THE CLASSIFICATION
OF CRITICALLY PREPERIODIC POLYNOMIALS
AS DYNAMICAL SYSTEMS

BEN BIELEFELD, YUVAL FISHER, AND JOHN HUBBARD

CONTENTS

1. INTRODUCTION

The object of this paper is to classify all polynomials p with the properties that all critical points of p are strictly preperiodic under iteration of p. We will also characterize the Julia sets of such polynomials.

In this section we will motivate the constructions of this paper by looking carefully at the quadratic case and relating it to the theory of kneading sequences.

Symbolic dynamics of unimodal mappings: kneading sequences. The combinatorial approach to the dynamics of mappings, which this paper develops, starts with kneading sequences, as developed in [MT] and [CE]. Choose $a < c < b \in \mathbb{R}$, and set $I = [a, b]$ and consider first unimodal maps, which we will take to mean continuous mappings $f: I \to I$ satisfying

(1) $f(a) = f(b) = b$;
(2) f is monotone decreasing (or increasing) on $[a, c]$;
(3) f is monotone increasing (or decreasing, respectively) on $[c, b]$.

We apologize to readers used to unimodal maps with maxima; monic polynomials are best adapted to our purposes. Thus typical unimodal mappings are elements of the quadratic family $p_c(x) = x^2 + c$ with the interval

$$I = [-(1 + \sqrt{1 - 4c})/2, (1 + \sqrt{1 - 4c})/2]$$

and $c = 0$.

In this case, the kneading sequence of a point $x \in I$ is the sequence $S_f(x) = (s_0, s_1, \ldots)$ with $s_i \in \{P, R, C\}$, determined by whether $f^i(x)$ is on the orientation-preserving (i.e., increasing) side of the "critical point" c, the orientation-reversing (i.e., decreasing) side of the "critical point," or on the "critical point," respectively.

The central questions are then:

1.1. When is x determined by the sequence $S_f(x)$?;

1.2. For what families of unimodal maps is f determined by $S_f(c)$? More specifically, is a quadratic polynomial determined by the symbolic sequence of the critical point?

These questions and their relatives are the key questions of this paper, and in fact of the whole subject; we answer them completely for preperiodic polynomials.

Example. Consider the graphs of unimodal maps in Figure 1.3, showing the forward orbit of the critical point. Both candidate mappings are unimodal

$$x_1 \qquad x_3 = x_4 \qquad x_2 \qquad\qquad x_1 \quad x_4 \quad x_3 = x_5 \qquad x_2$$
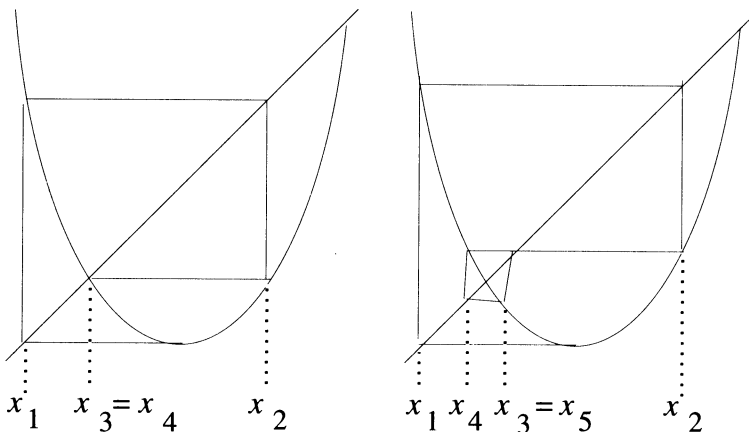
FIGURE 1.3. Two unimodal mappings with $x_{i+1} = f(x_i)$.

maps in which the critical points are strictly preperiodic. It is a corollary of the results presented later that preperiodic polynomials are uniquely determined by their kneading sequences. Thus, these graphs, which have the same kneading sequence, cannot both be realized as quadratic polynomials. We will also show how to explicitly compute the polynomial from its kneading sequence (showing that the graph on the left occurs as a quadratic polynomial, and the graph on the right does not).

Kneading sequences and external angles. Kneading sequences are only adapted to real polynomials; the appropriate extension of kneading sequences to complex polynomials is external angles. The extension to the complex case is not just a generalization, since even the real results have only been derived as a special case of the complex case. For instance, it is not known whether a mapping $|x|^p + c$ with $p > 1$ and preperiodic kneading sequence is determined by its kneading sequence, except when $p = 2, 4, 6, \ldots$, in which case the map is complex analytic.

If $p(z) = z^2 + c$ is a quadratic polynomial, the most important object for the dynamics of p is the set

$$K_p = \{z \in \mathbb{C} \mid \text{ the sequence } p^n(z) \text{ is bounded}\}.$$

Long before the study of unimodal mappings, Fatou [F1] discovered that, to a large extent, the behavior of the critical point under iteration controls the dynamics of the polynomial.

Theorem. *If the critical point* $0 \in K_p$, *then* K_p *is connected.*

If $0 \notin K_p$, *then* K_p *is a Cantor set.*

He also discovered the following relation between dynamics and conformal mapping. Let D be the open unit disc.

Theorem. (a) *There exists a unique analytic mapping* $\phi_p = z + c_1/z + \cdots$ *defined in a neighborhood of* ∞ *such that*

(1.4) $$p(\phi_p(z)) = \phi_p(z^2).$$

(b) *If* $0 \in K_p$, *then* ϕ_p *extends to an analytic isomorphism* $\mathbb{C} - \overline{D} \to \mathbb{C} - K_p$.

Example. For the polynomial $p(z) = z^2 - 2$ the function ϕ_p can be written explicitly:

$$\phi_p(z) = z + \frac{1}{z}.$$

This is the well-known conformal mapping for the complement of the interval $[-2, 2]$, and when you substitute $z = e^{2\pi i\theta}$ in the functional equation (1.4), you obtain the trigonometric identity $4\cos^2(2\pi\theta) - 2 = 2\cos 2(2\pi\theta)$.

The mapping ϕ_p is used to define external angles. By Carathéodory's theorem [DH1], when K_p is locally connected, the mapping ϕ_p extends continuously to a mapping $\overline{\phi}_p \colon \overline{\mathbb{C}} - D \to \overline{\mathbb{C}}$ (extending ϕ_p to the unit circle ∂D). We will focus on the case when K_p is locally connected. For example, when 0 is preperiodic under iteration by $p(z)$, this is the case (see [JM, Theorem 17.5]).

We denote the *external rays*

$$R_\theta = \overline{\phi}_p(\{re^{2\pi i\theta} | 1 \le r \le \infty\}).$$

Equation (1.4) implies that $p(R_\theta) = R_{2\theta}$. Each ray R_θ has a limit at the point $\overline{\phi}_p(e^{2\pi i\theta})$ of ∂K_p. We say that the external ray *lands* at this point. Each angle determines a unique ray and each ray lands at some point in K_p. However, several external rays may land at the same point of K_p.

Theorem. *Suppose p is a real quadratic polynomial with K_p connected, and that the ray R_θ lands at $x \in K_p \cap \mathbb{R}$. The kneading sequence $s_0, s_1, \ldots$ of x and the angle*

$$\theta = .\varepsilon_0\varepsilon_1\ldots = \sum_{i=0}^{\infty} \frac{\varepsilon_i}{2^{i+1}}$$

are related by the table

s_n \\ ε_n	0	1
P	0	1
R	1	0
C	*	*

*where * can be 0 or 1.*

This theorem can be used to compute the symbolic sequence from the external angle, or the external angle from the symbolic sequence once ε_0 is chosen; the two choices correspond to the ray in the upper half-plane if $\varepsilon_0 = 0$ and its complex conjugate in the lower half-plane if $\varepsilon_0 = 1$.

Example. This fact can be used to compute the arccosine inductively. For the polynomial $p(z) = z^2 - 2$, we saw that $\phi_p(z) = z + 1/z$, so that $\phi_p(e^{2\pi i\theta}) = 2\cos(2\pi\theta)$. If the symbolic sequence of $x = 2\cos(2\pi\theta)$ is $S_p(x) = (s_0, s_1, \ldots)$, and θ is written in base 2 according to the table above with first digit 0, then the branch of $\arccos(x/2)$ in $[0, \pi]$ is given by $\arccos(x/2) = 2\pi\theta$. Choosing the first digit to be 1 gives the branch in $[\pi, 2\pi]$. For instance, the symbolic sequence of $x = 1$ is $P, R, R, R, \ldots$, leading to $\theta = .001\overline{01}\ldots = 1/6$, so that $\arccos 1/2 = 2\pi/6 = \pi/3$, as it should.

External angles and the topology of Julia sets. Let $K \subset \mathbb{C}$ be compact, connected, simply connected and locally connected, and let $\phi_K: \mathbb{C} - \overline{D} \to \mathbb{C} - K$ be a conformal isomorphism. Then (by Carathéodory's theorem as above) ϕ_K extends to $\overline{\phi}_K: \mathbb{C} - D \to \mathbb{C}$; define $\gamma_K: \mathbb{T} = \mathbb{R}/\mathbb{Z} \to \partial K$ by $\gamma_K(t) = \overline{\phi}_K(e^{2\pi it})$. Let $\sim_K$ be the equivalence relation induced on $\mathbb{T}$; i.e., $t_1 \sim_K t_2$ if and only if $\gamma_K(t_1) = \gamma_K(t_2)$. The equivalence relation $\sim_K$ then determines the topology of the pair $(\mathbb{C}, K)$ by the following construction.
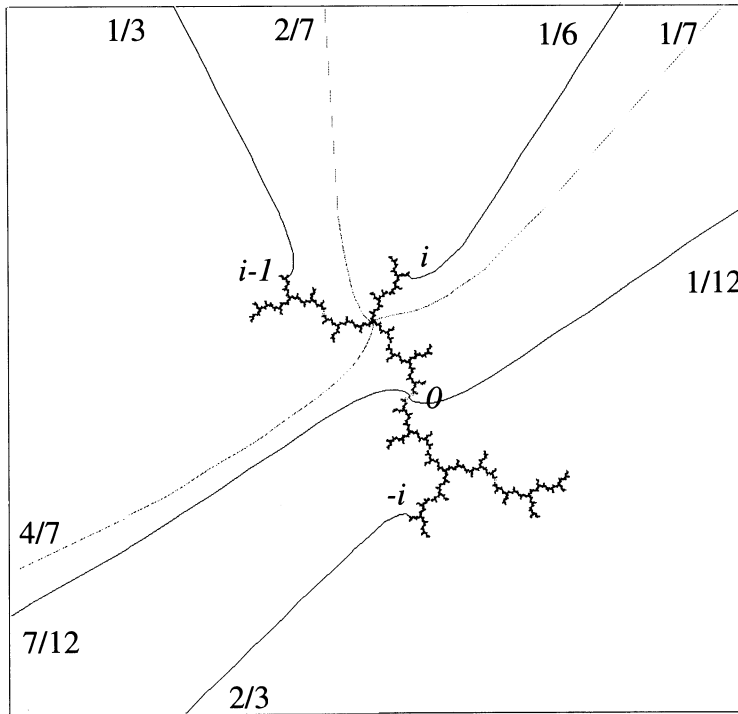
FIGURE 1.5. The set K_p for $p(z) = z^2 + i$ with external rays.

Consider the equivalence relation $\approx_K$ on $\mathbb{C}$ for which the nontrivial equivalence classes are exactly the convex hulls of sets

$$\{e^{2\pi i t} \mid t \in \text{ a nontrivial equivalence class under } \sim_K\}.$$

Proposition. *There exists a homeomorphism of pairs* $h\colon (\mathbb{C}, \overline{D})/\approx_K \to (\mathbb{C}, K)$ *extending* ϕ_K *on* $\mathbb{C} - \overline{D}$.

This is proved in [DH1]; the proof is easy, requiring mainly the Jordan curve theorem.

Motived by this proposition, our strategy will be to understand the equivalence relation $\sim_p$ induced as above by K_p.

Example. The polynomial $p(z) = z^2 + i$ is strictly preperiodic: the orbit of the unique critical point 0 is $(0, i, -1 + i, -i, -1 + i, -i, \ldots)$; but 0 is not periodic.

The set K_p shown in Figure 1.5 looks "chaotic," but we claim this is far from true, and in fact can easily be described combinatorially. First, it appears from the picture and is in fact true that there are unique rays landing at the two points of the cycle $\{-1 + i, -i\}$ of period 2. The angles of these rays must also be periodic of period 2 under angle doubling, and so must be $1/3$ and $2/3$, as indicated.

The ray landing at i must have angle one of the two halves of $1/3$, and the drawing shows it must be $1/6$. Now there are two rays landing at 0, with angles $1/12$ and $7/12$, the two halves of $1/6$. These curves together with $\{0\}$ divide $\mathbb{C}$ into two halves R, and P, where R is the half containing the critical value i.

This allows us to define "kneading sequences" for points of $z \in \partial K_p$: let $S_p(z) = (s_0, s_1, \ldots)$ where

$$s_n = \begin{cases} R & \text{if } p^n(z) \in R, \\ P & \text{if } p^n(z) \in P, \\ C & \text{if } p^n(z) = 0. \end{cases}$$

A very special case of Theorem I, stated in §2, asserts:

a point of ∂K_p is determined by its kneading sequence,

answering question 1.1 positively in this case. This also implies that the equivalence relation $\sim_p$ on $\mathbb{T}$ is in this case given by the following construction: cut $\mathbb{T}$ at $1/12$ and $7/12$, and label R and P the two halves, with $1/6 \in R$. Label the two angles $1/12$ and $7/12$ by C. Then an angle has a symbol for this partition under the angle-doubling map $t \mapsto 2t$. Then $\sim_p$ is given by

$t_1 \sim_p t_2$ *if and only if* $2^n t_1$ *has the same symbol as* $2^n t_2$ *for all* $n \geq 0$.

Thus for the polynomial $p(z) = z^2 + i$ the topology of K_p (and in fact the geometry up to affine maps of $\mathbb{C}$) *is completely determined by the number* $1/6$.

More generally, if for a quadratic polynomial p an external ray R_θ with θ rational lands at the critical value, then the equivalence relation $\sim_p$ is given by the "kneading sequence" with respect to the two halves of θ.

Multiple accesses to the critical value. The example $p(z) = z^2 + i$ shows that one should associate to a quadratic polynomial the angle of the ray landing at the critical value. But this example is simpler than the general preperiodic polynomial: there is only one ray landing at the critical value, and in general there may be several such rays.

Example. One of the eight roots of the equation $p_c^3(0) = -p_c^4(0)$ is approximately $-0.101096 + 0.956287i$, with Julia set shown in Figure 1.6. One of the rays landing at the critical value has angle $9/56$. But it is easy to check that $9/56$, $11/56$, and $15/56$ all have kneading sequence $RRP\overline{R}$ with respect to $9/56$, i.e., if the circle is cut at the two halves $9/112$ and $65/112$ of $9/56$ and R is the symbol of the half containing $9/56$. Thus the three rays with these angles land at the critical value.

It will be considerably easier to characterize polynomials together with a choice of a ray landing at the critical values, and in fact a bit more combinatorial information is needed in higher degrees. A polynomial with this extra information will be called a *marked polynomial*.
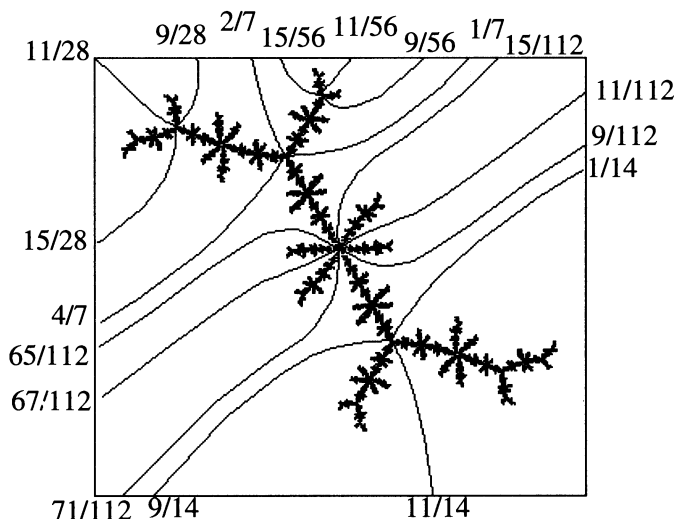
FIGURE 1.6. The Julia set for $p(z) \approx z^2 - 0.101096 + 0.956287i$ with external rays.

External angles and the Mandelbrot set. In one direction, we associate to each preperiodic polynomial the angles of the rays landing at the critical value. In the opposite direction, we are led to three questions:

1.7. Is there a polynomial corresponding to every angle?

1.8. Is it unique?

1.9. When do two angles correspond to the same polynomial?

Question 1.9 has already been answered; for quadratic polynomials there is a very nice way of answering questions 1.7 and 1.8, at least for rational angles, which leads to a description of the parameter space.

Consider the *Mandelbrot set* shown in Figure 1.10

$$M = \{c \in \mathbb{C} \,|\, 0 \in K_{p_c}\}.$$

It is known [DH1] that M is connected. More precisely, when $c \notin M$, then there is a disc D_R such that ϕ_c is defined on $\mathbb{C} - \overline{D}_R$, and $c \in \phi_c(\mathbb{C} - \overline{D}_R)$.
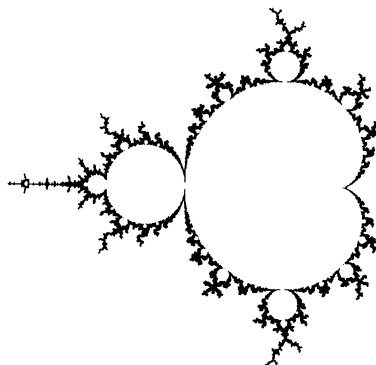
FIGURE 1.10. The boundary of the Mandelbrot set.

Then the mapping $c \mapsto \phi_c^{-1}(c)$ gives a conformal isomorphism $\mathbb{C} - M \to \mathbb{C} - \overline{D}$, the inverse of which we will denote ϕ_M. This conformal mapping ϕ_M allows us to speak of external angles of points in ∂M.

The self-referential nature of this formula accounts for many of the mysterious-sounding statements about M, such as the following.

The external ray at angle $1/6$ of M lands at the polynomial $z^2 + i$, the dynamics of which reflects the digits of $1/6 = .0\overline{01}$ in base two. Indeed, the orbit of the critical point, after one iteration, lands on a cycle of period two, and the digits, after the first, also repeat with period 2.

Now the explanation mentioned above is the following theorem.

Theorem. *For any rational angle θ, the external ray at angle θ of M lands at a point $c \in \partial M$.*

If θ is written in lowest terms has even denominator, then under p_c the critical point is strictly preperiodic, and the external ray at angle θ of K_{p_c} lands at the critical value c, and all strictly preperiodic polynomials arise in this fashion.

If θ written in lowest term has odd denominator, then p_c has a rationally indifferent periodic point, and all polynomials with rationally indifferent periodic points arise in this fashion.

This theorem answers questions 1.7 and 1.8 above, and further says:

If θ_1 and θ_2 are rational angles with even denominators, then the external rays of M at angles θ_1 and θ_2 land at the same point c if and only if the external rays of K_{p_c} at angle θ_1 and θ_2 both land at c.

Thus our answer to question 1.9 in terms of kneading sequences gives some insight into the equivalence relation $\sim_M$, (assuming M to be locally connected) and hence helps explain the topology of M.

This theorem also suggests a possible proof for question 1.7 above: given an angle θ, consider the external ray of M at angle θ, and try to prove that it lands. Then the landing point will be a good candidate for a corresponding polynomial. This strategy is carried out in [DH1], but works only for quadratic polynomials, or at least for polynomials $z^d + c$ with a unique critical point. If there is more than one critical point, the parameter space has dimension greater than one, and there is no good analog to ϕ_M.

This paper answers the analogs of questions 1.7, 1.8, 1.9 in higher degrees. This requires different methods, which even in degree 2 may be simpler than the methods in [DH1]; the basic tool is the *Thurston mapping*, a certain mapping from an appropriate Teichmüller space to itself.

A computer program. Although the reader might never notice it, this paper grew out of a computer program that actually computes preperiodic (and periodic) Thurston polynomials, carrying out the iteration of the Thurston mapping. The algorithm is given in §9, but the real quadratic case is very easy to describe, and we will give it here explicitly.

Let $s = (s_0, s_1, \dots)$ be a preperiodic kneading sequence for the critical point of a unimodal mapping f. Such a kneading sequence must satisfy various conditions; for our purposes we need to know that it starts $CRP \dots$.
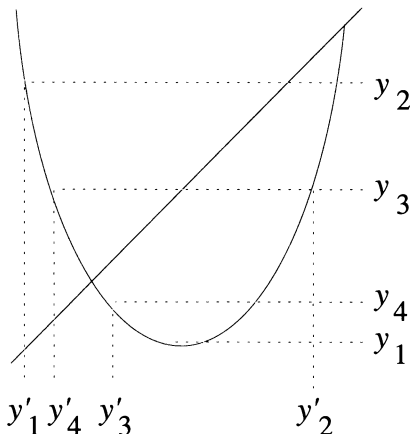
FIGURE 1.11. One iteration of the mapping σ_s.

Remark. Actually, this excludes two legal kneading sequences $CCCC\ldots$ and $CRCR\ldots$, which correspond to the polynomials x^2 and $x^2 - 1$.

We will describe a scheme to construct a quadratic polynomial $x^2 + c$ whose critical value has the same kneading sequence. Let x_1, $x_2 = f(x_1)$, $\ldots$ be the orbit of the critical value under f, with $x_1, \ldots, x_m$ distinct and $x_{m+1} = x_k$ for some $k \leq m$. Let Y be the set of finite sequences $y_1, y_2, \ldots, y_m$ of distinct points of $[-2, 2]$ in the same order as $x_1, \ldots, x_m$, and consider the map $\sigma_s : Y \to Y$ defined as follows.

Given $y = (y_1, y_2, \ldots) \in X$, let $\sigma_s(y) = y'$ where y'_i is the inverse image of y_{i+1} (or of y_k for $i = m$) by $x \mapsto x^2 + y_1$; i.e., $y'_i = \pm\sqrt{y_{i+1} - y_1}$, with the sign chosen positive if $s_i = P$ and negative if $s_i = R$. If $s_i = C$, then $y_{i+1} = y_1$ so that $y'_i = 0$ and there is no ambiguity. We leave it to the reader to check that the points of y' are distinct, belong to $[-2, 2]$ and are in the correct order.

Example. If we use the unimodal map on the right of Figure 1.3, the space Y consists of the sequences y_1, y_2, y_3, y_4 satisfying $-2 \leq y_1 < y_4 < y_3 < y_2 \leq 2$. The kneading sequence of the critical point is $CRP\overline{R}$, so the mapping is as shown in Figure 1.11.

The fundamental point of this construction is the following obvious statement.

If y is a fixed point of σ_s, then the polynomial $x^2 + y_1$ is a quadratic polynomial with the kneading sequence s.

So it is natural to ask whether σ_s has a fixed point, and this will follow (after modification) from the Brouwer fixed point theorem. Let $\overline{Y}$ be the closure of Y in $\mathbb{R}^m$. The map σ_s extends continuously to $\overline{Y}$. However, it has a trivial fixed point in $\overline{Y}$, namely $(0, \ldots, 0)$.

To avoid this fixed point, let $\overline{Y}_1$ be the subset where $y_2 - y_1 \geq 1$. Since

$$y'_2 - y'_1 = \sqrt{y_3 - y_1} + \sqrt{y_2 - y_1},$$

we see that $\overline{Y}_1$ is preserved by σ_s.

Remark. The signs above are due to the kneading sequence $s = CRP\ldots$, so that $y_2' = \sqrt{y_3 - y_1}$ and $y_1' = -\sqrt{y_2 - y_1}$.

Since $\overline{Y}_1$ is defined by linear equalities and inequalities, it is convex, hence homeomorphic to a ball. By the Brouwer fixed point theorem, σ_s must have a fixed point in $\overline{Y}_1$.

First, observe that using closures in order to find fixed points was necessary. Consider the unimodal mapping in the example above. We claimed that there is no quadratic polynomial topologically conjugate to that mapping. In fact, if σ_s is iterated in the space Y corresponding to that mapping, the points y_3 and y_4 coalesce, and the fixed point lies in the boundary of $\overline{Y}$, at a point corresponding to the function on the left of Figure 1.3.

The Brouwer fixed point theorem says nothing about uniqueness of fixed points, and there does not seem to be any easy way of proving that σ_s has a unique fixed point from this point of view. For instance, exactly the same program can be written using $|x|^p + c$, but as far as we know, nothing has been proved about convergence if p is not an even integer. Experimentally, the computer program appears to converge perfectly.

A very general philosophy is that to show a mapping is contracting, try to extend it to the complex domain. Complex manifolds carry intrinsic metrics (sometimes degenerate), like the Kobayashi and the Carathéodory metrics, with respect to which all analytic mappings are contracting. In the case of Thurston's mapping, the appropriate complexification of Y is a certain Teichmüller space, which carries a nondegenerate Kobayashi metric known as the Teichmüller metric. It is with respect to this metric that σ_s is contracting.

2. Definitions and Main Statement

To understand this paper the reader will need a background in complex analytic dynamics; we recommend [JM] and will use the results given there freely. The main tool in the classification above is *Thurston's topological characterization of rational functions* [DH2].

Some facts about the dynamics of polynomials. We recall here some results needed to state the main theorems.

Define the *critical set* Ω_f of a branched covering map $f: S^2 \to S^2$ by

$$\Omega_f = \{z \mid z \text{ is a branch point of } f\}$$

and the *postcritical set*

$$P_f = \bigcup_{n=1}^{\infty} f^{\circ n}(\Omega_f).$$

If p is a polynomial with $\mathrm{Card}(P_p) < \infty$, we follow Milnor's suggestion and call p a *Thurston polynomial*.

The *orbit* of a point z under p is the set $\{p^{\circ n}(z) \mid n = 0, 1, \ldots\}$. If $p^{\circ n}(x) = x$ for some $n > 0$ then x is called a *periodic point* of p, and if x is not periodic but $p^{\circ m}(x)$ is periodic, then x is said to be *preperiodic*.

When every critical point of p is preperiodic, we call p a preperiodic Thurston polynomial.

The *filled-in Julia set* is the compact set $K_p = \{z \in \overline{\mathbb{C}} \mid p^{\circ n}(z) \nrightarrow \infty\}$. The following theorem (see [JM], p. 94) will be central to the entire development.

Theorem 2.1. *If $p(z)$ is monic (i.e., a polynomial with leading coefficient 1) of degree d and P_p is bounded, then there exists a unique analytic homeomorphism*

$$\phi_p \colon \overline{\mathbb{C}} - \overline{D} \to \overline{\mathbb{C}} - K_p$$

tangent to the identity at ∞ (that is, $\phi_p(z) = z + c_0 + c_1/z + \cdots$) that satisfies

$$(2.2) \qquad f(\phi_p(z)) = \phi_p(z^d).$$

When p is not monic, we can still find a map ϕ_p tangent to the identity that conjugates $p(z)$ to az^d, where a is the leading coefficient of $p(z)$.

The Green's function $G_p \colon \mathbb{C} \to \mathbb{R}$ of K_p is the unique subharmonic function of $\mathbb{C}$ that vanishes on K_p and is harmonic on $\mathbb{C} - K_p$ with a logarithmic pole of the form $\log|z|$ at ∞. It satisfies $G_p = \log|\phi_p|$ on $\mathbb{C} - K_p$.

When $p(z)$ is a Thurston polynomial, K_p is locally connected [JM, Theorem 17.5], so that the mapping ϕ_p extends continuously to a mapping $\overline{\phi}_p \colon \overline{\mathbb{C}} - D \to \overline{\mathbb{C}}$ extending ϕ_p to $S^1 = \partial D$. We denote the *external rays*

$$R_\theta = \overline{\phi}_p(\{re^{2\pi i\theta} \mid 1 \le r \le \infty\}),$$

and if Θ is a finite set of angles we write

$$R_\Theta = \bigcup_{\theta \in \Theta} R_\theta.$$

Equation (2.2) implies that $p(R_\theta) = R_{d\theta}$. Each ray R_θ has a limit at the point $\overline{\phi}_p(e^{2\pi i\theta})$ of ∂K_p. We say that the external ray *lands* at this point. Each angle determines a unique ray and each ray lands at some point in K_p. However, several external rays may land at the same point of K_p.

If p is a preperiodic Thurston polynomial then $\partial K_p = K_p$. In particular, the postcritical set is contained in ∂K_p, and so at least one ray lands at each postcritical point. All the rays landing at periodic and preperiodic points in K_p (including the postcritical set) have rational angles.

Marked polynomials. It is easiest to classify marked polynomials and deal separately with the redundancy this introduces in the classification.

A *marked polynomial* is a polynomial with strictly preperiodic critical points that is monic and centered, together with the choice, for each critical point ω, of an external ray landing at $p(\omega)$. Define $\mathscr{P}_d$ as the set of all marked polynomials of degree d.

Throughout this paper we will let $\mathbb{T} \equiv \mathbb{R}/\mathbb{Z}$. The set $\mathbb{T}$ has a natural additive group structure that is equivalent to the multiplicative group structure on S^1. The circle $S^1 \in \mathbb{C}$ is better for geometry (convex hulls, etc.), but $\mathbb{T}$ is better

for arithmetic, and we will use them both, sometimes identifying them without comment.

Let $\theta \in \mathbb{T}$ and define the map $m_d(\theta): \mathbb{T} \to \mathbb{T}$ by $\theta \mapsto d\theta$ where d is some integer. If $\Theta \subset m_d^{-1}(\theta)$ and Θ contains at least two points then we call Θ a *d-preangle* of θ. A d-preangle Θ of θ is *periodic* of period k if $m_d^{\circ(k-1)}(\theta) \in \Theta$.

Define the mapping

$$PA: \mathscr{P}_d \to \text{ sets of } d\text{-preangles}$$

as follows. Let p be a monic centered polynomial, with critical points $\{\omega_1, \ldots, \omega_n\}$, and marked by the rays $\{R_{\theta_1}, \ldots, R_{\theta_n}\}$ landing at $p(\omega_1) \ldots, p(\omega_n)$. Set

$$PA(p; \theta_1, \ldots, \theta_n) = (\Theta_1, \ldots, \Theta_n),$$

where Θ_i is the d-preangle of θ_i made of the angles of external rays landing at ω_i that are inverse images of the ray R_{θ_i}.

Example 2.3. For the polynomial $z^2 + i$ above, there is only one ray landing at the critical value i that has angle $1/6$. So $(z^2 + i; 1/6)$ is a marked polynomial, and $PA(z^2 + i; 1/6) = \{\{1/12, 7/12\}\}$.

A more substantial example, shown in Figure 2.4, is provided by the polynomial $z^3 + c$ with $c \approx -.220330 + i1.186329$, where the third forward image of the critical point is fixed. In this case, there are two rays, with angles $19/72$ and $25/72$, that land at the critical value. So $(z^3 + c; 19/72)$ and $(z^3 + c; 25/72)$ are two markings of the same polynomial, and

$$PA(z^3 + c; 19/72) = \{\{19/216, 91/216, 163/216\}\};$$
$$PA(z^3 + c; 25/72) = \{\{25/216, 97/216, 169/216\}\}.$$

Conditions satisfied by marked polynomials. We wish to characterize the image of $\mathscr{P}_d$ under PA by combinatorial conditions on the preangles. For any marked polynomial $(p; \theta_1, \ldots, \theta_n)$ with

$$PA(p; \theta_1, \ldots, \theta_n) = \{\Theta_1, \ldots, \Theta_n\},$$

then the preangles $\{\Theta_1, \ldots, \Theta_n\}$ satisfy

(C.1) $\Theta_1, \ldots, \Theta_n$ are finite sets of rational angles (see [JM, p. 97]).

(C.2) $d - 1 = \sum_{i=1}^{n} (\text{Card}(\Theta_i) - 1)$.

(C.3) Each R_{Θ_i} partitions $\mathbb{C}$ into components, exactly one of which contains each R_{Θ_j} for $j \neq i$.

(C.4) None of the rays landing at critical points are periodic.

(C.5) Distinct R_{Θ_i} land at distinct points.

We will turn each of the conditions (C.1)–(C.5) into a combinatorial condition on the sets $\Theta_1, \ldots, \Theta_n$. The first four conditions have obvious interpretations, but condition (C.5) is more delicate.
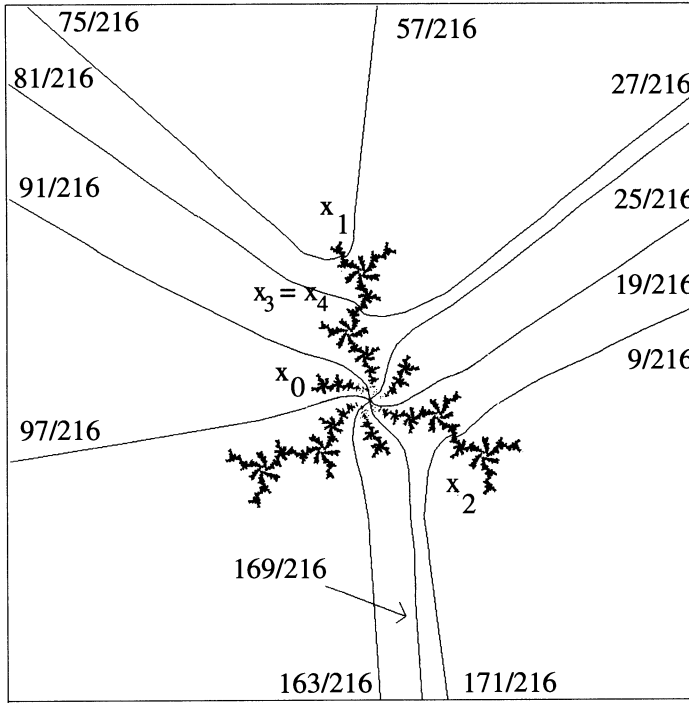
FIGURE 2.4. K_p for a cubic with six accesses to the critical point.

The equivalence relation $\sim_\Theta$. In this subsection we define an equivalence relation $\sim_\Theta$ that will characterize Julia sets of preperiodic Thurston polynomials.

Two nonempty closed subsets A, B of $\mathbb{T}$ (or S^1) are *unlinked* if there exist disjoint intervals I, $J \subset \mathbb{T}$ such that $A \subset I$ and $B \subset J$.

A collection $\{A_1, \ldots, A_n\}$ of closed subsets of S^1 is unlinked if they are pairwise unlinked. Figure 2.6 gives examples for the collection $\{\Theta_1, \Theta_2\}$.

We will find the following lemma useful on several occasions.

Lemma 2.5. *Two nonempty closed subsets A and B of S^1 are unlinked if and only if there exist disjoint closed connected subsets X, $Y \subset \overline{D}$ of the closed unit disk such that $A \subset X$ and $B \subset Y$.*

FIGURE 2.6. Linked and unlinked sets.

Proof. If A and B are unlinked, then we can take X and Y to be the convex hulls of intervals I and J in the definition.

Conversely, suppose X is contained in the component U of $\overline{D} - Y$. Then any two points $s_1, s_2 \in S^1 \cap U$ can be joined by an embedded arc $\alpha \subset U$ with interior in D. The arc divides $\overline{D}$ into two components, one of which contains Y, so s_1, s_2 can be connected in $U \cap S^1$. We have shown that $U \cap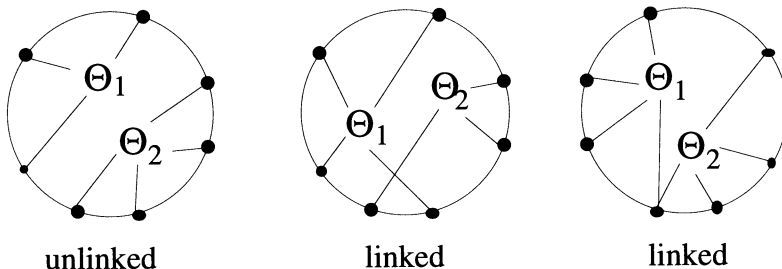 S^1$ is connected, and hence an open interval I', since B is nonempty. Let $I \subset I'$ be a closed interval containing A, and $J = S^1 - I'$; these satisfy the definition. □

Let $\Theta = \{\Theta_1, \ldots, \Theta_n\}$ be a collection of unlinked *d-preangles.*

Definition 2.7. We will say that two points $x, y \in \mathbb{T}$ are **Θ**-*unlinked* if the collection $\{x, y\}, \Theta_1, \ldots, \Theta_n$ is unlinked.

We will say that two points $x, y \in \mathbb{T}$ are **Θ**-*unlinkable* if there exist arbitrarily small perturbations x', y' of x, y such that $\{x', y'\}$ is **Θ**-unlinked.

Two angles will be called **Θ**-*related* if for all $m \geq 0$ the points $d^m x$ and $d^m y$ are **Θ**-unlinkable.

Let the equivalence relation $\sim_\Theta$ be the equivalence generated by $x \sim_\Theta y$ if x and y are **Θ**-related. This means that angles are $\sim_\Theta$-equivalent when they can be connected by a finite chain of **Θ**-related angles.

Remark. Being **Θ**-unlinked is an equivalence relation; an equivalence class is a finite union of open intervals. A nice way to visualize these equivalence classes is to consider the convex hull $\widehat{\Theta}_i$ of each preangle Θ_i; then for each component U of $\overline{D} - \bigcup_i \widehat{\Theta}_i$, the intersection $U \cap S^1$ is an equivalence class. See Figure 2.8.
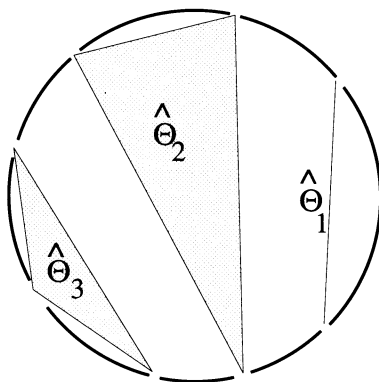
FIGURE 2.8. Unlinked equivalence classes.

Two angles are **Θ**-unlinkable when they belong to the closure of an equivalence class above; this is of course not an equivalence relation.

Example 2.9. Take $d = 2$ and $\Theta = \{\{1/12, 7/12\}\}$ as in Figure 1.5; then $1/7 \sim_\Theta 2/7 \sim_\Theta 4/7$. Indeed, since these three angles form a cycle under m_2, it is enough to show that $\{1/7, 2/7, 4/7\}$ is **Θ**-unlinked, which is true since $1/12 < 1/7 < 2/7 < 4/7 < 7/12$. We leave for the reader to check that for

$d = 3$ and both

$$\Theta = \{\{19/216, 91/216, 163/216\}\} \quad \text{and} \quad \Theta = \{\{25/216, 97/216, 169/216\}\},$$

we have $1/8 \sim_\Theta 3/8$.

A topological description of K_p when p is a preperiodic Thurston polynomial. The following theorem, which is one of our three main results, explains the central importance of the equivalence relation $\sim_\Theta$. Although $\sim_\Theta$ is an equivalence relation on $\mathbb{T}$, we will use the same symbol for the induced equivalence relation on S^1, and even on $\mathbb{C}$, with the understanding that all equivalence classes of numbers z with $|z| \neq 1$ are trivial.

Theorem I. *If $(p; \theta_1, \dots, \theta_n)$ is a marked critically preperiodic polynomial and $\Theta = PA(p; \theta_1, \dots, \theta_n)$, then the map $\phi_p: (\mathbb{C} - D, S^1) \to (\mathbb{C}, K_p)$ induces a homeomorphism*

$$((\mathbb{C} - D)/\sim_\Theta, S^1/\sim_\Theta) \to (\mathbb{C}, K_p).$$

The proof will be given in §3.

Remarks. 1. This result says that $\Theta = (\Theta_1, \Theta_2, \dots, \Theta_n)$ completely characterizes the topology of K_p, and even of the pair $(\mathbb{C}, K_p)$. A theorem of P. Jones and L. Carleson [CJ] goes further, and shows that Θ actually determines the pair $\mathbb{C}, K_p$ up to affine equivalence, in the sense that any homeomorphism $h: \mathbb{C} \to \mathbb{C}$ that is analytic on $\mathbb{C} - K_p$ is actually analytic on $\mathbb{C}$, hence is an affine map.

2. Theorem I goes a long way towards explaining when two rays in the complement of the Mandelbrot set land at the same point [DH1]. This will be explained in detail in a subsequent paper.

Preperiodic polynomial determining families of angles. A collection of d-preangles $\Theta = (\Theta_1, \dots, \Theta_n)$ is called a *preperiodic polynomial determining family of angles* or a *PPDFA* of degree d if the following conditions are satisfied.

    (1) All angles in $\Theta_1, \dots, \Theta_n$ are rational.
    (2) $\sum_{i=1}^n (\text{Card}(\Theta_i) - 1) = d - 1$.
    (3) $\Theta_1, \dots, \Theta_n$ are unlinked.
    (4) No $\theta \in \Theta_j$ is periodic for any j.
    (5) For all i, j with $i \neq j$ we have $\Theta_i \not\sim_\Theta \Theta_j$.

We denote the set of all PPDFA's by $\mathscr{A}_d$.

In particular, the preangles $\{\Theta_1, \dots, \Theta_n\}$ are disjoint by condition 3. Each of these conditions corresponds to a condition in (C.1)–(C.5). We are now in a position to prove the following result.

Proposition 2.10. *If $(p; \theta_1, \dots, \theta_1)$ is a marked polynomial, then*

$$PA(p; \theta_1, \dots, \theta_1) = \{\Theta_1, \dots, \Theta_n\}$$

is a PPDFA.

Proof. Let $PA(p; \theta_1, \dots, \theta_1) = (\Theta_1, \dots, \Theta_n)$. We must show the five conditions. Condition 1 follows from the fact that rays landing at preperiodic points

have rational angles [JM]. Condition 2 follows from the Riemann-Hurwitz formula, or in this case counting the zeros of the derivative of p. Condition 4 follows since the critical points of p are assumed to be strictly preperiodic, and condition 5 (the only delicate one) follows immediately from Theorem I.

We are left with condition 3, which can be shown as follows. Make a compact space $\overline{\overline{\mathbb{C}}}$ by adding circle S at ∞ in the obvious way; then the closure $\overline{\overline{R}}_\theta$ of R_θ in $\overline{\overline{\mathbb{C}}}$ intersects S exactly at $\theta \in S$. If $i \neq j$ then $\overline{\overline{R}}_{\Theta_i}$ and $\overline{\overline{R}}_{\Theta_j}$ satisfy the conditions of Lemma 2.5. □

Remark. We will of course not use Proposition 2.10 in the proof of Theorem I.

Example. We will exhibit the necessity of condition 5. Return to Example 2.3 above. It is a polynomial with a single critical point, but it is not clear why there could not be one just like it with two critical values, one at the end of the ray at angle $19/72$, and another at the end of the ray at angle $25/72$. If there were such a marked polynomial p, then

$$PA(p) = \{\{25/216, 169/216\}, \{91/216, 163/216\}\} = \Theta$$

is one of three possibilities. However, this is not a PPDFA; it fails condition 5 (and only condition 5). The reader may check $25/216 \sim_\Theta 91/216$.

The classification of preperiodic marked polynomials. Our main theorem is the following result.

Theorem II. *The map* $PA: \mathscr{P}_d \to \mathscr{A}_d$ *is a bijection.*

This theorem clearly gives a complete classification of marked polynomials in combinatorial terms. For instance, when $d = 2$, it reduces to the following result.

Corollary. *For every rational angle written in reduced form as p/q with q even, there exists a unique marked quadratic preperiodic polynomial such that the chosen ray to the critical value has angle p/q, and every marked quadratic preperiodic polynomial arises in this way.*

Proof. An element of $\mathscr{A}_2$ is a single 2-preangle by the second condition; it is easy to check that it must be of the form $\{p/(2q), (p+q)/(2q)\}$ with p/q as above. □

When two PPDFA's determine the same polynomial. When will two distinct PPDFA's determine the same polynomial? The redundancy in the description above is given by the following theorem.

Theorem III. *The polynomial determining families of angles* $\Theta = (\Theta_1, \ldots, \Theta_n)$ *and* $\Theta' = (\Theta'_1, \ldots, \Theta'_n)$ *determine the same polynomial if and only if the* Θ'_i *can be renumbered so that* $\Theta_1 \sim_\Theta \Theta'_1, \ldots, \Theta_n \sim_\Theta \Theta'_n$.

The proof can be found in §10. It is not a difficult proof, very much unlike the other cases where this sort of redundancy has been attacked, for instance when the matings of two polynomials lead to the same rational function (see [BW]).

Outline of the paper. Theorem I is proved in §3. The proof has two parts: seeing that rays with equivalent angles land at the same point, and that rays with nonequivalent angles land at distinct points.

The second part is essentially obvious, but requires a bit of care at the inverse images of the critical points; these difficulties are precisely the reason for the elaborate definitions of unlinked, unlinkable, etc. The first part is more substantial, and requires the expanding properties of the orbifold metric [JM]. This is not surprising: essentially the theorem says that a translation of a dynamical system to a shift mapping using a Markov partition is faithful; the obvious way of showing such things is to claim that the dynamical system is expanding.

The main part of Theorem II is to show the surjectivity of PA; this is done using Thurston's theorem, which is described in §4. Section 5 contains a careful analysis of the Thurston obstructions that can arise when a branched mapping mimics a polynomial.

Starting with a PPDFA Θ, we first construct a graph $S_\Theta \subset S^2$, and a mapping $f_\Theta : S_\Theta \to S_\Theta$, which we extend to a branched mapping $\overline{f}_\Theta : S^2 \to S^2$ (§7). This requires a criterion that allows maps on graphs in S^2 to be extended to S^2, which is given in §6.

In §8, we show that the obstructions described in §5 do not arise for the branched map f_Θ, which thus has no Thurston obstruction, and is equivalent to a polynomial p.

In §9 we show that p can be conjugated to a unique polynomial p_Θ, which is naturally marked. The end result of §§7, 8 and 9 is to construct a mapping $AP : \mathscr{A}_d \to \mathscr{P}_d$ such that $PA \circ AP = \text{id}$.

This is spelled out in §10, and we also show there that AP is surjective, which completes the proof of Theorem II. We also give the proof of Theorem III there.

The thesis of Alfredo Poirier at SUNY, Stony Brook, completes the description detailed in this paper by adding the extra information necessary to define a marked polynomial in the case of a periodic critical point.

3. WHEN EXTERNAL RAYS LAND AT THE SAME POINT

In this section we prove Theorem I. Our understanding of the theorem came from analyzing Levy cycles, which we discuss later. The proof of the theorem, however, does not use any of the techniques presented later in the paper.

If two rays land at the same point, their angles are equivalent. First we will see that if R_{θ_1} and R_{θ_2} land at the same point z, then $\theta_1 \sim_\Theta \theta_2$.

The proof will proceed by induction on

$$m(z) = \text{the small number } m \text{ such that the orbit}$$
$$\text{of } P^m(z) \text{ contains no critical point}.$$

Note that this number is finite since no critical point is periodic. First suppose that the forward orbit of z contains no critical points, so that $m(z) = 0$. Then because R_{θ_1} and R_{θ_2} land at the same point, Lemma 2.5 shows that θ_1 and θ_2 are Θ-unlinked (and thus Θ-unlinkable). Similarly, for any i, $d^i\theta_1$ and
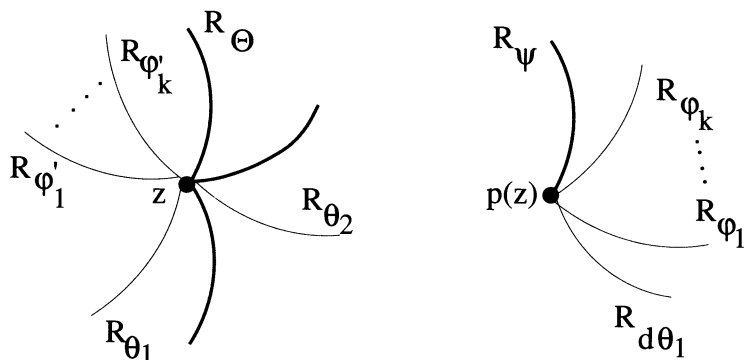
FIGURE 3.1. Pulling back Θ-related angles.

$d^i\theta_2$ are Θ-unlinkable, so that θ_1 and θ_2 are Θ-related and thus Θ-equivalent. This starts the induction.

Suppose by induction that the theorem is true of rays landing at points x with $m(x) < i$, and let z be a point with $m(z) = i$.

Suppose first that z is not critical. By induction, the rays at angles $d\theta_1$ and $d\theta_2$ are equivalent; i.e., there exist $\varphi_1, \ldots, \varphi_k$ such that in the sequence

$$d\theta_1, \varphi_1, \ldots, \varphi_k, d\theta_2$$

each angle is Θ-related to the next. Thus, there exist unique angles φ_i' such that $d\varphi_i' = \varphi_i$ and the rays $R_{\varphi_i'}$ land at z. Then in the sequence

$$\theta_1, \varphi_1', \ldots, \varphi_k', \theta_2,$$

each angle is Θ-related with its successor, which shows that $\theta_1 \sim_\Theta \theta_2$.

Now suppose that z is critical, corresponding to the d-preangle Θ_i of ψ. We will show that θ_1 is equivalent to one (and hence all) of the $\theta \in \Theta_i$; this will show $\theta_1 \sim_\Theta \theta_2$, by symmetry. The ray R_ψ lands at $p(z)$, and by induction ψ is equivalent to $d\theta_1$.

If $d\theta_1 = \psi$, the statement is clear. Otherwise, there exist $\varphi_1, \ldots, \varphi_k$ as above such that in the sequence

$$d\theta_1, \varphi_1, \ldots, \varphi_k, \psi$$

each angle is Θ-related to its successor. Now let φ_i' be the angles such that $d\varphi_i' = \varphi_i$ and the rays $R_{\varphi_i'}$ land at z in the component of $\mathbb{C} - R_{\Theta_i}$ containing R_{θ_1}. Then in the sequence

$$\theta_1, \varphi_1', \ldots, \varphi_k'$$

each is Θ-related to its successor, and the last one is Θ-related to both elements of Θ_i with corresponding rays in the boundary of that component. See Figure 3.1.

If two angles are equivalent, their rays land at the same point. Let X be a component of $\overline{\mathbb{C}} - \bigcup_i R_{\Theta_i}$, and construct a space $\overline{\overline{X}}$ by identifying points $x, y \in \partial X \subset \overline{\mathbb{C}}$ for which $p(x) = p(y)$.

Remark. The space $\overline{\overline{X}}$ is actually a Riemann surface, but we do not want to use this fact as we will use Lemma 3.2 in a setting where no analytic structure is available.

Lemma 3.2. *The space* $\overline{\overline{X}}$ *is homeomorphic to* S^2, *and the map* $p_X \colon \overline{\overline{X}} \to \overline{C}$ *induced by* p *is a homeomorphism.*

Proof. It is unnecessary and a bit fussy to deal with the point at infinity, so we eliminate it. Let $\overline{\overline{\infty}} \in \overline{\overline{X}}$ be the unique point corresponding to $\infty \in \overline{C}$; since $p_X^{-1}(\infty) = \overline{\overline{\infty}}$ we see that $p_X \colon \overline{\overline{X}} - \overline{\overline{\infty}} \to C$ is proper. But it is also clearly a local homeomorphism, so it is a covering map, and since C is simply connected and $\overline{\overline{X}}$ is connected, it is a homeomorphism. The lemma follows since $\overline{\overline{X}}$ is the one-point compactification of a space homeomorphic to C. $\square$

Let $T' \subset T$ be the set of postcritical angles, i.e.,

$$T' = \{ m_d^{\circ k}(\Theta_i) \mid k, i = 1, 2, \ldots \}$$

so that rays R_θ with $\theta \in T'$ land at one of the postcritical points P_p of p. Let $Y = \overline{C} - (R_{T'} - P_p)$, which is a connected set.

Lemma 3.3. *Any path* $\alpha \colon [0, 1] \to Y$ *lifts uniquely to a path* $\alpha' \colon [0, 1] \to \overline{\overline{X}}$ *such that* $p \circ \alpha' = \alpha$.

Proof. This follows from Lemma 3.2; the lift to $\overline{\overline{X}}$ does not intersect the identification locus (note that the critical points in ∂X are not in the identification locus). $\square$

Define a metric on K_p as follows. Choose $r > 0$, and let $Y_r = Y \cap \{G_p \leq r\}$, and

$$\delta(x, y) = \inf l(\alpha),$$

where α ranges over all paths from x to y in Y_r, and l is the length with respect to the orbifold metric on C [JM].

The map p is expanding in the orbifold metric. Therefore there exists $C > 1$ such that for any path $\alpha \in Y_r$, if α' is given as in Lemma 3.3, then $Cl(\alpha') \leq l(\alpha)$, since Y_r is a compact set; note also that α' is a path in $Y_{r/d} \subset Y_r$.

Suppose θ_1 is Θ-related to θ_2, and that R_{θ_1} lands at $x \in K_p$ and R_{θ_2} lands at y.

Then we see that

$$\delta(p(x), p(y)) \geq C\delta(x, y).$$

Indeed, let α be a length-minimizing curve joining $p(x)$ to $p(y)$ in Y_r. Then both x and y lie in the closure of some component X of $\overline{C} - \bigcup_i R_{\Theta_i}$, and the corresponding curve α' joins them. So

$$\delta(x, y) \leq l(\alpha') \leq C^{-1} l(\alpha).$$

Using this over and over, we see that if $x \neq y$, then $\delta(p^{\circ k}(x), p^{\circ k}(y))$ grows without bound, and will eventually be greater than the diameter of K_p.

Finally, if $\theta_1 \sim_\Theta \theta_2$ then there exists angles $\varphi_1, \ldots, \varphi_n$ such that each

successive pair of angles in the sequence $\theta_1, \varphi_1, \ldots, \varphi_n, \theta_2$ is Θ-related, and thus all the rays with these angles must land at the same point.

<div align="right">QED Theorem I</div>

4. Thurston's topological characterization of rational functions

In this section we describe a topological condition that is necessary and sufficient for a branched map to be equivalent, in a sense defined below, to a rational function. The reader should be forewarned that this section is technical, and it can be painlessly omitted in a first or second reading.

First we need some definitions. We will use $\mathbb{P}^1$ and S^2 interchangeably when the complex structure is relevant or not, respectively. All maps in this paper are understood to be orientation preserving.

A postcritically finite branched map $f: S^2 \to S^2$ will be called a *Thurston map*. Two Thurston maps f and g are *Thurston equivalent* if there are homeomorphisms θ_1 and θ_2 mapping $S^2 \to S^2$ such that the following diagram commutes:

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

where $\theta_1(P_f) = \theta_2(P_f) = P_g$, and θ_1 is isotopic to θ_2 relative to P_f (that is, θ_1 is homotopic to θ_2 through homeomorphisms that coincide on P_f).

We say that a simple closed curve $\gamma \subset S^2 - P_f$ is nonperipheral if each component of $S^2 = \{\gamma\}$ contains at least two points of P_f. A *multicurve* $\Gamma = \{\gamma_1, \ldots, \gamma_n\}$ is a set of simple, closed, disjoint, nonhomotopic, nonperipheral curves in $S^2 - P_f$.

We say a multicurve Γ is *f-stable* if for every $\gamma \in \Gamma$, every nonperipheral component of $f^{-1}(\gamma)$ is homotopic in $S^2 - P_f$ to a curve in Γ.

Let $\gamma_{i,j,\alpha}$ be the components of $f^{-1}(\gamma_j)$ homotopic to γ_i rel P_f, and $d_{i,j,\alpha}$ be the degree of the map $f|_{\gamma_{i,j,\alpha}} : \gamma_{i,j,\alpha} \to \gamma_j$.

The *Thurston linear transformation* $f_\Gamma \colon \mathbb{R}^\Gamma \to \mathbb{R}^\Gamma$ is defined as follows:

$$
f_\Gamma(\gamma_j) = \sum_{i,\alpha} \frac{1}{d_{ij\alpha}} \gamma_i \quad \text{where } d_{ij\alpha} = \deg f|_{\gamma_{ij\alpha}} : \gamma_{ij\alpha} \to \gamma_j.
$$

Since the matrix of f_Γ has nonnegative entries, its eigenvalue with largest modulus is real and positive; denote it $\lambda(f_\Gamma)$.

There is a function $\nu: P_f \to \{1, 2, 3, \ldots, \infty\}$ such that for all $x \in f^{-1}(y)$, $\nu(y)$ is an integer multiple of $\nu(x) \deg_x f$. Let ν_f be the smallest such function ν. We will say that the orbifold (see [JM]) is *hyperbolic* if its Euler characteristic

$$
2 - \sum_{x \in P_f} \left(1 - \frac{1}{\nu_f(x)}\right)
$$

is negative.

Remark. If $\text{Card}(P_f) > 4$ then this condition is always satisfied, because $\nu_f(x)$ will be at least 2 for $x \in P_f$.

The following theorem due to W. Thurston is proved in [DH2].

Theorem 4.1. *A postcritically finite branched map* $f: S^2 \to S^2$ *with hyperbolic orbifold is equivalent to a rational function if and only if for any f-stable multi-curve* Γ, $\lambda(f_\Gamma) < 1$. *In that case, the rational function is unique up to conjugation by an automorphism of the Riemann sphere* $\mathbb{P}^1$.

An f-stable multicurve Γ with $\lambda(f_\Gamma) \geq 1$ is called a *Thurston obstruction*.

In our specific case, if the orbifold is not hyperbolic, the branched map is equivalent to a polynomial anyway. Rather than treat this technical point, we refer the interested reader to [DH2] and restate the theorem is less generality. A branched map $f: S^2 \to S^2$ is said to be a *topological polynomial* if ∞ is a critical point and $f^{-1}(\infty) = \{\infty\}$, so that $\nu_f(\infty) = \infty$.

Remark. A topological polynomial will have hyperbolic orbifold if its postcritical set (excluding ∞) contains at least three points. The topological polynomials with nonhyperbolic orbifold are all equivalent to polynomials anyway. These are precisely the Tchebychev polynomials and the polynomials z^n for $n > 1$.

Theorem 4.2. *A postcritically finite topological polynomial* $f: S^2 \to S^2$ *is equivalent to a polynomial if and only if for any f-stable multicurve* Γ, $\lambda(f_\Gamma) < 1$. *In that case, the polynomial is unique up to conjugation by an affine transformation.*

We will apply this theorem to a branched map that we will generate from a PPDFA. Consider Figure 4.3, which shows points representing the postcritical set and a multicurve. It should be clear that the condition $\lambda(f_\Gamma) < 1$ is difficult to verify. In the next section we present a powerful simplification to the hypothesis of Theorem 4.2.
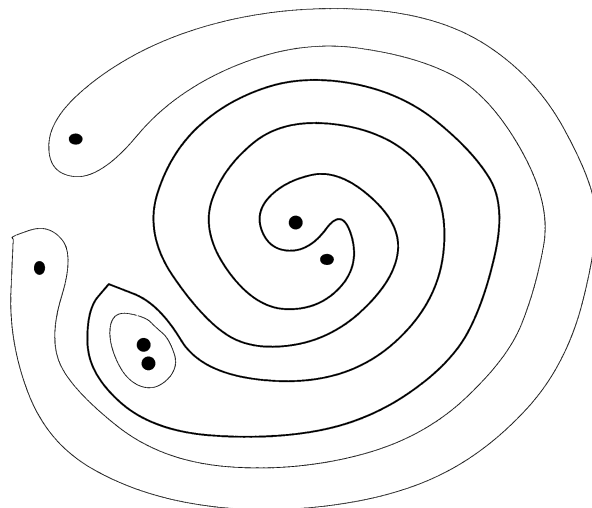

FIGURE 4.3. A complicated multicurve.

5. A CRITERION FOR THE EXISTENCE OF OBSTRUCTING MULTICURVES

In this section we prove that if a particular type of branched map is not equivalent to a polynomial, then we can find a special type of Thurston obstruction called a Levy cycle. This enormously simplifies proofs that certain types of branched maps are equivalent to polynomials since it is easy to show that in certain cases no Levy cycles exist. We know no general condition for the existence of an obstructing multicurve.

Lemma 5.1. *If f is a topological polynomial and $U \subset S^2$ homeomorphic to an open disk with $\infty \notin U$, then every component of $f^{-1}(U)$ is homeomorphic to a disk.*

Proof. If $f^{-1}(U)$ is not a disjoint union of disks, its complement Y has at least two components. Let Z be such a component that does not contain ∞. Then f maps Z surjectively to $X = S^2 - U$. Since

- $f: S^2 \to S^2$ is proper, so $f: Y \to X$ is proper, so $f: Z \to X$ is proper, so $f(Z)$ is closed in X;
- f is a topological polynomial, so f is open, so $f: Y \to X$ is open, so $f: Z \to X$ is open, so $f(Z)$ is open in X.

So Z contains an inverse image of ∞, which is a contradiction. □

Definition 5.2. Let f be a Thurston map, and let Γ be a Thurston obstruction. Suppose there exist curves $\{\gamma_0, \dots, \gamma_k = \gamma_0\} = \Lambda \subset \Gamma$ such that for each $i = 0, \dots, k-1$, γ_i is homotopic rel P_f to exactly one component γ' of $f^{-1}(\gamma_{i+1})$ and $f: \gamma' \to \gamma_{i+1}$ has degree 1. Then Λ is called a *Levy cycle*. See Figure 5.3.
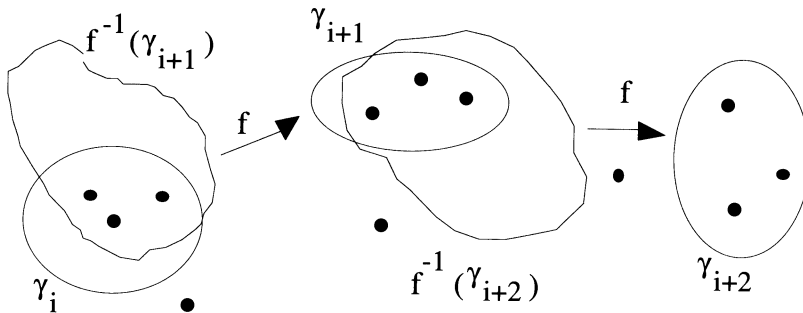
FIGURE 5.3. Part of a Levy Cycle. The dots represent points of P_f.

The following theorem shows that for any branched map, and in particular for a topological polynomial, a Levy cycle implies the existence of a Thurston obstruction.

Theorem 5.4. *If f is a rational function, then f cannot have a Levy cycle.*

Proof. Let d be the degree of f. First observe that $f^{-1}(P_f)$ strictly includes P_f, since $\Omega_f \subset f^{-1}(P_f)$. In fact,

$$\mathrm{Card}(f^{-1}(P_f)) = d\,\mathrm{Card}(P_f) - (2d - 2).$$

Then $P_f = f^{-1}(P_f)$ implies that $d = 1$ or $\mathrm{Card}(P_f) = 2$.

Suppose f has Levy cycle Λ. Then $\mathrm{Card}(P_f) \geq 4$ (otherwise the curves in the Levy cycle would be peripheral). Let $\gamma_1, \ldots, \gamma_n$ be the Poincaré geodesics on $\overline{\mathbb{C}} - P_f$ in the homotopy class of Λ. Then the component γ_i' of $f^{-1}(\gamma_{i+1})$ homotopic to γ_i has the same length as γ_{i+1} in the Poincaré metric of $\overline{\mathbb{C}} - f^{-1}(P_f)$. In the Poincaré metric of $\overline{\mathbb{C}} - P_f$, γ_i' is strictly shorter than γ_{i+1}, and since γ_i is the geodesic in the homotopy class of γ_i', it is shorter yet.

So each γ_i is strictly shorter than γ_{i+1}. This is a contradiction. $\square$

The converse of Theorem 5.4 is not true in general, but in the case of topological polynomials we have the following.

Theorem 5.5. *If the topological polynomial f has a Thurston obstruction Γ, then f has a Levy cycle $\Lambda \subset \Gamma$.*

Proof. This proof is based on ideas of Mary Rees. The idea is to take a Thurston obstruction Γ, and consider the mapping

$$P: \Gamma \to 2^{\Gamma}$$

to the power set of Γ, which associates to a curve γ the set of elements of Γ that are homotopic rel P_f to some component of $f^{-1}(\gamma)$. If P had images that were nonempty and disjoint, then clearly they would have to be singletons, and P would induce a permutation on Γ. A cycle of this permutation will then give the desired Levy cycle.

Unfortunately, the images of P are neither nonempty nor disjoint, in general. But we will show that there is a subset of *innermost* curves $\Gamma_i \subset \Gamma$ for which the induced map $P_i: \Gamma_i \to 2^{\Gamma_i}$ does have these properties. We will need to define *essential* and *negligible* curves along the way.

We assume without loss of generality that Γ is minimal in the sense that every sub-multicurve $\Gamma' \subset \Gamma$ that is still a Thurston obstruction is in fact Γ. This implies in particular that every $\gamma \in \Gamma$ is homotopic rel P_f to a component of an inverse image of some $\gamma' \in \Gamma$.

Write $\Gamma = \Gamma_e \cup \Gamma_n$, where $\Gamma_n = \{\gamma \in \Gamma \mid \|f_{\Gamma}^i(\gamma)\| \to 0$ as $i \to \infty\}$ is the set of *negligible* curves of Γ and $\Gamma_e = \Gamma - \Gamma_n$ consists of the *essential* curves. Since $\lambda(f_{\Gamma}) \geq 1$, Γ_e is not empty; otherwise we would have $\lim_{n\to\infty} f_{\Gamma}^n = 0$.

Lemma 5.6. *Let $A = (a_{ij}) \mid \mathbb{R}^n \to \mathbb{R}^n$ be a linear transformation with $a_{ij} \geq 0$ and suppose $\vec{x} = \sum a_i \vec{e}_i \in \mathbb{R}^n$ is a vector with positive entries such that $A^n \vec{x} \to 0$ as $n \to \infty$. Then $a_i \neq 0$ implies $A^n \vec{e}_i \to 0$ as $n \to \infty$.*

Proof. Since everything is positive, there can be no cancellation. $\square$

Lemma 5.6 implies that any nonperipheral component of the inverse image of a negligible curve is homotopic rel P_f to a negligible curve. On the other hand,

some component of the inverse image of every essential curve is homotopic rel P_f to an essential curve.

Indeed, write $f_\Gamma(\gamma) = \sum_{\delta \in P(\gamma)} a_\delta \delta$ where $a_\delta \geq 0$. So if γ is essential then not all of the $\delta \in P(\gamma)$ can be negligible; and if γ is negligible then all the δ must be negligible.

Therefore if we define $P_e : \Gamma_e \to 2^{\Gamma_e}$ by

$$P_e(\gamma) = P(\gamma) \cap \Gamma_e,$$

we see that the images of P_e are nonempty. This is half of our requirements.

For η a simple closed curve not containing ∞, we denote by $D(\eta)$ the component of $S^2 - \{\eta\}$ not containing ∞.

We call a curve $\gamma \in \Gamma_e$ *innermost* if $D(\gamma)$ contains no essential curves. Let $\Gamma_i = \{\gamma \in \Gamma_e \,|\, \gamma \text{ is innermost}\}$. There is always an innermost curve in Γ_e so Γ_i is not empty.

Proposition 5.7. *If $\gamma \in \Gamma_i$ then exactly one component of $f^{-1}(\gamma)$ is essential; moreover, this curve is innermost.*

Proof. First let us see that if $\gamma \in \Gamma_i$ is an innermost curve and a component γ' of $f^{-1}(\gamma)$ homotopic to an essential curve $\eta \in \Gamma_e$, then η is innermost. By Lemma 5.1 the components of $f^{-1}(D(\gamma))$ are all homeomorphic to disks; precisely one of these components U' is bounded by γ'. Moreover, one of the two components of $S^2 - \{\eta\}$ contains the same post critical points as U'; we call this component U. Clearly U is the bounded component of $S^2 - \{\eta\}$, since ∞ is not in U'.

We will show that U contains no essential curves of Γ. If there were such a curve $\beta \in \Gamma_e$, then β must separate some postcritical points in U from others. Also, β must, by minimality of Γ, be homotopic rel P_f to a component δ' of the inverse image of some curve $\delta \in \Gamma_e$, which of course cannot be in $D(\gamma)$ since γ is innermost. Hence δ' lies outside U', and cannot separate any postcritical points in U' from others. See Figure 5.8. This proves that η is innermost.
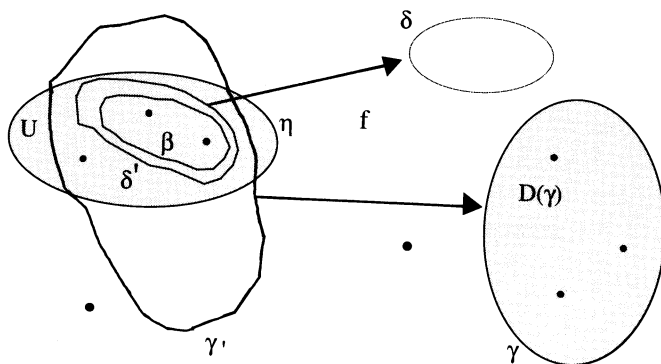

FIGURE 5.8. $\gamma \in \Gamma_i$ implies $\eta \in \Gamma_i$.

Remarks. It should be noted that the minimality condition was used crucially in the proof, and without it Proposition 5.7 is false. Also, minimality implies that $\Gamma_i = \Gamma_e$.

Lemma 5.9. *The images under P_e of distinct curves are nonempty, disjoint sets of curves.*

Proof. We have already seen that the images of P_e are nonempty. Now suppose $\gamma_1, \gamma_2 \in \Gamma_i$, $\gamma_1 \neq \gamma_2$ and $\eta \in P_e(\gamma_1) \cap P_e(\gamma_2)$. Let γ_1' and γ_2' be respective components of $f^{-1}(\gamma_1)$ and $f^{-1}(\gamma_2)$ that are homotopic rel P_f to $\eta \in \Gamma_i$. Let $X = P_f \cap D(\eta)$; then $f(X) \subset D(\gamma_1)$ and $f(X) \subset D(\gamma_2)$. But $D(\gamma_1)$ and $D(\gamma_2)$ are innermost and hence disjoint—a contradiction. $\square$

Since P_e is a map from a finite set into its power set whose images are nonempty and disjoint, the images must be singletons. This shows that exactly one component of $f^{-1}(\gamma)$ is essential.

QED Proposition 5.7

The map P_e induces a permutation that by minimality of Γ must be a cyclic permutation. Consider the cycle $\gamma_0 \mapsto \gamma_1 \mapsto \cdots \mapsto \gamma_n = \gamma_0$ arising from the action of P^{-1}.

Lemma 5.10. *The elements of the cycle $\gamma_0 \mapsto \gamma_1 \mapsto \cdots \mapsto \gamma_n = \gamma_0$ (in Γ_i) form a Levy cycle.*

Proof. We know that for each $i = 0, \ldots, n-1$, γ_i is the only essential curve homotopic to a component of $f^{-1}(\gamma_{i+1})$, so $f_\Gamma(\gamma_{i+1}) = b_i \gamma_i + \beta$ where $\beta \in \mathbb{R}^\Gamma$ has only negligible components. Since $f^{-1}(D(\gamma_{i+1}))$ is a union of disjoint disks that do not contain ∞, there is only one component γ_i' of $f^{-1}(\gamma_{i+1})$ homotopic to γ_i, so $b_i = 1/d_i$ where d_i is the degree of $f: \gamma_i' \to \gamma_{i+1}$.

We will be done if we can show that $d_i = 1$ for $i = 0, \ldots, n-1$.

Order

$$\Gamma = \{\gamma_0, \ldots, \gamma_{n-1}, \text{ negligible curves}\},$$

so the matrix of the Thurston linear transformation becomes

$$f_\Gamma = \begin{bmatrix} A & 0 \\ * & B \end{bmatrix}$$

where

$$A = \begin{bmatrix} 0 & b_0 & \cdots & 0 \\ 0 & 0 & \ddots & \\ 0 & & \ddots & b_{n-2} \\ b_{n-1} & & 0 & 0 \end{bmatrix},$$

$B^m \to 0$ as $m \to \infty$, and we know nothing about $*$.

The characteristic polynomial of f_Γ is

$$\det(\lambda I - A)\det(\lambda I - B) = (\lambda^n - b_0 \cdots b_{n-1})\det(\lambda I - B)$$

and since the roots of $\det(\lambda I - B)$ all have absolute value smaller than 1, we see that $b_0 \cdots b_{n-1} = 1$ since $\lambda(f_\Gamma) \geq 1$. This implies that $b_0 = b_1 = \cdots = b_{n-1} = 1$. □

<div align="right">QED Theorem 5.5</div>

Proposition 5.11. *Let* $\Lambda = \{\gamma_0, \ldots, \gamma_k = \gamma_0\}$ *be a Levy cycle, and let* γ_i' *be the component of* $f^{-1}(\gamma_{i+1})$ *homotopic ref* P_f *to* γ_i. *Then* $f\colon D(\gamma_i') \to D(\gamma_{i+1})$ *is a homeomorphism, and in particular* $D(\gamma_i')$ *cannot contain any critical points.*

Proof. The mapping $f\colon D(\gamma_i') \to D(\gamma_{i+1})$ is proper, hence has a degree that is equal to the degree of its restriction to the boundary, i.e., it has degree 1. □

Proposition 5.12. *Let* $\Lambda = \{\gamma_0, \ldots, \gamma_k = \gamma_0\}$ *be a Levy cycle consisting of innermost curves. Then the* $D(\gamma_i)$ *contain only periodic postcritical points of* f.

Proof. Since $\mathrm{Card}(P_f) < \infty$ every critical point must iterate onto a periodic cycle. if $x \in P_f$ is contained in $D(\gamma_i)$ for $\gamma_i \in \Lambda$, then $f(x) \in D(\gamma_{i+1})$. So $\bigcup D(\gamma_j)$ will contain all the iterates of x, and in particular, $\bigcup D(\gamma_j)$ will contain whatever periodic cycle x lands on or is part of. Suppose y is not periodic, but its image $z = f(y) \in D(\gamma)$ is periodic. Now for $\gamma \in \Lambda$, $f^{-1}(D(\gamma))$ consists of components homeomorphic to disks, only one of which has a boundary γ' homotopic rel P_f to a curve in Λ. By hypothesis, the curves of the Levy cycle are innermost, so only one curve can contain the periodic inverse image of z, and this curve must be homotopic to γ'. However, $y \notin \gamma'$ since $f\colon D(\gamma') \to D(\gamma)$ is one-to-one.

Since $y \notin D(\gamma')$ and y maps to z, we know that $y \notin \bigcup D(\gamma_j)$. So none of the preperiodic postcritical points are in $\bigcup D(\gamma_j)$. □

Corollary 5.13 (Berstein-Levy). *If* f *is a postcritically finite topological polynomial such that every critical point lands in a period cycle that contains a critical point then there is no Thurston obstruction.*

Proof. Suppose a Thurston obstruction exists. Since every periodic cycle contains a critical point, and since, by Proposition 5.12, the disks of the curves of the Levy cycle contain periodic points only, there is a curve γ_{j+1} such that the component of $f^{-1}(\gamma_{j+1})$ homotopic to γ_j contains a critical point; this contradicts Proposition 5.11. □

The following example illustrates the discussion above.

Example 5.14. Consider the pseudo-PPDFA

$$\Theta = \{\{25/216, 169/216\}, \{91/216, 163/216\}\} = \{\Theta_1, \Theta_2\}$$

given in Example 2.9. Let

$$T_\Theta = \{m_3^{\circ k}(\theta) \mid k \geq 0, \ \theta \in \Theta_1 \cup \Theta_2\}.$$

Consider the graph S (related but not quite the same as the spiders we will soon be seeing) that consists of the radial lines $\overline{e^{2\pi i \theta}, \infty}$, together with the segments

$$\overline{e^{2\pi i 25/216}, e^{2\pi i 169/216}} \quad \text{and} \quad \overline{e^{2\pi i 91/216}, e^{2\pi i 163/216}},$$
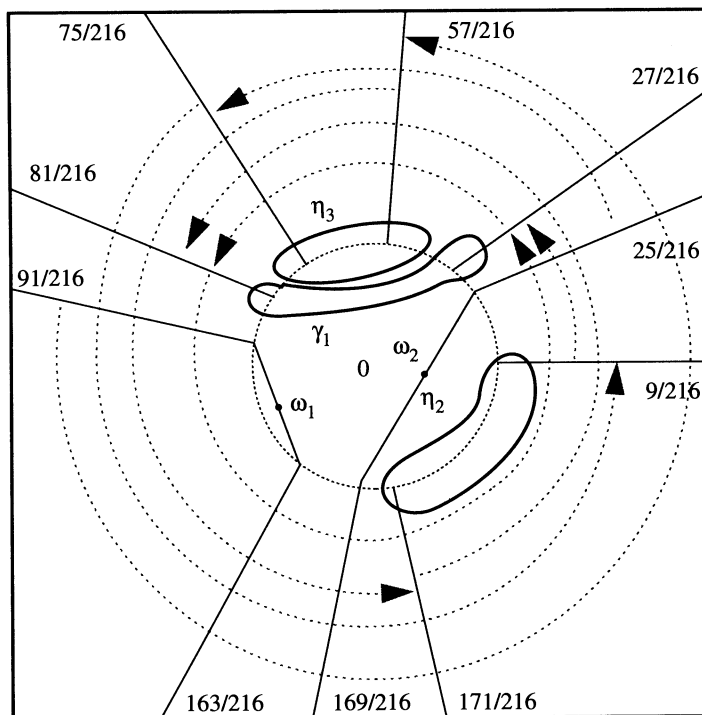
FIGURE 5.15. The set S with a Thurston obstruction.

and add points ω_1 and ω_2 in the middles of the segments above. It is clearly possible to map this graph S to itself by mapping each component of S.– $\{\omega_1, \omega_2\}$ with angle θ to the component with angle $m_3(\theta)$. This mapping extends to a branched mapping $S^2 \to S^2$, which is fairly easy to visualize. See Figure 5.15.

This mapping has a Thurston obstruction consisting of the curves marked $\Gamma = \{\gamma_1, \eta_2, \eta_3\}$. The inverse image of γ_1 consists of one curve homotopic to itself and one curve homotopic to η_2, the inverse image of η_2 consists of a curve homotopic to η_3 and a peripheral curve, and the inverse image of η_3 consists of two peripheral curves.

Thus only γ_1 is essential, and it itself forms a Levy cycle.

6. EXTENDING MAPS ON FINITE GRAPHS

In this section we discuss how maps on graphs in S^2 may be extended to maps of the whole sphere, giving a criterion for extendibility. While we are really dealing with $\overline{\mathbb{C}}$, we prefer to use the notation S^2 to emphasize the fact that the complex structure is not relevant for this construction. In §7, we construct a graph and a graph map that we will extend to the sphere using the tools in this section.

A fundamental observation that we will use is

Lemma 6.1. *If $g: S^1 \to S^1$ is an orientation-preserving homeomorphism, then there exists an orientation-preserving homeomorphic extension $\tilde{g}: \overline{D} \to \overline{D}$ such that $\tilde{g}|_{S^1} = g$. Moreover, $\tilde{g}$ is unique up to isotopy rel S^1.*

Proof (Alexander's Trick). We define $\tilde{g}(re^{i\theta}) = rg(e^{i\theta})$, which is clearly an orientation-preserving homeomorphism. We now show that if two homeomorphisms on $\overline{D}$ both agree on S^1, then they are isotopic. It is sufficient to show that if $f: \overline{D} \to \overline{D}$ is the identity on S^1 then it is isotopic to the identity on D. The isotopy $H: [0, 1] \times \overline{D} \to \overline{D}$ is given by

$$H(t, re^{i\theta}) = \begin{cases} \text{identity} & \text{on } \overline{D} - D_t, \\ tf(re^{i\theta}/t) & \text{on } D_t, \end{cases}$$

where $D_t = \{z \in \mathbb{C} : |z| < t\}$. $\quad\square$

Corollary 6.2. *If the orientation-preserving homeomorphisms $f, g: S^1 \to S^1$ are isotopic rel some finite number of points $X \subset S^1$, then the extensions $\tilde{g}, \tilde{f}: \overline{D} \to \overline{D}$ are isotopic rel X.*

Proof. It is clear that the isotopy between the extensions is given by the extensions of the isotopy between the maps f and g. $\quad\square$

A *finite graph* Γ is the quotient of a finite disjoint union of *arcs* arcs(Γ), which are sets homeomorphic to closed intervals, by an equivalence relation on the (finite) set of endpoints verts(Γ) (called *vertices*) of the arcs. A *finite embedded graph* is the image by a homeomorphism from a finite graph into S^2.

Let X_1 and X_2 be spaces homeomorphic to S^2, and let $\Gamma_1 \subset X_1$ and $\Gamma_2 \subset X_2$ be connected finite embedded graphs. Let $f: \Gamma_1 \to \Gamma_2$ be a continuous mapping, which is injective on arcs, and such that forward and inverse images of vertices are vertices. We call such a map a *graph map*. If a graph map f has an extension $\overline{f}: X_1 \to X_2$ that is an orientation-preserving branched map that is injective when restricted to each component of $X_1 - \Gamma_1$, then we call $\overline{f}$ a *regular extension*.

Corollary 6.3. *Let $f, g: \Gamma_1 \to \Gamma_2$ be graph maps coinciding on verts(Γ_1) such that for each $\gamma \in$ arcs(Γ_1) we have $f(\gamma) = g(\gamma)$, and suppose that f and g have regular extensions $\overline{f}, \overline{g}: X_1 \to X_2$. Then there is a map $\psi: X_1 \to X_1$ such that $\overline{f} = \overline{g} \circ \psi$ and ψ is a homeomorphism isotopic to the identity rel verts(Γ_1).*

Proof. Since Γ_1 is connected, each component U of $X_1 - \Gamma_1$ is homeomorphic to a disk. Since $\overline{f}$ and $\overline{g}$ are injective on U, the images $\overline{f}(U)$ and $\overline{g}(U)$ are homeomorphic to disks. The map ψ is defined on each component U by $(\overline{g}|_U)^{-1} \circ \overline{f}$ and on each arc γ by $(\overline{g}|_\gamma)^{-1} \circ \overline{f}$. Corollary 6.2 shows that $\psi|_{\overline{U}}$ is isotopic to the identity rel verts(Γ_1). The isotopies can clearly be chosen to coincide on the shared boundaries of the components, creating an isotopy on the whole sphere. $\quad\square$

For convenience in proving the next proposition, we will assume that both Γ_1 and Γ_2 are piecewise-linear, and that the graph map f preserves length near the vertices. If ∞ is in a graph then piecewise-linear means that an arc containing

∞ is straight when projected into $\mathbb{C}$. By possibly adding more vertices, we can assume that each arc is in fact a straight segment. The length-preserving properties of f are used only to minimize the notation.

For each vertex $x \in \mathrm{verts}(\Gamma_1)$, let D_x be the disk of radius r around x, we will assume that r is chosen sufficiently small so that D_x and D_y are disjoint when $x \neq y$; that no arcs enter D_x other than those leading to x; that these arcs are straight in D_x; and that f is length preserving there.

We can now define the sectors at each vertex x as the components of $D_x - \Gamma_1$, and extend f sector by sector to a map $\tilde{f}$ on each D_x. Let γ_1 and γ_2 be two arcs bounding a sector with γ_2 following γ_1 in counter-clockwise order. Choose θ_1, θ_2 arguments for these arcs satisfying $0 < \theta_2 - \theta_1 \leq 2\pi$ (if there is only one arc leading to x, then this condition forces us to take $\theta_2 = \theta_1 + 2\pi$). Let θ'_1 and θ'_2 be the arguments of the images of the γ_i by f. Let (ρ, θ) be the polar coordinates centered at x, and (ρ', θ') the polar coordinates centered at $f(x)$. The extension is given by $(\rho, \theta_1 + \theta) \mapsto (\rho', \theta'_1 + \theta')$ such that $\rho = \rho'$ and $\theta' = ((\theta'_2 - \theta'_1)/(\theta_2 - \theta_1))\theta$. (If f is not distance preserving, then ρ' will be some function of ρ.) This formula says that we map sectors in the domain in a counter-clockwise way onto sectors in the range.

Proposition 6.4. *The map* $f: \Gamma_1 \to \Gamma_2$ *has a regular extension if and only if for every vertex* $y \in \Gamma_2$ *and every component* U *of* $X_1 - \Gamma_1$, $\tilde{f}$ *is injective on*

$$\left(\bigcup_{x \in f^{-1}(y)} D_x \right) \cap U.$$

The extension $\overline{f}$ *may have critical points only at the vertices of* Γ_1.

Proof. We first further extend $\tilde{f}$ to neighborhoods T_a of the arcs $a \in \mathrm{arcs}(\Gamma_1)$. There can be no obstruction to this, since the extension at each vertex has been chosen to be orientation preserving. On each linear part of Γ_1 we extend f linearly by interpolation in an orientation-preserving way along small circular arcs in ∂D_x with $x \in \mathrm{verts}(\Gamma_1)$, as is suggested by Figure 6.5.

Let

$$\widetilde{\Gamma}_1 = \bigcup_{x \in \mathrm{verts}(\Gamma_1)} D_x \cup \bigcup_{a \in \mathrm{arcs}(\Gamma_1)} T_a.$$

Since Γ_1 is connected, the components of $X_1 - \Gamma_1$ are homeomorphic to disks. If for each component U the map $\tilde{f}$ is injection on $U \cap \partial\widetilde{\Gamma}_1$ then we can extend $\tilde{f}$ to all of U in an injective way using Lemma 6.1. This will yield a branched map $\overline{f}: X_1 \to X_2$ that coincides with f on Γ_1.

By our hypothesis, $\tilde{f}$ is injective on $\bigcup_{x \in \mathrm{verts}(\Gamma_1)} D_x \cap U$. Suppose $\tilde{f}$ were not injective on $U \cap \bigcup_{a \in \mathrm{arcs}(\Gamma_1)} T_a$. Then $\tilde{f}$ would not be injective on the neighborhoods in U of two arcs in Γ_1. This implies that $\tilde{f}$ would not be injective on the sectors that are in U at the end of these arcs. This contradicts the hypothesis. This gives the result in one direction.
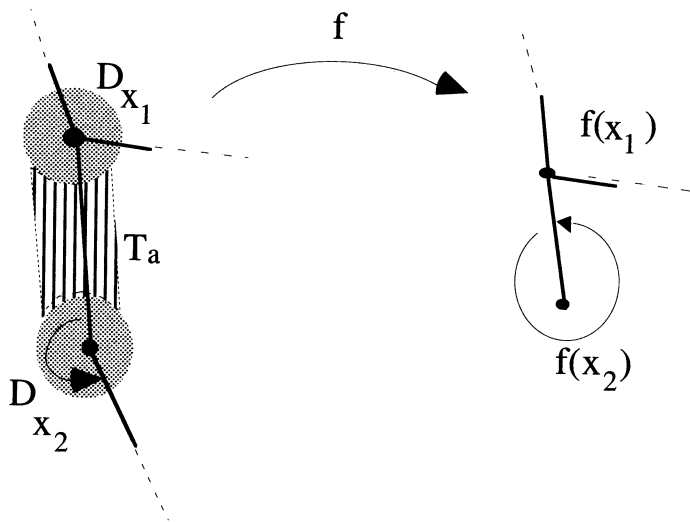
The converse is immediate.

FIGURE 6.5. Extending to a neighborhood of the graph.

Note that the map $\overline{f}$ can have critical points only at the vertices of the graph Γ_1, since it is locally injective elsewhere. The degree of the critical point can be computed from the number of times that the neighborhood D_x of a vertex x "wraps around" its image. □

Corollary 6.6. *Let* $\Gamma_1, \Gamma_2 \subset S^2$ *be connected graphs, and* $h: \Gamma_1 \to \Gamma_2$ *a homeomorphism. Then* h *extends to a homeomorphism* $\overline{h}: S^2 \to S^2$ *if and only if* h *preserves the circular order at all the vertices of* Γ_1.

Proof. There is only one vertex x' in $h^{-1}(x)$ for each vertex x of Γ_2, and the map induced on $D_{x'}$ is injective since h preserves circular order. □

7. DEFINING THE BRANCHED MAP ON S^2

In this section we construct a branched map $\overline{f}: S^2 \to S^2$ from a PPDFA $\Theta = \{\Theta_1, \ldots, \Theta_n\}$. The branched map is constructed to mimic the dynamical behavior of a polynomial that can be marked with the angles $d\Theta_1, \ldots, d\Theta_n$. The main complication is that angles may be related by $\sim_\Theta$, in which case the external rays with those angles land at the same point in the filled-in Julia set for the polynomial.

We will begin by defining the Θ-spider S_Θ, which is a graph corresponding to the union of the marked rays to critical points and their images, and a map $f_\Theta: S_\Theta \to S_\Theta$ that corresponds to the polynomial. We will then extend f_Θ to a Thurston mapping $\overline{f}_\Theta: S^2 \to S^2$.

Construction of S_Θ. If $E \subset \mathbb{T}$ is a finite subset, denote by $\mu(E)$ the center of mass

$$\mu(E) = \frac{1}{\text{Card}(E)} \sum_{\theta \in E} e^{2\pi i \theta}.$$

Consider the set

$$T = \{d^p\theta \mid \theta \in \Theta_i, \ p \geq 0, \ 0 \leq i \leq n\};$$

since the θ are rational, it is clear that T is a finite set.

For $\theta \in T$, let $[\theta]$ be the $\sim_{\Theta}$ equivalence class in T containing θ. Define the leg L_θ by the union of straight segments

$$L_\theta = \overline{0\mu([\theta])e^{2\pi i\theta}} \cup \overline{e^{2\pi i\theta}\infty},$$

and the Θ-spider

$$S_{\Theta} = \bigcup_{\theta \in T} L_\theta.$$

Construction of f_{Θ}. To define a map $f_{\Theta}: S_{\Theta} \to S_{\Theta}$ that maps L_θ to $L_{d\theta}$, we need to know that legs intersect only at their endpoints. Proposition 7.1 says that this is the case.

Proposition 7.1. *The equivalence classes of $\sim_{\Theta}$ are unlinked.*

The proof requires a number of lemmas of intrinsic interest.

Let $C_1, \ldots, C_d$ be the equivalence classes of the equivalence relation on $\mathbb{T} - \bigcup_{i=1}^n \Theta_i$ given by x equivalent to y when $\{x, y\}$ are Θ-unlinked. Note that two angles x and y are Θ-unlinkable precisely when they both belong to the closure of some class.

The proofs of Lemma 7.2(1)–(4) are left to the reader.

Lemma 7.2. (1) *Let $\{x, y\} \subset \overline{C}_i$ and $\{u, v\} \subset \overline{C}_j$, for $i \neq j$. Then if $\{x, y\} \cap \{u, v\} = \varnothing$ then $\{x, y\}$ and $\{u, v\}$ are unlinked.*

(2) *Let $\Psi \subset \mathbb{T}$ be a finite set. Then Ψ is Θ-unlinked if and only if $\Psi \subset C_i$ for some i.*

(3) *The number d of equivalence classes satisfies $d = 1 + \sum_{i=1}^n (\mathrm{Card}(\Theta_i) - 1)$.*

(4) *Each C_i is a finite union of disjoint open arcs of $\mathbb{T}$ with the following property: the set of boundary points of the arcs composing each C_i can be written in circular order as $\{x_0, y_0, x_1, y_1, \ldots, x_m = x_0, y_m = y_0\}$ such that for every i, the points x_i, y_i belong to Θ_{k_i} for some k_i, and $k_i \neq k_j$ if $i \neq j$. See Figure 7.2.*
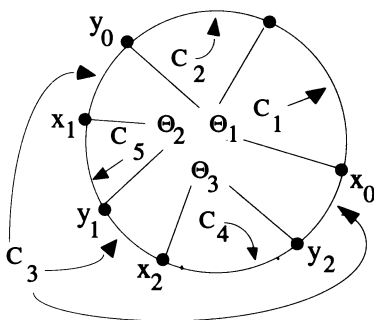

FIGURE 7.2. Labels on the boundary of C_3.

It follows from 7.2(4) that for each class C_i, we can construct a topological circle $\widehat{C}_i$ by identifying the points x_j and y_j. Since x_j and y_j are part

of the same preangle, the map $m_d \colon x \mapsto dx$ induces a continuous mapping $m_{d,i} \colon \widehat{C}_i \to \mathbb{T}$.

Lemma 7.3. *The map $m_{d,i}$ is an orientation-preserving homeomorphism.*

Proof. Clearly $m_{d,i}$ is a local orientation-preserving homeomorphism; since $\widehat{C}_i$ is compact it is proper, so it is a finite covering map, of some positive degree. Since $\bigcup \overline{C}_i = \mathbb{T}$, the sum of the degrees of the $m_{d,i}$ must be d. By Lemma 7.2(3), there are d classes C_i, so each $m_{d,i}$ must have degree 1. $\square$

Corollary 7.4. *The circular order of the points of the closure of an equivalence class $\overline{C}_i$ is preserved by $z \mapsto dz$.*

Note that the endpoints of the arcs of the C_i may be identified by $z \mapsto dz$, but this does not alter the circular order.

The following three lemmas say that we will be able to connect the legs with angles that are related by $\sim_{\Theta}$ without any legs crossing.

Lemma 7.5. *Suppose θ_1 and θ_2 are Θ-related, θ_1' and θ_2' are Θ-related, and that $\{\theta_1, \theta_2\} \not\sim_{\Theta} \{\theta_1', \theta_2'\}$. Then $\{\theta_1, \theta_2\}$ and $\{\theta_1', \theta_2'\}$ are unlinked.*

Proof. Suppose that $\{\theta_1, \theta_2\}$ and $\{\theta_1', \theta_2'\}$ are linked. Then the angles have circular order $\theta_1, \theta_1', \theta_2, \theta_2'$ in $\mathbb{T}$. Since $\theta_1 \not\sim_{\Theta} \theta_1'$, $\{\theta_1, \theta_2\} \cap \{\theta_1', \theta_2'\} = \varnothing$, so the four points are distinct.

For every m, $d^m \theta_1$ and $d^m \theta_2$ are Θ-unlinkable; hence both belong to the closure $\overline{C}_{i_m}$ of some class, and similarly $d^m \theta_1', d^m \theta_2' \in \overline{C}_{j_m}$. Since $\{\theta_1, \theta_2\}$ and $\{\theta_1', \theta_2'\}$ are linked and disjoint, Lemma 7.2(1) implies $i_0 = j_0$. But Corollary 7.4 implies that $d\theta_1, d\theta_1', d\theta_2, d\theta_2'$ have the same circular order as $\theta_1, \theta_1', \theta_2, \theta_2'$, and hence they are linked.

We now proceed by induction on m. If $i_m = j_m$ and $\{d^m \theta_1, d^m \theta_2\}$ and $\{d^m \theta_1', d^m \theta_2'\}$ are linked and disjoint, then as above we see that $i_{m+1} = j_{m+1}$. This implies that $\{\theta_1, \theta_2\} \sim_{\Theta} \{\theta_1', \theta_2'\}$, a contradiction. $\square$

Lemma 7.6. *Suppose $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are unlinked, and that $\{\theta_1, \theta_2\}$ and $\{\psi_2, \psi_3\}$ are unlinked. Then $\{\theta_1, \theta_1\}$ and $\{\psi_1, \psi_3\}$ are unlinked.*

Proof. This is easy to see. $\square$

Proof of Proposition 7.1. Let Q and Q' be distinct linked $\sim_{\Theta}$ equivalence classes. Then there exist $\{\theta_1, \theta_2\} \subset Q$ and $\{\theta_1', \theta_2'\} \subset Q'$ with $\{\theta_1, \theta_2\}$ and $\{\theta_1', \theta_2'\}$ linked. Now $\theta_1 \sim_{\Theta} \theta_2$ and $\theta_1' \sim_{\Theta} \theta_2'$ so there exist m, m' (with m, m' possibly 0) such that in the sequences $\theta_1, \psi_1, \ldots, \psi_m, \theta_2$ and

θ'_1, ψ'_1, ..., $\psi'_{m'}$, θ'_2 each angle is Θ-related to the next. Using Lemma 7.5, we know that $\{\theta_1, \psi_1\}$ and $\{\theta'_1, \psi'_1\}$ are unlinked. By using Lemma 7.6 inductively, we see that $\{\theta_1, \psi_1\}$ and $\{\theta'_1, \theta'_2\}$ are unlinked. Now we use Lemma 7.6 inductively again and see that $\{\theta_1, \theta_2\}$ and $\{\theta'_1, \theta'_2\}$ are unlinked, a contradiction.   □

Choose a map f on S_Θ by letting

$$f|_{L_\theta} : L_\theta \to L_{d\theta}$$

be any homeomorphism that fixes ∞; this can be done because if $\theta_1 \sim_\Theta \theta_2$ then $d\theta_1 \sim_\Theta d\theta_2$, and such an f satisfies $f(\mu[\theta]) = \mu[d\theta]$. It is possible to write down a formula for f on each leg L_θ, but since any homeomorphism will do fine, it would not be instructive to do so.

Extending f_Θ to $\overline{f}_\Theta$.

Lemma 7.7. *The set S_Θ is a finite connected graph, with vertices*

$$\{\mu([\theta]) \mid \theta \in T\} \cup \{\infty\}.$$

Proof. Since each leg contains ∞, S_Θ is connected. To show it is a graph with the given vertices, we must show that the L_θ intersect only at ∞ and $\mu([\theta])$. Clearly, outside the unit disk, the L_θ intersect only at ∞. One can easily show that if two finite subsets of the unit circle are unlinked, then their convex hulls do not intersect. Thus, the convex hulls of the points of each equivalence class in $T/\sim_\Theta$ do not intersect. For each equivalence class, the points $\mu([\theta])$ and the segments of $\overline{D} \cap L_\theta$ that intersect $\mu([\theta])$ are contained in these convex hulls, and hence they do not intersect.   □

Theorem 7.8. *The map $f : S \to S$ extends to a topological polynomial $\overline{f}_\Theta : S^2 \to S^2$ of degree*

$$d = 1 + \sum_{i=1}^{n} (\mathrm{Card}(\Theta_i) - 1)$$

with strictly preperiodic critical points $\mu([\Theta_i])$ and ∞ of local degree $\mathrm{Card}(\Theta_i)$ and d, respectively. Moreover, the Thurston equivalence class of $\overline{f}_\Theta$ is independent of the choice of extension or the choice of the homeomorphisms $f_\Theta|_{L_\theta}$.

Proof. First we will show that the map f_Θ satisfies the hypothesis of Proposition 6.4 of §6, and can therefore be extended to an orientation-preserving branched map whose only critical points are among the $\mu([\theta])$, $\theta \in T$ and ∞.

Let $\tilde{f}$ be an extension of f_Θ to small disks D_x of the vertices $x \in \mathrm{verts}(S_\Theta)$ as in §6. We need to show that for every vertex y in $f_\Theta(S_\Theta)$ and every component U of $S^2 - S_\Theta$, the map $\tilde{f}$ is injective on $(\bigcup_{x \in f^{-1}(y)} D_x) \cap U$.

In fact we will show that f_Θ is injective on the vertices of ∂U. Since, at each vertex $x \in \partial U$ except ∞, the intersection $U \cap D_x$ is a single sector, this implies the hypothesis at finite vertices. To do this we need the following lemma.

Lemma 7.9. *If* $d\theta_1 \sim_{\Theta} d\theta_2$, *and there exists* k *so that* θ_1 *and* θ_2 *are in* $\overline{C}_k$, *then* $\theta_1 \sim_{\Theta} \theta_2$.

Proof. By our hypothesis, there exist n angles $\varphi_1, \dots, \varphi_n$ (with n possibly zero) such that in the sequence $d\theta_1, \varphi_1, \varphi_2, \dots, \varphi_n, d\theta_2$, each angle is Θ-related to the next. Because m_d maps $\overline{C}_k$ surjectively onto $\mathbb{T}$ (Lemma 7.3), there also exist angles $\varphi'_1, \dots, \varphi'_n \in \overline{C}_k$ such that $d\varphi'_i = \varphi_i$. We now have that consecutive angles in the sequence $\theta_1, \varphi'_1, \varphi'_2, \dots, \varphi'_n, \theta_2$ are Θ-related. So $\theta_1 \sim_{\Theta} \theta_2$. $\square$

The intersection $\overline{U} \cap S^1$ is a subset of some $\overline{C}_l$ The finite vertices in ∂U are a subset of

$$\{\mu([\theta]) | \theta \in T \cap \overline{C}_l\}.$$

If two such vertices x, y map to the same vertex $f_{\Theta}(x)$, then the angles of the legs attached to x and y must map to the angles of the legs attached to $f_{\Theta}(x)$. The angles of the legs attached to $f_{\Theta}(x)$ are equivalent, and the angles attached to x and y are in the same $\overline{C}_l$, so these angles are equivalent by Lemma 7.9. Thus, we must have $x = y$.

In the case $x = \infty$ we have to see that $\tilde{f}$ is injective on $D_{\infty} \cap U$. At ∞, $\tilde{f}$ is given by $re^{i\theta} \mapsto re^{id\theta}$, the angular coordinates of the points in U near ∞ are all in some C_l, and by Lemma 7.3 $\tilde{f}$ will be injective on $D_{\infty} \cap U$.

Applying Proposition 6.4 constructs a branched map $\overline{f}_{\Theta}: S^2 \to S^2$ to which we will later apply Thurston's theorem.

By construction, the critical points of $\overline{f}_{\Theta}$ must be among the vertices of the graph S.

Let $x = \mu(E)$ be the vertex corresponding to the equivalence class E. If E contains no critical angles, i.e., $E \cap \Theta_i = \varnothing$ for all $\Theta_i \in \Theta$. Then $E \subset C_i$ (by Proposition 7.1 and Lemma 7.2(2)) for some i, and by Lemma 7.3 $m_d|_E$ is injective and preserves circular order. This means that the vertex x is not a critical point. See Figure 7.10.

If $\Theta_i \subset E$, then all the legs L_θ with $\theta \in \Theta_i$ map to a single leg, and the restriction of m_d to the angles of legs in any sector bounded by the L_θ with $\theta \in \Theta_i$ is also injective and order preserving. This means that $\overline{f}_{\Theta}$ maps such
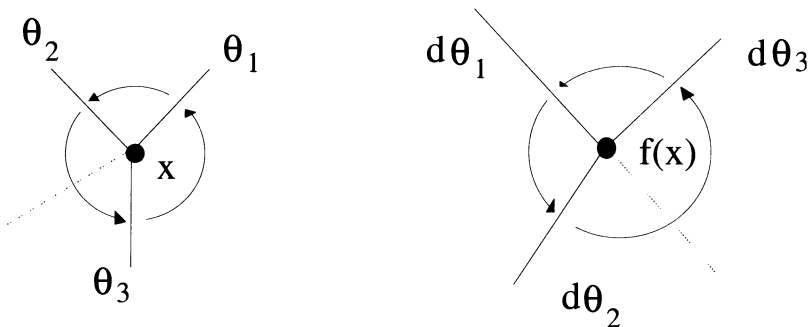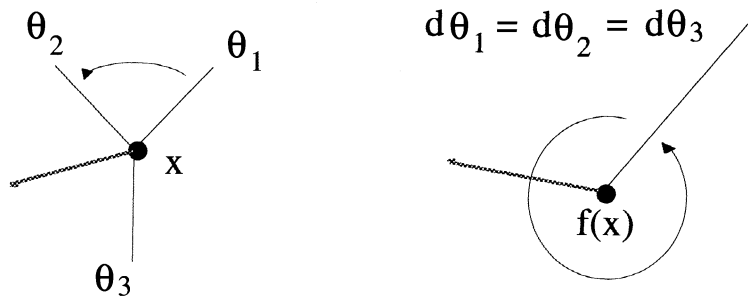

FIGURE 7.10. Local degree is one at x.

FIGURE 7.11. A critical point of degree 3.

a sector of D_x to the full disc $D_{f_\Theta(x)}$ with degree 1, and so x is a critical point with local degree $\text{Card}(\Theta_i)$. See Figure 7.11.

This covers all the vertices of the spider, except the point at ∞. Since the map on the legs is induced by $z \mapsto z^d$, the only way in which the degree could be lower than d is if a sector at infinity has an opening greater than $1/d$. This does not occur since m_d is injective on C_i.

The degree of $\overline{f}_\Theta$ is now immediate from the Riemann-Hurwitz formula.

Clearly the orbits of the critical points are finite, so $\overline{f}_\Theta$ is a Thurston mapping. It remains to show that no critical point is periodic, which is a bit more delicate than one might imagine.

Lemma 7.12. *If no Θ_j is periodic, then no $[\Theta_j]$ is periodic.*

Proof. Pick a fixed Θ_k. The orbit of Θ_k is eventually periodic, and the periodic part of the orbit contains no angle in any Θ_j, since these are assumed not to be periodic. Each angle in the periodic part of the orbit is therefore in one of the C_i, and these periodic angles will then have an itinerary $C_{k_1}, C_{k_2}, \ldots, C_{k_m}$ (up to cyclic permutation). Let Q be the set of rational angles that have a cyclic permutation of $C_{k_1}, C_{k_2}, \ldots, C_{k_m}$ as itinerary.

For any rational angle θ in the complement of Q and any $\phi \in Q$, we will show that θ is not Θ-related to ϕ. Thus $\theta \nsim_\Theta \phi$ for any $\phi \in Q$, and therefore the equivalence class of Θ_k is not periodic since Θ_k is not in Q.

Suppose there is a rational angle $\theta \notin Q$ and $\phi \in Q$ such that θ is Θ-related to ϕ. Then there is some smallest i such that $d^{i-1}\theta \notin Q$ and $d^i\theta = \psi \in Q$. Since $d^{i-1}\theta \in \partial C_{k_j}$ for some j, and since θ is rational, ψ is either preperiodic or periodic. If it is preperiodic there are two angles (the last preperiodic angle in the orbit of ψ and the last periodic angle in the orbit of ψ) in the same C_{k_l} that map to the same angle, contradicting Lemma 7.3. If ψ is periodic then we get a contradiction also since the inverse image of ψ in C_{k_j} is accounted for so we cannot have another inverse in ∂C_{k_j}. $\square$

Finally, that the Thurston equivalence class of $\overline{f}_\Theta$ is independent of the choice of extension or the choice of the homeomorphisms $f_\Theta|_{L_\theta}$ follows from

Corollary 6.3. If f_0 and f_1 are two such extensions, the diagram

$$
\begin{array}{ccc}
(S^2, P_{f_0}) & \xrightarrow{\ \psi\ } & (S^2, P_{f_1}) \\
f_0 \downarrow & & f_1 \downarrow \\
(S^2, P_{f_0}) & \xrightarrow{\text{identity}} & (S^2, P_{f_1})
\end{array}
$$

commutes and ψ is isotopic to the identity rel P_{f_0}.

<div align="right">QED Theorem 7.8</div>

We will use the notation f_Θ to refer to $\overline{f}_\Theta$ when there is no danger of confusion.

8. The topological polynomial f_Θ has no Thurston obstruction

This section is devoted to proving the following result.

Theorem 8.1. *The topological polynomial f_Θ has no Thurston obstruction.*

Proof. Assume that f_Θ has a Thurston obstruction. We know from Theorem 5.5 that if f_Θ has a Thurston obstruction then it has a Levy cycle $\Lambda = \{\gamma_0, \dots, \gamma_k = \gamma_0\}$. We will show that any curve $\gamma \in \Lambda$ can only intersect the L_{Θ_i} in a way that does not separate the points of $P_{f_0} \cap D(\gamma)$. This implies that the angles of the legs to the points contained in the disks of the Levy cycle must be Θ-equivalent, and hence the curves of the Levy cycle are peripheral, a contradiction.

Let $X \subset S^2$ be closed, let γ be a simple closed curve in $S^2 - X$, and let ρ be a union of closed arcs in S^2 with endpoints in X. Define the *geometric intersection number* by

$$
\rho \cdot \gamma = \min_{\gamma' \text{ isotopic to } \gamma \text{ rel } X} \mathrm{Card}(\rho \cap \gamma').
$$

We will say that ρ *intersects* γ *essentially* when $\rho \cdot \gamma \neq 0$.

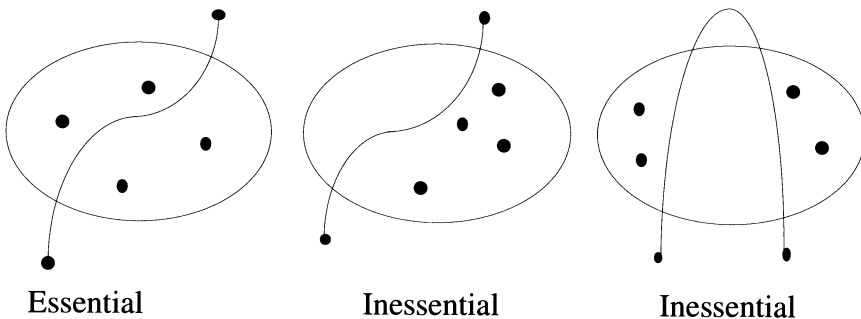| Essential	Inessential	Inessential

FIGURE 8.3. Essential and inessential intersections. The dots represent X.

Of course, $\rho \cdot \gamma$ depends on X, but in our application, X will always be $P_{f_\Theta} \cup \Omega_{f_\Theta}$ and it is safe not to mention it.

Our object is to show that $L_\theta \cdot \gamma = 0$, for any critical leg L_θ with $\theta \in \Theta_1$ and any curve $\gamma \in \Lambda$.

Lemma 8.5. *If* $\theta \in T$ *is a postcritical angle, then*

$$\sum_{\substack{\phi \in T \\ d\phi = \theta}} L_\phi \cdot \gamma_i \leq L_\theta \cdot \gamma_{i+1}.$$

Proof. By an isotopy, move γ_{i+1} rel X to be in minimal position with respect to L_θ, and let γ_i' be the component of $f_\Theta^{-1}(\gamma_{i+1})$ homotopic to γ_i. Then we see

$$\bigcup_{\substack{\phi \in T \\ d\phi = \theta}} L_\phi \cap \gamma_i' \subset (f_\Theta^{-1}(L_\theta \cap \gamma_{i+1})) \cap \gamma_i',$$

and since, by Definition 5.2, f_Θ is injective on γ_i' the result follows. $\square$

Corollary 8.6. *For a periodic* $x \in X$ *and* $\gamma_i \in \Lambda$, $L(f_\Theta^{\circ n}(x)) \cdot \gamma_{j+n} = L(x) \cdot \gamma_j$ *for all* j, n.

Proof. $L(x) \cdot \gamma_j \leq L(f_\Theta(x)) \cdot \gamma_{j+1} \leq L(f_\Theta^{\circ 2}(x)) \cdot \gamma_{j+2} \leq \cdots \leq L(x) \cdot \gamma_j$. $\square$

Lemma 8.7. *If* $x \in X$ *is strictly preperiodic then* $L(x) \cdot \gamma = 0$ *for any* $\gamma \in \Lambda$.

Proof. Let $x \in X$ be a strictly preperiodic point, and let $y = f_\Theta^{\circ k}$ be the first periodic point in its orbit. Let y' be the periodic inverse image of y, and let $y'' = f_\Theta^{\circ k-1}$.

Then for any $\gamma_{i+1} \in \Lambda$, we have

$$L(y') \cdot \gamma_i + L(y'') \cdot \gamma_i \leq L(y) \cdot \gamma_{i+1}$$

by Lemma 8.5, but $L(y') \cdot \gamma_i = L(y) \cdot \gamma_{i+1}$ by Corollary 8.6, so $L(y'') \cdot \gamma_i = 0$. $\square$

Note that

$$S^2 - \bigcup_{i=1}^{n} \bigcup_{\theta \in \Theta_1} L_\theta$$

consists of d simply connected components $\widetilde{C}_1, \ldots, \widetilde{C}_d$ we call *patches*, where $\widetilde{C}_i \cap S^1 = C_i$; in particular there are exactly d of them.

We are now in a position to finish the proof of the theorem. The previous lemma implies that the legs L_θ with $\theta \in \Theta_i$ for any i do not intersect any curve in the Levy cycle essentially. If $\gamma \in \Lambda$, then the orbit of the postcritical points in $D(\gamma)$ and the legs attached to them must have the same itinerary of patches. That is, if $x, y \in D(\gamma) \cap P_{f_\Theta}$, the legs $L(f_\Theta^{\circ j}(x)), L(f_\Theta^{\circ j}(y)) \in \widetilde{C}_{i_j}$ for some i_j. Since the legs $L(f_\Theta^{\circ j}(x)), L(f_\Theta^{\circ j}(y))$ are in the same patch, their angles are in the same equivalence class C_{i_j} for all j. This means that the angles of the legs $L(x), L(y)$ are Θ-equivalent. Since all legs with angles in the same Θ-equivalence class land at the same point by construction, the curves in Λ must be peripheral a contradiction. $\square$

QED Theorem 8.1

9. APPLYING THURSTON'S THEOREM TO f_Θ

In this section, we apply Theorem 4.2 to the branched map $\overline{f}_\Theta : S^2 \to S^2$ constructed from a PPDFA in §7. By Theorem 8.1, we see that there exists a

polynomial p that is Thurston equivalent to f_Θ; i.e., there exist homeomorphisms ϕ, $\phi'\colon S^2 \to \mathbb{P}^1$, isotopic rel P_{f_Θ} and such that the diagram

$$
\begin{array}{ccc}
(S^2, P_{f_\Theta}) & \xrightarrow{\;\phi'\;} & (\mathbb{P}^1, P_p) \\[4pt]
{\scriptstyle f_\Theta}\Big\downarrow & & {\scriptstyle p}\Big\downarrow \\[4pt]
(S^2, P_{f_\Theta}) & \xrightarrow{\;\phi\;} & (\mathbb{P}^1, P_p)
\end{array}
$$

commutes. We want to show that p is naturally a marked polynomial, and that the marking angles come from the original PPDFA Θ. This is not really difficult, but it is not enough to simply normalize p to be monic and centered. There are in general $d-1$ different conjugates of p that are monic and centered, and we need to distinguish one of these.

To do this, we need to show that the spider $\phi(S_\Theta)$ can be embedded in $\mathbb{P}^1$ so that the image is made up of external rays of p. In order to make this embedding unique, we will need to consider the Green's function [JM] G_p of K_p, and to choose for each leg L_θ of S_Θ a parametrization $\rho_\theta\colon [1, \infty] \to L_\theta$ mapping infinity to infinity.

Proposition 9.1. *There exists a unique embedding* $\Psi\colon S_\Theta \to \mathbb{P}^1$ *mapping* P_{f_Θ} *to* P_p *and isotopic rel* P_{f_Θ} *to* ϕ, *such that the image of each leg is an external ray, and*

$$
G_p(\Psi(\rho_\theta(t))) = \log t.
$$

Proof. The mapping Ψ will be a fixed point of a contracting mapping. The contraction will come from the fact that p is expanding in the orbifold metric. Unfortunately, p is strictly but not strongly expanding: the expansion factor tends to 1 near infinity; this requires setting up the space on which the contraction acts so that one never needs to measure distances near infinity. Consider the space $\mathscr{S}_\Theta$ of spider mappings

$$
\mathscr{S}_\Theta = \{\psi\colon S_\Theta \to \mathbb{P}^1 \mid \psi(P_{f_\Theta}) = P_p \text{ and } \psi \text{ isotopic to } \phi|_{S_\Theta} \text{ rel } P_{f_\Theta}\}
$$

and define a mapping $\tau_\Theta\colon \mathscr{S}_\Theta \to \mathscr{S}_\Theta$ as follows. For each θ, let us first find the leg $\tau(\psi(L_\theta))$. There exists a unique component X of $p^{-1}(\psi(L_{d\theta})) - \infty$ that contains $\phi'(\mu[\theta]) = \phi(\mu[\theta])$.

- If $\phi(\mu[\theta])$ is not a critical point the map p will be injective on X; we choose $\tau(\psi(L_\theta)) = X$.
- If $\phi(\mu[\theta])$ is a critical point, of local degree m, then X is the union of m legs, which have a natural order (not simply circular order) since exactly one component of the complement contains $\psi(\mu[d\theta])$. The same construction with ϕ instead of ψ has the same property, and $\phi'(L_\theta)$ has some position in that order. We choose $\tau(\psi(L_\theta))$ to be the leg that occupies the same position.

Define $\tau(\psi)(\rho_\theta(t)) \in p^{-1}(\rho_{d\theta}(t^d))$ to be the unique inverse image in $\tau(\psi(L_\theta))$.

We need to show that $\tau(\psi) \in \mathscr{S}_{\Theta}$. First observe that $\phi'|_{S_{\Theta}}$ and $\tau(\phi|_{S_{\Theta}})$ only differ by the parametrizations of the legs. Choose an isotopy ψ_s, $s \in [0, 1]$, with $\psi_0 = \phi|_{S_{\Theta}}$ to $\psi_1 = \psi$, and lift it to a ψ_s' starting at $\phi'|_{S_{\Theta}}$; i.e., ψ_s' is the spider mapping depending continuously on s, which satisfies $\psi_s \circ f_{\Theta} = p \circ \psi_s'$ and $\psi_0' = \phi'$. This is possible, and the isotopy ψ_s' is unique, by the curve lifting property of covering spaces. For all s, ψ_s' differs from $\tau(\psi_s)$ only by a parametrization of the legs, so $\tau(\psi)$ is in the correct isotopy class.

Unfortunately, there is no metric on $\mathscr{S}_{\Theta}$ for which τ is contracting; and we need to restrict to an invariant subset $\mathscr{S}_{\Theta}'$, given by the following conditions:

- on each leg L_{θ}, ψ maps the segment $\rho_{\theta}([2, \infty])$ to an external ray of p, with $G_p(\psi(\rho_{\theta}(t))) = \log t$.
- $G_p(\psi(\rho_{\theta}([1, 2]))) \leq \log 2$.

It is not quite obvious that $\mathscr{S}_{\Theta}'$ is nonempty, but we will leave this to the reader, as well as the easy verification that it is invariant. We also leave the reader to verify that any two elements of $\mathscr{S}_{\Theta}'$ are isotopic through elements of $\mathscr{S}_{\Theta}'$.

We will put a metric $d(\cdot, \cdot)$ on $\mathscr{S}_{\Theta}'$ as follows:

Choose ψ_0, $\psi_1 \mathscr{S}_{\Theta}'$. Let

$$\delta(\psi_0(\rho_{\theta}(t)), \psi_1(\rho_{\theta}(t)))$$
$$= \inf\{\text{the orbifold length of the curve } \psi_s(\rho_{\theta}(t)) \mid 0 \leq s \leq 1\},$$

where the infimum is taken over all isotopies ψ_s joining ψ_0 and ψ_1 through $\mathscr{S}_{\Theta}'$.

Then define the metric

$$d(\psi_0, \psi_1) = \sup_{\substack{\theta \in T \\ 1 \leq t < 2}} \delta(\psi_0(\rho_{\theta}(t)), \psi_1(\rho_{\theta}(t))).$$

An isotopy between ψ_0 and ψ_1 lifts to an isotopy between $\tau(\psi_0)$ and $\tau(\psi_1)$, and in the orbifold metric the lift of any curve α in the compact set $\{z | G_p(z) \leq \log 2\}$ is shorter than α by a fixed factor $K < 1$. Thus $\tau \colon \mathscr{S}_{\Theta}' \to \mathscr{S}_{\Theta}'$ is strongly contracting, and thus τ has a fixed point Ψ.

For any initial condition ψ we will have that $\tau^n(\psi(\rho_{\theta}(t))) \subset R_{\theta_n}$ for $t \geq 2^{1/d^n}$ and some θ_n. Since the distance between $\tau^n(\psi(\rho_{\theta}(2)))$ and $\Psi(\rho_{\theta}(2))$ goes to zero we must have that θ_n approaches some angle θ'. Therefore Ψ is a spider whose legs are external rays, and satisfies $G_p(\Psi(\rho_{\theta}(t))) = \log t$. □

Proposition 9.2. *There exists a unique affine mapping* $\alpha \colon \mathbb{C} \to \mathbb{C}$ *such that* $p_{\Theta} = \alpha \circ p \circ \alpha^{-1}$ *is centered and monic, and such that for all* $\theta \in T$, *we have that* $\alpha(\Psi)(L_{\theta})$ *is the external ray of* p_{Θ} *at angle* θ.

Proof. Either $0 \in T$, in which case $\Psi(L_0)$ is a distinguished fixed ray of p, or one sector at infinity for (S^2, S_{Θ}) contains the line $[1, \infty]$, and there is a unique distinguished fixed ray in the component of $\mathbb{C} - (K_p \cup \Psi(S_{\Theta}))$ corresponding to that sector.

There are d different affine mappings conjugating p to be monic and centered, since conjugation by a rotation of $2\pi/d$ leaves the polynomial monic. However, there is a unique affine mapping α such that $p_{\Theta} = \alpha \circ p \circ \alpha^{-1}$ is centered and monic, and such that the distinguished fixed ray is R_0 in $\overline{\mathbb{C}} - K_{p_{\Theta}}$.

The circular order of the external ray $\alpha \circ \Psi(L_{\theta})$ and the rays at angles k/d, $k = 0, \dots, d-1$, is the same as the circular order of θ and the elements k/d, $k = 0, \dots, d-1$, of $\mathbb{T}$. Hence the itinerary under p of a leg $\alpha \circ \Psi(L_{\theta})$ with respect to $\mathbb{C} - \bigcup_k R_{k/d}$ coincides with the itinerary of θ under m_d with respect to $\mathbb{T} - \bigcup_k k/d$. Since this itinerary is essentially the development of θ in base d, we have shown that $\alpha \circ \Psi(L_{\theta})$ is the external ray of p_{Θ} at angle θ.   □

10. Proofs of Theorems II and III

We are now in a position to prove the main theorems.

Proof of Theorem II. There is an obvious mapping PA, given in §2, which associates to a marked polynomial a collection of d-preangles. Let $\mathscr{A}$ be the set of PPDFA's.

Let $(p(z); \theta_1, \dots, \theta_n)$ be a marked polynomial of degree d, and

$$PA(p(z); \theta_1, \dots, \theta_n) = (\Theta_1, \dots, \Theta_n).$$

The final result of the previous three sections is to define the map $AP: \Theta \mapsto p_{\Theta}$ from $\mathscr{A}$ to $\mathscr{P}$, and it is clear from Proposition 9.2 that $PA(p_{\Theta}) = \Theta$, i.e., $PA \circ AP = \mathrm{id}$.

It remains to show that

Lemma 10.1. $AP \circ PA = \mathrm{id}$.

Proof. Let

$$PA(p(z); \theta_1, \dots, \theta_n) = \Theta,$$

and let

$$S_{(p; \theta_1, \dots, \theta_n)} = \bigcup_{\theta \in T} R_{\theta}.$$

There are homeomorphisms $h: S_{\Theta} \to S_{(p; \theta_1, \dots, \theta_n)}$ that map the leg L_{θ} to the external ray R_{θ} for all $\theta \in T_{\Theta}$; and they differ only by a parametrizations of the legs. Moreover, such a homeomorphism preserves circular order at the vertices, and by Corollary 6.6 can be extended to a homeomorphism $\overline{h}: S^2 \to \overline{\mathbb{C}}$. The map $h^{-1} \circ p \circ h: (S^2, S_{\Theta}) \to (S^2, S_{\Theta})$ is an extension of a spider map isotopic to f_{Θ} so by Theorem 7.8 $h^{-1} \circ p \circ h$ is equivalent to $\overline{f}_{\Theta}$. Since p is equivalent to any conjugate of itself, we have that p is equivalent to $\overline{f}_{\Theta}$ and thus $AP(\Theta) = (p(z); \theta_1', \dots, \theta_n')$. It follows from Proposition 9.2 that $\theta_i = \theta_i'$ for all $i = 1, \dots, n$.   □

QED Theorem II

Proof of Theorem III. Theorem III is an easy consequence of Theorems I and II.

Proof of Theorem III. In one direction, if p is a centered monic polynomial with critical points $\{\omega_1, \ldots, \omega_n\}$, and if $(\theta_1, \ldots, \theta_n)$ and $(\theta'_1, \ldots, \theta'_n)$ are two markings of p with R_{θ_i} and $R_{\theta'_i}$ both landing at $p(\omega_i)$, then the rays of Θ_i and Θ'_i also land at the same point and by Theorem I, we have $\Theta_i \sim_\Theta \Theta'_i$.

To show the converse, reorder Θ' so that $\Theta_1 \sim_\Theta \Theta'_1, \ldots, \Theta_n \sim_\Theta \Theta'_n$.

Let $(p; \theta_1, \ldots, \theta_n)$ and $(p'; \theta'_1, \ldots, \theta'_n)$ be marked polynomials that satisfy, respectively, $PA(p; \theta_1, \ldots, \theta_n) = \Theta$ and $PA(p'; \theta'_1, \ldots, \theta'_n) = \Theta'$. These exist by the surjectivity of PA. By I, the external rays of p with angles in Θ'_i all land at the same (critical) point as those with angles in Θ_i. Since the elements of Θ' are d-preangles, the external rays of p with angles in Θ'_i all map to a ray landing at the corresponding critical value. So we see that there is a marking $(\theta'_1, \ldots, \theta'_n)$ of p such that

$$\Theta' = PA(p; \theta'_1, \ldots, \theta'_n),$$

and the theorem follows from the injectivity of PA. $\square$

QED Theorem III

REFERENCES

[B] P. Blanchard, *Complex analytic dynamics on the Riemann sphere*, Bull. Amer. Math. Soc. (N.S.) **11** (1984), 85–141.

[CJ] L. Carleson and P. Jones, *On coefficient problems for univalent functions and conformal dimension*, Preprint.

[CE] P. Collet and J.-P. Eckmann, *Iterated maps on the interval as dynamical systems*, Birkhäuser-Boston, Boston, MA, 1980.

[DH1] A. Douady and J. H. Hubbard, *Systèmes dynamiques holomorphes.* I, II: *Itération des polynômes complexes*, Publ. Math. Orsay **84.02** (1984), **85.04** (1985).

[DH2] ____, *A proof of Thurston topological characterization of rational functions*, Institute Mittag-Leffler, Preprint (1984).

[F] O. Forster, *Riemann surfaces*, Springer-Verlag, Berlin-Heidelberg-New York, 1984.

[F1] P. Fatou, *Sur les solutions uniformes de certaines équations fonctionnelles*, C. R. Acad. Sci. Paris **143** (1906), 546–548.

[L] S. Levy, *Critically finite rational maps*, Ph.D. Thesis, Princeton Univ., 1985.

[JM] J. Milnor, *Dynamics in one complex variable: Introductory lectures*, SUNY Stony Brook Institute for Mathematical Sciences Preprints **5** (1990).

[MT] J. Milnor and W. Thurston, *Iterated maps of the interval*, Dynamical Systems (Maryland), J. C. Alexander, editor, Lecture Notes in Math., vol. 1342, Springer-Verlag, Berlin-Heidelberg-New York, 1988.

[TL] Tan Lei, *Accouplements des polynômes complexes*, Ph.D. Thesis, Université de Paris Sud, 1987.

[BW] B. Wittner, *On the bifurcation loci of rational maps of degree* 2, (to appear).

DEPARTMENT OF MATHEMATICS, STATE UNIVERSITY OF NEW YORK AT STONY BROOK, STONY BROOK, NEW YORK 11794
E-mail address: ben@math.sunysb.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA AT SAN DIEGO, SAN DIEGO, CALIFORNIA 92093
E-mail address: fisher@inls1.ucsd.edu

DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NEW YORK 14853
E-mail address: hubbard@mssun7.msi.cornell.edu