi : \mathcal{E} \to \mathbb{D}$, and left multiplication by f_j on $\mathcal{E}$ pushes forward to the action of f_j on $\mathbb{D}$.

Proposition 1.5.6 (Building a multishift). *Let $(\mathbb{D}, f)$ be as above, and suppose $\mathcal{E} \to \mathbb{D}$ is an embedding. Denote the image by Λ. Then there is a finite union of closed disks $K \subset \mathbb{D}$ containing Λ in the interior and each intersecting Λ, so that each $f_j(K)$ is contained in the interior of K and $f_j(K) \cap f_i(K)$ is empty when $i \neq j$. In particular, $(K - f(K), f)$ is a multislot with associated multishift $(K - \Lambda, f)$.*

Furthermore, any two K, L with this property give rise to multislots with a common splitting.

PROOF. For any $t > 0$ let K_t denote the closed t-neighborhood of Λ in the hyperbolic metric. Since $f_j(\Lambda) \subset \Lambda$, and each f_j is $(1-\epsilon)$-Lipshitz it follows that each $f_j(K_t)$ is contained in the interior of K_t.

We claim there is a positive t so that $f_i(K_t) \cap f_j(K_t)$ is empty for all $i \neq j$. For, if not, there's a fixed i, j and a sequence of pairs of points $p_t, q_t \in \Lambda$ with $d(f_i(p_t), f_j(q_t)) \leq 2t$. By compactness there's $p, q \in \Lambda$ with $f_i(p) = f_j(q)$, but this contradicts the fact that $\mathcal{E} \to \Lambda$ is an embedding.

Now, each component of K_t has diameter at least $2t$, so there are only finitely many components. Let U be a hole of biggest diameter in some component J of K_t. Since each $f_j(U)$ has strictly smaller diameter, and $f_j(K_t) \subset K_t$ it follows that $f_j(U)$ is disjoint from U.

Suppose $f_i(U)$ and $f_j(U)$ intersect for some i, j. Since $f_i(K_t)$ is disjoint from $f_j(K_t)$, without loss of generality we must have $f_i(J) \subset f_j(U)$; in other words, some component of some $f_i(K_t)$ is contained in some 'hole' of some $f_j(K_t)$. Since K_t has only finitely many components, there are only finitely many holes U that arise in this way. Since Λ is totally disonnected, there is $s < t$ for which every such U is contained in the 'big' component of $\mathbb{D} - K_s$. For such an s it follows that for the hole V in a component of K_s of largest diameter, the hole V and its images $f_j(V)$ are all disjoint. So we can fill in the hole V, and inductively all the smaller ones. At the end of this procedure we get K a finite union of disks with each $f_j(K)$ strictly contained in the interior of K, and different $f_j(K)$ disjoint. This proves the first part of the proposition.

For any two such subsets K, L as above their intersection $K \cap L$ is also a union of disks. Inductively throw away components of the intersection that don't contain any forward images. The result is obtained from both K and L by restriction, and the associated branched surfaces have a common splitting, by Lemma 1.5.5. $\square$

1.5.6. Extensions.

DEFINITION 1.5.7. Let (Ω, f) be a k-shift. An *extension* is an inclusion of (connected) Riemann surfaces $\Omega \to \Delta$ together with holomorphic maps $g_j : \Delta \to \Delta$ extending f_j.

We do *not* insist that (Δ, g) is a shift. The g_j don't need either to be injective nor to have disjoint image.

DEFINITION 1.5.8. Let $\Omega \to \Delta$ be an extension. The *precritical set* $K \subset \Delta$ is the set of k for which $v'(k) = 0$ for some word v in the g_j.

LEMMA 1.5.9 (Precritical limit is wandering). *Let $\Omega \to \Delta$ be an extension, and let $K \subset \Delta$ denote the precritical set. Let $\bar{K}$ and $\bar{\Omega}$ denote the closures of K and of Ω in Δ. Let k be a nontrivial limit point in $\bar{K}$. Then there is some infinite forward trajectory of k that never enters $\bar{\Omega}$.*

PROOF. Let $k_i \in K$ converge to k. For each i there is v_i with $v_i'(k_i) = 0$. Evidently, $w_i(k_i)$ is not in $\bar{\Omega}$ for any suffix w_i of v_i. Thus there is a sequence of longer and longer words w_i for which $w_i(k)$ is not in $\bar{\Omega}$. Now choose a subsequence of the w_i for which each word is a suffix of the next. $\square$

If $\Omega \to \mathbb{D}$ is the canonical realization provided by Theorem 1.3.2, then for any extension $\Omega \to \Delta$ there is a further (canonical) extension that factors through the realization $\Omega \to \mathbb{D}$. We call such an extension *real*.

LEMMA 1.5.10. *Let $\Omega \to \mathbb{D} \to \Delta$ be any real extension, and let $\pi : \tilde{\Delta} \to \Delta$ be the universal cover. Then there is a real extension $\Omega \to \mathbb{D} \to \tilde{\Delta}$, unique up to isomorphism, lifting the original extension.*

PROOF. For any continuous function $h : \Delta \to \Delta$ the composition $\tilde{\Delta} \to \Delta \to \Delta$ lifts to $\tilde{h} : \tilde{\Delta} \to \tilde{\Delta}$, uniquely up to composition with a deck transformation. Since $\mathbb{D}$ is simply-connected, it admits a lift to $\tilde{\Delta}$, and for any such lift, there is a unique lift of the maps g_j on Δ extending f_j that take the lift of $\mathbb{D}$ into itself. Thus the lift defines a real extension $\Omega \to \mathbb{D} \to \tilde{\Delta}$, unique up to composition with a deck transformation. □

Let Δ be a simply-connected real extension. There are three possibilities for Δ up to isomorphism: $\hat{\mathbb{C}}$, $\mathbb{C}$ or $\mathbb{H}$. We call the extension *spherical*, *Euclidean* and *hyperbolic* respectively. The case of a spherical or Euclidean extension require a special analysis and we'll return to them in § 1.5.7 and § 1.5.8, but for now let's restrict attention to hyperbolic extensions.

LEMMA 1.5.11 (Proper in the plane). *Let $\Omega \to \mathbb{D} \to \mathbb{H}$ be a hyperbolic extension, i.e. a simply-connected real extension for which $\mathbb{H}$ is the hyperbolic plane. Then for every $p \in \mathbb{H}$ there is a d so that $v(p) \in \mathbb{D}$ for all words v in the g_j of length at least d.*

PROOF. Every g_j has a fixed point in $\mathbb{D}$ where the derivative has norm strictly less than 1. By the Schwarz Lemma, every g_j is strictly length decreasing for the hyperbolic metric on $\mathbb{H}$. Thus, for every t, the g_j are uniformly contracting on the subset of points within hyperbolic distance t of Λ; the proof follows. □

COROLLARY 1.5.12 (Critical discrete preimages). *Let $\Omega \to \mathbb{D} \to \mathbb{H}$ be a hyperbolic extension. Then the precritical set K is discrete.*

CONSTRUCTION 1.5.13 (Extension by immersions). Let $\Omega \to \mathbb{D} \to \mathbb{H}$ be a hyperbolic extension. The precritical set K is discrete, is disjoint from $\mathbb{D}$, and is backward invariant by definition. Therefore the difference $\mathbb{H} - K$ is also an extension, and so is its universal cover. This universal cover of $\mathbb{H} - K$ is a hyperbolic extension with the additional property that the maps g_j are all holomorphic *immersions*.

DEFINITION 1.5.14. A hyperbolic extension $\Omega \to \mathbb{D} \to \mathbb{H}$ is of *immersion type* if the maps g_j extending the f_j are holomorphic immersions, and for any $p \in \Delta$ there is a d so that $v(p) \in \mathbb{D}$ for every word v of length at least d.

Construction 1.5.13 shows how to replace any hyperbolic extension by a canonical hyperbolic extension of immersion type.

1.5.7. Spherical extensions.

1.5.8. Euclidean extensions.

1.6. Topology of Teichmüller space

1.7. Moduli and MCG

CHAPTER 2

MCG of the plane minus a Cantor set

The most important and best studied example of a big mapping class group is the mapping class group of the plane minus a Cantor set. We denote this group $\Gamma_{\mathcal{C}}$, or sometimes just Γ for brevity.

Notation 2.0.1. If S is a surface, and A, B are subsets, we denote by $\mathrm{Homeo}^+(S, A; B)$ the subgroup of the group of orientation-preserving homeomorphisms of S that fix A setwise, and fix B pointwise. If B or A is empty, we write $\mathrm{Homeo}^+(S, A)$ or $\mathrm{Homeo}^+(S; B)$ respectively. These are all groups with the compact-open topology.

Thus, $\mathrm{Homeo}^+(\mathbb{C}, \mathcal{C})$ the group of orientation-preserving homeomorphisms of the plane fixing $\mathcal{C}$ setwise. We define

$$\Gamma_{\mathcal{C}} := \pi_0(\mathrm{Homeo}^+(\mathbb{C}, \mathcal{C}))$$

Equivalently, $\Gamma_{\mathcal{C}}$ is the quotient of $\mathrm{Homeo}^+(\mathbb{C}, \mathcal{C})$ by the normal subgroup of homeomorphisms of the plane isotopic to the identity rel. $\mathcal{C}$.

In this chapter we begin the systematic algebraic and geometric study of $\Gamma_{\mathcal{C}}$.

2.1. Embeddings of a Cantor set

Let $\mathcal{C}$ denote a Cantor set. For now we could think of this as an abstract topological space, e.g. as the infinite product $\{L, R\}^{\mathbb{N}}$, or as a concrete Cantor set, e.g. the middle-third Cantor set in $[0, 1]$. Each $\tau \in \mathcal{C}$ determines an element of the middle-third Cantor set as follows. Think of τ as an infinite word in L and R. Then let $\nu(\tau)$ be obtained from τ by replacing each L by 0 and each R by 2. The string $\cdot\nu(\tau)$ can then be read as the base three expansion of a number in $[0, 1]$.

Proposition 2.1.1 (Embeddings of $\mathcal{C}$). *For any connected oriented surface S, let $\mathrm{Emb}(\mathcal{C}, S)$ denote the space of embeddings $\mathcal{C} \to S$ with the compact-open topology, and for any n let $\mathrm{Emb}(n, S)$ denote the space of n ordered distinct points in S. If we pick three distinct points in $\mathbb{C}$, the forgetful map $\mathrm{Emb}(\mathcal{C}, S) \to \mathrm{Emb}(3, S)$ is a Serre fibration whose homotopy fiber has vanishing π_k for $k \neq 1$.*

Proof. Fix in advance a conformal structure on S.

Let X be a countable dense subset of $\mathcal{C}$, and for any integer n let X_n denote the first n elements of X in some enumeration. For any n there's a forgetful projection from $\mathrm{Emb}(\mathcal{C}, S)$ to the configuration space $\mathrm{Emb}(n, S)$ of n marked distinct points in S, and these projections are compatible with composition:

$$\mathrm{Emb}(\mathcal{C}, S) \to \mathrm{Emb}(n + 1, S) \to \mathrm{Emb}(n, S)$$

Note that for every n, the map $\mathrm{Emb}(n+1, S) \to \mathrm{Emb}(n, S)$ is an honest fibration whose fibers are copies of S minus n points. Likewise, for every $m > n$ the map $\mathrm{Emb}(m, S) \to$

$\mathrm{Emb}(n, S)$ is a fibration. Furthermore, when n is at least 3, every fiber admits a canonical hyperbolic metric coming from the conformal structure on S.

Now, any map $\alpha : S^k \to \mathrm{Emb}(\mathcal{C}, S)$ determines a family of maps $\alpha_n : S^k \to \mathrm{Emb}(n, S)$, and each α_{n+1} is a lift of α_n.

We show that the homotopy fiber has vanishing π_k for $k \neq 1$; a similar argument gives homtopy lifting for disks, and proves the Serre fibration property.

Let $k > 1$ and suppose for some $n \geq 3$ that $\alpha_n : S^k \to \mathrm{Emb}(n, S)$ extends to $\bar{\alpha}_n : D^{k+1} \to \mathrm{Emb}(n, S)$. Using the radial structure on D^{k+1} and the map $\bar{\alpha}_n$ we can canonically trivialize the pullback of the bundle over S^k and metrize it with the hyperbolic metric on the fiber over the origin. In this trivialization a lift $\alpha_{n+1} : S^k \to \mathrm{Emb}(n+1, S)$ is the same as a map from S^k to S minus n points; this fills in canonically (e.g. by heat flow) to a map from D^{k+1} to S minus n points (i.e. a lift $\bar{\alpha}_{n+1}$ of $\bar{\alpha}_n$).

In the compact open topology, finite sets of points that are sufficiently close in $\mathcal{C}$ stay close in compact families of embeddings. Heat flow of maps of spheres to negatively curved spaces are uniformly strictly energy decreasing. It follows that this canonical sequence of iterated lifts of homotopies converges to a map $\bar{\alpha} : D^{k+1} \to \mathrm{Emb}(\mathcal{C}, S)$. □

The group $\mathrm{Homeo}(\mathcal{C})$ of self-homeomorphisms of $\mathcal{C}$ acts freely on $\mathrm{Emb}(\mathcal{C}, S)$ by precomposition. Let $\mathrm{UEmb}(\mathcal{C}, S)$ denote the quotient. Then we have a fibration

$$\mathrm{Homeo}(\mathcal{C}) \to \mathrm{Emb}(\mathcal{C}, S) \to \mathrm{UEmb}(\mathcal{C}, S)$$

where $\mathrm{Homeo}(\mathcal{C})$ is totally disconnected.

It's reasonable to think of $\mathrm{Emb}(\mathcal{C}, S)$ and $\mathrm{UEmb}(\mathcal{C}, S)$ as marked and unmarked configuration spaces respectively, although one should be careful, as the following remark clarifies.

Remark 2.1.2. It is not true that $\mathrm{UEmb}(\mathcal{C}, S)$ is homeomorphic to the space of (unparameterized) Cantor sets in S with the Hausdorff metric. That's because a family of injective maps from $\mathcal{C}$ to S can limit to a non-injective map whose image is nevertheless homeomorphic to a Cantor set.

Now let's fix a specific Cantor set $\mathcal{C}$ in S, for instance, the middle-third Cantor set in the plane. Let $\mathrm{Homeo}^+(S)$ denote the group of self-homeomorphisms of S, and let $\mathrm{Homeo}(S; \mathcal{C})$ denote the subgroup fixing $\mathcal{C}$ *pointwise.*

Proposition 2.1.3 (Fibration). *For any connected oriented surface S there are fibrations*

$$\mathrm{Homeo}^+(S; \mathcal{C}) \to \mathrm{Homeo}^+(S) \to \mathrm{Emb}(\mathcal{C}, S)$$

and

$$\mathrm{Homeo}^+(S, \mathcal{C}) \to \mathrm{Homeo}^+(S) \to \mathrm{UEmb}(\mathcal{C}, S)$$

Proof. The group $\mathrm{Homeo}^+(S)$ acts on $\mathrm{Emb}(\mathcal{C}, S)$ by composition, and the kernel is evidently $\mathrm{Homeo}^+(S; \mathcal{C})$. The content of the proposition is that the action on $\mathrm{Emb}(\mathcal{C}, S)$ is transitive. This is the combination of two well-known facts: that any embedding of a Cantor set in a connected surface is tame (i.e. it is contained in a tamely embedded interval), and that every homeomorphism of a tame Cantor set extends to a homeomorphism of an open neighborhood.

The same argument applies to the action of $\mathrm{Homeo}^+(S)$ on $\mathrm{UEmb}(\mathcal{C}, S)$. □

REMARK 2.1.4. There are compactly-supported variants

$$\mathrm{Homeo}_c(S;\mathcal{C}) \to \mathrm{Homeo}_c(S) \to \mathrm{Emb}(\mathcal{C},S)$$

and

$$\mathrm{Homeo}_c(S,\mathcal{C}) \to \mathrm{Homeo}_c(S) \to \mathrm{UEmb}(\mathcal{C},S)$$

of Proposition 2.1.3, proved in the same way.

PROPOSITION 2.1.5 (Contractible components). *Let S be any connected oriented surface. Then the identity component of* $\mathrm{Homeo}^+(S;\mathcal{C})$ *(equivalently, the identity component of* $\mathrm{Homeo}^+(S;\mathcal{C})$*) is contractible.*

PROOF. First note that the identity component of $\mathrm{Homeo}^+(S,\mathcal{C})$ is evidently equal to the identity component of $\mathrm{Homeo}^+(S;\mathcal{C})$, and both are equal to the identity component of $\mathrm{Homeo}^+(S-\mathcal{C})$. This is a surface of hyperbolic type, and the usual proofs of the contractibility of its identity component hold. □

The long exact sequence of homotopy groups for the fibration in Proposition 2.1.3 thus breaks up into isomorphisms $\pi_k(\mathrm{Homeo}^+(S)) = \pi_k(\mathrm{Emb}(\mathcal{C},S))$ for $k>2$, and a 4-term sequence

$$0 \to \pi_1(\mathrm{Homeo}^+(S)) \to \pi_1(\mathrm{Emb}(\mathcal{C},S)) \to \pi_0(\mathrm{Homeo}^+(S;\mathcal{C})) \to \mathrm{MCG}(S) \to 0$$

where we use the identification $\pi_0(\mathrm{Homeo}^+(S))$ is equal to the mapping class group $\mathrm{MCG}(S)$ by definition, and the fact that $\pi_0(\mathrm{Emb}(\mathcal{C},S)) = 0$.

Likewise we have

$$0 \to \pi_1(\mathrm{Homeo}^+(S)) \to \pi_1(\mathrm{UEmb}(\mathcal{C},S)) \to \pi_0(\mathrm{Homeo}^+(S,\mathcal{C})) \to \mathrm{MCG}(S) \to 0$$

and the short exact sequence

$$0 \to \pi_0(\mathrm{Homeo}^+(S;\mathcal{C})) \to \pi_0(\mathrm{Homeo}^+(S,\mathcal{C})) \to \mathrm{Homeo}(\mathcal{C}) \to 0$$

Taking $S=\mathbb{C}$ we obtain a central extension

$$0 \to \mathbb{Z} \to \pi_1(\mathrm{UEmb}(\mathcal{C},\mathbb{C})) \to \Gamma_{\mathcal{C}} \to 0$$

In Proposition 2.2.2 we shall identify $\pi_1(\mathrm{UEmb}(\mathcal{C},\mathbb{C}))$ with the so-called 'Cantor braid group'.

2.2. Some closely related groups

The group $\Gamma_{\mathcal{C}}$ does not exist in isolation. It is closely related to two other big mapping class groups, namely $\widehat{\Gamma}_{\mathcal{C}}$, the mapping class group of the sphere minus a Cantor set, and $\mathcal{B}_{\mathcal{C}}$, the mapping class group of the disk minus a Cantor set. Explicitly, if we denote by $\mathrm{Homeo}^+(S^2,\mathcal{C})$ the group of orientation-preserving homeomorphisms of the sphere fixing $\mathcal{C}$ setwise, then

$$\widehat{\Gamma}_{\mathcal{C}} := \pi_0(\mathrm{Homeo}^+(S^2,\mathcal{C}))$$

and if we denote by $\mathrm{Homeo}_c(\mathbb{C},\mathcal{C})$ the group of compactly supported homeomorphisms of the plane fixing $\mathcal{C}$ setwise, then

$$\mathcal{B}_{\mathcal{C}} := \pi_0(\mathrm{Homeo}_c(\mathbb{C},\mathcal{C}))$$

We refer to this last group colloquially as the 'Cantor braid group'. Note that we can also write

$$\mathcal{B}_{\mathcal{C}} = \pi_0(\mathrm{Homeo}(\mathbb{D}, \mathcal{C}; \partial\mathbb{D}))$$

where $\mathrm{Homeo}(\mathbb{D}, \mathcal{C}; \partial\mathbb{D})$ denotes the group of homeomorphisms of the unit disk, fixing $\mathcal{C}$ setwise, and fixing $\partial\mathbb{D}$ pointwise.

These groups enjoy the following relations:

PROPOSITION 2.2.1 (Birman exact sequence). *There is a short exact sequence (the Birman 'point-pushing' sequence)*

$$\pi_1(S^2 - \mathcal{C}) \to \Gamma_{\mathcal{C}} \to \widehat{\Gamma}_{\mathcal{C}}$$

PROOF. A homeomorphism of the plane fixing $\mathcal{C}$ setwise extends canonically to a homeomorphism of the sphere fixing $\mathcal{C}$ setwise and ∞ pointwise. There is a map

$$\mathrm{Homeo}^+(S^2, \mathcal{C}) \to S^2 - \mathcal{C}$$

sending a homeomorphism ϕ to the point $\phi(\infty)$. This map is a fibration, and the fiber over ∞ is $\mathrm{Homeo}^+(S^2, \mathcal{C}; \infty) = \mathrm{Homeo}^+(\mathbb{C}, \mathcal{C})$. Now take the homotopy exact sequence of the fibration. Since the identity component of $\mathrm{Homeo}^+(S^2, \mathcal{C})$ is contractible by Proposition 2.1.5, the proposition follows. $\square$

PROPOSITION 2.2.2 (Central extension). *There is a central extension*

$$\mathbb{Z} \to \mathcal{B}_{\mathcal{C}} \to \Gamma_{\mathcal{C}}$$

whose kernel is generated by a Dehn twist around ∞.

PROOF. The standard argument for finite-type mapping class groups goes through without modification. $\square$

Notice by comparing short exact sequences that we obtain the identification $\mathcal{B}_{\mathcal{C}} = \pi_1(\mathrm{UEmb}(\mathcal{C}, \mathbb{C}))$. One can also see this from the fibrations (see Remark 2.1.4)

$$\mathrm{Homeo}_c(\mathbb{C}; \mathcal{C}) \to \mathrm{Homeo}_c(\mathbb{C}) \to \mathrm{Emb}(\mathcal{C}, \mathbb{C})$$

and

$$\mathrm{Homeo}_c(\mathbb{C}, \mathcal{C}) \to \mathrm{Homeo}_c(\mathbb{C}) \to \mathrm{UEmb}(\mathcal{C}, \mathbb{C})$$

and using the fact that $\mathrm{Homeo}_c(\mathbb{C})$ is contractible, which follows by the Alexander trick of coning to a point.

2.3. Orderability

LEMMA 2.3.1 (Center is twist). *The center of the Cantor Braid group is $\mathbb{Z}$, generated by a Dehn twist at infinity. The quotient is the mapping class group (in the usual sense) of the plane minus a Cantor set.*

THEOREM 2.3.2 (Orderability). *The mapping class group of the plane minus a Cantor set is circularly orderable, and the Cantor Braid group is consequently left orderable. The Pure Cantor Braid group is bi-orderable.*

PROOF. The sphere minus a Cantor set is hyperbolizable, and its universal cover is compactified by a circle at infinity. By the Birman exact sequence, the mapping class group of the plane minus a Cantor set may be represented by lifts of mapping classes of $S^2 - \mathcal{C}$ to the universal cover fixing a base point. There is a natural induced action on the circle at infinity by homeomorphisms. This action is faithful, and defines a circular order on the mapping class group of $\mathbb{C} - \mathcal{C}$, and a left order on $\mathcal{B}_{\mathcal{C}}$. □

2.4. Trees, Coarsenings and Pantscades

DEFINITION 2.4.1 (Coarsening). A *coarsening* of $\mathcal{C}$ is an isotopy class of a finite collection of disjoint closed disks in $\mathbb{C}$, each of which intersects $\mathcal{C}$, and such that $\mathcal{C}$ is contained in the union of the interiors of the disks.

If D is a coarsening, $|D|$ is the cardinality of the set. By abuse of notation, we let $\mathbb{D}$ denote the unique trivial coarsening consisting of a single disk (e.g. the unit disk) containing $\mathcal{C}$ in its interior.

DEFINITION 2.4.2 (Refinement and k-refinement). Let D and E be two coarsenings. We say D *refines* E, denoted $D \prec E$, if after an isotopy rel. $\mathcal{C}$ we can arrange that the disks of D are contained in the disks of E.

For any integer $k > 1$ we say D *k-refines* E, denoted $D \prec_k E$, if after an isotopy rel. $\mathcal{C}$, each disk of E contains exactly k disks of D.

The relation $\prec$ is transitive, but $\prec_k$ isn't. If $D \prec_k E$ then $|D| = k|E|$.

DEFINITION 2.4.3 (Pantscade). A *k-pantscade* (or just *pantscade* for short if k is understood) is an infinite inverse sequence of k-refinements,

$$\mathcal{C} \to \cdots \prec_k D_n \prec_k D_{n-1} \prec_k \cdots \prec_k D_2 \prec_k D_1 \prec_k \mathbb{D}$$

for which the Cantor set is in bijection with the inverse limit of π_0 of the components of the refinement.

An infinite inverse sequence of k-refinements determines a k-regular rooted tree T, and if this sequence is a pantscade, there is a natural bijection between $\mathcal{C}$ and the set of ends of this tree.

2.4.1. The braid group of n-coarsenings.

2.4.2. Rotation and Association.

2.5. Braiding

2.6. Ray graphs

2.6.1. Definition and examples.

DEFINITION 2.6.1. The *ray graph*, denoted $\mathcal{R}$, is the graph whose vertices are isotopy classes of properly embedded rays from $\mathcal{C}$ to infinity, and whose edges are pairs of rays that may be made disjoint by an isotopy, except possibly at a common endpoint.

For any finite subset X of $\mathbb{C}$ there is a ray graph $\mathcal{R}_X$. The isomorphism type of $\mathcal{R}_X$ depends only on the cardinality $n := |X|$ and we sometimes write $\mathcal{R}_n$ for $\mathcal{R}_X$ by abuse of notation if the particular isomorphism of some fixed $\mathcal{R}_n$ to $\mathcal{R}_X$ is unimportant.

EXAMPLE 2.6.2 ($|X| = 1, 2, 3$). $\mathcal{R}_1$ is a single point, and $\mathcal{R}_2$ is two points joined by an edge.

The double cover of the Riemann sphere branched over four points is a torus. The preimage of an isotopy class of ray is an isotopy class of simple closed curve through a marked point. Two rays are disjoint (adjacent in $\mathcal{R}_3$) if and only if their preimage loops intersect at most once or twice. The graph $\mathcal{R}_3$ is therefore quasi-isometric to the Farey graph. This graph has infinite diameter, and is quasi-isometric to a tree.

2.6.2. Coarsening.

CONSTRUCTION 2.6.3 (Coarsening map). A *coarsening* of $\mathcal{C}$ is a finite collection D of disjoint closed disks in $\mathbb{C}$, each of which intersects $\mathcal{C}$, and such that $\mathcal{C}$ is contained in the union of the interiors of the disks.

There is a quotient map $\pi_D : \mathbb{C} \to \mathbb{C}/D$ obtained by crushing every component of D to a point. Note that $\mathbb{C}/D$ is homeomorphic to $\mathbb{C}$. By abuse of notation we think of D as a finite subset of $\mathbb{C}/D$, and we let $\mathcal{R}_D$ denote the ray graph of this finite set.

There is a *coarsening map* $\pi_D : \mathcal{R} \to \mathcal{R}_D$ defined as follows. Given an embedded ray r, eliminate bigons of intersection of r with ∂D so that r meets ∂D efficiently. Then take r to the tail component of $r - D$; i.e. the connected component containing infinity. This map is well-defined on isotopy classes of rays.

LEMMA 2.6.4 (Coarsening map is Lipschitz). *The map π_D is surjective and 1-Lipschitz, and there are (many) isometric sections $\mathcal{R}_D \to \mathcal{R}$.*

PROOF. The first two claims follow from the definition.

There is a hyperbolic structure on $\mathbb{C} - \mathcal{C}$, and any coarsening D determines a canonical collection of disks — those bounded by hyerpbolic geodesics. Every ray in $\mathcal{R}_D$ determines a canonical geodesic ray in $\mathbb{C} - D$ meeting ∂D perpendicularly. There are countably many endpoints of such rays, and the countably many endpoints may be extended (arbirarily) to disjoint rays in D, thereby giving a section. $\square$

REMARK 2.6.5. Coarsening is a special case of subsurface projection.

REMARK 2.6.6. It makes sense to define ray graphs and coarsenings in much more generality. For example, there is a ray graph $\mathcal{R}_X$ whenever X is compact and totally disconnected, and there are coarsening maps π_D^X for suitable decompositions D containing X. We shall not develop this more general theory here.

2.6.3. Bavard's Hyperbolicity Theorem.

THEOREM 2.6.7 (Bavard's hyperbolicity theorem [1]). *The graph $\mathcal{R}$ is connected, hyperbolic, and has infinite diameter.*

We prove the first and third statements now, and defer the proof of hyperbolicity.

PROOF. If two rays start at different points in $\mathcal{C}$, then by general position we can arrange for them to intersect transversely, and then we can comb all but finitely many intersections to infinity. If two rays intersect in finitely many points, by cut and paste we can find a new ray that intersects each in at least one fewer points. Thus $\mathcal{R}$ is connected.

To see the diameter is infinite, let $\mathcal{C} \subset D$ be a coarsening with at least three components. The coarsening map $\pi_D : \mathcal{R} \to \mathcal{R}_D$ is surjective and 1-Lipschitz by Lemma 2.6.4, and since $|D| \geq 3$ the diameter of $\mathcal{R}_D$ (and therefore also $\mathcal{R}$) is infinite, by Example 2.6.2. $\square$

2.6.3.1. *Slim Unicorns.* Hyperbolicity can be proved using the slim unicorn method of Hensel–Przytycki–Webb [**21**] which is indifferent about the topology of the underlying surface, and even indifferent about whether it is of infinite type or not. We give a brief exposition of this argument; for details see [**21**] and [**1**].

The argument of HPW applies most directly to the *arc* complex, whose vertices are isotopy classes of arcs between pairs of marked or boundary points. Given a pair of oriented arcs α, β they construct a *path* $I'(\alpha, \beta)$ of intermediate arcs between α and β as follows. First, put α and β in minimal position. For each point $p \in \alpha \cup \beta$ form μ from two initial subpaths of α and β. If the result is embedded, it goes in $I'(\alpha, \beta)$; otherwise it doesn't.

This finite set of paths is *ordered* by how big an initial path they share with α, the bigger the earlier in the ordering. Then $I(\alpha, \beta)$ is the ordered path $\alpha, I'(\alpha, \beta), \beta$. Note that the result is *symmetric* in α and β, although both must be oriented.

LEMMA 2.6.8. *$I(\alpha, \beta)$ is a path; i.e. adjacent vertices have distance 1.*

PROOF. Let p and q be successive points on α associated to adjacent μ, ν in $I(\alpha, \beta)$. Write $\mu = \mu_\alpha \cup \mu_\beta$ and $\nu = \nu_\alpha \cup \nu_\beta$ where $\mu_\alpha \subset \nu_\alpha$ and $\nu_\beta \subset \mu_\beta$. Let γ be the segment of α between p and q so that $\nu_\alpha = \mu_\alpha \cup \gamma$, and let δ be the segment of β between p and q so that $\mu_\alpha = \nu_\alpha \cup \delta$.

We simply need to show that δ and γ are disjoint (except at their endpoints). But if not, we could take a rightmost point of intersection and build a new element of $I(\alpha, \beta)$ strictly between μ and ν. $\square$

PROPOSITION 2.6.9 (HPW [**21**], Lemma 3.3). *For any α, β, γ every point on $I(\alpha, \beta)$ is within distance 1 of $I(\beta, \gamma) \cup I(\gamma, \alpha)$.*

In fact, it turns out that one can find a triple of points on the three sides at pairwise distance at most 1.

PROOF. In other words, for every $\mu \in I(\alpha, \beta)$ there is a disjoint ν in $I(\beta, \gamma)$ or $I(\gamma, \alpha)$. This is trivial for $\mu = \alpha$ or β, so suppose $\mu \in I'(\alpha, \beta)$.

The curve μ is associated to a point $p \in \alpha \cup \beta$, and is the union of two initial segments α_μ, β_μ of α, β meeting only at the point p. Consider the curve γ; if it's disjoint from μ we're already done. Otherwise it meets one of α_μ, β_μ first at q. Then ν is the curve associated to q. $\square$

2.6.3.2. *Guess the geodesics.* From Proposition 2.6.9 hyperbolicity of the arc complex follows, by a theorem of Masur–Schleimer:

THEOREM 2.6.10 (Masur–Schleimer, Guess the Geodesics [**25**] Thm. 3.15). *Suppose X is a connected graph, and suppose there is a constant M, and for every pair of vertices x, y a connected subgraph $I(x, y) \subset X$ satisfying*

(1) if $d(x, y) \leq 1$ then $I(x, y)$ has diameter at most M; and

(2) for all x, y, z the graph $I(x,y)$ is contained in the M-neighborhood of $I(y,z) \cup I(z,x)$.

Then X is δ-hyperbolic for some $\delta(M)$.

Bavard applies [**21**] in the following way. She defines a *lasso* to be a properly embedded *line* ℓ in $\mathbb{R}^2 - \mathcal{C}$, dividing $\mathcal{C}$ nontrivially into two parts. There is a lasso complex $\mathcal{L}$ with isotopy classes of lassos for vertices, and those pairs with disjoint representatives as edges.

If ℓ is a lasso, any two rays disjoint from ℓ are distance at most 2 from each other. Conversely, if ℓ and ℓ' are disjoint non-isotopic lassos, there is a ray disjoint from both. It follows that $\mathcal{R}$ and $\mathcal{L}$ are quasi-isometric. Actually, it is possible to think of a ray as a kind of completion of a sequence of lassos as follows: any maximal clique in $\mathcal{L}$ determines a tree; noncompact ends of such a tree are rays, and every ray arises in this way.

The slim unicorns proof works without modification to prove hyperbolicity of the lasso graph. Hyperbolicity of the ray graph follows, completing the proof of Theorem 2.6.7.

2.7. Some examples of dynamics

Bavard [**1**] gives some elegant examples of mapping classes that act loxodromically on $\mathcal{R}$.

2.8. Gromov boundary of the Ray graph

Bavard–Walker [**4**] gave an elegant description of the Gromov boundary of $\mathcal{R}$. From their description certain dynamical conclusions follow, which we shall now explain.

2.8.1. Conical cover and the circle of geodesics. To discuss long rays it's convenient to take one of two equivalent points of view.

Let Ω denote the plane minus a Cantor set, and fix a complete hyperbolic structure. The *conical cover* $\widehat{\Omega}$ is the covering space associated to the $\mathbb{Z}$ subgroup of $\pi_1(\Omega)$ of loops around infinity. Thus $\widehat{\Omega}$ is conformally equivalent to the punctured open disk. By abuse of notation we refer to the puncture in $\widehat{\Omega}$ as infinity. There is a circle $S^1(\mathcal{G})$ compactifying the disk, which is in natural bijection with the set of complete geodesics in $\widehat{\Omega}$ with one endpoint at infinity. By lifting we may also identify this circle with the set of complete (not necessarily simple) geodesics in Ω with one end at infinity. As in Theorem 2.3.2 the mapping class group $\mathcal{M}(\Omega)$ acts on $S^1(\mathcal{G})$ by orientation-preserving homeomorphisms.

In this notation it makes sense to pick out a subset $S \subset S^1(\mathcal{G})$ that we call the *simple set*. This is precisely the set of complete *simple* geodesics in Ω with one end at infinity. This set decomposes into three subsets:

(1) Proper simple geodesics with one end at infinity and one end on the Cantor set; these are exactly the vertices in the ray graph. We call its elements *short rays*. This set is uncountable; we denote it R. Note that $\mathcal{M}(\Omega)$ acts transitively on R.
(2) Proper simple geodesics with both ends at infinity; these are exactly the oriented lassos. This set is countable; we denote it L. Note that $\mathcal{M}(\Omega)$ acts transitively on L.
(3) Non-proper simple geodesics with one end at infinity; Bavard–Walker call these *long rays*. We denote this set by X.

The main theorem in this section describes the structure of S.

THEOREM 2.8.1. *The following describes the structure of S and the action of $\mathcal{M}(\Omega)$ on it.*

(1) The simple set S is compact and its complement is dense.
(2) $R \cup X$ is a Cantor set, and is the unique minimal set for the action of $\mathcal{M}(\Omega)$ on $S^1(\mathcal{G})$.
(3) Every complementary interval to $R \cup X$ contains exactly one point of L, so that $\mathcal{M}(\Omega)$ acts transitively on this set of intervals.
(4) The boundary points of $R \cup X$ are in X, and consist precisely of long rays that spiral around a finite (simple) geodesic in Ω.
(5) Every long ray which does not spiral around a simple geodesic is accumulated on both sides by short rays.

PROOF. The complement of S is dense, for any long geodesic from infinity either already intersects itself, or can be perturbed slightly so that it runs into itself in the future. Evidently the first self-intersection point of a geodesic is stable under perturbation, so the complement of S is open. This proves (1).

Next we show that every orbit accumulates on a simple ray. To see this, let γ be an arbitrary geodesic from infinity. Then γ must cross some closed simple geodesic α separating part of the Cantor set from the other part. Take a sequence of mapping classes that shrinks α down to a single point p of the Cantor set; the images of γ accumulate on a simple ray from infinity to p. It follows that every minimal set must contain some simple ray, and therefore every simple ray; i.e. the whole set R. Since the closure $\overline{R}$ is invariant, it is the unique minimal set.

Now let's let $p \in S^1(\mathcal{G})$ be an arbitrary point corresponding to a self-intersecting geodesic γ. The part of γ up to its first intersection point is shaped like the letter ρ. Sliding the vertex of the ρ up or down defines an open interval I containing p in the complement of S which limits on one side to an oriented lasso, and on the other side to a long ray x spiralling around a finite simple geodesic α in Ω which is in the isotopy class of the loop of the ρ. The oriented lasso arises from exactly two complementary intervals this way; the other interval corresponds to ρs with the opposite orientation whose loop is isotopic to some simple geodesic α'. The geodesics α and α' cobound a once-punctured annulus, and the lasso is the 'punctured equator' of this punctured annulus.

Note that α, being a geodesic, must separate infinity from some of the Cantor set. Thus we can find a point in the Cantor set on the far side of α, and join it by a simple arc to α, and then drag this simple arc around α so it becomes a simple spiralling geodesic y. Perturb $x \cup y$ so that each of them spirals only finitely many times around α and then joins up with the other; the result is a short ray approximating x. Thus x is not isolated in $R \cup X$. This shows $R \cup X$ is perfect.

Since p and the interval I that contained it was arbitrary, we see that $R \cup X$ is perfect (and therefore a Cantor set) whose complementary intervals all contain a unique lasso. This proves (3) and (4).

Since boundary points of $R \cup X$ are accumulated by short rays on one side, and since non-boundary points of a Cantor set are accumulated by boundary points on both sides, it follows that every long ray which does not spiral around a simple geodesic is accumulated by short rays on both sides. This proves (5), and it also shows that $\overline{R} = R \cup X$. In particular, $R \cup X$ is the unique minimal set, proving (2). $\square$

2.8.2. Cliques. Following Bavard–Walker we define a 'simple graph' whose vertices are simple geodesics from infinity, two of which are joined by an edge if they are disjoint. Thus the simple graph extends the ray-lasso graph.

PROPOSITION 2.8.2. *The simple graph has one unbounded component called the* big component*, namely the component containing the ray-lasso graph. The inclusion of the ray-lasso graph into the big component is a quasi-isometry. The big component is contained in the 2-neighborhood of its image. Every other component is a clique and therefore has diameter one.*

PROOF. Say a long ray is *filling* if it is not disjoint from any short ray. Any long ray which is not filling is distance 1 from the ray-lasso subgraph. Let x be a filling long ray, and suppose x is disjoint from y and z. We claim y is disjoint from z. For, otherwise, we can make a bigon from initial segments of y and z by extending them from infinity until they intersect and throwing away arcs beyond the intersection. Such a bigon necessarily has Cantor set on both sides. Since x is disjoint from $y \cup z$ it is contained on one side of the bigon. Therefore we can join infinity to a point in the Cantor set on the other side to create a short ray disjoint from x, contrary to the hypothesis that x is filling. This proves the claim.

This claim shows that the big component is contained in the 2-neighborhood of its image, and also that every component disconnected from the big component is a clique. It remains to show that the inclusion of the ray-lasso graph into the big component is a quasi-isometry. $\square$

2.8.3. Gromov boundary as a quotient.

2.9. Inverse limits

The collection of coarsenings D of $\mathcal{C}$ up to isotopy rel. $\mathcal{C}$ is an inverse system. We write $D \prec E$ if D refines E; i.e. if E is a further coarsening of D, and $\pi_E^D : \mathcal{R}_D \to \mathcal{R}_E$.

LEMMA 2.9.1 (Inverse limit). *The graph $\mathcal{R}$ is the inverse limit of the $\mathcal{R}_D$. In other words, a compact graph Γ embeds in $\mathcal{R}$ if and only if it embeds in $\mathcal{R}_n$ for some n.*

PROOF. Lemma 2.6.4 explains how to produce isometric embeddings of $\mathcal{R}_n$ in $\mathcal{R}$ for any n. Conversely, let $r_1, \cdots, r_k$ be a finite collection of isotopy classes of rays in $\mathbb{C}$, and let Γ be the induced subgraph of $\mathcal{R}$ they span. We must find a coarsening D so that $\pi_D|\Gamma$ is an isomorphism.

Two rays are connected by an edge if and only if their geodesic representatives are interior disjoint. If two rays r_i, r_j intersect, this is witnessed by their restriction to some planar subsurface $S_{i,j}$ of finite type with totally geodesic boundary. If r_i and r_j don't intersect but are not isotopic, this is also witnessed by their restriction to some $S_{i,j}$ as above. We may enlarge each $S_{i,j}$ if necessary so that it contains ∞. Let S be the union of the $S_{i,j}$. Then $D := \mathbb{C} - S$ is a coarsening with the property that $r_i \cap S_{i,j}$ is contained in the tail of $r_i - D$ for all i, j. Thus $\pi_D|\Gamma$ is an isomorphism. $\square$

Each braid group $\mathcal{B}_D$ acts on $\mathcal{R}_D$ and also on the collection of all $\mathcal{R}_E$ where $D \prec E$. The Cantor braid group $\mathcal{B}_{\mathcal{C}}$ acts on the entire directed system.

If D and E are coarsenings, the stabilizer of E in $\mathcal{B}_D$ is denoted $\mathcal{B}_D^E$, and there is an induced homomorphism $\mathcal{B}_D^E \to \mathcal{B}_E$. By abuse of notation, we allow the case that $D = \mathcal{C}$ even though this is not strictly speaking a coarsening, so that the stabilizer of E in $\mathcal{B}_{\mathcal{C}}$ is denoted $\mathcal{B}_{\mathcal{C}}^E$, and there is an induced homomorphism $\mathcal{B}_{\mathcal{C}}^E \to \mathcal{B}_E$.

Construction 2.9.2 (Inverse limit). Let $\mathcal{A}$ be a directed set indexing a family D_α of coarsenings. When $\mathcal{A}$ has additional structure — for instance if $\mathcal{A}$ is invariant under a family of endomorphisms — there is a corresponding structure on the inverse limit of the braid groups $\mathcal{B}_{\mathbb{D}_\alpha}$, thought of as a subgroup of $\mathcal{B}_{\mathcal{C}}$.

2.10. Endomorphisms

Finite type mapping class groups are very rigid — they rarely admit any interesting automorphisms. Infinite type mapping class groups also have a tendency to be rigid. The situation looks very boring. But automorphisms are the wrong thing to look at: infinite type mapping class groups admit very interesting endomorphisms.

The distinction is largely moot in the finite type world, because finite type mapping class groups are Hopfian and (nearly) co-Hopfian. A group is *Hopfian* if every surjective endomorphism is an isomorphism, and *co-Hopfian* if every injective endomorphism is an isomorphism.

Finite Braid groups are Hopfian because they are finitely generated and residually finite. Furthermore, they are *almost* co-Hopfian by a theorem of Bell–Margalit [**6**]: every injective endomorphism is an isomorphism modulo the center (which is $\mathbb{Z}$ and generated by a Dehn twist at infinity).

Bell–Margalit prove that braid groups are boring in a very interesting way: they show that injective endomorphisms between finite type braid groups are geometric — they are induced by homeomorphisms between punctured disks. This geometric fact has its analog for Cantor braid groups, but the implications are very different, as we shall see in § 2.11

2.11. Shift-invariant Cantor braids

2.11.1. Tuning.

CHAPTER 3

Renormalization

3.1. Asymptotic self-similarity

3.2. Lei–Misiurewicz points

3.3. Cantor braids in the torus

CHAPTER 4

A Bestiary of Big Mapping Class Groups

Part 2

Examples and Applications

CHAPTER 5

The Shift Locus

In this chapter we study the space of *shift polynomials*, complex polynomials for which every critical point is in the attractive basin of infinity. Every such polynomial has a Cantor Julia set, and there is a natural representation from the fundamental group of the shift locus to the mapping class group of the plane minus a Cantor set. This is ground zero for the relationship between complex dynamics and big mapping class groups.

The point of view taken in this chapter was partly developed in the preprint [**3**] and there are substantial connections to the earlier pioneering work of DeMarco et. al. (see [**14**, **13**, **15**]) not to mention Branner–Hubbard [**8**, **9**], Blanchard–Devaney–Keen [**7**] and many others.

On the other hand, I believe that most of the theorems presented here are new (as far as I know), and for some of those which are not, the proofs are new or hold in greater generality than have appeared elsewhere.

5.1. Definition

For each integer d let Poly_d denote the space of complex polynomials of degree $\le d$. Taking coefficients gives a natural identification $\mathrm{Poly}_d = \mathbb{C}^{d+1}$. The affine group $z \to \alpha z+\beta$ gives a natural conjugation action on Poly_d. This action is generically free, and the reduced quotient is a complex variety birational to $\mathbb{C}^{d-1}$.

The *Shift Locus* $\mathcal{S}_d \subset \mathrm{Poly}_d$ is the subset of polynomials of degree exactly d for which every critical point is in the basin of attraction of ∞. We write $\mathcal{S}$ if d is understood.

LEMMA 5.1.1. *A polynomial f is in the Shift Locus if and only if J_f is a Cantor set and the dynamics of f on J_f is uniformly expanding.*

In particular, the action of f on J_f is conjugate to the shift on d letters, explaining the name. If f has degree 2, then f is in the Shift Locus if and only if J_f is a Cantor set. But for f of higher degree it's possible for J_f to be a Cantor set and at the same time contain some critical point.

EXAMPLE 5.1.2. The polynomial $f : z \to z^3 - 6z + 2\sqrt{2}$ is not in $\mathcal{S}_3$ since the critical point $\sqrt{2}$ is fixed by f. However, J_f is a Cantor set.

5.2. Surfaces out of Böttcher Paper

5.2.1. Böttcher coordinates. A polynomial f of positive degree fixes infinity, and if the degree is at least two, the derivative vanishes at infinity; one calls such a fixed point *superattractive*. Böttcher famously proved that holomorphic functions have trivial moduli near superattractive fixed points — the only invariant is the degree. This means there are canonical holomorphic coordinates near a superattractive fixed point, known as *Böttcher coordinates.*

THEOREM 5.2.1 (Böttcher coordinates). *Let $d > 1$ and suppose $f \in \mathcal{S}_d$. Then there are open neighborhoods U_f, U_B of infinity in $\hat{\mathbb{C}}$ and a holomorphic isomorphism $\psi_f : U_f \to U_B$ conjugating f on U_f to $z \to z^d$ on U_B. Furthermore, the map ψ_f is unique up to multiplication by a $(d-1)$st root of unity.*

PROOF. Let $f(z) = az^d + bz^{d-1} + \cdots$ where $a \neq 0$ but b might be. If we conjugate by $z \to \alpha z + \beta$ then we get a new map $\hat{f}$ given by

$$\hat{f}(z) = \alpha^{-1}(a(\alpha z+\beta)^d + b(\alpha z+\beta)^{d-1} + \cdots) - \alpha^{-1}\beta = a\alpha^{d-1}z^d + (ad\beta + b)\alpha^{d-2}z^{d-1} + \cdots$$

Setting $\alpha = a^{-1/(d-1)}$ and $\beta = -b/ad$ puts $\hat{f}$ in the form $z \to z^d + O(z^{d-2})$. The new $\hat{f}$ is the unique polynomial (up to a $(d-1)$st root of unity) holomorphically conjugate to $z \to z^d$ to second order at infinity. Hensel's Lemma gives a formal power series conjugating $\hat{f}$ to $z \to z^d$ at infinity, and this power series converges uniformly on an open neighborhood. $\square$

5.2.2. Analytic continuation. The attracting basin of ∞ under the map $z \to z^d$ is $\mathbb{C} - \mathbb{D}$, a punctured disk. If $f \in \mathcal{S}_d$ and $\psi_f : U_f \to U_B$ is a conjugacy between f and $z \to z^d$ near infinity, we can analytically continue ψ_f to any simply-connected open set $U \subset \hat{\mathbb{C}}$ satisfying

$$U_f \subset U \subset \hat{\mathbb{C}} - J_f$$

Conversely, we can analytically continue ψ_f^{-1} along any ray of constant argument from ∞ until we meet a precritical point; i.e. a point $p \in \mathbb{C}$ for which $(f^n)'(p) = 0$ for some n.

The union of these maximal open rays defines a canonical open dense subset $V_B \subset \mathbb{C}-\mathbb{D}$ which is the complement of countably many half-open segments of the form $\ell_{(v,\theta)} := \{z : \arg(z) = \theta, 1 < |z| \le e^v\}$ for some $(v,\theta) \in \mathbb{R}^+ \times \mathbb{R}/2\pi\mathbb{Z}$, where there are only finitely many $\ell_{(v,\theta)}$ for v greater than any given positive number. The set V_B is forward invariant under $z \to z^d$, and the map ψ_f^{-1} is injective on V_B and conjugates $z \to z^d$ to f. We refer to the ℓ as *semi-leaves.*

5.2.3. Böttcher models. For ℓ of the form $\ell_{(v,\theta)}$ let $\partial\ell$ denote the unique endpoint of ℓ; i.e. $\partial\ell_{(v,\theta)} = e^{v+i\theta}$. We may extend ψ_f^{-1} over the set $\partial\ell$, but it fails to be injective. If c is a critical point of f where $f'(c)$ vanishes to order k then there are $k+1$ semi-leaves ℓ with the same v, and with θ values that differ by a multiple of $2\pi/d$. These ℓ all map to the same semi-leaf under $z \to z^d$, and their endpoints all map to $\psi_f f(c) \in V_B$, the image of the critical *value* $f(c)$ associated to c.

We may build a Riemann surface Ω_f by cutting $\mathbb{C} - \mathbb{D}$ open along the semi-leaves ℓ, and then regluing the result in the unique way for which the result is compatible with the dynamics $\ell_{(v,\theta)} \to \ell_{(dv,d\theta)}$ and with the map ψ_f^{-1}. In particular, regluing preserves the v coordinate and changes the θ coordinate by some fixed multiple of $2\pi/d$. If ℓ_j are the $k+1$ semi-leaves whose endpoints all map to a precritical point c under ψ_f^{-1}, we can reorder them so that the indices j mod $k+1$ agree with the cyclic order of the arguments in $\mathbb{R}/2\pi\mathbb{Z}$. Cutting replaces each ℓ_j with two copies ℓ_j^L and ℓ_j^R which can be thought of as lying infinitesimally to the 'left' and the 'right' of ℓ_j in $\mathbb{C} - \mathbb{D}$. We reglue by identifying ℓ_j^R with ℓ_{j+1}^L.

LEMMA 5.2.2 (Böttcher inverse is isomorphism). *The map ψ_f^{-1} extends to an isomorphism from Ω_f to $\mathbb{C} - J_f$. The Riemann surface Ω_f has genus 0.*

PROOF. By construction, ψ_f^{-1} extends uniquely over Ω_f and conjugates $z \to z^d$ to f. The map $z \to z^d$ on Ω_f is critical exactly when f is and with the same degree and combinatorics, so ψ_f^{-1} is an immersion. Since it is proper it is a covering map, and since it is degree 1 near infinity, it is an isomorphism. □

COROLLARY 5.2.3 (No linking). *If ℓ and ℓ' are two families of semi-leaves whose endpoints correspond to precritical points c, c' of f then the arguments of ℓ and ℓ' do not link in $\mathbb{R}/2\pi\mathbb{Z}$.*

PROOF. If there were linking, the genus of Ω_f would be positive. □

We call the pair $(\Omega_f, z \to z^d)$ the *Böttcher model* for f.

5.2.4. Böttcher space and realization. We can now define the space $\mathcal{FS}_d$ of *formal d-shifts* to be certain dynamical systems of the form $(\Omega, z \to z^d)$ where Ω is obtained from $\mathbb{C} - \mathbb{D}$ by cut-and-paste along a backwards-invariant collection of semi-leaves $\{\ell\}$, and for which the map $z \to z^d$ has $d - 1$ critical points on Ω, counted with multiplicity.

Such an Ω is given by the following data:

1. a choice of $d-1$ semi-leaves (with multiplicity) corresponding to the critical values of $z \to z^d$;
2. for each critical value semi-leaf, a partition of its preimage leaves into subsets, where the subsets of cardinality greater than 1 correspond to the critical semi-leaves; and such that
3. no two critical subsets have pairs of arguments that link in $\mathbb{R}/2\pi\mathbb{Z}$.

Every Ω_f gives rise to data of this form by Corollary 5.2.3, and conversely this data gives rise to a canonical collection of pre-critical semi-leaves by pulling back the partitions under the dynamics. The condition of no linking propagates backwards, so the semi-leaves of Ω satisfy the conclusion of Corollary 5.2.3.

The space of formal d-shifts has the structure of a complex variety of dimension $d - 1$. Where the critical points are distinct, their values in $\mathbb{C} - \mathbb{D}$ define local holomorphic coordinates. Where $k \leq d - 1$ critical points coalesce, the local model is $\mathbb{C}^{d-k} \times (\mathbb{C}^k/S_k)$ where the symmetric group S_k acts on $\mathbb{C}^k$ by permuting coordinates. As is well-known, the quotient $\mathbb{C}^k/S_k$ is biholomorphic to $\mathbb{C}^k$; think of the map that takes a monic polynomial to its unordered set of roots.

THEOREM 5.2.4 (Formal shift is realizable). *Let $(\Omega, z \to z^d)$ be a formal d-shift. Then there is a (unique) $f \in \mathcal{S}_d$ up to conjugacy for which $\Omega = \Omega_f$.*

PROOF. Let $0 < v_0$ be smaller than the altitude of any critical point. Then the level set $v = v_0$ is a collection of circles that separates Ω. Since the semi-leaves satisfy no linking, the genus of Ω is zero, so each circle in the level set bounds a subset of Ω for which the inverse $z \to z^{1/d}$ has d well-defined branches making this subset into a shift. The Realization Theorem 1.3.2 says that there is a canonical Riemann surface structure on the union $\hat{\Omega} := \infty \cup \Omega \cup \mathcal{E}$ making it biholomorphic to $\hat{\mathbb{C}}$, so that $z \to z^d$ extends over $\hat{\Omega}$ to a proper holomorphic map of degree d with ∞ as a superattracting fixed point of degree d.

Such a map is necessarily a polynomial of degree d, and by construction it is in the Shift Locus. □

REMARK 5.2.5. Essentially the same theorem is proved as [**14**], Thm. 7.1, although the method of proof is a little different, and non-constructive. One advantage of our proof of the Realization Theorem 1.3.2 is that it finds a fixed point by a quickly converging algorithm; this is useful e.g. for computer implementation. We discuss this further in the sequel.

COROLLARY 5.2.6. *The Böttcher map* $\mathcal{S}_d \to \mathcal{FS}_d$ *taking* $(\hat{\mathbb{C}}, f)$ *to* $(\hat{\Omega}_f, z \to z^d)$ *is a biholomorphic isomorphism.*

5.2.5. Laminations. It is convenient to encode the cutting and pasting of semi-leaves by extra data. If ℓ is a set of semi-leaves with endpoints mapping to the same precritical point of f, we may extend them to closed segments in $\mathbb{C}$ by adding endpoints on the unit circle, and then join endpoints consecutive in the circular order by semicircular arcs in the unit disk. These arcs (and their extensions to a pair of semi-leaves) are called *leaves*.

The no linking property — i.e. Corollary 5.2.3 — implies that no two leaves will cross, so the union of the leaves are embedded. We call this structure the *Böttcher lamination* and denote it $\mathcal{L}$. Laminations of this sort were introduced by Thurston, and they are omnipresent in 1-dimensional complex dynamics. Some people add leaves that arise as limits so that the collection of all leaves is closed; we don't see an obvious advantage to this in the current context, but feel free to do so.

As far as I know, DeMarco and Pilgrim were the first to introduce laminations for the analysis of shift polynomials; see [**13, 15**].

5.2.6. Green coordinates and Squeezing. We have been implicitly using the coordinates coming from the Green's function $\log(z)$ on $\mathbb{C} - \mathbb{D}$. This function takes values in $\mathbb{R}^+ \times \mathbb{R}/2\pi\mathbb{Z}$. We call the first coordinate the *altitude* and the second coordinate the *argument*, and we continue to denote these coordinates v and θ respectively. The dynamical map $z \to z^d$ multiplies altitude and argument by d. There is a holomorphic 1-form $\varphi := d(v + i\theta)$ on $\mathbb{C} - \mathbb{D}$, and the map $z \to z^d$ pulls back φ to $d\varphi$.

Let f be in $\mathcal{S}_d$, and let Ω_f be obtained from $\mathbb{C} - \mathbb{D}$ by cut and paste. The coordinate v continues to make sense on Ω_f, but θ is only defined as an element of $\mathbb{R}/2\pi\mathbb{Z}[\frac{1}{d}]$. Nevertheless, φ makes sense on Ω_f, and it is still true that $z \to z^d$ pulls back φ to $d\varphi$.

There is a smooth (actually, holomorphic) vector field on $\mathbb{C} - \mathbb{D}$ which is given by $v\partial_v$ in v, θ coordinates. Flowing critical points and their preimages by this vector field defines a holomorphic vector field V on $\mathcal{S}_d$ whose flow defines a free proper action of $\mathbb{R}$ on $\mathcal{S}_d$. The Hausdorff dimension of J_f decreases[1] under this flow to zero.

The orbits of this flow are dynamical systems that are topologically conjugate but not holomorphically conjugate; the conjugacy is given by an f-invariant quasiconformal deformation of Ω which is a Teichmüller geodesic associated to the canonical holomorphic differential φ^2 on $\mathbb{C} - \mathbb{D}$. Deforming f under this flow is called *squeezing*.

[1] monotonically?

5.2.7. Examples in low degree.

EXAMPLE 5.2.7 (Mandelbrot is connected). A degree 2 polynomial up to conjugacy is of the form $f : z \to z^2 + c$. The set $\mathcal{M}$ of $c \in \mathbb{C}$ for which J_f is connected is the *Mandelbrot set*, first defined and studied by Brooks and Matelski [**10**]. The complement is identified with the shift locus $\mathcal{S}_2$ (up to conjugacy). From the Böttcher model we see that $\mathcal{S}_2$ is biholomorphic to a punctured disk, which uniformizes the complement of $\mathcal{M}$. Thus, $\mathcal{M}$ is connected with connected complement. See Figure 5.1.

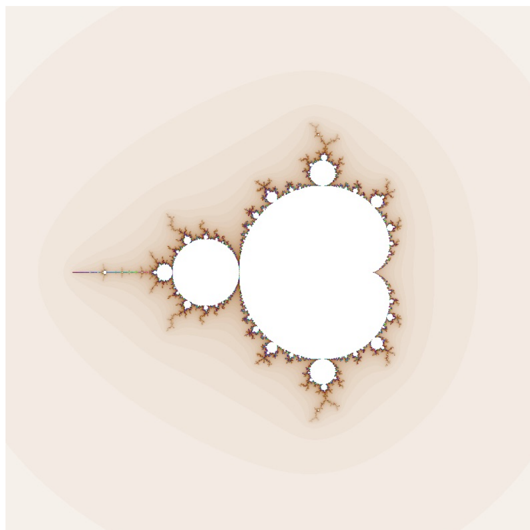

FIGURE 5.1. $\mathcal{M}$ is the complement of $\mathcal{S}_2$ in $\mathbb{C}$.

The connectivity of $\mathcal{M}$ was proved by Douady–Hubbard [**17**] using essentially this argument. They studied the Böttcher map $\mathbb{C} - \mathcal{M} \to \mathbb{C} - \mathbb{D}$ and proved that it was proper and an isomorphism near infinity, and therefore invertible. However, they did not directly construct an inverse (in particular, they did not prove any analog of the Realization Theorem).

EXAMPLE 5.2.8 (Multibrots and degenerately critical shifts). Let $\mathcal{S}_3^\Delta$ be the space of degree 3 polynomials in the shift locus for which f' has a single root of order 2, up to conjugacy. All such polynomials may be written in the form $z \to z^3 + c$. From the Böttcher model we see that $\mathcal{S}_3^\Delta$ is biholomorphic to a punctured disk. It uniformizes the complement of $\mathcal{M}_3^\Delta$, the space of degree 3 polynomials of the form $z \to z^3+c$ with connected Julia set. See Figure 5.2.

The same argument works for any d: if $\mathcal{S}_d^\Delta$ denotes the space of degree d shift polynomials of the form $z \to z^d + c$ then $\mathcal{S}_d^\Delta$ is biholomorphic to a punctured disk, and the complement $\mathcal{M}_d^\Delta$ in the c-plane is connected; these complements are called *Multibrot sets*.

EXAMPLE 5.2.9 (Real degree 3). Let $\mathcal{S}_3^+(\mathbb{R})$ denote the space of degree 3 shift polynomials of the form $z \to z^3 + az + b$ where a and b are real. See Figure 5.3. Not every real degree 3 polynomial is conjugate to one of this form; there is also the family $z \to -z^3 + az + b$ which we denote $\mathcal{S}_3^-(\mathbb{R})$; see Figure 5.4. From these figures it is evident that $\mathcal{S}_3^\pm(\mathbb{R})$ are

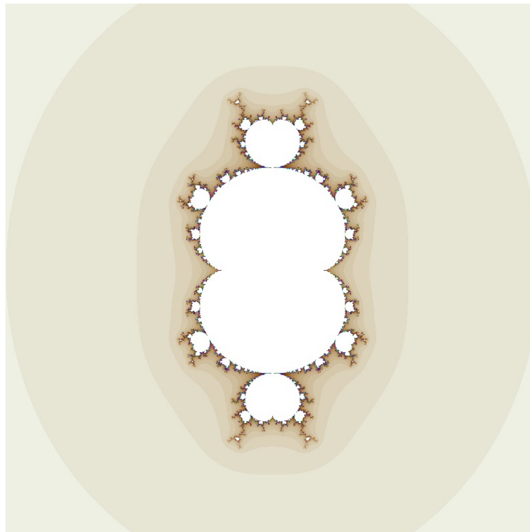

FIGURE 5.2. $\mathcal{M}_3^{\Delta}$ is the complement of $\mathcal{S}_3^{\Delta}$ in $\mathbb{C}$.

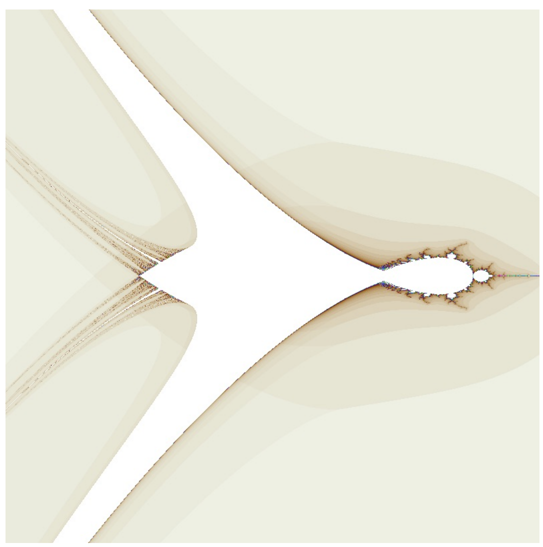

FIGURE 5.3. $\mathcal{M}_3^{+}(\mathbb{R})$ is the complement of $\mathcal{S}_3^{+}(\mathbb{R})$ in $\mathbb{R}^2$.

both homotopic to countably infinitely many proper rays, and that $\mathcal{M}_3^{\pm}(\mathbb{R})$ are not locally connected.

We may see this in the Böttcher model: a polynomial f is real if and only if it is invariant under complex conjugation. There is a Böttcher map to $z \to \pm z^3$ depending on the sign of the leading coefficient of f, and this map respects complex conjugation. Thus, the Böttcher lamination $\mathcal{L}$ corresponds to a real polynomial if and only if it is invariant under complex conjugation. We'll expand on this in § 5.6.

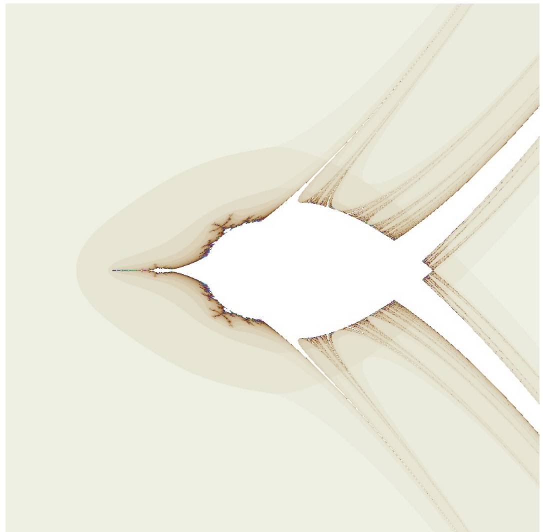

FIGURE 5.4. $\mathcal{M}_3^-(\mathbb{R})$ is the complement of $\mathcal{S}_3^-(\mathbb{R})$ in $\mathbb{R}^2$.

5.3. The topology of $\mathcal{S}_d$

In this section we give several different descriptions of $\mathcal{S}_d$: as a CW complex, as a limit of quasiprojective varieties, as a (singular) iterated fiber bundle, and so on. These descriptions are quite explicit. One application is that we can write down a simple (recursive) description of $\pi_1(\mathcal{S}_d)$.

5.3.1. A CW complex structure on $\mathcal{S}_d/\mathbb{R}$. Since $\mathcal{S}_d$ is homeomorphic (in fact, biholomorphic) to $\mathcal{FS}_d$, we can study the topology of the former by an analysis of the latter. This is very convenient, because there is a natural CW complex structure on the quotient $\mathcal{FS}_d/\mathbb{R}$ by the orbits of the squeezing flow.

A (formal) shift f in $\mathcal{FS}_d$ is encoded by a certain lamination $\mathcal{L}_f$ of the disk, together with an assignment of numbers (altitudes) to the leaves, and such that all but finitely many altitudes are bigger than any given positive ϵ; pictorially we denote this by extending a leaf of altitude v radially beyond the disk so that the endpoints have absolute value e^v.

In fact, we don't need the entire lamination $\mathcal{L}_f$ to recover f. Call a precritical leaf *intermediate* if its altitude is at least as big as some critical leaf; all other precritical leaves are *short*. Given a choice of critical leaves, the precritical leaves are uniquely determined by the no linking property, so technically all we need to keep track of are the critical leaves. However, the location of these critical leaves depends on continuous moduli (v and θ), and the combinatorics of $\mathcal{L}$ depends in a complicated (discontinuous) way on these moduli. Understanding the discontinuities in the combinatorics of $\mathcal{L}$ is tantamount to understanding the topology of $\mathcal{FS}_d$.

Let's suppose we're given $\mathcal{L}_f$. We restrict attention to the sublamination consisting only of the critical leaves and the intermediate precritical leaves. This is a finite set of leaves. We call this subset the *intermediate lamination*.

DEFINITION 5.3.1. Let $v_1 \le v_2 \le \cdots \le v_{d-1}$ be the critical altitudes. The *height* is the vector of *positive* integers $h_1, h_2, \cdots, h_{d-2}$ where h_j is the smallest integer *strictly* bigger than $\log(v_{d-1}/v_j)/\log d$.

The dynamics $z \to z^d$ multiplies altitudes v by d, so the height measures the number of orbits of the dynamical map separating the altitudes of a pair of critical points.

The number of leaves in the intermediate lamination is $\sum_j (1-d^{h_j})/(1-d)$ of which $d-1$ are critical. Height is invariant under the squeezing flow. There is a partial order $\prec$ on the set of heights given by the termwise order.

Among the set of values of v and θ achieved by the (semi)-leaves in the intermediate lamination there are equalities and inequalities. The *signature* of p, denoted $\sigma(p)$, is the information of the *order structure* on the set $\{v \text{ of semileaves}\}$, and the *circular order structure* on the set $\{\theta \text{ of semileaves}\} \cup 0$. We denote by $\mathcal{FS}_d(\sigma)$ the subspace of $\mathcal{FS}_d$ with a particular signature σ. We can and do think of σ as a refinement of h.

DEFINITION 5.3.2 (Primitive equality). A *primitive equality* between altitudes or arguments is one that is not a dynamical consequence of an equality between preimage leaves.

We denote the number of primitive equalities by $e(\sigma)$, or just e if σ is understood.

EXAMPLE 5.3.3.
(1) Two precritical leaves with the same image leaf will have equal altitudes; this is a non-primitive equality.
(2) If the argument of a critical semi-leaf is zero, it will have a preimage also with argument zero; the first equality is primitive but not the second.
(3) If two precritical leaves ℓ, ℓ' have the same altitude, the same is true of their forward images. If ℓ is critical but ℓ' is strictly precritical this is a primitive equality; otherwise not.

THEOREM 5.3.4 (CW complex structure). *For each σ in the image of $\mathcal{FS}_d$, the subset $\mathcal{FS}_d(\sigma)$ is invariant under the squeezing flow, and each is homeomorphic to a product $\mathbb{R} \times D^k$ where the first factor parameterizes squeezing orbits, and where D^k is an open disk of dimension $k = 2d-3-e$ for e the number of primitive equalities. Furthermore the quotients under the squeezing flow $\mathcal{FS}_d(\sigma)/\mathbb{R}$ are the open cells in a CW complex structure on $\mathcal{FS}_d/\mathbb{R}$.*

PROOF. The critical altitudes and arguments vary continuously in a stratum $\mathcal{FS}_d(\sigma)$, and each argument is either identically zero, or varies in $(0, 2\pi)$. The log of the ratio of any two critical altitudes is bounded from above by $\lceil\text{height}\rceil$ and is fixed by the action of the squeezing flow. Each primitive equality is either an equality of arguments or an equality of the ratio of altitudes (one is equal to the other times a power of d). Every orbit of the squeezing flow has a unique point for which the biggest critical altitude is equal to 1; therefore fixing $v_{\max} = 1$ gives a section from $\mathcal{FS}_d/\mathbb{R}$ to $\mathcal{FS}_d$. The altitude and argument coordinates on this section define a characteristic map from D^{2d-3-e} to $\mathcal{FS}_d(\sigma)/\mathbb{R}$ which evidently extends to the closed disk. $\square$

In particular, $\mathcal{FS}_d$ is homeomorphic to the product of $\mathbb{R}$ with a CW complex of dimension $2d-3$, and the subspace with height $\le n$ is homeomorphic to the product of $\mathbb{R}$ with a *finite* CW complex. Denote this complex X_d and denote its cells by $X_d(\sigma)$. Let $X_d(h)$ denote the subcomplex with heights bounded by h (in the partial order $\prec$).

5.3.2. Low dimensional cells of the dual complex. Since X_d is a manifold, and the cells are embedded balls, there is a *dual* CW structure Y_d, with cells of complementary dimension. The advantage of working with the dual complex is that we may compute π_1 from low dimensional cells (those of dimension ≤ 2) that correspond to cells of the original complex with low co-dimension. In other words, these correspond to signatures σ that are very *generic* — they satisfy at most two primitive equalities — corresponding pictures that are easy to draw and reason about. This leads to a natural groupoid presentation for π_1 whose generators and relations are very easy to understand and enumerate.

Let's give a precise description of the cells of Y_d of low codimension.

(1) Codimension 0: There are no primitive altitude or argument equalities; leaves are distinct and share no endpoints, arguments are all different from zero, and the altitudes are ordered with no equalities except that preimages of the same leaf have the same altitude.
(2) Codimension 1: The strata fall into three kinds:
    (a) a critical argument goes to zero;
    (b) a critical leaf increases or decreases altitude until its altitude is equal to that of another precritical leaf; or
    (c) one edge of a critical leaf bumps into an edge of a taller precritical leaf.
(3) Codimension 2: The strata fall into two kinds:
    (a) two codimension 1 conditions occur simultaneously; or
    (b) a critical point runs into a precritical point.

Dualizing gives the 2-skeleton of the dual complex, from which we can read off π_1.

5.3.3. Tuning. As everybody knows, baby Mandelbrots are born in cauliflowers [16]. Douady defined the operation of *tuning* ...

5.3.4. $X_d(1)$ is a $K(\mathcal{B}_d, 1)$. We denote by $X_d(1)$ the union of cells with height h a vector of 1s. This means exactly that the ratio of the critical altitudes is $< d$, or in other words that every critical leaf has altitude strictly bigger than every strictly precritical leaf.

THEOREM 5.3.5 ($X_d(1)$ is a configuration space). *There is a homeomorphism from $\mathbb{R} \times X_d(1)$ to the configuration space $\mathbb{C}\backslash(\mathbb{C}^d - \Delta)/S_d$ of unordered d-tuples of distinct complex numbers up to translation.*

This theorem can be proved directly by an elementary cut-and-paste argument. However we give a slightly more subtle argument which teases out the (asymptotic) holomorphic relationship between the two sides.

PROOF. Let Ω_f be the Riemann surface associated to a formal shift. Altitude defines a map $v : \Omega_f \to \mathbb{R}^+$, and we let $v_{\max}$ be the biggest critical altitude. Then all other critical altitudes are in the interval $(v_{\max}/d, v_{\max}]$. Let Ω_f^D be obtained from Ω_f by cutting out the subspace $v^{-1}(0, 1/v_{\max}]$ and gluing back d punctured round disks using θ coordinates to glue in their boundaries. Then Ω_f^D is conformally isomorphic to $\mathbb{C}$ minus d distinct complex numbers, and we get a map in one direction.

The squeezing flow acts on $X_d(1)$ by multiplying the critical v coordinates by a real number e^t. There is a *different* $\mathbb{R}^+$ action which *adds* t to the critical v coordinates. This is not well-defined on all of $\mathcal{S}_d$, but it is defined on $\mathbb{R} \times X_d(1)$. We call this action *additive squeezing*. It takes $\mathbb{R} \times X_d(1)$ properly inside itself.

Conversely, associated to any distinct d-tuple $\{z_j\}$ summing to zero we can form the polynomial $z \to z^d + a_2 z^{d-2} + \cdots$ with the z_j as roots. There is an action of $\mathbb{R}$ on the configuration space given by

$$\{z_1 \cdots z_d\} \to \{e^t z_1 \cdots e^t z_d\}$$

This $\mathbb{R}$ action is free, and for $t \gg 1$ its image lies in $\mathbb{R} \times X_d(1)$. Furthermore, the pushforward of this $\mathbb{R}$ action on the image converges (uniformly on compact subsets) to additive squeezing. Taking $v_{\max} \to \infty$ the compositions of these two maps converge to the identity on compact subsets and we can construct the desired homeomorphism as a limit. $\square$

Remark 5.3.6. Note that $(\mathbb{C}^d - \Delta)/S_d$ is isomorphic to the space of degree d monic polynomials with distinct roots and the quotient $\mathbb{C}\backslash(\mathbb{C}^d - \Delta)/S_d$ can be identified with the subspace of the form $z \to z^d + a_2 z^{d-2} + \cdots$. In this formalism, the homotopy class of

$$\mathbb{R} \times X_d(1) \subset \mathcal{S}_d \to \mathbb{C}\backslash(\mathbb{C}^d - \Delta)/S_d$$

is represented by the map that takes a shift polynomial f to the polynomial $f - z$ whose roots are the fixed points of f. This map is a holomorphic inclusion, and is homotopic (though not properly homotopic) to a homeomorphism.

For every compact subset K of $\mathbb{C}\backslash(\mathbb{C}^d - \Delta)/S_d$ there is an R so that for any $|\lambda| \geq R$ the map $f \to \lambda f$ embeds K holomorphically in $\mathbb{R} \times X_d(1) \subset \mathcal{S}_d$, and this map is a homotopy inverse to $f \to (f - z)$. For sufficiently big λ the Julia set of the polynomial λf is ϵ-close to $\{z_1, \cdots, z_d\}$ in the Hausdorff metric; this elegant observation is due to Oleg Ivrii [22]. Evidently any braiding of these roots in $\mathbb{C}$ may be accomplished by a loop in $X_d(1)$.

One definition of $\mathcal{B}_d$, the braid group on d strands, is that it is the fundamental group of $(\mathbb{C}^d - \Delta)/S_d$. It follows that $\pi_1(X_d(1))$ is isomorphic to $\mathcal{B}_d$. In fact, this isomorphism factors through the monodromy representation $\rho : \pi_1(\mathcal{S}_d) \to \mathcal{B}_{\mathcal{C}}$.

Corollary 5.3.7. *The monodromy representation $\rho : \pi_1(X_d(1)) \to \mathcal{B}_{\mathcal{C}}$ preserves a d-element coarsening D and the composition $\rho : \pi_1(X_d(1)) \to \mathcal{B}_{\mathcal{C}} \to \mathcal{B}_D$ is an isomorphism.*

5.3.5. Colored shift space. One key source of combinatorial complexity comes from the fact that there is no way to distinguish globally between different critical leaves. Because of this it's easier to work in a closely related space, the *colored shift space*.

Let $\Delta(f') \subset \mathcal{S}_d$ denote the collection of shift polynomials with a critical point of degree > 1, and let $\mathcal{N}_d := \mathcal{S}_d - \Delta(f')$ denote the complement. $\Delta(f')$ is the discriminant variety of the derivative f'. We call $\mathcal{N}_d$ the space of *nondegenerate shift polynomials*. There is a degree $(d-1)!$ regular cover $\hat{\mathcal{N}}_d$ of $\mathcal{N}_d$ whose elements correspond to nondegenerate shift polynomials together with a bijection (a *coloring*) of the critical points with the $\{1, \cdots, d-1\}$.

The *colored shift space* $\hat{\mathcal{S}}_d$ is the completion of $\hat{\mathcal{N}}_d$. It maps to $\mathcal{S}_d$ as a cover branched over $\Delta(f')$. An element of $\hat{\mathcal{S}}_d$ is a shift polynomial together with a bijection between the critical points and the elements in an equivalence relation on $\{1, \cdots d-1\}$ so that a critical point which is a root of f' of multiplicity j corresponds to a subset of cardinality j.

The discriminant $\Delta(f')$ is smooth and complex codimension 1 in $\mathcal{S}_d$ away from a subset of complex codimension 2. At a generic point on Δ two critical values coalesce. There are $\binom{d}{3}$ components of Δ corresponding to the distinct combinatorial ways the critical points can map to the critical values. These are parameterized by a choice of 3 preimages of a point in S^1 under the d-fold covering map.

In particular, $\pi_1(\mathcal{S}_d)$ is a quotient of $\pi_1(\mathcal{N}_d)$ by a subgroup normally generated by $\binom{d}{3}$ relations. These relations are of *braid type*.

5.3.6. $\hat{\mathcal{N}}_d$ as an iterated fibration. We want to argue that the space $\hat{\mathcal{N}}_d$ is an iterated fibration: since the critical points are (by fiat) distinct and labeled, we can just "put them in one at a time." Each successive choice is parameterized by the fiber of a tautological fibration, which is an explicit finite cover of a Riemann surface intermediate between $\mathbb{C}-\mathbb{D}$ and Ω_f, obtained by cut and paste along the preimages of some subset of the critical leaves.

We start with $\mathbb{C}-\mathbb{D}$. A choice of critical leaf can be made as follows. First choose a critical value, and then choose exactly two preimages (out of the d possibilities) to be the endpoints of the critical leaf.

The set of choices of pairs of preimages is a $\binom{d}{2}$-fold cover of the punctured disk. It's connected if $d = 2$ or $d = 3$ and disconnected otherwise. The choice of the first critical point determines a canonical Riemann surface Ω, a plane minus a Cantor set, the result of cut-and-paste and pullback as in § ??.

Next we choose the second critical point. Again, we first need to choose a critical value v_2. This critical value is a point in Ω. It's not completely arbitrary though: by fiat the critical value is not allowed to be equal to the first critical value we chose. So the second criticial value is a point in $\Omega - v_1$. Here's the important observation: the topology of the space of choices for v_2 is independent of the choice of v_1, and varies continuously as a function of it. In other words, there is a *fiber bundle* over the space of choices for c_1 (which is a finite cover of the punctured disk) and the fibers are homeomorphic to $\mathbb{C}-\mathcal{C}$. Having chosen v_2, choosing c_2 is equivalent to choosing a pair of preimages of v_2 in Ω, which again amounts to passing to a finite cover (connected if $d < 5$ but not otherwise). This procedure can be iterated.

5.3.7. Monodromy representation. To complete the description of $\hat{\mathcal{N}}_d$ as a fiber bundle we must give the monodromy representation; i.e. how the fundamental group of configurations for $v_1, \cdots, v_k$ acts on the k-th fiber Ω. Since Ω is a $K(\pi, 1)$, the homotopy type of the bundle is completely determined by the action of the fundamental group of the base on the fundamental group of the fiber, which is isomorphic to F_∞, the free group on countably infinitely many generators. The answer is rather simple and elegant: the monodromy of the generators is by powers of Dehn twists in a the (disjoint) loops arising in an iterated pants decomposition of Ω.

5.3.8. $\mathcal{S}_3$ as a link complement. Let's return to the special case of $X_3 := \mathcal{F}\mathcal{S}_3/\mathbb{R}$. This is a real 3-manifold. In degree 3, the height h is a single positive integer, and $X_3 = \cup_n X_3(n)$. Let's describe X_3 explicitly.

Let's call the two critical altitudes v_1 and v_2, normalized by squeezing so that $v_1 \le v_2 = 1$. The critical leaves ℓ_1 and ℓ_2 aren't distinguished globally, but at least when $v_1 < v_2$ we can label them unambiguously. Each critical leaf has two critical semi-leaves whose arguments differ by $2\pi/3$; we let θ_j be the 'bigger' of the two arguments in $\mathbb{R}/2\pi\mathbb{Z}$, in the sense that the other argument is equal to $\theta_j - 2\pi/3$. Since critical leaves don't intersect we have $\theta_1 \in [\theta_2 + 2\pi/3, \theta_2 + 4\pi/3]$.

LEMMA 5.3.8 (X_3 by cut and paste). *Let $v_1 : X_3 \to (0, 1]$ denote the shorter critical altitude, and for each $t \in (0, 1]$ let $\Theta(t) \subset X_3$ denote the preimage. Let Θ be the torus of pairs (θ_1, θ_2) where $\theta_2 \in \mathbb{R}/2\pi\mathbb{Z}$ and $\theta_1 \in [\theta_2 + 2\pi/3, \theta_2 + 4\pi/3]/\sim$.*
Then

1. *for $v_1 = 1$ the space $\Theta(1)$ is obtained from Θ by the relations $(\theta_1, \theta_2) \sim (\theta_2, \theta_1)$ and $(\theta_2 + 2\pi/3, \theta_2) \sim (\theta_2 + 4\pi/3, \theta_2)$;*
2. *for $1/3 < v_1 < 1$ the space $\Theta(v_1)$ is naturally identified with Θ by the critical arguments (θ_1, θ_2); and*
3. *for $3^{-1-n} < v_1 <= 3^{-n}$, $n > 0$, the space $\Theta(v_1)$ is obtained from Θ by successively cutting $\theta_2 =$ constant into intervals at the points $\theta_2/3^n + p2\pi/3^n$ and $\theta_2/3^n - 2\pi/3^{n+1} + p2\pi/3^n$ mod 2π for $0 \le p < 3^n$, and regluing these intervals up into disjoint circles.*

PROOF. This is completely formal, and comes from the CW complex structure on X_3. For each n the precritical leaves of ℓ_2 at depth n have arguments which are equal to $\theta_2/3^n$ and $(\theta_2 - 2\pi/3)/3^n$ mod $(2/3^n)\pi$. If we are in $X_3(n)$, equivalently if $3^{-1-n} < v_1 \le 3^{-n}$, then for each $m < n$ in turn, and for each fixed θ_2, we must cut the circle of θ_1 values $\theta_2 =$ constant along the precritical arguments at depth m and reglue, realizing the discontinuity of θ_2 when we push ℓ_1 over a bigger precritical leaf of ℓ_2. $\square$

THEOREM 5.3.9 (X_3 as a union of iterated cable complements). *Let L_1 denote the right handed trefoil in S^3. Let L_n be obtained from L_{n-1} by iterated cables as follows:*

1. *The components fall into* types, *indexed by powers of 2. The trefoil is of type 1.*
2. *The link L_n is obtained from L_{n-1} as follows: for every component K of L_{n-1} of type 2^k we keep the same component in L_n but change its type to 2^{k+1}, and we also add a $(-1, 2^k)$ cable of type 1.*

Let's think of each L_n as being contained in successive tubular neighborhoods $N(L_{n-1})$ of the previous L_{n-1}.
Then with this notation, $X_3(n) = S^3 - N(L_n)$, and $X_3 = S^3 - \cap N(L_n)$.

PROOF. This theorem is really just a restatement of Lemma 5.3.8 in 3-manifold terminology. Recall the notation Θ to denote the torus with parameters (θ_1, θ_2) where $\theta_2 \in \mathbb{R}/2\pi\mathbb{Z}$ and $\theta_1 \in [\theta_2 + 2\pi/3, \theta_2 + 4\pi/3]/\sim$. We think of Θ as being obtained from a pair of trapezes in the square $[0, 2\pi]^2$. The quotient $(\theta_1, \theta_2) \sim (\theta_2, \theta_1)$ identifies these two trapezes with one, and under the natural identifications this trapeze glues up to a Möbius band whose core is double-covered by the circle $\theta_1 = \theta_2 + \pi$, and whose boundary $\theta_1 = \theta_2 + 2\pi/3$ triple-covers its image. If we think of S^3 as the join $S^1 * S^1$ of these two singular circles, then the circles become the Hopf link, and the complement of $(\Theta/\sim)$ is (a neighborhood of) a $(3, 2)$ torus link — i.e. a right-handed trefoil. Thus $X_3(1)$ is homeomorphic to S^3 minus a (neighborhood of a) trefoil. The circles $\theta_2 =$ constant are meridians of the trefoil, and the circles $\theta_1 =$ constant are longitudes.

Inductively the meridians of L_n correspond to circles of length $2^k \cdot 3^{-n}\pi$ for type 2^k, obtained from segments in the circles $\theta_2 =$ constant by cutting along precritical arguments, gluing up the smallest segments into circles of type 1, and the longer segments into a single circle of type 2^{k+1}. As we wind around the θ_2 circle, the precritical arugments also wind positively, but each is 3 times as slow as the previous one. Thus for type 2^k it gives rise to a core of type 2^{k+1} and a cable $(-1, 2^k)$ torus knot of type 1. Figure 5.5 depicts the links L_j for $j = 1, 2, 3$. □

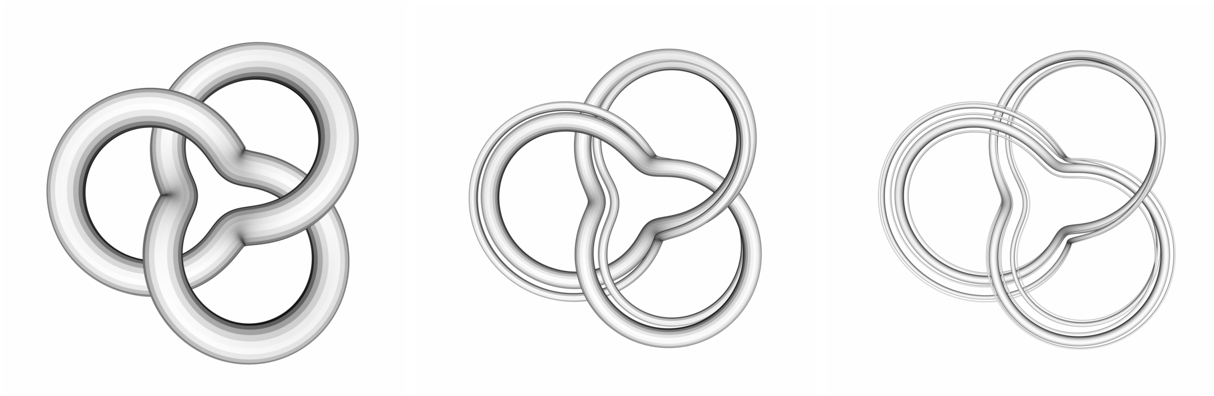

FIGURE 5.5. $X_3(j)$ for $j = 1, 2, 3$ are the complements of the links L_k in S^3.

5.3.9. Higher degree. We're now in a position to give a completely explicit description of $\hat{\mathcal{S}}_d$ and $\mathcal{S}_d$, and associated CW complexes $\hat{X}_d$ and X_d.

The shift space $\mathcal{S}_d$ is a quotient of $\hat{\mathcal{S}}_d$ by the symmetric group S_{d-1}. The squeezing flow commutes with this action, and we can write $\hat{\mathcal{S}}_d = \hat{X}_d \times \mathbb{R}$ where $\hat{X}_d$ is a locally finite CW complex mapping cellularly to X_d.

5.3.9.1. *The altitude plane.* The altitudes v_j, being the logarithms of the absolute value of the critical Böttcher coordinates, are well-defined positive functions on $\hat{\mathcal{S}}_d$, and their logarithms are the coordinates of a surjection $\log_d v : \hat{\mathcal{S}}_d \to \mathbb{R}^{d-1}$.

Note that we use $\log_d$ — i.e. the logarithm in base d — so that the natural stratification by singularities of $\log v$ is concentrated on hyperplanes with integral coordinates. Anyway, we suppress the subscript in the sequel.

The squeezing flow factors through $\log v$, and translates $\mathbb{R}^{d-1}$ in the direction of the vector $(1, 1, \cdots, 1)$. It's convenient to normalize coordinates so that $\log v : \hat{X}_d \to \mathbb{R}^{d-2}$ maps to the subspace with coefficients summing to zero (equivalently, the product of the v_j is 1). We call $\mathbb{R}^{d-2}$ the *altitude plane*.

In these coordinates we may identify $\mathbb{R}^{d-2}$ as the root space of the Lie algebra $\mathfrak{sl}_{(d-1)}$ and S_{d-1} as the Weyl group. Passing from $\hat{X}_d$ to X_d corresponds to quotienting out $\mathbb{R}^{d-2}$ by the Weyl group, so we can think of $\log v$ as a map from X_d to a Weyl chamber $v_1 \le v_2 \le \cdots \le v_{d-1}$ whose faces have been *orbifolded*; i.e. they are decorated by the stabilizer subgroup.

5.3.9.2. *Triangulation by walls.* Let Φ^* be the dual root lattice, and let H be the system of hyperplanes ('walls') defined by $\phi^* \in \mathbb{Z}$ for some (dual) root ϕ^*. The hyperplanes in H fall into $(d-1)(d-2)/2$ distinct parallelism classes, and each hyperplane corresponds to an equality of the form $\log v_j - \log v_k = n$ for some $j > k$ and some integer n. We can write such a hyperplane $H_{j,k}(n)$; the $H_{j+1,j}(0)$ cut out the walls of the Weyl chamber.

The hyperplanes $H_{j,k}(n)$ cut up the Weyl chamber $v_1 \leq \cdots v_{d-1}$ into simplices, the cells in a CW decomposition Z_d, and the map $\log v : \hat{X}_d \to Z$ is cellular. Each simplex of Z_d is defined by a finite number of equalities of the form $\log v_j - \log v_k = n$ and inequalities of the form $n < \log v_k - \log v_j < n + 1$. The preimage of each cell is evidently a subcomplex of $\hat{X}_d$.

5.3.9.3. *Fibers of singular tori.* Over each open simplex σ of Z the preimage in $\hat{X}_d$ is a product of σ with a space that may be thought of as a mildly singular flat $(d-2)$-torus. This latter space is itself a $(d-2)$-fold iterated fiber bundle. If we denote the total space E_1 we can write $F_1 \to E_1 \to E_2$ and successively $F_j \to E_j \to E_{j+1}$ until we get to $F_{d-2} \to E_{d-2} \to B$. The base space B is always a circle of length 2π which we may canonically identify with $\mathbb{R}/2\pi\mathbb{Z}$ parameterized by the θ_{d-1} coordinate. Likewise, each F_j is 'parameterized' by the θ_j coordinate, and is obtained from a circle of length $2j\pi/d$ by iterated operations we call *pinch* and *cut*. These operations occur when passing through the walls $H_{j,k}(n)$.

When no $\log v_k - \log v_j$ is an integer but $\log v_{d-1} - \log v_j < n$, the jth fiber F_j is a disjoint union of circles each with length of the form $2p\pi/d^n$ for various integers p, and summing to $2j\pi/d$.

In some family where $t := \log v_k - \log v_j$ increases from $n - \epsilon$ to $n + \epsilon$ some components of the fiber F_j *pinch* at $t = n$ to real cacti (i.e. a tree of circles glued at isolated points), and when $t > n$ these cacti are *cut* apart into their constituent circles.

5.3.9.4. *Wall complements and recursive factorization.* We've seen that when we pass through a wall $H_{j,k}(n)$ the fibers of $\log v$ become disconnected. The 'big' wall of the Weyl chamber is $H_{d-2,d-1}(0)$, i.e. the one where $v_{d-2} = v_{d-1}$. Its preimage disconnects $\hat{X}_d$ (and X_d for that matter); the complement is where v_{d-1} is the unique biggest altitude. This complement has $\lfloor d/2 \rfloor$ components, corresponding to the difference in the two arguments of the tallest critical leaf ℓ_{d-1}; this difference is of the form $p2\pi/d$ for some $0 < p < d$ and there is no distinction between p and $d - p$. Denote these components $X'_d(p)$.

LEMMA 5.3.10 (Wall complements are fiber products). *There is a natural homeomorphism from $X'_d(p)$ to the fiber product $X'_{p+1}(1) \times_{S^1} X'_{d-p+1}(1)$ respecting the combinatorial structure, where the S^1 action on either side acts fiberwise (as above) by rotating each circle in the base space at constant speed.*

Thus, for instance, $X'_4(2) = X'_3(1) \times_{S^1} X'_3(1) = X_3 \times_{S^1} X_3$.

PROOF. All the combinatorial complexity of each X_d comes from the way in which fibers pinch and break as we pass through hyperplanes $\phi \in \mathbb{Z}$. At the level of laminations, this occurs when some critical leaf collides with a precritical leaf at the same height. But in the stratum $X'_d(p)$ the big critical leaf ℓ_{d-1} separates $p - 1$ of the smaller critical leaves from the other $d - p - 1$ of them, and inductively, the preimages of the big critical leaf separate $p-1$ of the smaller critical leaves from the preimages of the other $d-p-1$ of them. There are no interactions, and the two factors decompose into an independent $X'_{p+1}(1)$ and $X'_{d-p+1}(1)$. The circle action on either side comes from rotating the same θ_{d-1} argument, so we obtain an identification of the given stratum with the fiber product. $\square$

It follows recursively that there is only one really 'new' component of X_d for each d, namely $X'_d(1)$, where the arguments of the biggest critical leaf differ by exactly $2\pi/d$.

There is an analog of Lemma 5.3.10 for every $H_{j,k}(n)$ ...

5.3.10. $\mathcal{S}_d$ is a $K(\pi,1)$. From the explicit description of X_d we can immediately deduce:

THEOREM 5.3.11 ($\mathcal{S}_d$ is a $K(\pi,1)$).

PROOF. We first show that $\hat{\mathcal{S}}_d$ is a $K(\pi,1)$. This follows from the iterated singular fibration picture for $\hat{X}_d$, which should be thought of as a kind of 'horotorus slice' of $\hat{\mathcal{S}}_d$. Away from the singular hyperplanes the space $\hat{X}_d$ looks like a bundle over $\mathbb{R}^{d-1}$ with fibers which are disjoint unions of Euclidean tori, and the metric is locally flat. Passing through a critical hyperplane causes the fibers to pinch along parallel subspaces, which are then cut open. At the level of the universal cover this can be accomplished by taking a collection of parallel codimension two Euclidean subspaces, and for each cutting out the product with a ray and inserting a new Euclidean halfspace. The codimension two subspaces become singular in the new metric with cone angle 3π.

Another topologically equivalent but less obviously geometric way to see this is to consider the iterated fibration structure on $\hat{\mathcal{N}}_d$ discussed in § 5.3.6 and apply the long exact sequence in homotopy, since the fibers are all surfaces of infinite type.

Passing from $\hat{\mathcal{S}}_d$ to $\mathcal{S}_d$ is achieved by quotienting by the action of the symmetric group. This acts in the Euclidean space model in the standard way — i.e. by identifying the Euclidean space with $\mathbb{C}^{d-1}$, and quotienting out by the permutation action of S_{d-1} on the coordinates. Topologically, the permutation action is the same before and after taking branched covers and lifting to the universal cover: the singular locus for the permutation action is a union of linear subspaces, and when we take double branched covers these are covered by Euclidean spaces in the same topological configuration.

The usual observation that the quotient of $\mathbb{C}^{d-1}$ by S_{d-1} is homeomorphic (actually biholomorphic!) to $\mathbb{C}^{d-1}$ completes the proof. $\square$

5.4. Algebraic geometry

The Shift space $\mathcal{S}_d$ embeds naturally as an open subset of $\mathbb{C}^{d-1}$, and is therefore a complex analytic manifold in a natural way. However, this embedding is highly transcendental. It turns out that there is a more algebraic way to see $\mathcal{S}_d$, at least up to homeomorphism. In this section we'll see that each $X_d(h)\times\mathbb{R}$ is homeomorphic to a quasiprojective variety, and these quasiprojective varieties arise in a natural way as subsets from an algebraic compactification (com-cacti-fication?) of shift space by *cactus shifts*.

5.4.1. Cactus shifts. Spaces of rational maps can be naturally compactified by dynamics on nodal curves called *cacti*. This is analogous to the Deligne-Mumford compactification of moduli spaces of Riemann surfaces. These cacti are typically not compact, nor are their moduli. However the space of cacti is naturally exhibited as an increasing union of pieces, each of which is a quasiprojective variety, and (although reducible) each finite union can be deformed (not canonically) to an irreducible quasiprojective variety homeomorphic to some $X_d(h)\times\mathbb{R}$.

Let's start with the definition of a cactus.

DEFINITION 5.4.1 (Cactus). A *cactus* K is a locally finite rooted tree of $\mathbb{CP}^1$s. This means that

(1) there is a rooted tree T, and
(2) for each vertex $v \in T$ there is a copy of $\mathbb{CP}^1$ which we denote $\mathbb{CP}^1_v$, and
(3) for each edge $e \in T$ with parent v there is a distinguished value $z(e) \in \mathbb{CP}^1_v - \infty$ with distinct edges e_j beginning at v corresponding to distinct $z(e_j)$ in $\mathbb{CP}^1_v$.

We think of the union of $\mathbb{CP}^1$s as being glued together into a (locally finite) nodal curve, where for each edge e with parent v and child w we attach $\infty \in \mathbb{CP}^1_w$ to $z(e) \in \mathbb{CP}^1_v$.

So much for cacti. Now let's talk about dynamics.

DEFINITION 5.4.2 (Cactus polynomial). Let K be a cactus with associated tree T. A *cactus polynomial* p *of degree* d is the following data:

(1) for every vertex v there is a polynomial p_v of the form $p_v : z \to z^n + a_2 z^{n-2} + a_3 z^{n-3} + \cdots a_n$ (i.e. the polynomial is in Böttcher form);
(2) there is a simplicial map of T. It collapses each root edge to the root vertex, and is injective on every other edge.
(3) if v is the root vertex, the degree of p_v is d and the distinguished values in $\mathbb{CP}^1_v$ are the roots of p_v;
(4) for every vertex v the sum of the degrees of the p_w over the preimages w of v is equal to d; and
(5) for every vertex v and every preimage w the distinguished values in $\mathbb{CP}^1_w$ are exactly the preimages of the distinguished values in $\mathbb{CP}^1_v$ under p_w.

A critical point of some p_w is *genuine* if it is not equal to a distinguished value. A cactus polynomial of degree d is a *shift* if it has exactly $d-1$ genuine critical points, counted with multiplicity.

The p_v define a map $p : K \to K$ as follows. If e is an edge with parent v and child w then we think of p_w as a map from $\mathbb{CP}^1_w$ to $\mathbb{CP}^1_v$, so that by bullet (4) the map p takes distinguished values to distinguished values. In the special case that v is the root vertex, we let $p|\mathbb{CP}^1_v$ be the constant map to $\infty \in \mathbb{CP}^1_v$.

With this convention p determines an honest map from K to itself. The dynamics is very simple: ∞ in the root node is fixed, and every other point converges to ∞ in finite time. More interesting is the action of p on the space of ends of K, which is (except in some very degenerate cases) a Cantor set on which p acts as a quotient of a shift; this Cantor set is, if you like, the Julia set of p.

Every cactus polynomial of degree d has at most $d-1$ critical points counted with multiplicity. We may *extend* a cactus polynomial by adding a new root vertex, an edge from the new root to the old, putting the unique distinguished value in the new root $\mathbb{CP}^1$ at 0, and defining p_v on the new root $\mathbb{CP}^1_v$ to be the polynomial $z \to z^d$. A cactus polynomial is *reduced* if there are at least two distinguished values in the root $\mathbb{CP}^1$, equivalently if for v the root vertex, the polynomial has at least two roots. Every cactus polynomial is obtained from a reduced one in a unique way by iterated extension, except for the unique cactus polynomial in any fixed degree whose underlying tree is an infinite ray.

The *critical subcactus* is the cactus associated to the minimal (finite) subtree of T containing all the genuine critical points.

The cactus polynomials as we define them above are the *nondegenerate* ones. Define the class of *generalized cacti* by allowing the possibility that more than one child $\mathbb{CP}^1$ shares the same distinguished value in a parent $\mathbb{CP}^1$.

For a degenerate cactus polynomial if p is a distinguished value of multiplicity n we also think of it as a genuine critical point of multiplicity $n-1$.

Definition 5.4.3 (height). The *height* of a cactus polynomial is the vector whose components are equal with multiplicity to 1 plus the combinatorial distance (measured in number of edges) from the genuine critical points to the root. If a cactus polynomial is not a shift, this vector is padded by a $*$ symbol to have $d-1$ entries.

For a generalized cactus polynomial we add 1 to the distance for those genuine critical points that occur as distinguished values of multiplicity at least 2.

We define a partial order on heights componentwise, except that $*$ is considered incomparable. With this convention,

Lemma 5.4.4. *For any h the set of generalized cactus polynomials of degree d and height $\le h$ is a (highly reducible) quasiprojective variety of (complex) dimension $d-1$, and the subset of nondegenerate cactus polynomials is an open dense subvariety.*

Proof. The key point is that once we have the finite subset of the cactus containing all the genuine critical points, the remainder of the cactus is completely determined. Each genuine critical point contributes one degree of freedom, and the total number is $d-1$. Since the height is bounded there are only finitely many combinatorial possibilities for the critical subcactus. ☐

In the sequel we shall prove the following theorem:

Theorem 5.4.5 (Quasiprojective). *Let $\mathcal{CS}_d(h)$ denote the variety of generalized cactus shifts of degree d and height $\le h$. Then $\mathcal{CS}_d(h)$ can be deformed to a nonsingular quasiprojective variety homeomorphic to $X_d(h) \times \mathbb{R}$.*

5.4.2. Examples in low degree. Let's work out some explicit examples in low degree. We restrict attention to reduced cacti, so that the root $\mathbb{CP}^1$ must contain at least one critical point.

Example 5.4.6 (Degree 2). A reduced cactus K of degree 2 has root polynomial $p_v : z \to z^2 + c$ for some nonzero c. The distinguished points are $\pm\sqrt{c}$ and at each of these points the subcactus is a copy of K (whose underlying tree is, by induction, a two-valent rooted tree), and every child polynomial p_w is the identity map $z \to z$. Thus every reduced cactus of degree 2 is a shift, and the space of reduced cacti is $\mathbb{C}^*$.

Example 5.4.7 (Degree 3). A reduced cactus K of degree 3 has root polynomial $p_v : z \to z^3 + az + b$ where at least one of a or b is nonzero. The height is 1 if and only if p_v has three distinct roots; equivalently if and only if its discriminant $-4a^3 - 27b^2$ is nonzero. Let Δ be the discriminant locus in $\mathbb{C}^2$ with coordinates a, b.

If the discriminant of p_v is zero, p_v has two distinct roots α and β where β is a double root and $\alpha + 2\beta = 0$ so that $\alpha = -\beta/2$. These two distinguished values are associated to two child vertices x and y where p_x is the degree one map $p_x : z \to z$ and p_y is a degree two map $p_y : z \to z^2 + c$. The distinguished values in $\mathbb{CP}^1_y$ are the preimages of α and β under p_y. The height is 2 if and only if α and β are regular values, equivalently if c is not equal to α or β. Thus we can think of the space of reduced cacti of degree 3 and height exactly 2 as a bundle over $\Delta - 0$ whose fibers are copies of $\mathbb{C}$ − two points.

If the discriminant of p_v is zero there are two ways for p_y to be degenerate: either $p_y : z \to z^2 + \alpha$ or $p_y : z \to z^2 + \beta$. In either case there are three distinguished points in $\mathbb{CP}^1_y$. In the first case these are at 0 and at $\pm\sqrt{\beta - \alpha}$, and in the second case these are at 0 and at $\pm\sqrt{\alpha - \beta}$. Two of these distinguished points are attached to a $\mathbb{CP}^1$ which maps by degree 1, and zero is attached to a $\mathbb{CP}^1$ which maps by a degree 2 map $z \to z^2 + d$. This special $\mathbb{CP}^1$ has distinguished points, at the square roots of $-d$ and $\pm\sqrt{\alpha - \beta} - d$ in the first case and $\pm\sqrt{\beta - \alpha} - d$ in the second case. The height is 3 if and only if these square roots are all distinct.

It's clear how to continue this pattern. The cactus maps of a particular height are the complement of a certain discriminant, and those of the next height are a bundle over this discriminant with fiber a copy of the complement of the distinguished points at the previous level. Over every point in $\Delta - 0$ the union of the shift cacti is itself an infinite cactus ...

5.4.3. Proof of Theorem 5.4.5. We now give the proof of Theorem 5.4.5.

Proof. In fact, the correspondence between $\mathcal{CS}_d$ and $\mathcal{FS}_d$ is quite direct. A reduced cactus K admits a pair of measured singular foliations defined as follows. For the root vertex v the foliations of $\mathbb{CP}^1_v$ are the preimage under p_v of the (singular) foliations of $\mathbb{CP}^1$ by lines of constant argument, and circles of constant absolute value; and then for every non root vertex w which is a child of v we define the foliations on $\mathbb{CP}^1_w$ to be the preimage of the foliations on $\mathbb{CP}^1_v$ under p_w. These foliations are singular on each $\mathbb{CP}^1$ at infinity, at the distinguished points, and at the genuine critical points. If we remove infinity and the distinguished points we get a pair of foliations of an open planar surface that can be completed to compact surfaces by adding circle leaves as boundaries. These foliations give altitude and argument functions on each subsurface. Reparameterize log of the altitude function by a homeomorphism $\mathbb{R} \to (0, 1)$ (for instance, the normal cumulative distribution function is a natural choice) and glue these subsurfaces together with the added boundary. The result is an infinite type surface Ω together with a single self-map $p : \Omega \to \Omega$ which is evidently a formal shift; reversing this procedure takes a formal shift to a cactus shift providing none of the critical values have log altitude equal to an integer.

This shows a few things: firstly that each end of each component of the space of cactus maps is a product, homeomorphic to the corresponding end of the component it attaches to, and that the result of gluing the components together by this product structure is homeomorphic to $\mathcal{FS}_d$. It remains to show that this topological gluing can be accomplished algebraically by a deformation. $\square$

5.5. Monodromy

5.5.1. Representations to braid groups. For f in the Shift locus, the action of f on J_f is uniformly expanding, and the dynamics is structurally stable in a neighborhood of J_f. We therefore get natural representations of $\pi_1(\mathcal{S}_d)$ to various braid groups. Here is a (non-exhaustive) list:

(1) the representation to $\mathcal{B}_\mathcal{C}$ given by the braiding of J_f in $\mathbb{C}$;
(2) the representation to $\mathcal{B}_{d^n}$ given by the braiding of the periodic orbits of J_f of period dividing n in $\mathbb{C}$;
(3) for every periodic point p of period dividing n, the kernel of $\pi_1(\mathcal{S}_d) \to \mathcal{B}_{d^n} \to S_{d^n}$ admits a representation to $\mathcal{M}(T^2 - \mathcal{C})$ given by the braiding of $(J_f - p)/\langle f^n \rangle$ in the torus $(\mathbb{C} - p)/\langle f^n \rangle$; ...

5.5.2. The braid group of d^n-coarsenings.

5.5.3. The monodromy representation is injective. We are now in a position to prove:

THEOREM 5.5.1 (Monodromy is injective). *The monodromy representation $\rho : \pi_1(\mathcal{S}_d) \to \mathcal{B}_\mathcal{C}$ is injective.*

5.5.4. Landing rays and the action on the Ray graph. For a shift polynomial f the *landing rays* are proper rays from J_f to infinity which are the images of the vertical segments $\theta =$ constant in Böttcher coordinates containing no critical or precritical leaf. As we move around in shift space, the isotopy class of a landing ray will change when it passes over a critical leaf. Nevertheless there is a coarse map from the universal cover of $\mathcal{S}_d$ to the Cantor ray graph $\mathcal{R}_\mathcal{C}$.

Of course, this is an orbit map for the monodromy action of $\mathcal{B}_\mathcal{C}$. But the point is we don't need to *choose* an orbit — the complex structure picks one out for us canonically, and there is a close connection between the holomorphic geometry of a family in $\mathcal{S}_d$ and the metric geometry of $\mathcal{R}_\mathcal{C}$.

5.6. Real Shift locus

A real degree d polynomial $f(z) = a_0 z^d + a_1 z^{d-1} + \cdots + a_d$ is always conjugate to a polynomial of the form $z^d + O(z^{d-2})$, but this polynomial is not necessarily real, because a_0 may not have a real $(d-1)$st root. If d is even, then of course any real a_0 has a (unique) real $(d-1)$st root; if d is odd, then either a_0 or $-a_0$ has such a root. Thus for d even it makes sense to define $\mathcal{S}_d(\mathbb{R})$ to be the space of real polynomials of the form $z \to z^d + O(z^{d-2})$, whereas for d odd we define $\mathcal{S}_d^{\pm}(\mathbb{R})$ to be the space of real polynomials of the form $z \to \pm z^d + O(z^{d-2})$ respectively. We denote the image of these subspaces in $\mathcal{FS}$ by $\mathcal{FS}_d(\mathbb{R})$ for d even and $\mathcal{FS}_d^{\pm}(\mathbb{R})$ for d odd.

The Böttcher map commutes with complex conjugation, because this is evidently true near infinity. A polynomial is real if and only if it is invariant under complex conjugation. It follows that a formal shift is in $\mathcal{FS}_d(\mathbb{R})$ if and only if the lamination $\mathcal{L}$ is invariant under complex conjugation. This means we can read off the topology of $\mathcal{FS}_d(\mathbb{R})$ from the combinatorics of symmetric laminations in the Böttcher model.

5.6.1. Examples.

EXAMPLE 5.6.1 (Degree 2). Let $\mathbb{R}$ be the space of degree two polynomials $z \to z^2 + c$ with c real. $\mathcal{M}_2(\mathbb{R})$ is equal to the interval $[-2, 1/4]$, and $\mathcal{S}_2(\mathbb{R})$ is the union of two open rays.

Example 5.6.2 (Degree 3).

5.6.2. Local connectivity.

CHAPTER 6

Linear IFS

For $1 \leq j \leq k$ pick complex numbers α_j, β_j with $0 < |\alpha_j| < 1$ and define $f_j : \mathbb{C} \to \mathbb{C}$ by $f_j(z) = \alpha_j z + \beta_j$. The dynamical system $(\mathbb{C}, f_j)$ is usually called a (complex linear) iterated function system, or IFS for short.

Because the generators are uniform contractions, there is a continuous endpoint map $e : \mathcal{E} \to \mathbb{C}$ with image the attractor Λ. We say $(\mathbb{C}, f_j)$ is *Schottky* if $e : \mathcal{E} \to \Lambda$ is a homeomorphism.

CHAPTER 7

Holomorphic foliations

7.1. Holomorphic vector fields

Let V be a complex surface and let X be a holomorphic vector field on V. The integral curves of X define a one (complex) dimensional singular foliation $\mathcal{F}$. This foliation is singular at the zeroes of X which form a subvariety $\Sigma \subset V$.

Let Ω be the open subset of the projective space $PH^0(V;\Theta)$ consisting of vector fields for which $\mathcal{F}$ contains an exceptional minimal set.

CHAPTER 8

Schottky groups

8.1. Handlebodies

CHAPTER 9

Roots

Acknowledgments

Danny Calegari was supported by NSF grant DMS 1405466.

Bibliography

[1] J. Bavard, *Hyperbolicité du graphe des rayons et quasi-morphismes sur un gros groupe modulaire*, Geom. Topol. **20** (2016), no. 1, 491–535

[2] J. Bavard, D. Calegari and A. Walker, *Notes on the Cantor Mapping Class Group*, preprint, 2015

[3] J. Bavard, D. Calegari, Y-M. He, S. Koch and A. Walker, *Building Cantor Julia sets and the topology of the shift locus*, preprint, in preparation

[4] J. Bavard and A. Walker, *The Gromov boundary of the ray graph*, TAMS, **370** (2018), no. 11, 7647–7678

[5] J. Bavard and A. Walker, *Two simultaneous actions of big mapping class groups*, preprint, `https://arxiv.org/abs/1806.10272`

[6] R. Bell and D. Margalit, *Braid groups and the co-Hopfian property*, J. Algebra **303** (2006), 275–294

[7] P. Blanchard, R. Devaney and L. Keen, *The dynamics of complex polynomials and automorphisms of the shift*, Invent. Math. **104** (1991), no. 3, 545–580

[8] B. Branner and J. Hubbard, *The iteration of cubic polynomials. I. The global topology of parameter space*, Acta Math. **160** (1988), no. 3-4, 143–206

[9] B. Branner and J. Hubbard, *The iteration of cubic polynomials. II. Patterns and parapatterns*, Acta Math. **169** (1992), no. 3-4, 229–325

[10] R. Brooks and J. Matelski, *The dynamics of 2-generator subgroups of* $\mathrm{PSL}(2,\mathbb{C})$, Riemann surfaces and related topics: Proc. 1978 Stony Brook Conference, 65–71 Ann. of Math. Stud. **97**, Princeton Univ. Press, Princeton N.J., 1981

[11] D. Calegari, *Foliations and the geometry of 3-manifolds*, Oxford Mathematical Monographs. Oxford University Press, Oxford, 2007.

[12] D. Calegari, S. Koch and A. Walker, *Roots, Schottky semigroups and a proof of Bandt's Conjecture*, ETDS, **37** (2017), no. 8, 2487–2555

[13] L. DeMarco, *Combinatorics and topology of the shift locus*, in Conformal Dynamics and Hyperbolic Geometry, AMS Contemp. Math. **573** (2012), 35–48

[14] L. DeMarco and C. McMullen, *Trees and the dynamics of polynomials*, Ann. Sci. Éc. Norm. Supér. (4) **41** (2008), no. 3, 337–382

[15] L. DeMarco and K. Pilgrim, *Polynomial basins of infinity*, Geom. Func. Anal. **21** (2011), 920–950

[16] A. Douady, X. Buff, R. Devaney and P. Sentenac, *Baby Mandelbrot sets are born in cauliflowers*, The Mandelbrot set, theme and variations, 19–36, LMS. Lec. Note Ser. **274**, Cambridge Univ. Press, Cambridge, 2000

[17] A. Douady and J. Hubbard, *Itération des polynômes quadratiques complexes*, C. R. Acad. Sci. Paris Sér. I Math. **294** (1982), no. 3, 123–126

[18] D. Eisenbud and W. Neumann, *Three-dimensional link theory and invariants of plane curve singularities*, Ann. Math. Stud. **110**, Princeton Univ. Press, Princeton N.J., 1985

[19] B. Farb and D. Margalit, *A Primer on Mapping Class Groups*, Princeton University Press

[20] D. Gabai and U. Oertel, *Essential laminations in 3-manifolds*, Ann. of Math. (2) **130** (1989), no. 1, 41–73

[21] S. Hensel, P. Przytycki, and R. Webb, 1-*slim triangles and uniform hyperbolicity for arc graphs and curve graphs*, JEMS **17** (2015), 755–762

[22] O. Ivrii, *Approximation by Cantor Sets Julia Sets*, preprint, `http://www.its.caltech.edu/ ivrii/CantorApproximation.pdf`

[23] T. Lei, *Similarity between the Mandelbrot set and Julia sets*, Comm. Math. Phys **134** (1990), no. 3, 587–617

[24] A. Malcev, *On the faithful representation of infinite groups by matrices*, Mat. Sb. **8** (1940), 405–422

[25] H. Masur and S. Schleimer, *The geometry of the disk complex*, Jour. AMS **26** no. 1, (2013), 1–62

[26] C. McMullen, *Complex Dynamics and Renormalization*, Ann. Math. Stud. **135**, Princeton Univ. press, Princeton N.J., 1994