FUNDAMENTA
MATHEMATICAE
203 (2009)

Critical portraits for postcritically finite polynomials

by

Alfredo Poirier (Lima)

Abstract. We extend the work of Bielefeld, Fisher and Hubbard on critical portraits to arbitrary postcritically finite polynomials. This gives the classification of such polynomials as dynamical systems in terms of their external ray behavior.

1. Introduction. The subject of this paper is the classification (as dynamical systems) of postcritically finite polynomials. Essentially this is done by refining ideas originally introduced by Bielefeld, Fisher and Hubbard [BFH].

The main ideas presented here represent part of the author's Ph.D. thesis at SUNY Stony Brook. This work has previously circulated as a preprint [P2] in the IMS Stony Brook series. This paper contains a new exposition, which means that some proofs had to be reorganized. The results, however, are exactly the same.

We start with a degree $d \geq 2$ polynomial $f : \mathbb{C} \to \mathbb{C}$. If the critical set $\Omega_f = \{z : \deg_z f \geq 1\} = \{z : f'(z) = 0\}$ has finite orbit $\mathcal{O}(\Omega_f) = \bigcup_{n=0} f^{\circ n}(\Omega_f)$, then f is called *postcritically finite*. In this case, the Julia set and the filled-in Julia set of f are connected and locally connected. Furthermore, the unbounded Fatou component is conformally equivalent to the standard punctured disk. (For this and other elementary facts, we refer the reader to [Mi] or [CG].)

A great deal of the dynamics turns out to be conjugated to the simple model $z \mapsto z^d$. Therefore, it can be understood by the ray behavior in the basin of attraction of ∞. However, in the bounded part of the plane, the endpoints of several rays usually collide, and the topological model for the filled-in Julia set is a pinched disk with lot of identifications (cf. [D]). In a sense, to classify postcritically finite polynomials is to pick in a skillful way selected data from which all possible identifications can be recovered.

2000 *Mathematics Subject Classification*: Primary 37F10; Secondary 37E15.
Key words and phrases: complex dynamical systems, postcritically finite polynomials.

DOI: 10.4064/fm203-2-2

Still, there is an extra bit of information needed to achieve that goal. For a postcritically finite polynomial the closures of all bounded Fatou components are topological disks that eventually map onto loops of periodic components. In turn, a periodic Fatou component is related to a periodic critical cycle, indicating a well defined center.

Next we take a bounded Fatou component U and select a boundary point $p \in \partial U$. If we order counterclockwise the external rays $R_{\theta_1}, \ldots, R_{\theta_k}$ landing at p, the complex plane splits into k pieces. Whenever U fits inside the region determined by R_{θ_1} and R_{θ_2} ($\theta_2 = \theta_1$ in case a single ray lands at p), the argument θ_1 becomes by definition the (*left*) *supporting argument* of the Fatou component U at p, and R_{θ_1} the *supporting ray*. Likewise, we can define right supporting rays, but we will seldom use them. An argument supports at most one component. Even more, by definition, given a Fatou component U, at every point $p \in \partial U$ lands a supporting ray for U. It follows from the basic theory of Hubbard trees (see [DH1] or [P1]) that in this postcritically finite setting points that belong to the boundary of several bounded components are eventually periodic. There are only a finite number of periodic points where this happens and the concept is really needed. The arguments of the rays landing at such points are rational as they are absorbed by periodic orbits.

We prolong an external ray R_θ supporting a Fatou component $U(\omega)$ up to its center ω through an internal ray and call the resulting set the *extended ray E_θ with argument θ*. The *internal rays* are the preimages of the radial segments under the coordinate with ω corresponding to 0.

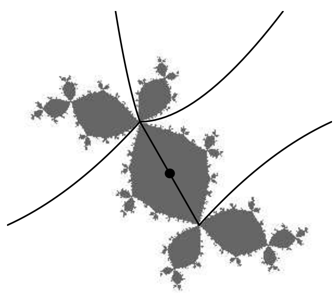

Fig. 1.1

Critically marked polynomials. Given a degree $d \geq 2$ postcritically finite polynomial f, we assign to every critical point a finite subset of $\mathbb{Q}/\mathbb{Z}$ and construct a *critically marked polynomial* (f, Θ), where $\Theta = (\mathcal{F}, \mathcal{J})$, the *marking*, is actually a pair of families of subsets of the circle. Here $\mathcal{F} = \{\mathcal{F}_1, \ldots, \mathcal{F}_l\}$ is a finite family of subsets where each $\mathcal{F}_k$ stands for a finite set of arguments associated with the critical point $\omega_k^{\mathcal{F}}$ in the Fatou set, whereas $\mathcal{J} = \{\mathcal{J}_1, \dots, \mathcal{J}_n\}$ is composed of sets $\mathcal{J}_i$ that collect the arguments of some rays landing at a given Julia set critical point. In each case the size of the set equals the local degree at the critical point. We remark once and for all that the marking may not be unique. Also note that half the marking finishes up empty if there are no critical points in the Fatou or in the Julia set to work with.

Construction of $\mathcal{F}_k$. We first consider the case where the Fatou critical point $\omega = \omega_k^{\mathcal{F}}$ is periodic. Let $\omega \mapsto f(\omega) \mapsto \cdots \mapsto f^{\circ n}(\omega) = \omega$ be a critical cycle of period n and total degree $m > 1$. We construct the associated set $\mathcal{F}_k$ for every critical element ω_k in the cycle simultaneously. Denote by d_ω the local degree of f at ω_k. Pick any periodic point $p_\omega \in \partial U(\omega)$ of period dividing n—which is not critical since it is periodic and belongs to the Julia set—and consider the left supporting ray R_θ for this component at p_ω. This choice, in a natural way, determines periodic supporting rays all along the Fatou components in the cycle. Moreover, they are all of period n. Given the supporting periodic ray R_θ, we track down the d_ω supporting rays of $U(\omega)$ that are inverse images of $f(R_\theta) = R_{d\theta}$. The set of arguments of these rays is by definition $\mathcal{F}_k$. Keeping in mind that a preferred periodic supporting ray has already been chosen, we repeat the procedure for all critical points in the cycle. When the cycle has total degree m, there are $m-1$ different ways to accomplish the selection. If $\mathcal{F}_k$ is the set subordinated to the periodic critical point ω_k, there is one argument in $\mathcal{F}_k$ which is periodic (namely θ as above), the so called *preferred supporting argument associated with* ω_k. By definition, the period of ω_k and of the preferred supporting argument are equal.

Now we consider the case of a non-periodic Fatou critical point ω_k of degree d_ω. There is a minimal $n > 0$ for which $\omega' = f^{\circ n}(\omega)$ is also critical. Recursively we assume that ω' is marked and has associated a preferred supporting ray R_θ (at the beginning only periodic critical elements have). Then $(f^{\circ n})^{-1}(R_\theta)$ contains many rays, but only d_ω among them support this Fatou component $U(\omega)$. Again, the set of arguments of these rays is $\mathcal{F}_k$. We choose one, and call it the preferred supporting argument associated with ω. We continue this process for all Fatou critical points until exhausted.

Construction of $\mathcal{J}_i$. Given a Julia critical point $c = c_k^J$ of degree $d_i > 1$, we proceed according to whether its forward orbit is critical point free or not. If the forward orbit of c contains no further critical point, then for some θ, usually non-unique, the ray R_θ lands at the critical value $f(c)$. Now $f^{-1}(R_\theta)$ consists of d different rays, and among them exactly d_k land at c. As usual, define $\mathcal{J}_i$ to be the set of arguments of those latter rays. For future reference we choose a preferred one.

Otherwise, the Julia critical point c will reach in n steps another critical point, to which we assume we have already associated a preferred ray R_θ. In the n-fold iterated inverse $(f^{\circ n})^{-1}(R_\theta)$ we will discover precisely d_i rays landing at c. We group these arguments in $\mathcal{J}_i$ and select a preferred argument.

Summing up, we see that the construction is in several steps. First we complete the choice the closest to a periodic orbit, and from there we proceed backwards. At some instances we will face decisions that would affect the subsequent marking. When this happens, we will be performing a *hierarchic selection*. At this stage of the exposition we encourage the reader to examine the examples contained in the next section.

The importance of the above construction is hinted at by the next result. (Compare also Theorem 1.3 where uniqueness issues are discussed.)

THEOREM 1.1. *Every centered monic postcritically finite polynomial has a critical marking. This marking determines the polynomial in the sense that two centered monic polynomials whose critical markings agree are the same.*

The combinatorics of critically marked polynomials. In order to analyze the properties satisfied by $\Theta = (\mathcal{F}, \mathcal{J})$, we introduce combinatorial notation.

A set $\Lambda \subset \mathbb{T} = \mathbb{R}/\mathbb{Z}$ is a *degree d preargument set* if $d\Lambda = \{d\lambda : \lambda \in \Lambda\}$ is a singleton. For practical reasons it is convenient to assume always that Λ contains at least two elements. If all values in Λ are rational, we call Λ a *rational preargument set*. As an illustration we single out each individual collection $\mathcal{F}_k, \mathcal{J}_i$ as presented before.

For a family $\Lambda = \{\Lambda_1, \dots, \Lambda_n\}$ of subsets of the circle $\mathbb{T}$, we define the *family union* as $\Lambda^\cup = \bigcup \Lambda_i$, and say that $\lambda \in \Lambda^\cup$ *participates in the family*. We also write $\Lambda^\cup_{\mathbf{per}}$ for the subset of Λ whose members are periodic under multiplication by d.

Hierarchic families. A family $\Lambda = \{\Lambda_1, \dots, \Lambda_n\}$ is called *hierarchic* if whenever $\lambda, \lambda' \in \Lambda^\cup$ and, for some k, we have $d^i\lambda, d^j\lambda' \in \Lambda_k$ with $i, j > 0$, then $d^i\lambda = d^j\lambda'$.

This is like retaining a preferred element λ_k in each Λ_k. Think of them as gates: as soon as $\lambda \in \Lambda_l$ is such that $d^i\lambda \in \Lambda_k$ for $i > 0$, then $d^i\lambda$ is the preferred member of Λ_k chosen before.

Linkage relations. Two subsets $\mathcal{T}, \mathcal{S}$ of the circle are said to be *unlinked* if they belong to disjoint connected subsets of $\mathbb{T}$, or equivalently, if $\mathcal{S}$ is contained in one component of $\mathbb{T} - \mathcal{T}$. In particular, these sets are disjoint. If we identify $\mathbb{T}$ with the boundary of the unit disk, an analogous condition is that the convex hulls of $\mathcal{T}$ and of $\mathcal{S}$ are disjoint. When $\mathcal{T}$ and $\mathcal{S}$ are not unlinked, then either $T \cap \mathcal{S} \neq \emptyset$ or there are $t_1, t_2 \in \mathcal{T}$ and $s_1, s_2 \in \mathcal{S}$ that can be displayed cyclically as $t_1 < s_1 < t_2 < s_2 < t_1$ (if this is so, they are *linked*). More generally, a family $\Lambda = \{\Lambda_1, \dots \Lambda_n\}$ is *unlinked* if its members are pairwise unlinked. Alternatively, each Λ_i is contained in a connected component of $\mathbb{T} - \Lambda_j$ for all $j \neq i$.

Our definitions are highly motivated by the dynamics of external rays for a polynomial map. For example, if we choose a finite number of different points, and for each we pick the arguments of the rays landing there, then we obtain an unlinked family, since otherwise the external rays involved will cross. The same applies to a set of arguments supporting different Fatou components. However, when we compare groups of arguments landing at a point with a set of arguments supporting a component, we can anticipate minor problems (see Example 2.7). Anyhow, it is not difficult to see—because of the consistent way the supporting rays are collected—that rays sharing their landing point are "almost" unlinked to rays supporting a component.

Consider two families $\mathcal{F} = \{\mathcal{F}_1, \dots, \mathcal{F}_n\}$ and $\mathcal{J} = \{\mathcal{J}_1, \dots, \mathcal{J}_m\}$. We say that $\mathcal{J}$ *is weakly unlinked to* $\mathcal{F}$ *on the right* if we can choose arbitrarily small $\varepsilon > 0$ so that the family $\{\mathcal{F}_1, \dots, \mathcal{F}_n, \mathcal{J}_1 - \varepsilon, \dots, \mathcal{J}_m - \varepsilon\}$ is unlinked. (Here $\mathcal{J}_i - \varepsilon = \{\lambda - \varepsilon : \lambda \in \mathcal{J}_i\}$.) In particular, each family by itself is unlinked. Note that we may even work with empty families.

Formal critical portraits. Consider a pair of families $\mathcal{F} = \{\mathcal{F}_1, \dots, \mathcal{F}_n\}$ and $\mathcal{J} = \{\mathcal{J}_1, \dots, \mathcal{J}_m\}$ of rational degree d preargument sets. We say that $\Theta = (\mathcal{F}, \mathcal{J})$ is a *degree* d *formal critical portrait* if the following conditions are fulfilled:

(c1) $d - 1 = \sum(|\mathcal{F}_k| - 1) + \sum(|\mathcal{J}_k| - 1)$,
(c2) $\mathcal{J}$ is weakly unlinked to $\mathcal{F}$ on the right,
(c3) each family is hierarchic,
(c4) for any γ that participates in $\mathcal{F}$, some periodic forward iterate $d^i\gamma$ also participates in $\mathcal{F}$,
(c5) no θ that participates in $\mathcal{J}$ is periodic.

These conditions represent the minimal abstract requirements imposed on a marking. Condition (c1) says that we are choosing the correct number of critical elements. Condition (c2) tells us that the rays and extended rays, once the last ones support a component, determine sectors that do not conflict with each other. Condition (c3) asks for dynamically preferred rays. Condition (c4) favors the fact that arguments in $\mathcal{F}$ are associated with Fatou critical points. Condition (c5) means that Julia set critical points are non-periodic.

Unfortunately, some formal critical portraits are unrelated to polynomials (see Example 5.15). To state necessary and sufficient conditions for a marking to arise from a postcritically finite polynomial, we study the partitions of the unit circle they determine.

Given $\Theta = (\mathcal{F}, \mathcal{J})$ as above, we form a partition $\mathcal{P} = \{A_1, \ldots, A_d\}$ of the circle minus the finite set of family members $\Theta^{\cup}$ by means of an equivalence relation. We say that two angles t, s not in $\Theta^{\cup} = \mathcal{F}^{\cup} \cup \mathcal{J}^{\cup}$ are *unlink equivalent* if they belong to the same connected component of $\mathbb{T} - \mathcal{F}_j$ and $\mathbb{T} - \mathcal{J}_i$ for all possible choices of i, j. It is trivial to check that each resulting equivalence class is a finite union of open intervals with total length $1/d$. (A detailed analysis is carried out in Section 4.)

Each element $A_i \in \mathcal{P}$ is a finite union $A_i = \bigcup (x_j, y_j)$ of open connected intervals. We also define the sets $A_i^+ = \bigcup [x_j, y_j)$ and $A_i^- = \bigcup (x_j, y_j]$. This time we get two full partitions $\mathcal{P}^+ = \{A_1^+, \ldots, A_d^+\}$ and $\mathcal{P}^- = \{A_1^-, \ldots, A_d^-\}$. As every angle t belongs to only one A_k^+, we recognize in $a_\Theta^+(t) = A_k^+$ its *right address*. In a similar fashion there is a left address $a_\Theta^-(t)$. We associate to every angle a *right symbol sequence* $s_\Theta^+(t) = (a_\Theta^+(t), a_\Theta^+(dt), a_\Theta^+(d^2t), \ldots)$ and a left one $s_\Theta^-(t) = (a_\Theta^-(t), a_\Theta^-(dt), a_\Theta^-(d^2t), \ldots)$. For all but a countable number of arguments (the angles present in the families and their iterated inverses) left and right symbol sequences agree.

Admissible critical portraits. A formal critical portrait $\Theta = (\mathcal{F}, \mathcal{J})$ is *admissible* if two extra conditions are satisfied.

(c6) Suppose γ and t are periodic with the same period and the same symbol sequences (i.e. $s_\Theta^+(\gamma) = s_\Theta^+(t)$). If γ participates in $\mathcal{F}$, then $\gamma = t$.

(c7) Let $\theta \in \mathcal{J}_l$ and $\theta' \in \mathcal{J}_k$. Take $i \geq 0$. If $s_\Theta^-(d^i\theta) = s_\Theta^-(\theta')$ then $d^i\theta = \theta'$.

Proposition 1.2. *If (f, Θ) is a postcritically finite marked polynomial, then Θ is an admissible critical portrait.*

For the time being, a brief comment regarding the significance of these two new conditions is in order. As is usual in this kind of situation, angles sharing a symbol sequence are supposed to land at the same point (cf. Corollary 5.9; see also Section 7 devoted to a complete study of when two or more rays land together). Under this proviso, condition (c6) explains why the arguments in $\mathcal{F}^{\cup}_{\mathbf{per}}$, and just them, are the periodic ones expected to support Fatou components. On the other hand, condition (c7) reinforces the idea that different elements of $\mathcal{J}$ are associated with different critical points, chosen, of course, according to a hierarchic scheme. In any case, conditions (c1)–(c7) represent a finite amount of information to be tested.

Now we are ready to state a structural result for postcritically finite marked polynomials.

Theorem 1.3. *Let $\Theta = (\mathcal{F}, \mathcal{J})$ be a degree d admissible critical portrait. Then there is a unique monic centered postcritically finite polynomial f with marking (f, Θ).*

2. Examples. We illustrate with examples the concepts involved so far. This will help us anticipate possible complications.

EXAMPLE 2.1 (The rabbit; see Figure 1.1). In the dynamical system determined by iteration of the polynomial $f(z) = z^2 + c$ with c approximately $-0.12 + 0.74i$, popularly known as the *rabbit*, the critical point $z = 0$ has period 3 under iteration. This tells us that $f^{\circ 3}$ restricted to the critical component is a degree 2 cover of itself. So, for $f^{\circ 3}$ there is a fixed point on the boundary of the Fatou critical component. It is well known that the rays $R_{1/7}, R_{2/7}, R_{4/7}$ land at this point, which happens to be fixed also for f. Among them only $R_{4/7}$ supports the main component. We seek for the other ray that supports this component and maps to $R_{1/7} = f(R_{4/7})$, which can only be $R_{1/14}$. The marking $\mathcal{F} = \{\mathcal{F}_1\}$ and $\mathcal{J} = \emptyset$, with $\mathcal{F}_1 = \{4/7, 1/14\}$, is thus unique.

EXAMPLE 2.2 (The Ulam–von Neumann map). This is a strictly preperiodic case. The critical orbit for $f(z) = z^2 - 2$ is given by $0 \mapsto -2 \mapsto 2 \mapsto 2 \mapsto \cdots$. Only the external ray $R_{1/2}$ lands at -2, so the marking for the critical point $z = 0$ should include both $R_{1/4}$ and $R_{3/4}$. The allowed marking is $\mathcal{F} = \emptyset$ and $\mathcal{J} = \{\mathcal{J}_1\}$, where $\mathcal{J}_1 = \{1/4, 3/4\}$.

EXAMPLE 2.3 (Preperiodic case: two possible choices; see Figure 2.1). Consider the degree 2 polynomial $f(z) = z^2 + c$, where c near -1.54 is the only negative root of the equation $c^3 + 2c^2 + 2c + 2 = 0$. In this case the critical orbit is confined to $0 \mapsto c \mapsto c^2+c \mapsto -(c^2+c) \mapsto -(c^2+c)$. The rays

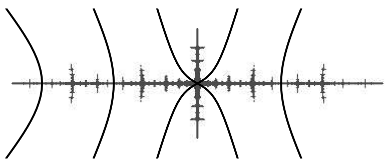

Fig. 2.1

$R_{1/3}, R_{2/3}$ both land at the fixed point $-(c^2 + c)$ and are swapped by f. We will proceed to track the rays backwards. At $c^2 + c$, the rays $R_{1/6}, R_{5/6}$ land, and at the critical value c, both $R_{5/12}$ and $R_{7/12}$ land. Therefore, for each hierarchic choice we must get a different marking: set $\mathcal{J}_1 = \{5/24, 17/24\}$ if you like $R_{5/12}$, or $\mathcal{J}_1 = \{7/24, 19/24\}$ if you prefer $R_{7/12}$. The marking will be $\mathcal{F} = \emptyset$ and $\mathcal{J} = \{\mathcal{J}_1\}$. In either case, we can infer from the marking that the critical point takes three iterations to become periodic. The exact period, however, cannot be read immediately from the data (cf. Corollaries 5.11 and 7.8).

EXAMPLE 2.4 (A non-trivial cycle; see Figure 2.2). For $f(z) = z^3 - 3/2z$ the critical points satisfy $z^2 = 1/2$ and are interchanged by f. In each critical Fatou component the map $f^{\circ 2}$ is a degree 4 (the product of local degrees of the cycle members) cover of itself. Thus, at the boundary of each component there are up to three choices of points of period 2. One of these, namely $z = 0$, belongs to both of them. The rays landing here are $R_{1/4}$ and $R_{3/4}$, each supporting one and only one Fatou component. The period 2 rays that support A, the "rightmost" component, are $R_{3/4}, R_{7/8}, R_{1/8}$, while their images $R_{1/4}, R_{5/8}, R_{3/8}$ support B, the other. The choice should be made simultaneously.

This polynomial has three markings, each of type $\mathcal{F} = \{\mathcal{F}_A, \mathcal{F}_B\}, \mathcal{J} = \emptyset$. The periodic supporting rays are listed on the left.

Component A	$\mathcal{F}_A$
$R_{3/4}$	$\{3/4, 1/12\}$
$R_{7/8}$	$\{7/8, 5/24\}$
$R_{1/8}$	$\{1/8, 19/24\}$

Component B	$\mathcal{F}_B$
$R_{1/4}$	$\{1/4, 7/12\}$
$R_{5/8}$	$\{5/8, 7/24\}$
$R_{3/8}$	$\{3/8, 17/24\}$

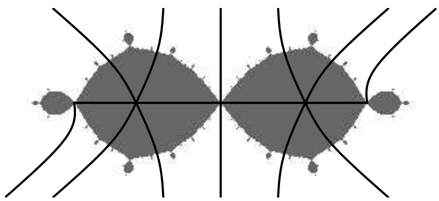

Fig. 2.2

It is important to ask why not take $\mathcal{F}_A = \{3/4, 1/12\}, \mathcal{F}_B = \{3/8, 17/24\}$ as marking. This is forbidden by the rules since 3/4 and 3/8 do not belong to the same cycle. A strong motive for this interdiction is given next.

Example 2.5 (Bad choice, wrong polynomial; cf. Figure 2.3). There is a postcritically finite polynomial related to $\mathcal{F} = \{\mathcal{F}_A, \mathcal{F}_B\}$, $\mathcal{J} = \emptyset$ where $\mathcal{F}_A = \{3/4, 1/12\}$ and $\mathcal{F}_B = \{3/8, 17/24\}$, but it is not the one from Example 2.4.

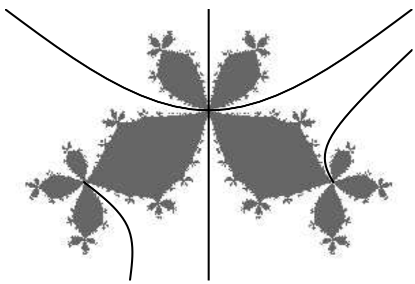

Fig. 2.3

For the polynomial $f(z) = z^3 + az + b$ (where $a = -0.75$ and $b \approx 0.66i$), the rays $R_{1/8}, R_{1/4}, R_{3/8}, R_{3/4}$ land at a fixed point, which pins together the boundaries of the four periodic Fatou components. These components are associated pairwise in cycles, so we are looking at a couple of disjoint period 2 cycles. Only $R_{3/4}$ and $R_{3/8}$ support critical components. In conclusion, this polynomial admits a single marking.

EXAMPLE 2.6 (Hierarchic choice). Consider the map $f(z) = \sqrt{2}(z^2-1)^2$ with critical points $z = 0, \pm 1$ (even if it is not a monic polynomial, the theory adapts effortlessly after applying a positive real dilatation). The critical points quickly fall into the stable sequence $\pm 1 \mapsto 0 \mapsto \sqrt{2} \mapsto \sqrt{2}$. At the fixed point $z = \sqrt{2}$ only R_0 lands; at $z = 0$, the rays $R_{1/4}, R_{3/4}$; at $z = 1$, $R_{1/16}, R_{3/16}, R_{13/16}, R_{15/16}$; and at $z = -1$, $R_{5/16}, R_{7/16}, R_{9/16}, R_{11/16}$. The marking is not unique since it depends upon the choice at $z = 0$. However, it will always fit the pattern $\mathcal{F} = \emptyset$ plus $\mathcal{J} = \{\mathcal{J}_0, \mathcal{J}_{-1}, \mathcal{J}_1\}$.

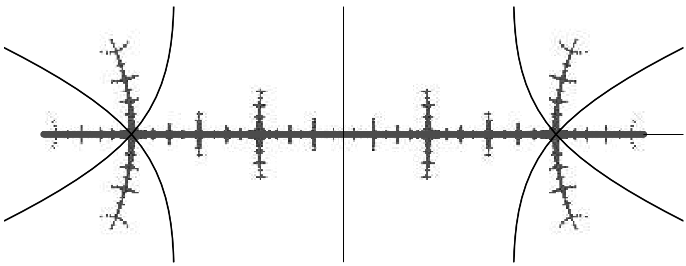

Fig. 2.4

$\mathcal{J}_0$	Preferred ray at $z = 0$	$\mathcal{J}_1$	$\mathcal{J}_{-1}$
$\{1/4, 3/4\}$	$R_{1/4}$	$\{1/16, 13/16\}$	$\{5/16, 9/16\}$
$\{1/4, 3/4\}$	$R_{3/4}$	$\{3/16, 15/16\}$	$\{7/16, 11/16\}$

In the following example we take a closer look at condition (c2).

EXAMPLE 2.7 (Badly mixed case, see Figure 2.5). Consider $f(z) = c(z^5 + 3z^4 + 3z^3 + z^2)$, where c is approximately 4.36. It has two Fatou critical components: one (on the right) fixed of degree 2, and the other (the big one on the left) preperiodic of degree 3, absorbed by the first in one iteration. The boundaries of these components share a point, which happens to be critical. The image of this Julia set critical point is the only fixed point in the boundary of the fixed Fatou component. Just the ray R_0 lands there. The rays $R_{1/5}, R_{4/5}$ land at the Julia critical point, each supporting one of the two critical Fatou components. Now, $R_{4/5}$ is the one that supports the fixed component, while $R_{1/5}$ relates to the other. Also, R_0 must have two inverses supporting the fixed component (they are $R_0, R_{4/5}$), and three sup-

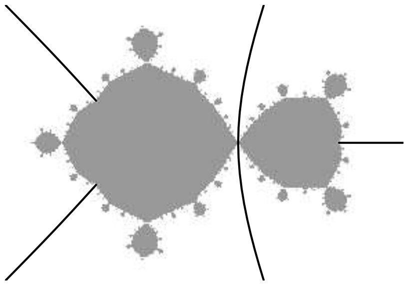

Fig. 2.5

porting the preperiodic one (this time $R_{1/5}, R_{2/5}, R_{3/5}$). Thus, there is only one possible marking $\mathcal{F} = \{\{0, 4/5\}, \{1/5, 2/5, 3/5\}\}, \mathcal{J} = \{\{1/5, 2/5\}\}$ that puts together all these facts. This time there are arguments which belong to one family and also to the other. Of course, when something like this occurs, they are strictly preperiodic.

In order to deal with a formal critical portrait, for $\varepsilon > 0$ small, the three sets $\{0, 4/5\}, \{1/5, 2/5, 3/5\}, \{1/5 - \varepsilon, 4/5 - \varepsilon\}$ should be unlinked; an obvious fact.

Example 2.8 (Several critical cycles; see Figure 2.2). Consider the degree 9 polynomial $f \circ f$ where f is the polynomial given in Example 2.4. The filled-in Julia set as well as the external rays remain unchanged. This time, however, we have two fixed Fatou components of degree 4 (instead of two of period 2). Each one pulls in one iteration another critical component. The cycles are independent, and the choice of a marking must reflect this fact. Nevertheless, they determine the choice of marking in the components they attract. Let us denote by A, B the fixed components and by A', B' the ones they absorb. The marking should be $\mathcal{F} = \{\mathcal{F}_A, \mathcal{F}_{A'}, \mathcal{F}_B, \mathcal{F}_{B'}\}$ together with $\mathcal{J} = \emptyset$.

Component A	$\mathcal{F}_A$	$\mathcal{F}_{A'}$
$R_{3/4}$	$\{3/4, 62/72, 6/72, 14/72\}$	$\{30/72, 38/72\}$
$R_{7/8}$	$\{7/8, 7/72, 15/72, 55/72\}$	$\{31/72, 39/72\}$
$R_{1/8}$	$\{1/8, 17/72, 57/72, 65/72\}$	$\{41/72, 33/72\}$

Component B	$\mathcal{F}_B$	$\mathcal{F}_{B'}$
$R_{1/4}$	$\{1/4, 26/72, 42/72, 50/72\}$	$\{66/72, 2/72\}$
$R_{5/8}$	$\{5/8, 53/72, 21/72, 29/72\}$	$\{3/72, 67/72\}$
$R_{3/8}$	$\{3/8, 43/72, 51/72, 19/72\}$	$\{5/72, 69/72\}$

In all, we have nine different markings at our disposal.

Example 2.9 (Why we should have separated families). Finally, we clarify why we should not consolidate the families together into one.

Let $\mathcal{A} = \{0, 1/3\}$ and $\mathcal{B} = \{5/9, 8/9\}$. The polynomial $f(z) = z^3 + Az + B$, where $A = 2.25, B \approx -0.43i$, has marking $\mathcal{F} = \{\mathcal{A}, \mathcal{B}\}, \mathcal{J} = \emptyset$, and at the same time $f(z) = z^3 + A'z + B'$, where $A' \approx 2.18, B' \approx -0.39i$, admits $\mathcal{F} = \{\mathcal{A}\}, \mathcal{J} = \{\mathcal{B}\}$ as portrait.

One critical point is fixed and the other one is not. The difference is that the second point eventually maps to the center of the fixed component in the first case and to the supporting fixed point in the boundary in the other.

Fig. 2.6. Just about the same marking

3. A conceptual overview. Conceptually our work is divided in three blocks. In the first we deal with necessary conditions, that is, we prove that the marking of a postcritically finite polynomial satisfies conditions (c1)–(c7) above. In the second, from the combinatorial data we build a map that will turn out to be equivalent to a polynomial in a sense to be specified shortly. In the last we verify that the postcritically finite polynomial obtained admits the original marking.

Near ∞ the dynamics of any polynomial is completely understood, being in essence that of $z \mapsto z^d$. All works because the dynamics of rays behaves as multiplication by d of the arguments. We dedicate Section 4 to the study of this simplified dynamics using portraits. In Section 5 we relate those ideas to the dynamical plane. The outcome is that we deduce that the marking of our polynomial is admissible, as expected.

We also study the statical layout of the problem. This is done by sketching external rays as they approach the unit disk. However, there will always be rays landing at the same Julia set point or external rays pushing their way to the center of a Fatou component, and our topological model should reflect directly or indirectly, but nonetheless skilfully, these facts. Thus, Section 6 focuses on how to delineate inside the unit disk a compatible set of identifications that in the end will efficiently balance the dynamical and statical interplay.

Instead of seeking for all identifications determined by the Julia set in order to get a so called lamination (cf. [K]), we benefit from the fact that we are working with a combinatorially finite amount of information and proceed differently. First we decide which arguments are meaningful. Then we group them in equivalence classes to reflect our wish that the corresponding rays land together (see Section 7). We will also learn that these sets are unlinked, so we can draw the terminal points of the rays without any leg crossing for different classes. Here we followed [BFH].

But there are also Fatou components to study. We group likewise a subset of arguments in supporting sets (this is worked out in Section 8). As these classes are essentially unlinked to the previous ones, we will be able to allocate inside the regions points capable to represent Fatou centers. We join this reference point to the Julia vertices associated to the corresponding angles. (All of this is done before, in Section 6, in an abstract setting.)

Up to this point we have a web but no dynamics whatsoever. If the Julia and Fatou angles fit into an invariant set, then this dynamics can be extended naturally to the graph. And then, as our previous construction slices $\mathbb{C}$ into relatively small pieces, we can extend the web dynamics, patch by patch, to the whole plane (see Section 9). This is the branched covering we look for. In there we will retain all the vertices as special points and apply Thurston's theorem on topological characterization of rational maps. Bear in mind that Thurston's theorem only concerns postcritical points, so our rays, internal or external, honest or extended, only play a marginal role in the discussion. Nevertheless, we will not delete them but instead think of them as a bonus item that would help us understand the global picture.

In Section 10 we discuss why we can apply Theorem 3.2 to this topological polynomial in order to get a unique postcritically finite polynomial up to conjugation. However, we still do not know why this polynomial admits the expected marking. The answer is given in Section 12 and is a result of the precautions taken all over the place together with supplementary details squeezed from Thurston's theorem in Section 11.

To conclude, we recall several facts concerning Thurston's topological characterization of rational maps, with emphasis on polynomials.

Let $f : S^2 \to S^2$ be an orientation preserving branched covering map of the sphere. The set $\Omega_f = \{z : \deg_z f > 1\}$ is named the *critical set* of f. The *postcritical set* of f is the set $P(\Omega_f) = \bigcup_{n=1}^{\infty} f^{\circ n}(\Omega_f)$ of forward iterates of the critical set. Whenever $P(\Omega_f)$ is finite, we say that f is *postcritically finite*. Every time there is a reference point ∞ for which $f^{-1}(\infty) = \{\infty\}$, we talk about a *topological polynomial*. In what follows, we always take f to be postcritically finite.

Two marked branched maps (f, Ω_f) and (g, Ω_g) are *Thurston equivalent* if there are isotopic homeomorphisms $\phi_0, \phi_1 : (S^2, P(\Omega_f)) \to (S^2, P(\Omega_g))$ that preserve the postcritical sets (that is, the isotopy itself maps $P(\Omega_f)$ to $P(\Omega_g)$), so that the diagram

$$\begin{tikzcd} (S^2, P(\Omega_f)) \arrow[r, "f"] \arrow[d, "\phi_1"'] & (S^2, P(\Omega_f)) \arrow[d, "\phi_0"] \\ (S^2, P(\Omega_g)) \arrow[r, "g"] & (S^2, P(\Omega_g)) \end{tikzcd}$$

commutes.

For the branched map (f, Ω_f), a simple closed curve $\gamma \subset S^2 - P(\Omega_f)$ is *non-peripheral* if each component of $S^2 - \gamma$ contains at least two points in $P(\Omega_f)$. A *multicurve* $\Gamma = \{\gamma_1, \ldots, \gamma_k\}$ is a set of simple, closed, disjoint, non-homotopic, non-peripheral curves in $S^2 - P(\Omega_f)$. A multicurve Γ is *stable* when, for all $\gamma \in \Gamma$, every non-peripheral component of $f^{-1}(\gamma)$ is homotopic—relative to $P(\Omega_f)$—to a curve in Γ.

Whenever inside a stable multicurve Γ we can find a cyclically arranged collection $\gamma_0, \ldots, \gamma_{k-1}, \gamma_k = \gamma_0 \in \Gamma$ such that γ_i is homotopic relative to $P(\Omega_f)$ to exactly one component γ_j of $f^{-1}(\gamma_{i+1})$, and additionally $f : \gamma_j \to \gamma_{i+1}$ has degree 1, we will have a *Levy cycle*. Its importance is made clear by the next two theorems and the interpretation below.

THEOREM 3.1. *Suppose a topological polynomial f admits $\gamma_0, \ldots, \gamma_{k-1}$ as a Levy cycle. Let M_i be the intersection of $P(\Omega_f)$ with the bounded component of $S^2 - \gamma_i$. Then M_i consists of periodic non-critical points. Furthermore, we necessarily have $M_{i+1} = f(M_i)$.*

THEOREM 3.2. *A marked topological polynomial is Thurston equivalent to an actual postcritically finite rational map if and only if it admits no Levy cycle. If this is the case, all these rational maps are affine conjugated to the same postcritically finite polynomial.*

The proof of both results, based on Thurston's theorem as stated in [DH2], can be found in [BFH].

What is important is that *for postcritically finite topological polynomials, only misidentification of non-critical cycles can lead to an obstruction.*

4. Basic combinatorics. In this section we describe the basic properties of formal critical portraits. We manipulate unlinked families and take a closer look at the dynamical properties that can be derived. A central role is played by multiplication by d, the dynamics that mimics in $\mathbb{T}$ what iteration of a polynomial does with rays near ∞.

In the next section we will transfer everything to the dynamical plane, concretely to a neighborhood of the Julia set. Among other things, we study the physical layout of rays and extended rays, and as a reward we establish the admissibility of the critical portrait of a postcritically finite polynomial.

In Section 6 we return to weakly unlinked families to reconstruct some static features.

In the double family $\Theta = (\mathcal{F} = \{\mathcal{F}_j\}_{j=1}^n, \mathcal{J} = \{\mathcal{J}_i\}_{i=1}^m)$, suppose $\mathcal{J}$ is weakly unlinked to $\mathcal{F}$ on the right; so for the moment we only bother with condition (c2). Set $\mathcal{F}^\cup = \bigcup_{j=1}^n \mathcal{F}_j$, $\mathcal{J}^\cup = \bigcup_{i=1}^m \mathcal{J}_i$ and $\Theta^\cup = \mathcal{F}^\cup \cup \mathcal{J}^\cup$. The *weight* of Θ is by definition $w(\Theta) = 1 + \sum_{j=1}^n (|\mathcal{F}_j| - 1) + \sum_{i=1}^m (|\mathcal{J}_i| - 1)$.

Call $\theta_1, \theta_2 \notin \Theta^\cup$ *unlinked equivalent* (with Θ as a prefix when needed) if they belong to the same connected component of $\mathbb{T} - \mathcal{J}_i$ and of $\mathbb{T} - \mathcal{F}_j$ for all possible choices of $\mathcal{J}_i$ and $\mathcal{F}_j$. Whenever $\mathcal{J}_i$ or $\mathcal{F}_j$ are singletons their complements are always connected, and they do not contribute to the global weight.

We first concentrate on a single family $\Lambda = \{\Lambda_l\}_{l=1}^n$ known to be unlinked. In this case, the weight of Λ is just $w(\Lambda) = 1 + \sum_{j=1}^n (|\Lambda_j| - 1)$. The following result is key to describe the number and shape of Λ-unlinked classes.

LEMMA 4.1. *If $\Lambda = \{\Lambda_k\}_{k=1}^n$ is an unlinked family, then the number of unlinked classes equals the weight of Λ. Moreover, we may write each class A_k as the union of open intervals as $A_k = (\theta_0, \theta_1) \cup \cdots \cup (\theta_{2p_k-2}, \theta_{2p_k-1})$ with subscripts modulo $2p_k$ and respecting the cyclic order. Additionally, every θ_l participates in Λ and, in the notation above, successive $\theta_{2j-1}, \theta_{2j}$ belong to the same Λ_i.*

Proof. The proof will be by induction on the number of basic sets in Λ. If Λ consists of a single member Λ_1, the unlinked classes are the connected components of $\mathbb{T} - \Lambda_1$, exactly $|\Lambda_1| = 1 + (|\Lambda_1| - 1)$ of them. Everything else is clear.

Suppose the result holds for the family $\widetilde{\Lambda} = \{\Lambda_1, \ldots, \Lambda_{n-1}\}$. In particular, let $\widetilde{A}_1, \ldots, \widetilde{A}_d$, where $d = 1 + \sum_{k=1}^{n-1} (|\Lambda_k| - 1)$, be the $\widetilde{\Lambda}$-equivalence classes. Now append Λ_n to $\widetilde{\Lambda}$ so that the enlarged family remains unlinked. This can only happen if Λ_n is contained in a connected component of $\mathbb{T} - \Lambda_i$ for $i = 1, \ldots, n-1$. But this property is exactly the one we require to recognize all elements in Λ_n unlinked equivalent relative to $\widetilde{\Lambda}$. That is, the whole of Λ_n can be placed inside one $\widetilde{A}_i$. The new Λ_n subdivides $\widetilde{A}_i$ into exactly $|\Lambda_n|$ subsets in a standard way: first sort $\Lambda_n = \{\lambda_1, \ldots, \lambda_l\}$ in cyclic order, and then replace the class $\widetilde{A}_i$ by the l non-empty disjoint open sets $\widetilde{A}_i \cap (\lambda_j, \lambda_{j+1})$. The rest is evident. ∎

The next lemma is of interest in itself and justifies the names.

LEMMA 4.2. *If $\Lambda = \{\Lambda_k\}_{k=1}^n$ is an unlinked family, then the Λ-unlinked equivalence classes are unlinked.*

Proof. Suppose the pairs $\theta_i, \theta_i' \in A_i$ and $\theta_j, \theta_j' \in A_j$ can be arranged as $\theta_i < \theta_j < \theta_i' < \theta_j' < \theta_i$ in cyclic order. As they are linked, our goal is reduced to confirm the equality $A_i = A_j$. By definition of unlinked equivalent classes, any Λ_k which does not fit completely in (θ_i, θ_i') appears in (θ_i', θ_i), the opposite interval. But this property of Λ_k also holds with (θ_j, θ_j') and (θ_j, θ_j'). Hence, any fixed Λ_k is contained in one of the four intervals $(\theta_i, \theta_j), (\theta_j, \theta_i'), (\theta_i', \theta_j'), (\theta_j', \theta_i)$. This explains why all four arguments are unlinked equivalent and we should have $A_i = A_j$. ∎

We are interested in formal critical portraits, so we must allow one family to be weakly unlinked to the other. This is handled right away.

LEMMA 4.3. *In $\Theta = (\mathcal{F} = \{\mathcal{F}_j\}_{j=1}^n, \mathcal{J} = \{\mathcal{J}_i\}_{i=1}^m)$ let $\mathcal{J}$ be weakly unlinked to $\mathcal{F}$ on the right. Then there is an unlinked family $\Lambda = \{\Lambda_1, \ldots, \Lambda_k\}$ of the same weight as Θ for which the Λ-classes are equal to the Θ-classes. Moreover, each Λ_l is a union of some basic constituents of Θ.*

Proof. We supply only a sketch of the proof, details are left to the reader. First send all elements in $\mathcal{J}$ that do not intersect any $\mathcal{F}_j$ to $\mathcal{F}$. If there is nothing left in $\mathcal{J}$, we are done. When $\mathcal{J}_1$ intersects $\mathcal{F}_1, \ldots, \mathcal{F}_k$, disregard $\mathcal{J}_1$ as well as $\mathcal{F}_1, \ldots, \mathcal{F}_k$ and replace them by $\Lambda_1 = \mathcal{J}_1 \cup \mathcal{F}_1 \cup \cdots \cup \mathcal{F}_k$. It is clear that the unlinked classes do not change whatsoever. Notice that the weak linkage relation implies that all $\mathcal{J}_1 \cap \mathcal{F}_i$ are different singletons. In this way, a straightforward count gives $|\Lambda_1| = |\mathcal{J}_1| + |\mathcal{F}_1| + \cdots + |\mathcal{F}_k| - k$. The weights therefore match. The general claim follows by induction. ∎

We next require that each member be a preargument set and describe the resulting classes. First we rewrite this as in Lemma 4.1.

LEMMA 4.4. *Let $\Lambda = \{\Lambda_k\}_{k=1}^n$ be an unlinked family of d-prearguments. If we write each class A_k as a union $A_k = (\theta_0, \theta_1) \cup \cdots \cup (\theta_{2p_k-2}, \theta_{2p_k-1})$ of intervals with subscripts modulo $2p_k$ and respecting the cyclic order, then necessarily $d\theta_{2j-1} = d\theta_{2j}$. As a consequence, the total Lebesgue measure of each class is a positive multiple of $1/d$.*

Proof. Imitating Lemma 4.1 everything is identical up to where we can conclude that θ_{2j-1} and θ_{2j} belong to the same Λ_k. However, since Λ_k is a preargument set, we also get $d\theta_{2j-1} = d\theta_{2j}$.

As $d\theta_{2j-1} = d\theta_{2j}$ implies $\theta_{2j} = \theta_{2j-1} + l/d$ for some l, the span of the closed interval $[\theta_{2j-1}, \theta_{2j}]$ is a multiple of $1/d$. Thus, the class A_k, the complement of a union of such intervals, has the correct measure. ∎

Now we are prepared to describe the classes of interest.

LEMMA 4.5. *If Θ is a formal critical portrait, then there are d unlinked equivalence classes. Moreover, we may present each unlinked class A as the union of finitely many intervals as $A = (\theta_0, \theta_1) \cup \cdots \cup (\theta_{2p-2}, \theta_{2p-1})$ with subscripts modulo $2p$ and respecting the cyclic order. Additionally, all θ_k participate in Θ and, in the notation above, they satisfy $d\theta_{2i-1} = d\theta_{2i}$. The total length of each equivalence class is exactly $1/d$.*

Proof. Reviewing the merging that took place in Lemma 4.3 we assume without loss of generality that $\mathcal{F} \cup \mathcal{J}$ is an unlinked family of d-preargument sets, of weight d because of condition (c1). According to Lemma 4.4 each of the resulting d non-empty open unlinked sets measures at least $1/d$, and so the exact size is $1/d$ since there are d of them. ∎

Each $A_i \in \mathcal{P}$ is a finite union $A_i = \bigcup_j (\theta_{2j}, \theta_{2j+1})$ of open connected intervals. We also define $A_i^+ = \bigcup_j [\theta_{2j}, \theta_{2j+1})$ and $A_i^- = \bigcup_j (\theta_{2j}, \theta_{2j+1}]$. This time we get two full partitions $\mathcal{P}^+ = \{A_1^+, \ldots, A_d^+\}$ and $\mathcal{P}^- = \{A_1^-, \ldots, A_d^-\}$ of $\mathbb{T}$.

For convenience and to avoid confusion, multiplication by d will also be referred to as m_d. This will prove useful when thinking of it as a function, specially while taking inverse images.

COROLLARY 4.6. *Each $A \in \mathcal{P}$ is mapped bijectively onto the complement of a finite set by m_d. In turn, each $A^\pm$ is mapped bijectively by m_d onto the whole unit circle. Furthermore, the restriction of m_d to each of those sets preserves the cyclic order.*

Proof. This is an easy consequence of $d\theta_{2j-1} = d\theta_{2j}$. ∎

REMARK 4.7. In the description $A_i = \bigcup (\theta_{2j}, \theta_{2j+1})$ an equality $d\theta_{2j} = d\theta_{2j'}$ can only be a consequence of $\theta_{2j} = \theta_{2j'}$. This is true because m_d is injective in $A_i^\pm$.

As every angle θ belongs to only one A_k^+, its *right address* $a_\Theta^+(\theta) = A_k^+$ is well defined. In a similar fashion we assign a left address $a_\Theta^-(\theta)$. When $a_\Theta^+(\theta) = a_\Theta^-(\theta)$, we write $a_\Theta(\theta)$. We associate to θ a *right symbol sequence* $s_\Theta^+(\theta) = (a_\Theta^+(\theta), a_\Theta^+(d\theta), \ldots)$ and a *left* one $s_\Theta^-(\theta) = (a_\Theta^-(\theta), a_\Theta^-(d\theta), \ldots)$. For all but a countable number of arguments (namely the angles present in the families and their iterated inverses), left and right symbol sequences agree.

REMARK 4.8. If we take θ, θ' in $\mathcal{J}_k$ and λ is an argument whose left address $a_\Theta^-(\lambda)$ agrees with that of θ, then $\lambda \in (\theta', \theta]$ by definition. Similarly, if θ, θ' are in $\mathcal{F}_k$ and λ satisfies $a_\Theta^+(\lambda) = a_\Theta^+(\theta)$, we can infer $\lambda \in [\theta, \theta')$. (There is nothing special about $\mathcal{J}$ or $\mathcal{F}$ in this formulation, but this is how these properties will be used later.)

LEMMA 4.9. *Let θ, θ' be such that $a_\Theta^+(d^i\theta) = a_\Theta^+(d^i\theta')$ for $i = 0, \ldots, n-1$. If $d^n\theta = d^n\theta'$, then $\theta = \theta'$. The same is true when we consider left symbol sequences instead.*

Proof. Everything is set to pass from $d^n\theta = d^n\theta'$ to $d^{n-1}\theta = d^{n-1}\theta'$ and so on. ∎

Warning. This last result is not necessarily true when we compare left with right symbol sequences. From $a^+_\Theta(\theta) = a^-_\Theta(\theta')$ and $d\theta = d\theta'$ we cannot deduce $\theta = \theta'$. For example, in the Ulam–von Neumann map (Example 2.2) we get $a^+_\Theta(1/4) = a^-_\Theta(3/4)$ and at the same time both of these different arguments turn equal after doubling.

Refinements. We are ready to introduce dynamically defined refinements. When $A_0, A_1, \ldots \in \mathcal{P}^-$, write $U_{A_0,\ldots,A_n} = \{\theta \in \mathbb{T} : d^i\theta \in A_i,\ i = 0, \ldots, n\}$. Each set measures $1/d^{n+1}$ as can be verified from $dU_{A_0,\ldots,A_n} = U_{A_1,\ldots,A_n}$ by induction. Also define $U_{A_0,A_1,\ldots} = \bigcap_{n=0}^{\infty} \mathrm{cl}(U_{A_0,\ldots,A_n})$. This last set, being a nested intersection of non-empty compact sets, is non-empty. It is easy to see that $s^\pm_\Theta(\theta) = (A_0, A_1, \ldots)$ implies $\theta \in U_{A_0,A_1,\ldots}$. We conclude that for any given $A_0, A_1, \ldots \in \mathcal{P}$ there is an argument which has either left or right symbol sequence $(A_0, A_1, \ldots)$. The same can be done when taking $\mathcal{P}^+$ or $\mathcal{P}$ in place of $\mathcal{P}^-$.

Lemma 4.10. *For each $n \geq 0$, the family $\{U_{A_0,A_1,\ldots,A_n}\}$ is unlinked.*

Proof. We work by induction on n. For $n = 0$, we are only rephrasing Lemma 4.2. Now suppose that $\theta, \lambda \in U_{A_0,\ldots,A_n}$ and $\theta', \lambda' \in U_{A'_0,\ldots,A'_n}$ can be arranged as $\theta, \theta', \lambda, \lambda', \theta$ in cyclic order. This means first $\theta, \lambda \in U_{A_0} = A_0$ and $\theta', \lambda' \in U_{A'_0} = A'_0$, and then, by Lemma 4.2, that A_0 and A'_0 are the same unlink class. Now, Corollary 4.6 shows that the order $d\theta, d\theta', d\lambda, d\lambda', d\theta$ is suitable for those values. Since $d\theta, d\lambda \in U_{A_1,\ldots,A_n}$ as well as $d\theta', d\lambda' \in U_{A'_1,\ldots,A'_n}$, the inductive hypothesis guarantees $A_i = A'_i$ also for $i = 1, \ldots, n$. ∎

Corollary 4.11. *If $s^-_\Theta(\theta_1) = s^-_\Theta(\theta_2)$, $s^-_\Theta(\psi_1) = s^-_\Theta(\psi_2)$ but $s^-_\Theta(\theta_1) \neq s^-_\Theta(\psi_1)$, then the sets $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are unlinked.*

Proof. There must be a first $m \geq 0$ such that $a^-_\Theta(d^m\theta_1) \neq a^-_\Theta(d^m\psi_1)$. Apply the last lemma with this m. ∎

Lemma 4.12. *There is a finite bound on the number of arguments* (*right or left*) *that can share a given symbol sequence.*

Proof. Consider the full orbit of both families $\Lambda = \mathcal{O}(\mathcal{F}^\cup) \cup \mathcal{O}(\mathcal{J}^\cup)$, which is obviously invariant. At the end $|\Lambda|$—the cardinality of Λ—will be the bound.

First we claim that the number of connected components of $U_{A_0,\ldots,A_n} - \Lambda$ is less than or equal to $|\Lambda|$. Once more we resort to induction. For $n = 0$, this is plain and simple because trivially we get $\mathbb{T} - \bigcup_{A\in\mathcal{P}} A \subset \Lambda$. Next suppose $U_{A_1,\ldots,A_n} - \Lambda = \bigcup_{i=1}^{k} I_i$, where each I_i is a connected open interval and $k \leq |\Lambda|$. By construction, every set $A_0 \cap m_d^{-1}(I_i)$ is contained in a component of $\mathbb{T} - \Lambda$ and therefore is connected. The auxiliary claim follows.

Assume now that $\theta_0, \ldots, \theta_N$, where $N = |\Lambda|$, share a symbol sequence. Let $\tau > 0$ be such that the distance between two of these values is bigger than τ. Take n subject to $1/d^n < \tau$. Then the angles listed above will belong to the closure of a connected component of $U_{A_0,\ldots,A_n} - \Lambda$, where $A_i = a_\Theta(d^i\theta_0) = \cdots = a_\Theta(d^i\theta_N)$. As the number of elements $(N+1)$ is bigger than the number of intervals (at most N, as proved before), two elements are forced to fall in the same box. This is impossible because a connected subset of $U_{A_0,\ldots,A_n}$ measures at most $1/d^n$ while different elements ought to be at least τ units apart. ∎

Given a symbol sequence $s = (A_0, A_1, \ldots)$, its *shift* is $\sigma(s) = (A_1, A_2, \ldots)$. In the context of dynamics we always have $\sigma(s_\Theta(\theta)) = s_\Theta(d\theta)$.

LEMMA 4.13. *If θ is periodic, then the period of $s_\Theta(\theta)$ divides the period of θ.*

Proof. From $d^n\theta = \theta$ we get $\sigma^n(s_\Theta(\theta)) = s_\Theta(d^n\theta) = s_\Theta(\theta)$. ∎

LEMMA 4.14. *Periodic symbol sequences only happen for periodic elements.*

Proof. Let m be the period of $s_\Theta(\theta)$. If θ is irrational, then $\theta, d^m\theta, d^{2m}\theta, \ldots$ is an infinite collection of arguments sharing a symbol sequence. This goes against Lemma 4.12.

Let θ be preperiodic. For example, suppose $d^{k-1}\theta$ is not periodic but $d^k\theta$ is. Let n be the period $d^k\theta$. This tells us that the $(k-1)$th and $(n+k-1)$th entries of $s_\Theta(\theta)$ are equal, meaning that $d^{k-1}\theta$ and $d^{n+k-1}\theta$ belong to the same $A \in \mathcal{P}$ (or $A^\pm$). However, these two angles are different (simply because one is periodic, while the other is not) but they map under m_d to the same argument. This contradicts the injectivity stated in Corollary 4.6. ∎

LEMMA 4.15. *Suppose θ, θ' share the same periodic left* (*respectively right*) *symbol sequence. Then θ, θ' have the same period.*

Proof. It is a trivial application of Corollary 4.6 that multiplication by any power of d is order preserving in the set of elements with identical symbol sequence. Let k be the exact period of θ. For contradiction take a periodic θ' that shares with θ its symbol sequence but not the period. We have $d^k\theta' \neq \theta'$. Then the finite orbit of θ' under multiplication by d^k, call it $\mathcal{O}_k(\theta')$, involves more than one point. Taking into account the periods of these elements, the cyclic order of $\{\theta\} \cup \mathcal{O}_k(\theta')$ is not preserved under multiplication by d^k. This is impossible because all those elements have equal symbol sequences. ∎

Assigning a symbol sequence is a semicontinuous operation in the sense described below.

LEMMA 4.16. *Suppose γ_n is a clockwise sequence that converges to γ. Then there is an N such that for all $n \geq N$ we have $a^+_\Theta(\gamma_n) = a^+_\Theta(\gamma)$.*

Proof. Choose $\varepsilon > 0$ subject to $[\gamma, \gamma+\varepsilon) \subset A^+ = a^+_\Theta(\gamma)$. Then combining cyclic order and continuity we find that $\gamma_n \in [\gamma, \gamma+\varepsilon) \subset A^+$. But this is equivalent to $a^+_\Theta(\gamma_n) = a^+_\Theta(\gamma) = A^+$. ∎

COROLLARY 4.17. *Let γ_n be a clockwise sequence that converges to γ. Then for every M there is an N such that for $n \geq N$ the symbol sequences $s^+_\Theta(\gamma_n)$ and $s^+_\Theta(\gamma)$ agree up to the Mth position.*

Proof. We apply the last lemma to the sequences $d^i\gamma_n$, where $i = 0, \ldots, M$, in order to conclude equality of the addresses $a^+_\Theta(d^i\gamma_n) = a^+_\Theta(d^i\gamma)$ if n is large enough. But this is the same as claiming that the symbol sequences agree up to that place. ∎

5. The induced partition in the dynamical plane. In this section we introduce the induced partition of the Julia set relative to the marking. As this partition is Markov, the admissibility of the critical marking of a postcritically finite polynomial follows easily.

Let f be a postcritically finite polynomial with Θ as marking. In analogy to the partition $\mathcal{P}$ of the circle in which only arguments in $\Theta^\cup = \mathcal{F}^\cup \cup \mathcal{J}^\cup$ were missing, we build a partition of the dynamical plane off the rays with argument in $\mathcal{J}^\cup$ and extended rays with values in $\mathcal{F}^\cup$.

To simplify the exposition, we introduce more notation. For $\Lambda \subset \mathbb{T}$, denote by $\overline{\mathcal{R}}(\Lambda)$ the collection of external rays with arguments in Λ together with their landing points. Also, when Λ is a set of supporting arguments for Fatou components, we denote by $\bar{\mathcal{E}}(\Lambda)$ the union of the corresponding extended rays and the center of the component.

REMARK 5.1. Whenever $\mathcal{F}_i = \{\theta_1, \ldots, \theta_k\}$ is a Fatou marked constituent, the extended ray set $\bar{\mathcal{E}}(\mathcal{F}_i)$ forms a small web. Most of the points here are of Fatou type, the exceptions being the landing point of the rays R_{θ_j}, exactly k points as they support the component at different locations. If R_λ lands where R_{θ_1} does, then two things can happen: either the rays are equal (and so are their arguments), or they are different, in which case, due to the definition, the ray R_λ must surround the Fatou component, pressing λ to skip $\theta_2, \ldots, \theta_k$. To take care of both eventualities we can write $\lambda \in (\theta_k, \theta_1]$. This will be useful while picturing the layout.

DEFINITION 5.2. Two points $z_1, z_2 \in \mathbb{C} - \overline{\mathcal{R}}(\mathcal{J}^\cup) - \bar{\mathcal{E}}(\mathcal{F}^\cup)$ are *unlinked equivalent* if they belong to the same connected component of $\mathbb{C} - \overline{\mathcal{R}}(\mathcal{J}_i)$ and of $\mathbb{C} - \bar{\mathcal{E}}(\mathcal{F}_j)$ for all possible choices $\mathcal{F}_j, \mathcal{J}_i$ in the marking.

If $\mathcal{F}$ is non-empty, we pick $\mathcal{F}_i \in \mathcal{F}$. We then cut open $\mathbb{C}$ along $\bar{\mathcal{E}}(\mathcal{F}_i)$. The boundary of each of the $|\mathcal{F}_i|$ closed sectors so determined is given by

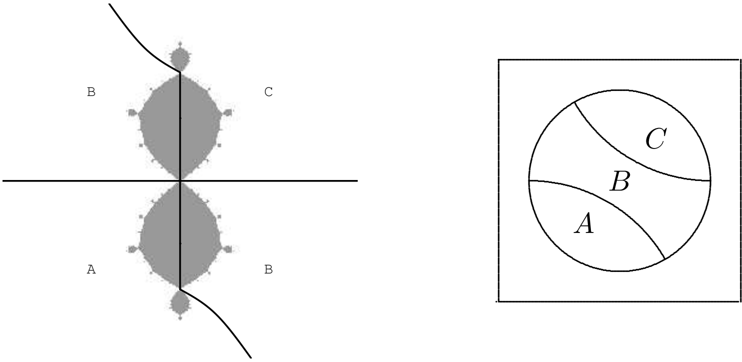

Fig. 5.1. The degree 3 critically marked polynomial $f(z) = z^3 + 1.5z$ with admissible critical portrait $\Theta = (\mathcal{F} = \{\{0, 1/3\}, \{1/2, 5/6\}\}, \mathcal{J} = \emptyset)$ determines a partition of the dynamical plane. However, this time the pieces are not connected open sets. The arguments 0 and 1/2 share the same left symbol sequence in the circle, while the rays R_0 and $R_{1/2}$ land at the same location in the dynamical plane.

two extended rays that have the same image under f. As a consequence, because f is a holomorphic map, every sector covers evenly its image, $\mathbb{C}$ in this case. Therefore, we can assign a well defined weight to each portion. Those values when taken together add up to d, the full degree of f. Most boundary points belong to the Fatou set, with the exception of the places where the supporting rays land. If θ, θ' are consecutive in $\mathcal{F}_i$, we are talking about the sector trapped between the extended rays E_θ and $E_{\theta'}$, sector that contains a ray R_λ if and only if $\lambda \in (\theta, \theta')$. As guaranteed by construction and explained in Remark 5.1, no ray with argument $\lambda \in [\theta, \theta']$ other than θ itself can share with R_θ its landing place, because otherwise λ will be the supporting argument instead of θ. Regarding θ', there might be plenty of rays colliding with it, but their arguments should range between θ and θ' (see, once again, Remark 5.1).

Take another $\mathcal{F}_j$ and again cut open. It is clear that $\bar{\mathcal{E}}(\mathcal{F}_i)$ fits in the interior of a surviving sector because the family $\mathcal{F}$ is unlinked. The worst scenario is when some of these new rays land at an already occupied Julia boundary spot, but then, as a simple argument shows, this ray will bounce to the inside and meet the center of the determining component. (Check Figure 5.1.) Still, the closed regions thus determined will be connected, the shared Julia set point pinning everything together. Furthermore, for each cutting step that we manage to accomplish, a single sector disappears and is replaced by $|\mathcal{F}_j|$ others. When we are done, we pass to the next $\mathcal{F}_i$ and the same ideas apply. At the end we are left with $1 + \sum(|\mathcal{F}_i| - 1)$ pieces, each of which covers the dynamical plane evenly.

When we run out of Fatou parts, we turn to the Julia critical ones. If the first we chose, say $\mathcal{J}_1$, results unlinked with the Fatou family, then $\overline{\mathcal{R}}(\mathcal{J}_1)$

fits inside one of the previous sectors and everything is fine: we cut open and we finish with sectors that cover $\mathbb{C}$ nicely. All this happens because, as usual, the newly added boundary lines are external rays that merge at the Julia critical point and fold under iteration.

This brings us to the exceptional case: not everything is nice in the definition, and it may occur that one or several $\mathcal{F}_k$ touch $\mathcal{J}_i$, so that really $\{\mathcal{F}_j\}_{j=1}^n$ and $\mathcal{J}_i - \varepsilon$ are unlinked. By Remark 5.1, the whole web $\overline{\mathcal{R}}(\mathcal{J}_i)$ can be installed in a single sector. However, this time at least one of the rays with argument in $\mathcal{J}_i$ is already part of the boundary of the sector. We do not bother and cut along the rays. Notice that after cutting, each problematic ray will belong to the boundary of only one of the resulting pieces. (Compare Example 2.7.)

Now let us recapitulate: we are left with $1 + \sum(|\mathcal{F}_j| - 1) + \sum(|\mathcal{J}_i| - 1)$ sectors, each "covering" evenly the complex plane at least once. As the sum above is by construction exactly d, it is impossible that any sector covers $\mathbb{C}$ more than once.

By inspecting the circle at infinity we derive additional properties. As suggested in the previous paragraphs, there are exactly d equivalence classes. Next, either an external ray is completely contained in an equivalence class or is disjoint from it. Furthermore, two rays $R_\theta, R_{\theta'}$ belong to the same class if and only if their arguments θ, θ' belong to a common $A \in \mathcal{P}$. Thus, these equivalence classes are in canonical correspondence with the elements of the partition $\mathcal{P}$ defined in the last section. For $A \in \mathcal{P}$ we denote by $\mathcal{U}_A$ the associated equivalence class in the dynamical plane. Each equivalence class is by definition a finite union of unbounded open sets. (Note that if two arguments are in the same connected component of some $A \in \mathcal{P}$, then the respective rays will be contained in the same connected region in the dynamical plane. The converse might fail.) We summarize as follows.

LEMMA 5.3. *Each of the sectors $\mathcal{U}_A$ is mapped bijectively by f into the complement of a finite number of rays and extended rays.* ■

The following technical lemma is the basis of all subsequent arguments. Let $\widetilde{\mathbb{C}}^f$ be defined as the surface with boundary obtained by cutting open the complex plane along external rays with argument in $\mathcal{O}(d\mathcal{J}^\cup)$ and extended rays with argument in $\mathcal{O}(d\mathcal{F}^\cup)$ and, in addition, by deleting the forward image of all Fatou critical points.

It should be clear what is meant here by the Fatou set $\widetilde{F}(f)$ or Julia set $\widetilde{J}(f)$.

LEMMA 5.4. *The surface $\widetilde{\mathbb{C}}^f$ is connected. The same holds for the "Julia set" $\widetilde{J}(f)$.*

Proof. It is enough to prove the second part since we can retract any Fatou point (other than a center) to the Julia set by pushing along internal or external rays.

Neither cutting external rays open nor removing centers of components disconnects the Julia set. If we pick a single extended ray supporting a Fatou component and cut open, still the Julia set remains connected. This is because the boundaries of Fatou components are Jordan curves, so that if in the modified $\widetilde{J}(f)$ an extended marked ray obstructs our way, we circumvent the obstacle by taking the alternative route.

The rest is easy. Due to the hierarchic selection of rays, each value in $\mathcal{O}(d\mathcal{F}^{\cup})$ supports a different Fatou component. As we are considering only a finite number of rays, the extended net they form does not disconnect the Julia set. This is the only place in this work where we cannot avoid the hierarchic selection of extended rays. ∎

By how we have dissected the surface, it is possible to define in $\widetilde{\mathbb{C}}^f$ (for each $A \in \mathcal{P}$) a continuous inverse branch f_A^{-1} onto $\mathrm{cl}(\mathcal{U}_A)$. From this we derive several topological properties.

LEMMA 5.5. *Both* $\mathrm{cl}(\mathcal{U}_A)$ *and its intersection* $J_A = J(f) \cap \mathrm{cl}(\mathcal{U}_A)$ *with the Julia set are pathwise connected.*

Proof. Both are the image under the continuous map f_A^{-1} of a pathwise connected set. ∎

LEMMA 5.6. *Let* $\gamma : [0,1] \to \mathbb{C}$ *be a path that crosses neither rays with argument in* $\mathcal{O}(d\mathcal{J}^{\cup})$ *nor extended rays with argument in* $\mathcal{O}(d\mathcal{F}^{\cup})$. *Suppose further that the image of* γ *is disjoint from the forward image of all Fatou critical points. If* γ *contains an interior point disjoint from those rays*, *then* γ *can be lifted in a unique way within any* $\mathrm{cl}(\mathcal{U}_A)$.

Proof. With an interior point $\gamma(t_0) \in \mathcal{U}_A$ we are able to determine the appropriate f_A^{-1}. ∎

We can go a step beyond and take the regions determined by the n-fold inverse images of those rays and external rays. Or, alternatively, we can dynamically define sets $\mathcal{U}_{A_0,\ldots,A_n}$ in concordance with Section 4. Intrinsically we have $z \in \mathcal{U}_{A_0,\ldots,A_n}$ if and only if $f^{\circ i}(z) \in \mathcal{U}_{A_i}$ for $i = 0, \ldots, n$. For an outsider standing close to ∞ the analogy is clear: a ray R_θ stays in $\mathcal{U}_{A_0,\ldots,A_n}$ if and only if $\theta \in U_{A_0,\ldots,A_n}$. Even if the set $\mathcal{U}_{A_0,\ldots,A_n}$ tends to split into several portions, their closures remain tight together.

LEMMA 5.7. *Given symbols* $A_0, \ldots, A_n \in \mathcal{P}$, *the closure* $\mathrm{cl}(\mathcal{U}_{A_0,\ldots,A_n})$ *of the region and its restriction to the Julia set* $J_{A_0,\ldots,A_n} = J(f) \cap \mathrm{cl}(\mathcal{U}_{A_0,\ldots,A_n})$ *are pathwise connected.*

Proof. We apply induction starting with Lemma 5.5. Suppose that $\mathrm{cl}(\mathcal{U}_{A_1,\ldots,A_n})$ is pathwise connected. Pick $z, z' \in J_{A_0,\ldots,A_n}$. Join $f(z), f(z')$ by an arc in $J_{A_1,\ldots,A_n}$ that avoids crossing an extended ray with argument in $\mathcal{O}(d\mathcal{F}^{\cup})$. If we lift this path within $\mathrm{cl}(\mathcal{U}_{A_0})$, we connect z to z'. ■

We are mostly concerned with the effect of this partition on the Julia set. Write $J_{A_0,A_1,\ldots} = \bigcap_{n=0}^{\infty} J_{A_0,\ldots,A_n}$. Because $J(f)$ is locally connected, the external ray R_θ lands somewhere in $J_{s^+_\Theta(\theta)} \cap J_{s^-_\Theta(\theta)}$. We now ask how big any $J_{s_\Theta(\theta)}$ could be.

LEMMA 5.8. *For any $S = (A_0, A_1, \ldots)$, the set J_S is a singleton.*

Proof (cf. [GM, Lemma 4.2]). We make use of the Thurston orbifold metric associated with f. We work in the surface with boundary equal to the disjoint union of all $\widetilde{\mathcal{U}}_A$ defined as $\mathrm{cl}(\mathcal{U}_A)$ cut open along all marked rays, extended rays and their forward images, and with the orbit of all Fatou critical points also removed. Define the distance $\varrho(z, z')$ between two points as the infimum of the lengths of smooth paths joining z and z' within $\widetilde{\mathcal{U}}_A$ (or ∞ if they belong to different components). If z and z' are in the same subset $J_{A_0,A_1} \subset J(f)$, then any path between $f(z)$ and $f(z')$ in $\widetilde{\mathcal{U}}_{A_1}$ can be lifted back to a unique path from z to z' inside $\widetilde{\mathcal{U}}_{A_0}$ (see also Lemma 5.6). Since the orbifold metric is locally strictly expanding, a compactness argument shows that restricted to the Julia set there is a universal constant $C > 1$ independent of A_i such that $\varrho(f(z), f(z')) \geq C\varrho(z, z')$. Therefore the inverse map $f^{-1}_{A_0} : J_A \to J_{A_0,A}$ contracts lengths by at least $1/C$. Hence, the iterated images $f^{-1}_{A_0} \circ \cdots \circ f^{-1}_{A_n}(J_{S_{A_{n+1}}})$ have diameter less than some initial constant divided by C^n. Taking limits with n gives us the desired unique point. ■

COROLLARY 5.9. *If two arguments share the symbol sequence, the associated rays land at the same point.*

Proof. If $s_\Theta(\theta)$ is a (right or left) symbol sequence for θ, then R_θ lands at $J_{s_\Theta(\theta)}$. ■

COROLLARY 5.10. *For any $(A_0, A_1, \ldots)$ we have $f(J_{A_0,A_1\ldots}) = J_{A_1,\ldots}$.*

Proof. For some θ either its right or left symbol sequence is $(A_0, A_1, \ldots)$. As the ray R_θ lands at the unique point in $J_{A_0,A_1\ldots}$, after mapping by f we obtain the result. ■

COROLLARY 5.11. *If $S = (A_0, A_1, \ldots)$ is periodic of period m, then the unique element of J_S is periodic of period dividing m.*

Proof. This follows directly from Lemma 5.8 and Corollaries 5.9 and 5.10. In fact, the period will be exactly m, but this piece of information is not straightforward and we will have to wait two sections. ■

We are now in a position to finally confirm that honest critical portraits satisfy the two additional conditions stated in the definition of admissible portraits.

LEMMA 5.12. *Let* $\theta \in \mathcal{J}_l$ *and* $\theta' \in \mathcal{J}_k$. *If* $s_\Theta^-(d^i\theta) = s_\Theta^-(\theta')$ *then* $d^i\theta = \theta'$.

Proof. Lemma 5.8 shows that the rays with argument $d^i\theta$ and θ' land at the same Julia set point, which is critical due to $\theta' \in \mathcal{J}_k$. If $i = 0$, the collections $\mathcal{J}_l$ and $\mathcal{J}_k$ are associated with the same critical point and *a fortiori* must agree. As different members of $\mathcal{J}_k$ assume different left addresses, we are limited to $\theta = \theta'$. When $i > 0$, it is enough to remember the hierarchic choice of rays. ∎

LEMMA 5.13. *If* λ *has the same period as* $\gamma \in \mathcal{F}_{\mathbf{per}}^{\cup}$, *then* $s_\Theta^+(\gamma) = s_\Theta^+(\lambda)$ *implies* $\gamma = \lambda$.

Proof. Both rays R_γ and R_λ can be found in the closure of the sector $\mathcal{U}_A$ where $A = a_\Theta^+(\lambda) = a_\Theta^+(\gamma)$. Because of Lemma 5.8, the rays R_γ and R_λ land together. As stated in Remark 5.1, these two facts occur simultaneously only when $\gamma = \lambda$. ∎

PROPOSITION 5.14. *The marking of a postcritically finite polynomial is an admissible portrait.* ∎

EXAMPLE 5.15 (A formal critical portrait unrelated to postcritically finite polynomials). The degree 4 strictly preperiodic formal critical portrait with $\mathcal{F} = \emptyset$ and $\mathcal{J} = \{\{3/60, 18/60\}, \{19/60, 34/60\}, \{1/60, 46/60\}\}$ does not come from the marking of a polynomial. (Compare condition (c7), here $s_\Theta^-(19/60) = s_\Theta^-(46/60)$).

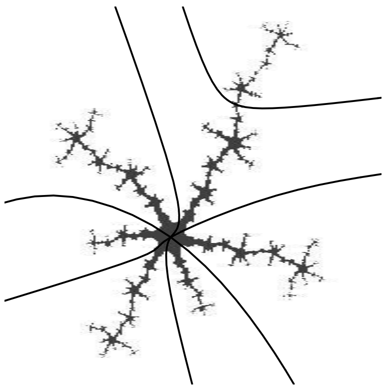

Fig. 5.2. The Julia set for $f(z) = z^4 + Az^2 + Bz + C$ with rays $R_{k/60}$ shown for $k = 1, 3, 18, 19, 31, 34, 46, 49$. (Here $A \approx 0.38 - 0.56i$, $B \approx 0.30 + 0.03i$ and $C \approx 0.49 + 0.93i$.)

In case a degree 4 polynomial realizes this critical portrait, there should likewise exist Julia set critical points $\omega_1 \neq \omega_2$ tied to $\{19/60, 34/60\}$ and

$\{1/60, 46/60\}$, respectively. But these values satisfy $s_{\Theta}^{-}(19/60) = s_{\Theta}^{-}(46/60)$, so from Corollary 5.9 we expect $\omega_1 = \omega_2$. Therefore something is wrong, and this polynomial does not carry three degree 2 critical points but only two: one of local degree 3, the other of degree 2.

If you are thinking about the polynomial of the figure, first note that at one critical point not only the rays of argument $19/60, 34/60, 1/60, 46/60$ land, but also those of argument $49/60$ and $31/60$. In short, the correct way to mark that object is as $\mathcal{J} = \{\{3/60, 18/60\}, \{19/60, 34/60, 49/60\}\}$.

6. Statical considerations that arise from weak linkage relations. The time has come to take a closer look at the abstract features hidden behind the linkage relations of two given families. Here we work with a Julia family $\mathcal{J}^* = \{\mathcal{J}_1^*, \ldots, \mathcal{J}_m^*\}$ assumed to be weakly unlinked on the right to $\mathcal{F}^* = \{\mathcal{F}_1^*, \ldots, \mathcal{F}_l^*\}$, a Fatou family.

We want to image each $\mathcal{J}_i^*$ as the arguments of rays landing at the same point, while every $\mathcal{F}_j^*$ should be related to rays that support the same Fatou component. We are expecting to gain the necessary insight into the statical layout, and, reasonably enough, no dynamics is imposed whatsoever. Slight variations of this kind of sets have been tested in other contexts, for example in the study of periodic orbit portraits (cf. [GM]).

Along this section we will impose $\mathcal{F}^{*\cup} \subset \mathcal{J}^{*\cup}$ as a technical condition. This property has a simple interpretation: all rays land somewhere but not all rays are expected to support a component.

In practice we can take a finite set and sort it into "landing" groups, each a $\mathcal{J}_i^*$. The resulting partition is $\mathcal{J}^*$. Automatically $\mathcal{J}^{*\cup}$ reconstructs the original set. Then we redistribute a subset of $\mathcal{J}^{*\cup}$ (into $\mathcal{F}^*$) so that $\mathcal{J}^*$ finishes up weakly unlinked to $\mathcal{F}^*$ on the right.

Divide $\mathbb{T} - \mathcal{J}^{*\cup}$ into unlinked classes but only with respect to $\mathcal{J}^*$. The output is a total of $w(\mathcal{J}^*) = 1 + \sum(|\mathcal{J}_i^*| - 1)$ equivalence $\mathcal{J}^*$-classes, as shown in Section 4. On the other hand, the original weak linkage relation between $\mathcal{J}^*$ and $\mathcal{F}^*$ implies that any $\mathcal{F}_j^* \in \mathcal{F}^*$ when perturbed a little to $\mathcal{F}_j^* + \varepsilon$ fits inside a unique $\mathcal{J}^*$-class.

Let us extrapolate these ideas to the plane. There we will be dealing first with the disconnection of $\mathbb{C}$ determined by rays coming together. Indeed, in the truly polynomial case, we are left with $w(\mathcal{J}^*) = 1 + \sum(|\mathcal{J}_i^*| - 1)$ connected regions with boundary, one for each basic $\mathcal{J}^*$-class.

In the abstract setting we have to prove that it is possible to achieve a similar layout. For each individual $\mathcal{J}_i^*$ consider the convex hull (in $\overline{\mathbb{D}}$) of $J_i = \{e^{2\pi i\theta} : \theta \in \mathcal{J}_i^*\}$ together with its barycenter v_i (or $v_{[\theta]}$ or whatever mnemonic allows us to remember its background). Since $\mathcal{J}^*$ is an unlinked family, the just mentioned convex hulls are disjoint, and each v_i can be joined to all the points in J_i through line segments without interference.

Whenever $\mathcal{J}_i^*$ is a singleton, everything relates to a single boundary point and the closed disk remains in one piece; fact fully coherent with the count $|\mathcal{J}_i^*|-1=0$. However, for every non-trivial $\mathcal{J}_i^*$, one of the simply connected regions is broken into $|\mathcal{J}_i^*|$ new ones by an equal number of segments, a contribution consistent with the usual weight $|\mathcal{J}_i^*|-1$ associated with $\mathcal{J}_i^*$. At this point, stage one is through.

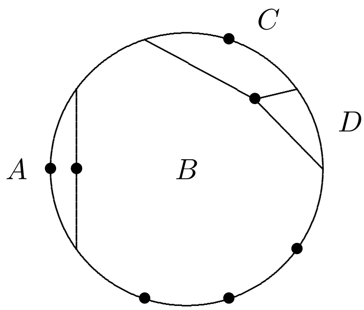

Fig. 6.1. Act One. Only the seven barycenters are marked.

We will analyze an example step by step to reinforce some ideas. For $\mathcal{J}^* = \{\{0/10, 1/10, 3/10\}, \{1/5\}, \{2/5, 3/5\}, \{1/2\}, \{7/10\}, \{4/5\}, \{9/10\}\}$, the resulting diagram is illustrated in Figure 6.1. The total degree is $w(\mathcal{J}^*)$ $= 4$, so there are four classes. Each is composed by the angles lying on the boundary of the regions shown in Figure 6.1. Call those regions A, B, C, D for the sake of discussion.

To make the example more complex we take compatible Fatou collections. For instance, let $\mathcal{F}^* = \{\{1/10\}, \{3/10, 6/10\}, \{7/10\}, \{8/10, 9/10\}\}$. We remark that it is OK not to mention all arguments since not all angles are supposed to support a component. The resulting family is weakly unlinked on the left to $\mathcal{J}^*$: the singleton $\{1/10\}+\varepsilon$ fits in the class named C, while $\{3/10, 6/10\}+\varepsilon, \{7/10\}+\varepsilon, \{8/10, 9/10\}+\varepsilon$ are all to be found inside B.

Let us say something about the most elemental case: $\mathcal{F}_1^* = \{1/10\}$ should be related to Sector C. We can describe the boundary of that region as follows: there is a circular arc running from $e^{2\pi i/10}$ to $e^{6\pi i/10}$ with a mark at $e^{4\pi i/10}$, and a wedge placed at $v_{[1/10]} = v_{[3/10]}$. The presence of this wedge is a sign that at least two arguments are identified. Outside Sector C there can be more identifications, but we do not bother to find out. We are analyzing the sector that, cyclically, starts at the angle associated with 1/10, so we conclude that among the marked arguments in $\mathcal{J}^{\cup *} \cap (1/10, 3/10]$, only 3/10 is identified with 1/10. Notice that in the postcritically finite polynomial case, we will be talking instead about an infinite wedge along the rays $R_{1/10}$ and $R_{3/10}$, together with the ray $R_{2/10}$ running across searching to land

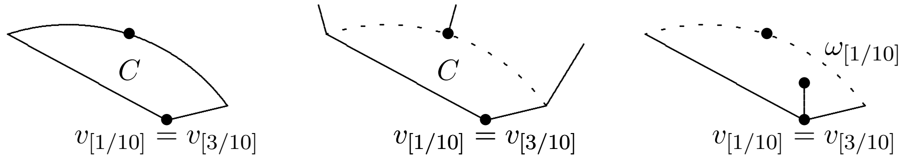

Fig. 6.2. The circular arcs can even be ignored.

elsewhere. Now we have all the ingredients to simulate the presence of the supporting ray $E_{1/10}$: the Fatou component it supports stands between $R_{1/10}$ and $R_{3/10}$. (Note that whether or not $R_{2/10}$, also in Sector C, supports this or another component is irrelevant.)

It is clear what we should do now with Sector B. We must first be able to allocate three different Fatou points $\omega_{[3/10]} = \omega_{[6/10]}, \omega_{[7/10]}, \omega_{[8/10]} = \omega_{[9/10]}$ inside. Each should be joined (imitating internal rays) to the "landing" points of the associated arguments. Still, it is not yet obvious how we can satisfactorily fulfill this task without crossings.

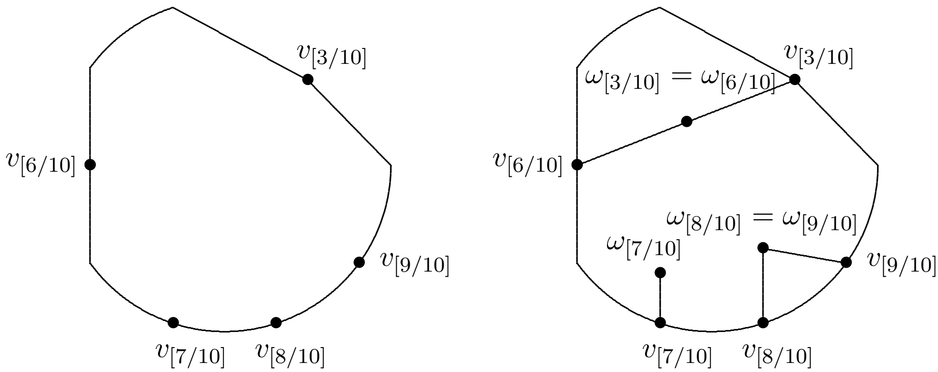

Fig. 6.3. From the inside, this looks much like a disk with boundary.

We stand inside Sector B and translate the features of the configuration to the unit closed disk. Inspecting the points on the boundary, we realize that $v_{[3/10]}, v_{[6/10]}, v_{[7/10]}, v_{[8/10]}, v_{[9/10]}$ appear in succession. We proceed to join the points supposed to support the same components. This can be done successfully since there are only $\mathcal{F}^{*\cup}$ members in sight and the $\mathcal{F}$ family is internally unlinked.

Two concluding remarks. First, it is not accidental that everything boils down to members listed in the $\mathcal{F}$ family. Notice that among $1/10 + \varepsilon$ and $3/10 + \varepsilon$ only the latter bears any relationship to Sector B. It is now clear why we refuse to work with $v_{[4/10]}$ and prefer $v_{[6/10]}$ instead (even if they are the same point). Second, for singletons we must push the barycenter to the inside: here $v_{[\theta]}$ is committed to be a Julia point while $\omega_{[\theta]}$, Fatou.

As remarked earlier, each $\mathcal{F}_k + \varepsilon$ can be thought of as immersed in an unlinked $\mathcal{J}^*$-class. Hence, we can picture them within the corresponding subsector in $\overline{\mathbb{D}}$. Let us watch the example in perspective. If we stand inside one of the simply connected regions and scan the boundary of this topological open disk, we detect two kinds of anomalies: there are isolated marked points (corresponding to a singleton $\mathcal{J}_i^*$) and "quoins" that start at $e^{2\pi i\lambda}$, then travel in the direction of $v_i = v_{[\lambda]} = v_{[\theta]}$ and finally hit $\partial\overline{\mathbb{D}}$ again at $e^{2\pi i\theta}$. This last case happens when $\lambda, \theta \in \mathcal{J}_i^*$ are different and $\theta + \varepsilon$ belongs to the correct $\mathcal{J}^*$-class. All those tents and isolated points are always disjoint. However, the interior of the region together with its visual boundary remit us to a closed topological disk, in which the order of the points is given by their subscripts. This implies that we can straighten things up so as to identify the wedge at $v_{[\theta]}$ with $e^{2\pi i\theta}$. The rest is easy: for any $\mathcal{F}_j^*$ we draw an interior point that represents $\omega_{\mathcal{F}_j^*}$ and then connect it with the vertices $v_{[\theta]}$, for $\theta \in \mathcal{F}_j^*$, as was done previously with the other families. In a hypothetical round disk this can be done by linear segments, in the actual topological disk they bend. As a matter of consistency, even if $\mathcal{F}_j^*$ is the singleton $\{\theta\}$, the vertex $\omega_{\mathcal{F}_j^*}$ should be driven slightly to the interior of the sector so as not to be mistaken with $v_{[\theta]}$. This must be done since one belongs to the Fatou set and the other to the Julia set. We must tell them apart.

As a final step, we join all points $e^{2\pi i\theta}$, for $\theta \in \mathcal{J}^{*\cup}$, to ∞ via a radial ray. These play the role of the unbounded part of the external rays. However, there is no conceptual need for the circular arcs in the unit circle anymore and it is better to ignore them.

For each $\theta \in \mathcal{J}^{*\cup}$, the external web ray comes from ∞ to $e^{2\pi i\theta}$. This point, however, is not the theoretical "landing point" for this ray: it should continue inward until it reaches $v_{[\theta]}$. In this description we are not at the mercy of $v_{[\theta]}$ sitting or not at $e^{2\pi i\theta}$.

7. The landing equivalence. For angles whose left or right symbol sequences coincide, Corollary 5.9 places the landing points of the associated rays together. However, there are noteworthy exceptions to the converse: different members of $\mathcal{J}_i \in \mathcal{J}$ are related to the same critical place. In this section we will conclude that all irregularities are a direct consequence of this phenomenon.

DEFINITION 7.1. The *landing relation* generated by an admissible critical portrait $\Theta = (\mathcal{F}, \mathcal{J})$, denoted $\sim_l$, is the smallest equivalence relation in $\mathbb{T}$ such that if one of the following two conditions hold, then $s \sim_l t$:

(l1) $s_\Theta^-(s) = s_\Theta^-(t)$,
(l2) there is j so that $a_\Theta^-(d^i s) = a_\Theta^-(d^i t)$ for all $i < j$ and $\{d^j s, d^j t\} \subset \mathcal{J}_k$ for some k.

Our primary goal is to prove that this equivalence relation captures the essence of two rays landing at the same point (see Proposition 7.7). In the meanwhile it is key to learn to recognize by inspection when two arguments are $\sim_l$-related.

We claim that an irrational or a periodic argument can only be related to arguments with identical left symbol sequence. In fact, those values and their forward orbits cannot belong to $\mathcal{J}_k$ because $\mathcal{J}^\cup$ only contains strictly preperiodic elements. Therefore, all special cases are related to the fact that future iterates share a symbol sequence with a marked "Julia" element.

For example, conditions (c7) and (c3) combined guarantee that whenever $\theta_i \in \mathcal{J}_i$ $(i = 0, 1)$, then $\theta_0 \sim_l \theta_1$ if and only if $\mathcal{J}_0 = \mathcal{J}_1$.

In practice, two arguments θ_0, θ_n are $\sim_l$-related through a sequence $\theta_0, \theta_1, \ldots, \theta_n$ where one argument is related to the next via condition (l1) or (l2) above.

We leave to the reader the task of showing that only a finite number of arguments can belong to a given equivalence class.

The following result shows that this new relation has dynamical meaning.

Lemma 7.2. *If $\theta \sim_l \theta'$ then $d\theta \sim_l d\theta'$.*

Proof. Note that $s^-_\Theta(\theta) = s^-_\Theta(\theta')$ implies $s^-_\Theta(d\theta) = s^-_\Theta(d\theta')$ by shifting. Now, if (l2) is satisfied with $j = 0$, then $\theta, \theta' \in \mathcal{J}_k$ and we get $s^-_\Theta(d\theta) = s^-_\Theta(d\theta')$ thanks to $d\theta = d\theta'$. Finally, when θ, θ' and their iterates share addresses up to the point when they reach $\mathcal{J}_k$, then the same is true for $d\theta, d\theta'$. ∎

There is a partial converse to Lemma 7.2.

Lemma 7.3. *Suppose $a^-_\Theta(\theta) = a^-_\Theta(\psi)$. Then $\theta \sim_l \psi$ if and only if $d\theta \sim_l d\psi$.*

Proof. In part, this was already done in Lemma 7.2 (despite the initial address).

In the other direction, we first restrict to the main two cases. Whenever $s^-_\Theta(d\theta) = s^-_\Theta(d\psi)$, the concordance of initial address $a^-_\Theta(\theta) = a^-_\Theta(\psi)$ implies the equality $s^-_\Theta(\theta) = s^-_\Theta(\psi)$. Otherwise, there is $j \geq 0$ such that $a^-_\Theta(d^i d\theta) = a^-_\Theta(d^i d\psi)$ for $i < j$ and $\{d^j d\theta, d^j d\psi\} \subset \mathcal{J}_k$. But all this can be rewritten as $a^-_\Theta(d^i\theta) = a^-_\Theta(d^i\psi)$ for $0 < i < j+1$, and $\{d^{j+1}\theta, d^{j+1}\psi\} \subset \mathcal{J}_k$, which together with $a^-_\Theta(\theta) = a^-_\Theta(\psi)$ gives us $\theta \sim_l \psi$.

Consider now a sequence $\lambda_0 = d\theta, \lambda_1, \ldots, \lambda_n = d\psi$ so that each element is related to the next either by (l1) or by (l2). Since multiplication by d is surjective when restricted to a fixed sector, we may write $\lambda_i = d\gamma_i$ where $a^-_\Theta(\gamma_i) = a^-_\Theta(\theta)$. The result is now a consequence of what was argued in the last paragraph taking one step at a time. ∎

Let (f, Θ) be a postcritically finite marked polynomial. Next we show in several lemmas how the $\sim_l$-classes effectively characterize the arguments of rays landing together.

Lemma 7.4. *Suppose the rays $R_\theta, R_{\theta'}$ both land at z. If z is non-critical, then θ, θ' belong to the same left sector* (*in fact*, $a^-_\Theta(\theta) = a^-_\Theta(\theta')$).

Proof. If z is not the landing point of a ray with argument in $\mathcal{F}^\cup$, then z is interior to some region $\mathcal{U}_A$. We then have $A = a^-_\Theta(\theta) = a^-_\Theta(\theta')$.

Otherwise, let $R_{\theta_1}, \dots, R_{\theta_k}$ be the rays that participate in $\mathcal{F}$ and land at z. Around z we consider local segments of these rays together with the internal rays that join z to the k centers of the different Fatou components they support. By definition of supporting argument, to every external ray follows in the cyclic order an internal segment pointing inward. Hence, this configuration divides the neighborhood of z into $2k$ consecutive regions. As every other region is contained in $\mathcal{U}_A$, where $A = a^-_\Theta(\theta_1) = \dots = a^-_\Theta(\theta_k)$, the result is established. ∎

Lemma 7.5. *Suppose R_θ lands at a critical point ω. Then for some $\theta' \in \mathcal{J}_\omega$ we have $a^-_\Theta(\theta) = a^-_\Theta(\theta')$.*

Proof. If θ is not itself a member of $\mathcal{J}_\omega$, the external ray R_θ is contained within some $\mathcal{U}_{a^-_\Theta(\theta')}$. ∎

Corollary 7.6. *Suppose θ and θ' share the left address. Then $R_\theta, R_{\theta'}$ land together if and only if $R_{d\theta}, R_{d\theta'}$ do.*

Proof. Of course, there is nothing mysterious about R_θ and $R_{\theta'}$ landing together forcing their images $R_{d\theta}$ and $R_{d\theta'}$ to do the same.

In the opposite direction, suppose $R_{d\theta}$ and $R_{d\theta'}$ land together. Let R_θ land at a point z of local degree $m \geq 1$. Then as many as m of the preimages of $R_{d\theta'}$ would also land at z: arrange them as $R_{\psi_1}, \dots, R_{\psi_m}$. However, in view of the last two results, we must have $a^-_\Theta(\theta') = a^-_\Theta(\theta) = a^-_\Theta(\psi_i)$ for some i. But as we also have $d\psi_i = d\theta'$, Corollary 4.6 implies $\psi_i = \theta'$. ∎

Proposition 7.7. *Let (f, Θ) be a marked postcritically finite polynomial. Then R_θ and $R_{\theta'}$ land at the same point if and only if $\theta \sim_l \theta'$.*

Proof. Assume first $\theta \sim_l \theta'$. If $s^-_\Theta(\theta) = s^-_\Theta(\theta')$, then the rays $R_\theta, R_{\theta'}$ land at the same point according to Corollary 5.9. Else, it is enough to consider θ, θ' related by (12). In this way, we necessarily have $a^-_\Theta(d^i\theta) = a^-_\Theta(d^i\theta')$ for $i < j$ and $\{d^j\theta, d^j\theta'\} \subset \mathcal{J}_k$ for some k. By definition, the rays $R_{d^j\theta}, R_{d^j\theta'}$ land at the same critical point ω_k. The result now becomes clear after successive applications of Corollary 7.6.

Conversely, suppose $R_\theta, R_{\theta'}$ land at z. There is a minimal $k \geq 0$ such that $f^{\circ k}(z)$ neither is critical nor contains a critical point in its forward orbit. We will prove by induction on k that $\theta \sim_l \theta'$. Let $f^{\circ n}(z)$ be non-critical for all

$n \geq 0$ (this is the starting case $k = 0$). As for all $n \geq 0$ the rays $R_{d^n\theta}, R_{d^n\theta'}$ land at the same non-critical point, Lemma 7.4 gives $a^-_\Theta(d^n\theta) = a^-_\Theta(d^n\theta')$, which proves the equality $s^-_\Theta(\theta) = s^-_\Theta(\theta')$. Next, suppose $d\theta \sim_l d\theta'$ (this is the inductive hypothesis). If z is regular, we use Lemma 7.4 again; in case z is critical, we refer the reader to Lemma 7.5. Either way the conclusion is $\theta \sim_l \theta'$. ∎

Corollary 7.8. *If $A_0, A_1, \ldots$ is a periodic sequence of period m, then the unique point in $J_{A_0,A_1,\ldots}$ has period m.*

Proof. In Corollary 5.11 we discussed why the period divides m. If the period were smaller, then two different periodic angles, and therefore non-$\sim_l$-equivalent values, will land at the same Julia set point. This contradicts Proposition 7.7. ∎

As landing classes ideally represent the physical landing spot of rays, the least we should demand from them is to be pairwise unlinked. The reader should be warned that from now on Θ is no longer the marking of a postcritically finite polynomial but merely an admissible critical portrait.

In the next results, we will be talking about two pairs $\theta_1 \sim_l \theta_2$ and $\psi_1 \sim_l \psi_2$. Whenever we add $\theta_1 \nsim_l \psi_1$ to this hypothesis, a dichotomy appears: the sets $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are either linked or unlinked. This is so because equality between one of the thetas and one of the psis will imply $\theta_1 \sim_l \psi_1$ as well.

Lemma 7.9. *Let $\theta_1 \sim_l \theta_2$ be such that $a^-_\Theta(\theta_1) = a^-_\Theta(\theta_2)$. Suppose also $\psi_1, \psi_2 \in \mathcal{J}_k$. Then $\theta_1 \nsim_l \psi_1$ implies that $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are unlinked.*

Proof. The extra condition $\theta_1 \nsim_l \psi_1$ tells us that $\{\theta_1, \theta_2\}$ and $\mathcal{J}_k$ are disjoint. However, the precise translation of $a^-_\Theta(\theta_1) = a^-_\Theta(\theta_2)$ is that $\theta_1 - \varepsilon$ and $\theta_2 - \varepsilon$ share an equivalence class. So, by definition, those arguments belong to the same connected component of $\mathbb{T} - \mathcal{J}_k$. This together with the initial remark is enough to confirm the proposed cyclic order. ∎

Lemma 7.10. *Suppose $s^-_\Theta(\theta_1) = s^-_\Theta(\theta_2)$ and that ψ_1, ψ_2 are related as in condition* (l2)*. If $\theta_1 \nsim_l \psi_1$, then $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are unlinked.*

Proof. By contradiction suppose that $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are linked. If we have $a^-_\Theta(d^i\psi_1) = a^-_\Theta(d^i\psi_2)$ for $i < j$ and $\{d^j\psi_1, d^j\psi_2\} \subset \mathcal{J}_k$, then Lemma 4.10 guarantees $a^-_\Theta(d^i\theta_1) = a^-_\Theta(d^i\psi_1)$ for $i < j$ as well. In view of this, we can apply Corollary 4.6 to conclude that $\{d^j\theta_1, d^j\theta_2\}$ and $\{d^j\psi_1, d^j\psi_2\}$ are also linked, in clear contradiction to Lemma 7.9, $\{d^j\psi_1, d^j\psi_2\}$ being part of $\mathcal{J}_k$. ∎

Corollary 7.11. *Let θ_1, θ_2 be related by condition* (l1) *and ψ_1, ψ_2 be related either by* (l1) *or* (l2)*. If $\theta_1 \nsim_l \psi_1$ then $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are unlinked.*

Proof. In fact, if ψ_1, ψ_2 share a symbol sequence this was already established in Corollary 4.11. If ψ_1, ψ_2 are related by condition (l2), this is just Lemma 7.10. ∎

LEMMA 7.12. *Let θ_1, θ_2 and ψ_1, ψ_2 be two couples related by condition* (l2) *and such that $\theta_1 \nsim_l \psi_1$. Then $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are unlinked.*

Proof. If $a^-_\Theta(\theta_1) = a^-_\Theta(\theta_2)$ and $a^-_\Theta(\psi_1) = a^-_\Theta(\psi_2)$, there are just two options: either $a^-_\Theta(\theta_1) \neq a^-_\Theta(\psi_1)$, in which case $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are unlinked by Corollary 4.11, or, by default, the four arguments belong to the same left class and we can analyze what is going on with $d\theta_1, d\theta_2, d\psi_1, d\psi_2$ instead.

Therefore, we may assume without loss of generality $a^-_\Theta(\psi_1) \neq a^-_\Theta(\psi_2)$. This in the context means $\psi_1, \psi_2 \in \mathcal{J}_k$. If also $a^-_\Theta(\theta_1) = a^-_\Theta(\theta_2)$, then we are in the setting of Lemma 7.9, and everything is fine. Else, we have $\theta_1, \theta_2 \in \mathcal{J}_l$. Now, $l = k$ implies $\theta_1 \sim_l \psi_1$, which is not the case, so we are left with $l \neq k$. But now the rest follows from the fact that $\mathcal{J}_l$ and $\mathcal{J}_k$ are unlinked due to condition (c2) in the definition of formal critical portrait. ∎

COROLLARY 7.13. *Let θ_1, θ_2 be related either by condition* (l1) *or* (l2). *Suppose $\psi_1 \sim_l \psi_2$ but $\theta_1 \nsim_l \psi_1$. Then $\{\theta_1, \theta_2\}$ and $\{\psi_1, \psi_2\}$ are unlinked.*

Proof. Let $\lambda_0 = \psi_1, \lambda_1, \ldots, \lambda_n = \psi_2$ be a sequence of arguments each $\sim_l$-related to the next by condition (l1) or (l2). Then Lemma 7.12 implies that $\{\lambda_i, \lambda_{i+1}\}$ and $\{\theta_1, \theta_2\}$ are unlinked. But this only says that λ_i and λ_{i+1} belong to the same connected component of $\mathbb{T} - \{\theta_1, \theta_2\}$. This situates all λ_i in the same interval, which means that $\{\psi_1 = \lambda_0, \lambda_1, \ldots, \lambda_n = \psi_2\}$ and $\{\theta_1, \theta_2\}$ are unlinked. ∎

PROPOSITION 7.14. *Different landing equivalent classes are unlinked.*

Proof. Use the technique of the last corollary. ∎

Finally, it is important to postulate a stronger version of condition (c2).

LEMMA 7.15. *Any $\sim_l$-class is weakly unlinked to $\mathcal{F}_k \in \mathcal{F}$ on the right.*

Proof. Call the landing class Λ. Let $\lambda_0 \in \Lambda$ and take γ_1, γ_2 consecutive in $\mathcal{F}_k$ so that $\lambda_0 \in (\gamma_1, \gamma_2]$. It is enough to prove that whenever λ_0 and λ_1 are related either by condition (l1) or by (l2), then necessarily $\lambda_1 \in (\gamma_1, \gamma_2]$. If $a^-_\Theta(\lambda_0) = a^-_\Theta(\lambda_1)$, this is a consequence of the definition (λ_0 and λ_1 belong to the same connected component of the complement of $\mathcal{F}_k$). Otherwise, we must have $\lambda_0, \lambda_1 \in \mathcal{J}_i$ for some i, and here we refer to condition (c2) in the definition of critical portrait since $\mathcal{J}_i$ is weakly unlinked on the right to $\mathcal{F}_k$. ∎

8. Special arguments that support the same component. In the last section we have recovered which rays land at the same point by screening the combinatorial data. If we intend to apply the machinery from Section 6, we need to have a clear picture of when two extended rays support the same component. In this generality this is a difficult problem. However, if we think of rays landing at periodic Fatou components, and among them we focus on those that eventually map to the preferred supporting elements, everything simplifies a great deal.

Until stated otherwise, once again Θ is the marking of a postcritically finite polynomial.

LEMMA 8.1. *Let $\gamma \in \mathcal{F}_i$. Then the extended ray E_λ supports the same Fatou component as E_γ if and only if $E_{d\lambda}$ and $E_{d\gamma}$ support the same Fatou component and, additionally, there is $\gamma' \in \mathcal{F}_i$ for which $a^+_\Theta(\gamma') = a^+_\Theta(\lambda)$.*

Proof. If E_λ supports the same Fatou component as E_γ, then E_λ belongs to the closure of one of the $|\mathcal{F}_i|$ sectors whose boundaries are determined by the extended rays with argument in $\mathcal{F}_i$. This implies $a^+_\Theta(\gamma') = a^+_\Theta(\lambda)$ for some $\gamma' \in \mathcal{F}_i$. Clearly $E_{d\lambda}$ and $E_{d\gamma}$ support the same Fatou component.

Conversely, if $E_{d\lambda}$ and $E_{d\gamma}$ support the same Fatou component and there is $\gamma' \in \mathcal{F}_i$ such that $a^+_\Theta(\gamma') = a^+_\Theta(\lambda)$, the inverse $f^{-1}_{a^+_\Theta(\lambda)}$ sends the wedge formed by $E_{d\gamma}$ and $E_{d\gamma'}$ to the wedge formed by E_γ and $E_{\gamma'}$, with center the critical point. As E_γ and $E_{\gamma'}$ by definition support the same component, the same is true with E_λ. ∎

LEMMA 8.2. *Take $\gamma \in \mathcal{O}(\mathcal{F}^\cup_{\mathbf{per}})$ that supports no critical Fatou component. Then E_λ supports the same component as E_γ if and only if $E_{d\lambda}$ and $E_{d\gamma}$ support the same Fatou component and $a^+_\Theta(\gamma) = a^+_\Theta(\lambda)$.*

Proof. The proof is almost identical to the one above. If the rays originally support a non-critical component, they should be in the interior of some $\mathcal{U}_A$; this A can only be $a^+_\Theta(\gamma) = a^+_\Theta(\lambda)$.

In the other direction, this time E_λ and E_γ are the images under $f^{-1}_{a^+_\Theta(\lambda)}$ of the rays $E_{d\lambda}$ and $E_{d\gamma}$. ∎

Let $\gamma \in \mathcal{O}(\mathcal{F}^\cup_{\mathbf{per}})$ be of period k. Then necessarily the Fatou component that E_γ supports has also period k. We are interested, as the lemmas indicate, in learning when an argument λ that eventually maps to γ supports jointly with the latter a component. Of course, a necessary condition is $d^{nk}\lambda = \gamma$ for some n.

LEMMA 8.3. *Let γ and λ be such that $\gamma \in \mathcal{O}(\mathcal{F}^\cup_{\mathbf{per}})$ is of period k and $d^{nk}\lambda = \gamma$ for some n. Then R_λ supports the same component as R_γ if and only if for each $i < nk$ either $a^+_\Theta(d^i\gamma) = a^+_\Theta(d^i\lambda)$, or $d^i\gamma$ belongs to some $\mathcal{F}_{\alpha_i}$ where we can find $\gamma' \in \mathcal{F}_{\alpha_i}$ so that $a^+_\Theta(d^i\lambda) = a^+_\Theta(\gamma')$.*

Proof. This is plain induction started by the previous two lemmas. ∎

Special arguments. All the above justifies a new abstract concept. Let $\Theta = (\mathcal{F}, \mathcal{J})$ be an admissible critical portrait. To every $\gamma \in \mathcal{O}(\mathcal{F}^{\cup}_{\mathbf{per}})$ we associate a (periodic) sequence of sets by first putting

$$\mathrm{Sp}(\gamma, 0) = \begin{cases} \{a^+_\Theta(\gamma') : \gamma' \in \mathcal{F}_\alpha\} & \text{if } \gamma \in \mathcal{F}_\alpha \text{ for some } \alpha, \\ \{a^+_\Theta(\gamma)\} & \text{otherwise,} \end{cases}$$

and later writing $\mathrm{Sp}(\gamma, j) = \mathrm{Sp}(d^j\gamma, 0)$.

Let $\gamma \in \mathcal{O}(\mathcal{F}^{\cup}_{\mathbf{per}})$ be of period $k = k(\gamma)$. We call λ a *special argument for* γ if there is an $n \geq 0$ such that $a^+_\Theta(d^i\lambda) \in \mathrm{Sp}(\gamma, i)$ for all $i < nk$, and $d^{nk}\lambda = \gamma$. In case both θ, θ' are special arguments for $\gamma \in \mathcal{O}(\mathcal{F}^{\cup}_{\mathbf{per}})$ we write $\theta \sim_\gamma \theta'$.

In the language of special arguments Lemma 8.3 reads:

PROPOSITION 8.4. *Let (f, Θ) be a postcritically finite marked polynomial. If θ is a special argument for $\gamma \in \mathcal{F}^{\cup}_{\mathbf{per}}$, then R_θ and R_γ support the same component.* ∎

The next lemma explains why an argument is special at most for one marked element.

LEMMA 8.5. *If λ is a special argument for both γ, γ' then $\gamma = \gamma'$.*

Proof. If n is a multiple of $k(\gamma)k(\gamma')$ large enough, then by definition we have $s^+_\Theta(\gamma) = \sigma^n s^+_\Theta(\lambda) = s^+_\Theta(\gamma')$. Now, on account of the hierarchic relation, interpreted as the presence of a dynamically preferred element, there is m so that $d^m\gamma$ belongs to some $\mathcal{F}_\alpha \in \mathcal{F}$. However, since we trivially have $s^+_\Theta(d^m\gamma) = \sigma^m s^+_\Theta(\gamma) = \sigma^m s^+_\Theta(\gamma') = s^+_\Theta(d^m\gamma')$, condition (c6) of admissibility implies $d^m\gamma = d^m\gamma'$, which immediately leads us toward $\gamma = \gamma'$ as both arguments are periodic. ∎

REMARK 8.6. If $\theta \sim_\gamma \theta'$ and $s^+_\Theta(\theta) = s^+_\Theta(\theta')$, we infer $\theta = \theta'$ from the definition of $\sim_\gamma$, condition (c6) and Lemma 4.9.

It is important to notice that multiplication by d is compatible with the relations defined so far.

LEMMA 8.7. *If $\lambda_1 \sim_\gamma \lambda_2$ then $d\lambda_1 \sim_{d\gamma} d\lambda_2$.*

Proof. For some k we have $\gamma = d^k\lambda_1 = d^k\lambda_2$. Thus we also have $d\gamma = d^k d\lambda_1 = d^k d\lambda_2$, and the result is a consequence of the definition of $\sim_{d\gamma}$. ∎

The next result comes as preparation for the proof of Theorem 1.3. Its significance when translated to the context of actual postcritically finite polynomials is that inverse images of marked periodic supporting rays can be found as close as you wish to the starting one (this is obvious in the context of polynomial dynamics as we are in the subhyperbolic case).

LEMMA 8.8. *Let $\Theta = (\mathcal{F}, \mathcal{J})$ be an admissible critical portrait. If $\gamma \in \mathcal{F}^{\cup}$ is periodic, then there exists arbitrarily small $\varepsilon > 0$ such that $\gamma+\varepsilon$ is a special argument for γ.*

Proof. Let $\gamma \in \mathcal{F}_i \in \mathcal{F}$ be of period k. Select $\gamma_0 \in \mathcal{F}_i$ different from γ. Remark 4.8 asserts that $a^+_\Theta(\theta) = a^+_\Theta(\gamma)$ implies $\theta \in [\gamma, \gamma_0)$. Let *mult* stand for multiplication by d^k when restricted to arguments that share with γ the first k addresses. As *mult* is a composition of bijective maps, it is injective and surjective in its own right. Also note the equality $mult(\gamma) = \gamma$. Let γ_1 be the only inverse image of γ_0 under *mult*. Then, as we have $a^+_\Theta(\gamma_1) = a^+_\Theta(\gamma)$ by construction, Remark 4.8 tells us that $\gamma, \gamma_1, \gamma_0$ are already cyclically ordered. If we continue in this fashion and inductively define γ_{n+1} as the inverse of γ_n, the fact that multiplication by d^k respects the cyclic order for sequences sharing their first k symbols allows us to conclude that γ_n is a decreasing sequence of special arguments for γ. Take a sequence converging to λ. As Corollary 4.17 implies $s^+_\Theta(\lambda) = s^+_\Theta(\gamma)$, from condition (c6) we conclude $\lambda = \gamma$. Because the original sequence is decreasing, the approximation has the correct type. ∎

Now we work out linkage issues.

LEMMA 8.9. *Suppose $\psi_1 \sim_\gamma \psi_2$. If $\gamma \notin \mathcal{F}_k$, then $\{\psi_1, \psi_2\}$ and $\mathcal{F}_k$ are unlinked.*

Proof. If $a^+_\Theta(\psi_1) = a^+_\Theta(\psi_2)$, this was observed in Remark 5.8. Otherwise the set $\mathrm{Sp}(\gamma, 0)$ is not a singleton, which happens only when γ belongs to some $\mathcal{F}_i$. Of course, we get $\mathcal{F}_i \neq \mathcal{F}_k$, as can be derived from the extra hypothesis. By definition of special argument there are two different $\gamma_1, \gamma_2 \in \mathcal{F}_i$ such that $a^+_\Theta(\psi_1) = a^+_\Theta(\gamma_1)$ and $a^+_\Theta(\psi_2) = a^+_\Theta(\gamma_2)$. By the particular case already verified, the arguments γ_j and ψ_j belong to the same connected component of $\mathbb{T} - \mathcal{F}_k$, the only one that contains $\mathcal{F}_i$ itself. ∎

Now we confirm that two different supporting sets are unlinked.

LEMMA 8.10. *Suppose $\theta_1 \sim_{\gamma_1} \gamma_1$ and $\theta_2 \sim_{\gamma_2} \gamma_2$, where $\gamma_1 \neq \gamma_2$. Then $\{\theta_1, \gamma_1\}$ and $\{\theta_2, \gamma_2\}$ are unlinked.*

Proof. As the two sets cannot have a common member, we suppose that $\{\theta_1, \gamma_1\}$ and $\{\theta_2, \gamma_2\}$ are linked and argue towards a contradiction. The strategy is simple: we will prove that $\{d\theta_1, d\gamma_1\}$ and $\{d\theta_2, d\gamma_2\}$ are still linked. Note that this is impossible, since for j large enough we have $d^j\theta_i = d^j\gamma_i$, and both "linked" sets turn out to be singletons.

So suppose the two sets are linked. First we try the case $a^+_\Theta(\theta_1) = a^+_\Theta(\gamma_1)$ and $a^+_\Theta(\theta_2) = a^+_\Theta(\gamma_2)$. By the very definition of right address, all four symbols agree and the desired fact, that $\{d\theta_1, d\gamma_1\}$ and $\{d\theta_2, d\gamma_2\}$ remain linked, is stated as Corollary 4.6.

Thus, we may further suppose $a^+_\Theta(\theta_1) \neq a^+_\Theta(\gamma_1)$. As $\mathrm{Sp}(\gamma_1, 0)$ is not a singleton, we should have $\gamma_1 \in \mathcal{F}_1$ for some member of $\mathcal{F}$. Moreover, there is $\gamma'_1 \in \mathcal{F}_1$ subject to $a^+_\Theta(\theta_1) = a^+_\Theta(\gamma'_1)$. However, the last lemma explains why $\{\gamma_1, \gamma'_1\} \subset \mathcal{F}_1$ is to be found entirely within a connected component of $\mathbb{T} - \{\theta_2, \gamma_2\}$. In this way $\{\theta_1, \gamma'_1\}$ and $\{\theta_2, \gamma_2\}$ are still linked. Now two things may happen.

If $a^+_\Theta(\theta_2) = a^+_\Theta(\gamma_2)$, then $a^+_\Theta(\theta_1) = a^+_\Theta(\gamma'_1)$ implies that $\{d\theta_1, d\gamma'_1\}$ and $\{d\theta_2, d\gamma_2\}$ are linked. However, as $\mathcal{F}_1$ is a d-preargument set we can replace $d\gamma'_1$ by $d\gamma_1$ and we are done.

Else we have $a^+_\Theta(\theta_2) \neq a^+_\Theta(\gamma_2)$. Reasoning as above we produce $\mathcal{F}_2$ in $\mathcal{F}$ from which we can take $\gamma_2, \gamma'_2 \in \mathcal{F}_2$ so that $a^+_\Theta(\theta_2) = a^+_\Theta(\gamma'_2)$ and, most important, the sets $\{\theta_1, \gamma'_1\}$ and $\{\theta_2, \gamma'_2\}$ are linked. From the equalities $a^+_\Theta(\theta_1) = a^+_\Theta(\gamma'_1)$, $a^+_\Theta(\theta_2) = a^+_\Theta(\gamma'_2)$, we conclude that $\{d\theta_1, d\gamma'_1\}$ and $\{d\theta_2, d\gamma'_2\}$ are linked. Therefore we are done because those two sets are actually equal to $\{d\theta_1, d\gamma'_1\}$ and $\{d\theta_2, d\gamma_2\}$. ∎

Corollary 8.11. *Suppose $\theta_1 \sim_{\gamma_1} \psi_1$ and $\theta_2 \sim_{\gamma_2} \psi_2$, where $\gamma_1 \neq \gamma_2$. Then $\{\theta_1, \psi_1\}$ and $\{\theta_2, \psi_2\}$ are unlinked.*

Proof. As these sets are disjoint, we assume for contradiction that they are linked. Then γ_1 shares a component of $\mathbb{T} - \{\gamma_2, \psi_2\}$, say, with ψ_1. This means that ψ_1 can be replaced by γ_1 in the linked pair. Applying a similar reasoning we change ψ_2 to γ_2. The bottom line is that we have managed to prove that $\{\theta_1, \gamma_1\}$ and $\{\theta_2, \gamma_2\}$ are linked, in contradiction to Lemma 8.10. ∎

Corollary 8.12. *Different supporting classes are unlinked.* ∎

To bring this discussion to an end we show that landing classes and supporting packages do not conflict with each other.

Lemma 8.13. *Fix a periodic $\gamma \in \mathcal{O}(\mathcal{F}^\cup)$ and let Λ be a $\sim_l$ equivalence class. Then Λ is weakly unlinked on the right to any finite subset of $\{\theta : \theta \sim_\gamma \gamma\}$.*

Proof. Suppose $\gamma \in \mathcal{F}_\gamma \in \mathcal{F}$. We will prove by induction on n that any $\sim_l$-class Λ is weakly unlinked on the right to $\Psi_n(\gamma') = \{\theta \sim_{\gamma'} \gamma' : d^n\theta \in \mathcal{F}_\gamma\}$ (here the periodic γ' belongs to the same cycle as γ, and satisfies $d^n\gamma' = \gamma$). The general case is immediately established. For $n = 0$, this is Lemma 7.15.

In general, we pick θ_1, θ_2 related by condition (l1) or (l2) in Definition 7.1 and assume for contradiction that $\{\theta_1, \theta_2\}$ is not weakly unlinked on the right to $\{\psi_1, \psi_2\} \subset \Psi_n(\gamma')$.

If $a^+_\Theta(\psi_1) \neq a^+_\Theta(\psi_2)$, then $\mathrm{Sp}(\gamma', 0)$ is not a singleton and we must have $\gamma' \in \mathcal{F}_k$ for some k. Furthermore, there are $\psi'_1, \psi'_2 \in \mathcal{F}_k$ related to the previous ones by $a^+_\Theta(\psi_1) = a^+_\Theta(\psi'_1)$ and $a^+_\Theta(\psi_2) = a^+_\Theta(\psi'_2)$. Lemma 7.15 then says that ψ'_1, ψ'_2 are in the same connected component of $\mathbb{T} - \{\theta_1 - \varepsilon, \theta_2 - \varepsilon\}$ (for small $\varepsilon > 0$), and so, the doubleton $\{\theta_1 - \varepsilon, \theta_2 - \varepsilon\}$ must be linked either to $\{\psi_1, \psi_1'\}$ or to $\{\psi_2, \psi_2'\}$ (both being subsets of $\Psi_n(\gamma')$ as soon as $n \geq 1$). Thus, eventually everything is reduced to the case $a^+_\Theta(\psi_1) = a^+_\Theta(\psi_2)$.

So let $a^+_\Theta(\psi_1) = a^+_\Theta(\psi_2)$. We claim that it is impossible to have $\theta_1 = \psi_1$ and $\theta_2 = \psi_2$ simultaneously. In fact, otherwise all four right addresses are the same, and by pushing forward with Corollary 4.6 we would deduce that $\{d\theta_1, d\theta_2\}$ is not weakly unlinked to $\{d\psi_1, d\psi_2\} \subset \Psi_{n-1}(d\gamma')$ on the right, conflicting with the inductive hypothesis. Therefore we may suppose further $\theta_1 \in (\psi_1, \psi_2)$, and $\theta_2 \in (\psi_1, \psi_2]$ as well. If $a^-_\Theta(\theta_1) = a^-_\Theta(\theta_2)$, then $a^+_\Theta(\theta_1 - \varepsilon/d) = a^+_\Theta(\theta_2 - \varepsilon/d) = a^+_\Theta(\psi_1) = a^+_\Theta(\psi_2)$ for $\varepsilon > 0$ small enough. But then, in view of Corollary 4.6, we easily deduce that $\{d\theta_1, d\theta_2\}$ is not weakly unlinked to $\{d\psi_1, d\psi_2\} \subset \Psi_{n-1}(d\gamma')$ on the right, challenging once more the inductive hypothesis. Otherwise, as θ_1, θ_2 must be related by condition (l2), we only have to deal with $\theta_1, \theta_2 \in \mathcal{J}_i$, a possibility ruled out by Remark 5.8. ∎

9. Webs. In this section we construct certain graphs and web maps associated to the portrait Θ. In order to gain generality we reshape $\mathcal{F}$ into a new $\mathcal{F}^*$ so that it conveys information about all the potential Fatou components involved. Likewise $\mathcal{J}^*$ will sort the participating members of into $\sim_l$ collections.

Fix a finite invariant set Γ of special arguments. We extend $\mathcal{F}$ to a (natural) full partition of $\Gamma \cup \mathcal{O}(\mathcal{F}^\cup)$ as follows. For $\mathcal{F}_i \in \mathcal{F}$ with no intrinsic periodic elements we write $\mathcal{F}^*_\theta = \mathcal{F}_i$ for all $\theta \in \mathcal{F}_i$. If $\theta \in \mathcal{O}(\mathcal{F}^\cup) - \mathcal{F}^\cup$ is non-periodic, then set $\mathcal{F}^*_\theta = \{\theta\}$. Finally, if $\theta \in \mathcal{O}(\mathcal{F}^\cup)$ is periodic write $\mathcal{F}^*_\theta = \{\gamma \in \Gamma \cup \mathcal{O}(\mathcal{F}^\cup) : \gamma \sim_\theta \theta\}$. It is clear that $\mathcal{F}^* = \{\mathcal{F}^*_\theta : \theta \in \Gamma \cup \mathcal{O}(\mathcal{F}^\cup)\}$ extends $\mathcal{F}$.

Next, we can also sort $\mathcal{O}(\mathcal{F}^\cup) \cup \mathcal{O}(\mathcal{J}^\cup) \cup \Gamma$ into $\sim_l$-classes to form $\mathcal{J}^* = \{\mathcal{J}^*_1, \ldots, \mathcal{J}^*_m\}$. That is, for $\theta \in \mathcal{O}(\mathcal{F}^\cup) \cup \mathcal{O}(\mathcal{J}^\cup) \cup \Gamma$ we write $\mathcal{J}^*_\theta = \{\lambda \in \mathcal{O}(\mathcal{F}^\cup) \cup \mathcal{O}(\mathcal{J}^\cup) \cup \Gamma : \lambda \sim_l \theta\}$. In the postcritically finite case this is how we group rays that land at the same point.

Proposition 9.1. *Let $\Theta = (\mathcal{F}, \mathcal{J})$ be an admissible critical portrait and Γ a finite invariant set of special arguments. Construct $\Theta^* = (\mathcal{F}^*, \mathcal{J}^*)$ as above. Then $\mathcal{J}^*$ is weakly unlinked to $\mathcal{F}^*$ on the right.*

Also, for every $\mathcal{J}^*_i$ there is $\mathcal{J}^*_{i'}$ so that $d\mathcal{J}^*_i \subset \mathcal{J}^*_{i'}$. In the same fashion, for every $\mathcal{F}^*_j$ we can find $\mathcal{F}^*_{j'}$ subject to $d\mathcal{F}^*_j \subset \mathcal{F}^*_{j'}$.

Proof. The first part is a synopsis of well established results (Proposition 7.14, Corollary 8.12 and Lemma 8.13). The second follows from Lemmas 7.2 and 8.7 (for the Julia and Fatou cases, respectively). ∎

With $\Theta^* = (\mathcal{F}^*, \mathcal{J}^*)$ as combinatorial data, the rest of this section is devoted to the construction of a topological polynomial of degree d using the graphs of Section 6 as scaffolds.

We first define an abstract 1-dimensional finite graph. The vertices will be of three types: a single special point called infinity and referred to as ∞; for each $\mathcal{J}_i^*$ a Julia type vertex denoted by $v_{[\theta]}$ whenever $\theta \in \mathcal{J}_i^*$; and for every $\mathcal{F}_j^*$ a Fatou type vertex w_j. The edges come in two different brands: "external" and "internal". As for the external edges, there will be one for each argument listed in $\mathcal{J}^{*\cup}$: for every $\theta \in \mathcal{J}^{*\cup}$ connect ∞ to $v_{[\theta]}$ along a *web ray* $\mathcal{R}_\theta$. Also, for $\theta \in \mathcal{F}_j^*$ we join each Julia vertex $v_{[\theta]}$ to w_j by an internal segment $\mathcal{I}_\theta^{\omega_j}$. Even if there are several ways to cross from ∞ to $v_{[\theta]}$ (indeed, as many as the cardinality of $\mathcal{J}_\theta^*$), there is at most one route from ω_j to $v_{[\theta]}$. This is because if $\theta, \theta' \in \mathcal{F}_j^*$ are different, then θ and θ' cannot belong to the same Julia class $\mathcal{J}_i^*$ due to weak linkage considerations. This graph is the *abstract web associated with* $\Theta^* = (\mathcal{J}^*, \mathcal{F}^*)$. We write $\mathcal{W}(\Theta^*)$ for future reference.

The second step is to embed the web in the complex plane in such a way that the order around each vertex is natural. Of course, in Section 6 this was done in the greatest of generalities, and there is no need to repeat the construction from scratch. What is important to notice, however, is that the construction is (for practical purposes) unique.

Definition 9.2. An embedding of $\mathcal{W}(\Theta^*)$ in $\widehat{\mathbb{C}}$ that matches ∞ with ∞ is an *embedded web* if the following conditions hold at the vertices.

- Around ∞ the order of the external rays $\mathcal{R}_\theta$ is clockwise (for $\theta \in \mathcal{J}^{*\cup}$). (This must be so because at finite places the observed order is reversed.)
- Around each ω_j the order of the internal segments $\mathcal{I}_\theta^{\omega_j}$ is that of $\mathcal{F}_j^*$.
- Around $v_{[\theta]}$, display $\mathcal{J}_\theta = \{\theta, \theta', \ldots, \theta\}$ cyclically. If $\theta \in \mathcal{F}^{*\cup}$ (so that θ belongs to $\mathcal{F}_j^*$ and $\mathcal{R}_\theta$ is due to support ω_j), then $\mathcal{I}_\theta^{\omega_i}$ is intercalated between $\mathcal{R}_\theta$ and $\mathcal{R}_{\theta'}$. Otherwise, when $\theta \notin \mathcal{F}^{*\cup}$, the ray $\mathcal{R}_{\theta'}$ goes after $\mathcal{R}_\theta$.

Lemma 9.3. *Any Θ^* as above determines an embedded web $\mathcal{W}(\Theta^*)$. The construction is unique in the sense that any label preserving homeomorphism between two embedded webs can be extended to all $\widehat{\mathbb{C}}$.*

Proof. The existence part was worked out in Section 6. For uniqueness, the key point is to learn how to extend the homeomorphism patch by patch. But the complement of each web is a union of topological disks that can be appended to the graph $\mathcal{W}(\Theta^*)$ by homotopic attaching functions. The result now follows from the fact that a homeomorphism of the boundary of a disk extends to a homeomorphism of the interior. For other details, we refer the reader to Lemma 9.5. ∎

Remark 9.4. We have used the fact that *any homeomorphism $\varphi : S^1 \to S^1$ can be extended to a homeomorphism of the closed disk* by means of the formula $\varphi(re^{2\pi i\theta}) = r\varphi(e^{2\pi i\theta})$.

Additionally, two such extensions are isotopic relative to the boundary. In fact, if $\Phi_0, \Phi_1 : \overline{\mathbb{D}} \to \overline{\mathbb{D}}$ agree on the unit circle, it is enough to consider $\Phi_0 \circ H$ where H is the isotopy between $\Phi_0^{-1} \circ \Phi_1$ and $\mathrm{id}_{\overline{\mathbb{D}}}$ given by

$$H(t, re^{2\pi i\theta}) = \begin{cases} re^{2\pi i\theta} & \text{if } r \geq t, \\ t\Phi_0^{-1} \circ \Phi_1(re^{2\pi i\theta}/t) & \text{otherwise.} \end{cases}$$

Moreover, when φ_0, φ_1 are homeomorphisms isotopic relative to $X \subset \partial\mathbb{D}$, then the extensions leave X intact, as the formula shows.

With Θ^* as structural data we begin to guess how f_Θ, the topological polynomial we are aiming at, would look like: it will map ∞ to ∞, send $v_{[\theta]}$ to $v_{[d\theta]}$ and take w_j to $w_{j'}$ where j' is the unique index for which $d\mathcal{F}_j^* \subset \mathcal{F}_{j'}^*$. As for the edges, it will map homeomorphically $\mathcal{R}_\theta$, which comes from ∞ to $v_{[\theta]}$, onto $\mathcal{R}_{d\theta}$, which runs from ∞ to $v_{d\theta}$, and further it will pair $I_\theta^{w_j}$, the edge from w_j to $v_{[\theta]}$ where $\theta \in \mathcal{F}_j^*$, with $I_{d\theta}^{w_{j'}}$, when $d\theta \in \mathcal{F}_{j'}^*$.

LEMMA 9.5. *For any Θ^* there is a well defined map $f_{\Theta^*} : \mathcal{W}(\Theta^*) \to \mathcal{W}(\Theta^*)$ that sends homeomorphically internal and external edges indexed by θ to the corresponding edges indexed by $d\theta$. Also, any two extensions are isotopic relative to the vertices.*

Proof. At the vertices everything is well and uniquely defined as they always map to the correct location. The map between edges is unique modulo homeomorphisms of the segments involved. It is trivial to check that any two directed homeomorphisms between segments are isotopic. ∎

A *patch* is a connected component of the complement of an embedded web. By its very definition, a patch is bounded by a sequence of edges whose corners are graph vertices. As the web map is to be extended patch by patch, the first thing to study is the function defined on the graph.

The map $f_{\Theta^*} : \mathcal{W}(\Theta^*) \to \mathcal{W}(\Theta^*)$ need not be injective at the boundary of a patch. This is inevitable since the fate of two rays with arguments listed in a marked member of the original family is to reach a critical point and fold under iteration. Even at regular Julia vertices we may find minor problems (see Lemma 9.7).

Each patch $\mathcal{P}$ is *a priori* canonically associated with an unlinked Θ^*-class. As a matter of fact, it can be thought of as immersed in one of the original Θ-unlinked classes, say $A(\mathcal{P})$.

LEMMA 9.6. *The web map $f_{\Theta^*} : \mathcal{W}(\Theta^*) \to \mathcal{W}(\Theta^*)$ is one-to-one when restricted to the Fatou vertices of a given patch.*

Proof. Take two Fatou vertices ω_1, ω_2 in the boundary a given patch. Each ω_i is associated to some $\mathcal{F}_i^* \in \mathcal{F}^*$. As a patch is naturally related to a Θ^*-class, Lemma 4.1 shows the existence of elements $\theta_i \in \mathcal{F}_i^*$ subject to $\theta_i + \varepsilon \in A(\mathcal{P})$. In this way we obtain arguments θ_1, θ_2 with a common right address and supporting ω_1, ω_2, respectively. If $f(\omega_1) = f(\omega_2)$, then $d\mathcal{F}_1^*, d\mathcal{F}_2^*$ are related to the Fatou vertex $f(\omega_1)$. Hence, both are subsets of a supporting $\mathcal{F}_3^*$. In sum, we must have $d\mathcal{F}_1^*, d\mathcal{F}_2^* \subset \mathcal{F}_3^*$.

Notice first that unless $\mathcal{F}_1^* = \mathcal{F}_2^*$, at least one of these two sets does not contain a periodic element. In fact, if $\mathcal{F}_1$ carries a periodic element, then in the periodic cycle that $\mathcal{F}_1^*$ sets in motion, $\mathcal{F}_3^*$ admits a single predecessor, which we already know to be $\mathcal{F}_1^*$. This excludes periodic elements from $\mathcal{F}_2^*$.

Therefore, we can assume without loss of generality that $\mathcal{F}_2^*$ contains no periodic elements. Since special arguments are only attached to periodic orbits, we recognize in $\mathcal{F}_2^*$ a subset of $\mathcal{O}(\mathcal{F}^{\cup})$, and conclude by the same cause that $d\mathcal{F}_2^*$ is a singleton. Its only member, say λ, can only be the preferred member of $\mathcal{F}_3^*$ whenever the hierarchic relation is in effect. We get $d\theta_2 = \lambda$.

If $\mathcal{F}_1^*$ contains no periodic element either, again $d\mathcal{F}_1^*$ equals $\{\lambda\}$. Therefore we must have $d\theta_1 = d\theta_2$, which jointly with $a_\Theta^+(\theta_1) = a_\Theta^+(\theta_2)$ implies $\theta_1 = \theta_2$ by Corollary 4.6. We are working with partitions, so $\mathcal{F}_1^* \cap \mathcal{F}_2^* \neq \emptyset$ forces $\mathcal{F}_1^* = \mathcal{F}_2^*$.

Finally, for contradiction, suppose $\tau \in \mathcal{F}_1^*$ is periodic. Then $d\tau \in \mathcal{F}_3^*$ is also periodic, and therefore must be λ as this is the preferred member of $\mathcal{F}_3^*$. On the other hand, θ_1 itself must be a special argument for τ. Here we are left with two possibilities, either $a_\Theta^+(\theta_1) = a_\Theta^+(\tau)$ or $a_\Theta^+(\theta_1) \neq a_\Theta^+(\tau)$. If $a_\Theta^+(\theta_1) = a_\Theta^+(\tau)$, transitivity leads us toward $a_\Theta^+(\theta_2) = a_\Theta^+(\tau)$, which together with $d\theta_2 = \lambda = d\tau$ implies $\theta_2 = \tau$; this is impossible because $\theta_2 \in \mathcal{F}_2^*$ is not periodic. Otherwise, if $a_\Theta^+(\theta_1) \neq a_\Theta^+(\tau)$, the very definition of special argument indicates $a_\Theta^+(\theta_1) \in \mathrm{Sp}(\tau, 0) - \{a_\Theta^+(\tau)\}$. This corroborates the existence of $\mathcal{F}_\tau \in \mathcal{F}$, in the primitive marking $\Theta = (\mathcal{F}, \mathcal{J})$, from which we can extract $\widetilde{\theta}_1 \in \mathcal{F}_\tau \subset \mathcal{F}_1^*$ subject to $a_\Theta^+(\widetilde{\theta}_1) = a_\Theta^+(\theta_1)$. Furthermore, the original $\mathcal{F}_\tau$ is a preargument set, and this implies $d\widetilde{\theta}_1 = d\tau = \lambda$. In brief, $\widetilde{\theta}_1$ and θ_2 have the same right address and both map to τ, hence by Corollary 4.6 they agree. As before, a common element in $\mathcal{F}_1^*$ and $\mathcal{F}_2^*$ is a contradiction so ω_1 and ω_2 must be the same Fatou vertex. ∎

For Julia type vertices, the situation even if not trivial is still easy to describe.

LEMMA 9.7. *Suppose $v_{[\theta_1]}, v_{[\theta_2]}$ are different Julia vertices in a given patch. Then $f(v_{[\theta_1]}) = f(v_{[\theta_2]})$ if and only if we can find $\tau_1, \tau_2 \in \mathcal{F}_\alpha \in \mathcal{F}$ (in the original marking) so that $\tau_i \sim_l \theta_i$.*

Proof. We first work out the easy part. Suppose there exist τ_1, τ_2 as above. As $\mathcal{F}_\alpha \in \mathcal{F}$ is a preargument set, we obtain $d\tau_1 = d\tau_2$. From this we get

$$f(v_{[\theta_1]}) = f(v_{[\tau_1]}) = v_{[d\tau_1]} = v_{[d\tau_2]} = f(v_{[\tau_2]}) = f(v_{[\theta_2]}),$$

and we are done.

Conversely, when $f(v_{[\theta_1]}) = f(v_{[\theta_2]})$ for $v_{[\theta_1]} \neq v_{[\theta_2]}$, then both $d\theta_1 \sim_l d\theta_2$ and $\theta_1 \nsim_l \theta_2$. Next we rule out a trivial possibility. If $a^-_\Theta(\theta_1) = a^-_\Theta(\theta_2) = A(\mathcal{P})$, then the simultaneous occurrence of $d\theta_1 \sim_l d\theta_2$ and $\theta_1 \nsim_l \theta_2$ is incompatible with Lemma 7.3. Therefore, at least one of these angles, say θ_1, has no associate with left address $A(\mathcal{P})$. If we replace θ_1 by an equivalent argument, we can also assume that $\mathcal{R}_{\theta_1}$ is visible within the patch. In that case we necessarily get $a^+_\Theta(\theta_1) = A(\mathcal{P})$ since otherwise $\mathcal{R}_{\theta_1}$ will be out of sight. We work with $a^-_\Theta(\theta_1) \neq a^+_\Theta(\theta_1) = A(\mathcal{P})$ from now on.

If we apply Lemma 4.5 to Θ, the original marking, there is $\widetilde{\theta}_1 \in \Theta^\cup$ for which $[\widetilde{\theta}_1, \theta_1]$ is a connected piece of the complement of the Θ-class labeled $A(\mathcal{P})$. This implies in particular $d\theta_1 = d\widetilde{\theta}_1$ and $a^-_\Theta(\widetilde{\theta}_1) = A(\mathcal{P})$, and also that θ_1 and $\widetilde{\theta}_1$ are marked elements. Three things can happen: both belong to the same $\mathcal{J}_i \in \mathcal{J}$, both to the same $\mathcal{F}_j \in \mathcal{F}$, or one to $\mathcal{J}_i$ and the other to $\mathcal{F}_j$. In the last case it follows easily from weak linkage considerations that θ_1 is in $\mathcal{F}_j$ while $\widetilde{\theta}_1$ is in $\mathcal{J}_i$. Also notice that the first of the three cases is impossible since $\theta_1, \widetilde{\theta}_1 \in \mathcal{J}_i$ implies that those arguments are $\sim_l$-equivalent, which is not the case as $a^-_\Theta(\widetilde{\theta}_1) = A(\mathcal{P})$. We write $\tau_1 = \theta_1$. For $\widetilde{\theta}_1$ we note that either it is equal to an element $\tau_2 \in \mathcal{F}_j$ (in the second case), or is $\sim_l$-equivalent to some $\tau_2 \in \mathcal{F}_j$ (in the weak linkage case).

An identical argument proves the existence of $\widetilde{\theta}_2 \in (\theta_1, \widetilde{\theta}_1]$ subject to $a^-_\Theta(\widetilde{\theta}_2) = A(\mathcal{P})$ and $d\theta_2 = d\widetilde{\theta}_2$. Hence we get $d\widetilde{\theta}_1 \sim_l d\widetilde{\theta}_2$ together with $a^-_\Theta(\widetilde{\theta}_1) = a^-_\Theta(\widetilde{\theta}_2)$, which implies $\widetilde{\theta}_1 \sim_l \widetilde{\theta}_2$. If $\widetilde{\theta}_2 \sim_l \theta_2$, then we are done, because it follows that θ_2 is already equivalent to τ_2 as defined before. If $\widetilde{\theta}_2$ is not equivalent to θ_2, then $[\widetilde{\theta}_2, \theta_2] \subset (\theta_1, \widetilde{\theta}_1)$ because $[\widetilde{\theta}_1, \theta_1]$ and $[\widetilde{\theta}_2, \theta_2]$ are different connected pieces of the complement of $A(\mathcal{P})$. However, in this case θ_1, θ_2 belong to different components of the complement of the common landing class of $\widetilde{\theta}_1$ and $\widetilde{\theta}_2$. By definition this means that $v_{[\theta_1]}$ and $v_{[\theta_2]}$ cannot appear in the same patch, a contradiction. ∎

Next we describe the technical way to cut open the complex plane in order to rescue the patches. To be more specific, we will be talking about two-way cuts and three-way cuts.

DEFINITION 9.8. Suppose $\theta, \theta' \in \mathcal{J}^{*\cup}$ are $\sim_l$-related. Then the path that comes from ∞ to $v_{[\theta]} = v_{[\theta']}$ along the web ray $\mathcal{R}_\theta$ and retreats through the web ray $\mathcal{R}_{\theta'}$ is the *two-way cut* $C(\theta, \theta')$.

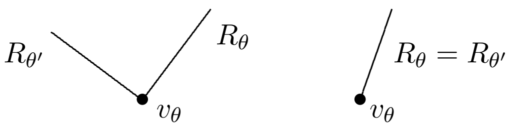

Fig. 9.1. The only two possible two-way cuts

Suppose $\theta, \widetilde{\theta} \in \mathcal{J}^{*\cup}$ are $\sim_l$-related. Suppose further that $\widetilde{\theta} \in \mathcal{F}_i^*$. Take $\theta' \in \mathcal{F}_i^*$ such that $\theta \leq \widetilde{\theta} \leq \theta'$ (here $\theta = \theta'$ implies $\widetilde{\theta} = \theta'$). The path that comes from ∞ to $v_{[\theta]}$, next pays a visit to ω_i through the internal ray $\mathcal{I}_{\widetilde{\theta}}$, continues via $\mathcal{I}_{\theta'}$ to $v_{[\theta']}$, and finally returns to ∞ using $\mathcal{R}_{\theta'}$ is the *three-way cut* $C(\theta, \widetilde{\theta}, \theta')$.

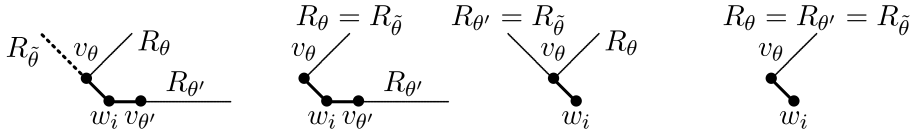

Fig. 9.2. The three-way cuts, four in total

If we think of a cut as a function, the *support* is its image.

In practice, a cut is determined by the web rays and the finite vertices it contains. In this way, we write $v(C(\theta, \theta')) = \{v_{[\theta]}\} = \{v_{[\theta']}\}$ to refer to the single finite vertex of a two-way cut.

For three-way cuts the set $v(C(\theta, \widetilde{\theta}, \theta')) = \{v_{[\theta]}, \omega_{\widetilde{\theta}}, v_{[\widetilde{\theta}]}\}$ always contains at least one Julia and one Fatou vertex. Hence, a three-way cut has either two or three elements. Anyhow, it is easy to distinguish one case from the other.

Lemma 9.9. *A three-way cut $C(\theta, \widetilde{\theta}, \theta')$ contains two vertices if and only if $\widetilde{\theta} = \theta'$.*

Proof. Suppose $v_{[\theta]} = v_{[\theta']}$, which happens to be true only if $\theta \sim_l \theta'$. In this case, the set $\mathcal{J}^*_{\widetilde{\theta}}$ contains aside from θ and $\widetilde{\theta}$, also θ'. However, the Fatou member $\mathcal{F}^*_{\widetilde{\theta}}$ by definition includes both $\widetilde{\theta}$ and θ'. This is impossible if we expect the family $\mathcal{J}^*$ to be weakly unlinked to $\mathcal{F}^*$, unless $\widetilde{\theta} = \theta'$.

If $\widetilde{\theta} = \theta'$, then $v_{[\widetilde{\theta}]} = v_{[\theta']}$, and we are talking about the same Julia vertex. ∎

On the other hand, two different cuts will intersect at a place different from ∞ if and only if they have a finite vertex in common.

Returning to embedded webs, their complements can be described as a union of simply connected patches whose visual boundaries are a succession of cyclically ordered disjoint cuts. (The *visual boundary* is the image of the oriented unit circle in a standard uniformizing coordinate.)

Lemma 9.10. *A succession $C_1, \ldots, C_n$ of cyclic disjoint cuts determines a simply connected region of the plane whose visual boundary* (*traversed in order*) *is $C_1 \cup \cdots \cup C_n$.*

Proof. Get rid of the excesses as suggested by Figures 9.1 and 9.2. ∎

Any map φ between the support of C and the support of C' will be called a *cut map* provided C', thought of as a function, is the composition of C and φ (so $C' = \varphi(C)$).

LEMMA 9.11. *Let $C_1, \ldots, C_n$ and $C'_1, \ldots, C'_n$ be cyclically ordered collections of disjoint cuts. Suppose φ is such that every C'_i factorizes as the composition of C_i and φ. Then φ can be extended to a continuous function that maps homeomorphically the interior of the regions bounded by the cuts. Any two such maps are isotopic relative to the vertices of the cuts.*

Proof. Indeed, the hypothesis guarantees that for the two domains we have at hand a well defined injective map between the visual boundaries. Extend this homeomorphism to the patch as you wish.

Besides, maps between cuts are isotopic relative to the set of vertices. This implies that any two compatible maps between the edges are isotopic. Now we can revert to Remark 9.4 and conclude that any two extensions are also isotopic. ∎

As a corollary we deduce that f_{Θ^*} can be extended, patch by patch, to a branched self covering map of the Riemann sphere.

THEOREM 9.12. *Any web map can be extended to a topological postcritically finite polynomial $f_{\Theta^*} : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ of degree d. The branched locus Ω_f of f_{Θ^*} consists of ∞ plus the vertices $v_{\mathcal{J}_i}$ and $\omega_{\mathcal{F}_k}$ where $\mathcal{J}_i \in \mathcal{J}, \mathcal{F}_k \in \mathcal{F}$ belong to the initial portrait. Furthermore, for the local degrees we have* $\deg(v_{\mathcal{J}_i}) = |\mathcal{J}_i|$ *and* $\deg(\omega_{\mathcal{F}_k}) = |\mathcal{F}_k|$. *Any two such regular extensions are isotopically equivalent.*

Proof. In order to deal with a unique extension (built patch by patch), we must comply with the conditions of Lemma 9.11. So suppose $f_{\Theta^*}(C_1) = f_{\Theta^*}(C_2)$ for two cuts C_1, C_2 associated to the same patch. This leads us to vertices $v_i \in C_i$ for which $f_{\Theta^*}(v_1) = f_{\Theta^*}(v_2)$. If $v_1 = v_2$, then $C_1 \cap C_2 \neq \emptyset$, and we have $C_1 = C_2$. Otherwise, if $v_1 \neq v_2$, the situation is described by Lemma 9.7, and necessarily v_1, v_2 belong to the same three-way cut, which can only be $C_1 = C_2$.

Therefore, after pasting together a finite number of patches, we wind up with a branched cover of a sphere. By construction we have $f^{-1}(\infty) = \{\infty\}$, so we get a topological polynomial. To determine its degree, it is enough to look at the circle at infinity, where everything behaves as multiplication by d.

For the degrees at finite places, at the moment we only have $\deg(v_{\mathcal{J}_i}) \geq |\mathcal{J}_i|$ and $\deg(\omega_{\mathcal{F}_k}) \geq |\mathcal{F}_k|$ due to the folding that is taking place. Therefore

$$\sum(\deg(v_{\mathcal{J}_i}) - 1) + \sum(\deg(\omega_{\mathcal{F}_k}) - 1) \geq \sum(|\mathcal{J}_i)| - 1) + \sum(|\mathcal{F}_k| - 1).$$

However, when we add 1 to both sides, the left hand member turns smaller than d, the total degree of the topological polynomial, while the right one becomes d, the weight of Θ, the original portrait. Hence equality must hold everywhere. As a consequence, there is no room for extra critical points.

That the polynomial is postcritically finite can be determined by trailing the vertex dynamics of the branched locus. ∎

10. The associated polynomial. Theorem 9.12 explains why, no matter how many special arguments we add to $\Theta^{*\cup}$, the end result, the topological polynomial, will be in practice the same. A good idea is to enlarge the family by shielding any periodic $\lambda \in \mathcal{F}^{*\cup}$ with a special argument $\lambda + \varepsilon$ in such a way that the interval $(\lambda, \lambda + \varepsilon)$ does not intersect the forward orbit of any member of the portrait (see Lemma 8.8). *We will assume this up to the end.*

If we start with any invariant set of special arguments, it is time to show that the postcritically finite polynomial constructed in the last section does not allow a Levy cycle. And so, as we will be unable to supply a Thurston obstruction, we end up with a postcritically finite polynomial.

So assume $\gamma_0 \mapsto \gamma_1 = f(\gamma_0) \mapsto \cdots \mapsto \gamma_{i+1} = f(\gamma_i) \mapsto \gamma_n = f(\gamma_{n-1})$, where γ_n and γ_0 are homotopic relative to the postcritical set $P(\Omega_f)$, describes a Levy cycle. Given a discrete invariant set M—preferably one containing Ω_f—we can suppose, after applying a small perturbation, that all γ are disjoint from M. For us, this M will be the set of vertices in the web.

Given a segment ℓ in the embedded web $\mathcal{W}(\Theta^*)$, we define the *incidence of* γ_i *in* ℓ, denoted $\mathrm{in}(\ell, \gamma_i)$, to be the minimal number of intersection points of ℓ and a closed curve γ' that is homotopic to γ_i relative to M, the preferred set.

The next elementary result reduces the scope of action of a potential Levy cycle (cf. [BFH, Section 8]).

Lemma 10.1. *For any edge $\ell \in \mathcal{W} = \mathcal{W}(\Theta^*)$, we have*

$$\sum_{\{\ell' \in \mathcal{W} : f(\ell') = \ell\}} \mathrm{in}(\ell', \gamma_i) \le \mathrm{in}(\ell, \gamma_{i+1}).$$

Proof. We start by picking γ_{i+1} in the same isotopy class relative to M so that it intersects ℓ exactly $\mathrm{in}(\ell, \gamma_{i+1})$ times. Now, sorting things out in such a way that $\gamma_{i+1} = f(\gamma_i)$, each intersection point of ℓ' and γ_i accounts for a new different intersection point of $\ell = f(\ell')$ and $\gamma_{i+1} = f(\gamma_i)$. Thus, we obtain the crude estimate

$$\sum_{\{\ell' \in \mathcal{W} : f(\ell') = \ell\}} |\ell' \cap \gamma_i| \le |\ell \cap \gamma_{i+1}| = \mathrm{in}(\ell, \gamma_{i+1}),$$

and the result follows due to the trivial relation $|\ell' \cap \gamma_i| \ge \mathrm{in}(\ell', \gamma_i)$. ∎

Corollary 10.2. *For any edge* $\ell \subset \mathcal{W}$ *we have* $\mathrm{in}(\ell, \gamma_i) \leq \mathrm{in}(f(\ell), \gamma_{i+1})$.

Proof. Note that on the one hand we have

$$\mathrm{in}(\ell, \gamma_i) \leq \sum_{\{\ell' \in \mathcal{W}\,:\, f(\ell') = f(\ell)\}} \mathrm{in}(\ell', \gamma_i)$$

since ℓ is one of the $\ell' \in \mathcal{W}$ subject to $f(\ell') = f(\ell)$, while on the other, the last lemma gives

$$\sum_{\{\ell' \in \mathcal{W}\,:\, f(\ell') = f(\ell)\}} \mathrm{in}(\ell', \gamma_i) \leq \mathrm{in}(f(\ell), \gamma_{i+1}).$$

Of course, both estimates together imply $\mathrm{in}(\ell, \gamma_i) \leq \mathrm{in}(f(\ell), \gamma_{i+1})$. ∎

Corollary 10.3. *If* ℓ *is a periodic edge of* $\mathcal{W}$, *then* $\mathrm{in}(\ell, \gamma_i) = \mathrm{in}(f(\ell), \gamma_{i+1})$.

Proof. By taking high multiples of the periods of ℓ and of the cycle, we can assume both periods to be n. Going around the full cycle gives

$$\mathrm{in}(\ell, \gamma_i) \leq \mathrm{in}(f(\ell), \gamma_{i+1}) \leq \cdots \leq \mathrm{in}(f^{\circ n}(\ell), \gamma_{i+n}) = \mathrm{in}(\ell, \gamma_i),$$

and we are done. ∎

Corollary 10.4. *If* ℓ *is not periodic while* $f(\ell)$ *is, then* $\mathrm{in}(\ell, \gamma_i) = 0$.

Proof. As $f(\ell)$ is periodic, there must be a periodic $\widetilde{\ell}$ so that $f(\ell) = f(\widetilde{\ell})$. This $\widetilde{\ell}$ is different from ℓ because one edge is periodic while the other is not. By limiting the inequality in Lemma 10.1 to $\ell, \widetilde{\ell}$ we are left with

$$\mathrm{in}(\ell, \gamma_i) + \mathrm{in}(\widetilde{\ell}, \gamma_i) \leq \mathrm{in}(f(\ell), \gamma_{i+1}).$$

However, Corollary 10.3 guarantees $\mathrm{in}(\widetilde{\ell}, \gamma_i) = \mathrm{in}(f(\ell), \gamma_{i+1})$, yielding $\mathrm{in}(\ell, \gamma_i) \leq 0$, a relation that implies equality as intersection numbers are always non-negative. ∎

The intersection of a closed curve γ and a subset X of the embedded web is said to be *essential* if any γ' isotopic to γ relative to M intersects X.

Proposition 10.5. *The only web edges that can essentially intersect a Levy cycle are periodic.*

Proof. We pick a preperiodic edge ℓ and show $\mathrm{in}(\ell, \gamma_i) = 0$. As every edge ℓ is eventually periodic, there is a unique $n \geq 0$ so that $f^{\circ n+1}(\ell)$ is periodic while $f^{\circ n}(\ell)$ is not. Then successive applications of Corollary 10.2 imply the chain of inequalities $0 \leq \mathrm{in}(\ell, \gamma_i) \leq \mathrm{in}(f^{\circ n}(\ell), \gamma_{i+n})$. However, the definition of n together with Corollary 10.4 implies $\mathrm{in}(f^{\circ n}(\ell), \gamma_{i+n}) = 0$ as well. ∎

This implies that by small pushes a Levy cycle can be forced never to cross a preperiodic web edge. Let us take a closer look at similar curves.

Lemma 10.6. *Suppose γ is a Jordan curve disjoint from M that meets the following criteria.*

- *All vertices in M which belong to the interior of γ are periodic and non-critical.*
- *The curve γ does not intersect a preperiodic web edge essentially.*

If $v_\theta, v_{\theta'} \in M$ (supposed to be the landing points of the periodic web rays $\mathcal{R}_\theta, \mathcal{R}_{\theta'}$, respectively) are interior to γ, then $a_\Theta^-(\theta) = a_\Theta^-(\theta')$.

Proof. Take different values $\lambda, \lambda' \in \mathcal{J}_k$, for a given k. Then the rays $\mathcal{R}_\lambda, \mathcal{R}_{\lambda'}$ divide the complex plane in two. If the vertices $v_\theta, v_{\theta'}$ each belong to one of these domains, it is clear that either $\mathcal{R}_\lambda$ or $\mathcal{R}_{\lambda'}$ cuts γ in an essential way. This is impossible since these rays are preperiodic. We conclude that θ and θ' are in the same connected component of $\mathbb{R}/\mathbb{Z} - \mathcal{J}_k$.

Now let λ, λ' be different members of $\mathcal{F}_j$. If both values are preperiodic, then working with the extended rays $\mathcal{E}_\lambda, \mathcal{E}_{\lambda'}$, the argument of the last paragraph applies, and again θ, θ' should belong to the same connected component of $\mathbb{R}/\mathbb{Z} - \{\lambda, \lambda'\}$. Finally, suppose λ is periodic (this automatically implies λ' is not). By assumption, there is $\varepsilon > 0$ for which $\lambda + \varepsilon$ is a special argument for λ and such that the orbit of $\Theta^\cup$ never catches up with the interval $(\lambda, \lambda + \varepsilon)$. From this, as $\lambda + \varepsilon$ and λ' are preperiodic, we derive once more that θ, θ' are in the same connected component of $\mathbb{R}/\mathbb{Z} - \{\lambda + \varepsilon, \lambda'\}$. As a matter of fact, θ, θ' belong either to the directed open arc $(\lambda + \varepsilon, \lambda')$ or to $(\lambda', \lambda + \varepsilon)$. (Both θ and θ' are periodic, which is not the case for λ' nor $\lambda + \varepsilon$; this is why the end points of the intervals are out of the picture.) However, we also know that $\varepsilon > 0$ was chosen so that $\theta, \theta' \notin (\lambda, \lambda + \varepsilon)$. We conclude that for $\eta > 0$ small enough, the values $\theta - \eta$ and $\theta' - \eta$ are contained in the same connected component of $\mathbb{R}/\mathbb{Z} - \mathcal{F}_j$.

All these facts together account for $a_\Theta^-(\theta) = a_\Theta^-(\theta')$. ∎

Proposition 10.7. *Let $\widehat{f}_\Theta : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ be a regular extension of the web map. Then $\widehat{f}_\Theta$ admits no Levy cycles.*

Proof. Suppose for contradiction that $\widehat{f}_\Theta$ admits a Levy cycle $\gamma_0, \ldots \ldots, \gamma_{n-1}$. When γ_i surrounds v_θ and $v_{\theta'}$, Lemma 10.6 gives $a_\Theta^-(\theta) = a_\Theta^-(\theta')$. According to Theorem 3.1 there is another element γ_{i+1} in this Levy cycle that encloses $v_{d\theta}, v_{d\theta'}$. After iterated applications of Lemma 10.6 we get $s_\Theta^-(\theta) = s_\Theta^-(\theta')$. But this is the same as $\theta \sim_l \theta'$, which in turn implies that v_θ and $v_{\theta'}$ are equal. In sum, there is a unique marked point encircled by γ_i, so there is no Levy cycle at all. ∎

Theorem 10.8. *Let $\Theta = (\mathcal{F}, \mathcal{J})$ be an admissible critical portrait. There is a unique (up to conjugation) polynomial f_Θ which is Thurston equivalent to $\widehat{f}_\Theta$. Here $\widehat{f}_\Theta$ is any regular extension of the web map.* ∎

11. Thurston's theorem revisited. Every time two postcritically finite branched covers $f, g : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ are Thurston equivalent, several practical problems appear. The way to tackle them is through the general theory of ramified and unramified covering transformations. We start by recalling standard facts.

Branched covers $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ of the Riemann sphere admit a well defined degree d. Let Ω_f be the branch locus of f. Fix a set M containing at least the critical values $f(\Omega_f)$. By its very definition, we are left with an honest degree d covering map $f : \widehat{\mathbb{C}} - f^{-1}(M) \to \widehat{\mathbb{C}} - M$ to which we can apply the theory of coverings and liftings (see for example [Mu]).

In particular, suppose there is a continuous function $g : X \to \widehat{\mathbb{C}} - M$ subject to $g(x_0) = f(z_0)$. Then the question of whether there is a continuous $\widetilde{g} : X \to \widehat{\mathbb{C}} - f^{-1}(M)$ with $\widetilde{g}(x_0) = z_0$ and such that g factors as $g = f \circ \widetilde{g}$ is named the *lifting problem* for g and is traditionally referred to by the diagram

$$\begin{tikzcd} & (\widehat{\mathbb{C}} - f^{-1}(M), z_0) \arrow[d, "f"] \\ (X, x_0) \arrow[ur, dashed, "\widetilde{g}"] \arrow[r, "g"] & (\widehat{\mathbb{C}} - M, f(z_0)). \end{tikzcd}$$

The broken line suggests that $\widetilde{g}$ is the expected unique solution. In fact, it is well known that when X is pathwise connected, under extremely mild conditions *the lifting problem admits a solution* (*which happens to be unique*) *if and only if the image of the fundamental groups under the induced functions are related as* $g_*(\pi(X, x_0)) \subset f_*(\pi(\widehat{\mathbb{C}} - f^{-1}(M), z_0))$ (cf. [Mu, Lemma 79.1]).

LEMMA 11.1. *Suppose $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ is a branched cover of the Riemann sphere. Let M be a set that includes the critical values $f(\Omega_f)$. If φ is homotopic to the identity relative to M, then there is a unique continuous ψ such that the diagram*

$$\begin{tikzcd} \widehat{\mathbb{C}} \arrow[r, "f"] \arrow[d, "\psi"'] & \widehat{\mathbb{C}} \arrow[d, "\varphi"] \\ \widehat{\mathbb{C}} \arrow[r, "f"] & \widehat{\mathbb{C}} \end{tikzcd}$$

commutes. Furthermore, this ψ is homotopic to the identity relative to $f^{-1}(M)$.

Proof. The uniqueness part is a consequence of the fact that ψ is a solution to a lifting problem for $\varphi \circ f$. (Reference points can be obtained readily. Fix a point $z_0 \in f^{-1}(f(\Omega_f))$. As φ is an isotopy, there is a path $\gamma : [0, 1] \to \mathbb{C} - \Omega_f$ from $f(z_0)$ to $\varphi(f(z_0))$. As this path avoids the critical set, we can lift it back by f^{-1} to a path σ starting at $\sigma(0) = z_0$. For the lifting problem, make sure to specify $\psi(z_0) = \sigma(1)$.)

To prove existence and the extra properties we reshape everything into a more suitable lifting problem

$$\begin{tikzcd}
 & (\widehat{\mathbb{C}} - f^{-1}(M), z_0) \arrow[d, "f"] \\
((\widehat{\mathbb{C}} - f^{-1}(M)) \times [0,1], (z_0, 0)) \arrow[ur, dashed, "\Psi"] \arrow[r, "\Phi\circ(f,\mathrm{id}_{[0,1]})"] & (\widehat{\mathbb{C}} - M, f(z_0))
\end{tikzcd}$$

where Φ is the homotopy modulo M that joins the identity (at time $t = 0$) to φ (when $t = 1$), and takes as reference any $(z_0, 0)$ since the relation $\Phi(f(z_0), 0) = f(z_0)$ is always satisfied.

As $[0, 1]$ is contractible, the fundamental group of the Cartesian product $(\widehat{\mathbb{C}} - f^{-1}(M)) \times [0, 1]$ can be identified with that of the first factor frozen at $t = 0$. Therefore, to push forward we evaluate $\Phi(f(z), t)$ at time 0 and get $f(z)$. As a consequence, $(\Phi \circ (f, \mathrm{id}_{[0,1]}))_*(\pi_1(\widehat{\mathbb{C}} - f^{-1}(M)) \times [0, 1], (z_0, 0)))$ equals $f_*(\pi_1(\widehat{\mathbb{C}} - f^{-1}(M)), z_0)$. The result amounts now to the general lifting lemma stated previous to Lemma 11.1. ∎

LEMMA 11.2. *Suppose $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ is a branched cover of the Riemann sphere. Let M be a set that includes the critical values $f(\Omega_f)$. If φ is a homeomorphism isotopic to the identity relative to M, then the unique continuous ψ that makes the diagram*

$$\begin{tikzcd}
\widehat{\mathbb{C}} \arrow[r, "f"] \arrow[d, "\psi"'] & \widehat{\mathbb{C}} \arrow[d, "\varphi"] \\
\widehat{\mathbb{C}} \arrow[r, "f"] & \widehat{\mathbb{C}}
\end{tikzcd}$$

commute is a homeomorphism isotopic to the identity relative to $f^{-1}(M)$.

Proof. If φ is such a homeomorphism, the two diagrams

$$\begin{tikzcd}
\widehat{\mathbb{C}} \arrow[r, "f"] \arrow[d, "\psi"'] & \widehat{\mathbb{C}} \arrow[d, "\varphi"] \\
\widehat{\mathbb{C}} \arrow[r, "f"] & \widehat{\mathbb{C}}
\end{tikzcd}
\qquad
\begin{tikzcd}
\widehat{\mathbb{C}} \arrow[r, "f"] \arrow[d, "\tau"'] & \widehat{\mathbb{C}} \arrow[d, "\varphi^{-1}"] \\
\widehat{\mathbb{C}} \arrow[r, "f"] & \widehat{\mathbb{C}}
\end{tikzcd}$$

can be filled in a compatible way in accordance with Lemma 11.1. Writing this out, we find $\varphi \circ f = f \circ \psi$ as well as $\varphi^{-1} \circ f = f \circ \tau$, relations that lead us toward

$$f \circ \mathrm{id}_{\widehat{\mathbb{C}}} = f = \varphi \circ f \circ \tau = f \circ \psi \circ \tau.$$

From this we deduce that $\psi \circ \tau$ solves the problem in Lemma 11.1 for the identity thought of as homotopic to the identity. Since the identity at both sides makes the diagram commutative, by uniqueness we must have $\psi \circ \tau = \mathrm{id}_{\widehat{\mathbb{C}}}$. This implies that ψ and τ are inverses of each other. Omitted details are left to the reader.

The rest of the properties are derived from what we have just learned: the homotopy between $\mathrm{id}_{\widehat{\mathbb{C}}}$ and ψ runs along invertible maps. ∎

As an application, we see how in a commutative diagram such as

$$\begin{tikzcd} (S^2, P(\Omega_f)) \arrow[r, "f"] \arrow[d, "\psi_0"'] & (S^2, P(\Omega_f)) \arrow[d, "\phi_0"] \\ (S^2, P(\Omega_g)) \arrow[r, "g"] & (S^2, P(\Omega_g)), \end{tikzcd}$$

where ϕ_0, ψ_0 are homeomorphims, we can replace ϕ_0 by any isotopic ϕ_1 by choosing an *ad-hoc* ψ_1 instead of ψ_0. In fact, as $g = \phi_0 \circ f \circ \psi_0^{-1}$, Lemma 11.2 tells us that both problems

$$\begin{tikzcd} \widehat{\mathbb{C}} \arrow[r, "f"] \arrow[d, "\psi_1"'] & \widehat{\mathbb{C}} \arrow[d, "\phi_1"] \\ \widehat{\mathbb{C}} \arrow[r, "g"] & \widehat{\mathbb{C}} \end{tikzcd} \qquad \begin{tikzcd} \widehat{\mathbb{C}} \arrow[r, "f"] \arrow[d, "\psi_0^{-1}\psi_1"'] & \widehat{\mathbb{C}} \arrow[d, "\phi_0^{-1}\phi_1"] \\ \widehat{\mathbb{C}} \arrow[r, "f"] & \widehat{\mathbb{C}} \end{tikzcd}$$

can be solved simultaneously. This means that as $\phi_0^{-1}\phi_1$ is isotopic to the identity relative to the appropriate set, the same is true for $\psi_0^{-1}\psi_1$.

In our problem, the topological polynomial we have constructed "admits" the correct marking if we work with artificial rays. The remaining task is to authenticate Θ as a marking for f. Thus, in some sense we have to prove that the chief member of the class borrows properties from the rough model.

As it is too much to ask from f, the truly analytical polynomial, and f_Θ, its topological model, to be identical, we fit everything into a commutative diagram

$$\begin{tikzcd} \widehat{\mathbb{C}} \arrow[r, "f_\Theta"] \arrow[d, "\psi"'] & \widehat{\mathbb{C}} \arrow[d, "\mathrm{id}"] \\ \widehat{\mathbb{C}} \arrow[r, "f"] & \widehat{\mathbb{C}} \end{tikzcd}$$

where ψ is isotopic to the identity modulo the postcritical set.

The second small improvement for f_Θ that we can handle on theoretical grounds alone is to assume that the critical sets of f and f_Θ are the same, so that ψ is also the identity in the critical set.

When this is accomplished, we conclude in harmony with Lemma 11.1 that ψ is isotopic to the identity not only relative to the postcritical set but also relative to the critical set. (There is a danger here: to be the identity in a set does not imply that the isotopy fixes this set pointwise. For an example see Example 12.1.)

This raises a new possibility. Start with $\varphi_0 = \mathrm{id}$ and inductively define new homeomorphisms φ_n isotopic to the identity modulo the postcritical set framed in the Thurston type diagram

$$\begin{tikzcd}
\widehat{\mathbb{C}} \arrow[r, "f_\Theta"] \arrow[d, "\varphi_{n+1}"'] & \widehat{\mathbb{C}} \arrow[d, "\varphi_n"] \\
\widehat{\mathbb{C}} \arrow[r, "f"] & \widehat{\mathbb{C}}.
\end{tikzcd}$$

Whenever W is a f_Θ-invariant set that displays some analytic features (there is no reason to hide we are thinking about the web $\mathcal{W}(\Theta)$), the successive embeddings $\varphi_{n+1}(W)$ are actual subsets of $f^{-1}(\varphi_n(W))$, as an easy diagram chase shows. This is good because holomorphic preimages tend to improve the analytical features and dissipate the non-analytical ones.

12. The marking of the associated polynomial. We already know about the existence of a unique postcritically finite polynomial associated with the admissible critical portrait $\Theta = (\mathcal{F}, \mathcal{J})$. Our final task is to find a suitable global chart where this polynomial admits Θ as marking. The existence of such a coordinate system is not as obvious as it may seem, and we will still have to cope with minor complications.

Example 12.1. For $\Theta = (\mathcal{F}, \mathcal{J})$ with $\mathcal{F} = \{0, 1/3, 2/3\}, \mathcal{J} = \emptyset$, imagine $\widehat{f}(z) = z^3$ as a "topological polynomial" in a web $\mathcal{W}(\Theta)$ with vertices $V = \{0, 1, e^{2\pi i/3}, e^{4\pi i/3}\}$ and extended rays $\mathcal{E}_{k/3} = \{re^{2k\pi i/3} : r \in [0, \infty)\}$ for $k = 0, 1, 2$. Theorem 10.8 claims that this branched covering is equivalent to a unique polynomial: doubtlessly $f(z) = z^3$.

Consider the two homeomorphisms

$$\psi_0(r^3 e^{2\pi i\theta}) = \begin{cases} r^3 e^{2\pi i\theta} & \text{if } r \le 3; \\ r^3 e^{2\pi i[\theta + \frac{3}{2}(\frac{\ln r - \ln 3}{\ln 4 - \ln 3})]} & \text{if } 3 \le r \le 4; \\ r^3 e^{2\pi i[\theta + 3/2]} & \text{if } 4 \le r, \end{cases}$$

$$\psi_1(re^{2\pi i\theta}) = \begin{cases} re^{2\pi i\theta} & \text{if } r \le 3; \\ re^{2\pi i[\theta + \frac{1}{2}(\frac{\ln r - \ln 3}{\ln 4 - \ln 3})]} & \text{if } 3 \le r \le 4; \\ re^{2\pi i[\theta + 1/2]} & \text{if } 4 \le r \end{cases}$$

inserted in the commutative diagram

$$\begin{tikzcd}(\mathbb{C}, V) \arrow[r, "\widehat{f}"] \arrow[d, "\psi_1"'] & (\mathbb{C}, V) \arrow[d, "\psi_0"] \\ (\mathbb{C}, V) \arrow[r, "f"] & (\mathbb{C}, V).\end{tikzcd}$$

Even if ψ_0 is isotopic to the identity in the postcritical set and is actually the identity near ∞, it is not true that ψ_0 is isotopic to the identity modulo a huge neighborhood of ∞. This always represents a source of misunderstanding.

Here we explain what is really going on. The map ψ_0 performs a "Dehn twist" of 3/2 turns far away from ∞. Thus, the web $\mathcal{W}(\Theta)$ itself is coiled that much. By this we mean that when keeping track of the image $\psi_0(\mathcal{R}_0)$ of the web ray $\mathcal{R}_0$, we start at the actual ray R_0 for a while, then move in counterclockwise direction until we have completed 3/2 turns and, finally, continue our way to ∞ along the track of the ray $R_{1/2}$. Something similar occurs with all other elements.

Now, when pulling back $\psi_0(\mathcal{W}(\Theta))$ by f^{-1}, the resulting embedded web $\psi_1(\mathcal{W}(\Theta))$ has a different intermediate behavior which translates, in the end, into a non-standard labeling near ∞ (however, the configuration as a whole remains isotopic to the optimal choice). For example, the image web ray $\psi_1(\mathcal{R}_0)$ travels for a while in the direction of the actual R_0 ray, then twists 1/2 turns, and finally continues happily imitating $R_{1/2}$.

The situation gets even more delicate for successive lifts of the web rays associated with $\theta = 0$. In these cases, near ∞ they will be respectively confused with fragments of $R_{1/2}, R_{1/6}, R_{1/18}, \ldots$. Of course, aside from its impertinence, there is nothing wrong about this. But we rather handle the correct identifications to avoid distractions.

To propose a possible way out, we note that ψ_0 and ψ_1 agree near ∞, say for $|z| \geq \alpha$ with α large. If we remove the set $\{z : |z| \geq \alpha\}$ from the complex plane, then ψ_0, ψ_1 will not be isotopic in this new Riemann surface with boundary, since they will differ by one spin around the circle $|z| = \alpha$. This is hardly a surprise because the difference in 360 degrees can be measured by comparing the embedded web with its lift. True, in this particular example, the embedding $\psi_0(\mathcal{W})$ was not the wisest choice: it spirals 3/2 turns too much. When we lift back the web, this mismatch will be divided by the degree of the polynomial (3 in this case). Thus, the "difference in twist" (which can always be measured) allows us to state the equation-like relation

$$\text{twist} - \text{twist}/d = \text{difference in twist},$$

where d is the degree of the polynomial and *difference in twist* is the relative twist of the web ray $\psi_1(\mathcal{R}_0)$ (in the lifted web) with respect to the original $\psi_0(\mathcal{R}_0)$. The formula hints that any strange behavior is to be blamed on a Dehn twist in a neighborhood of Fatou points. This is indeed the case as we will establish shortly.

From the admissible critical portrait $\Theta = (\mathcal{F}, \mathcal{J})$ we have constructed a unique (up to affine conjugation) postcritically finite polynomial f of degree d (which we take for granted to be monic and centered). Also, the standard commutative diagram holds for any regular extension f_Θ. By replacing f_Θ by $\psi_0 \circ f_\Theta \circ \psi_0^{-1}$ and ψ_1 by $\psi_1 \circ \psi_0^{-1}$, we may assume that ψ_0 is the identity. As explained in the last section, there is no problem to take f and f_Θ with the same critical set and to assume ψ_1 to be the identity also in Ω_f.

For each periodic Fatou point $\omega \in \Omega_f$ let ϕ_ω denote a fixed Böttcher coordinate associated with ω. (For notational convenience we include ∞ in the critical set.) Given $r < 1$ write $N_r(\omega) = \{z \in U(\omega) : |\phi_\omega(z)| < r\}$. For each preperiodic Fatou point $\omega \in \mathcal{O}(\Omega_f)$, inductively let $N_r(\omega)$ be the component of $f^{-1}(N_r(f(\omega)))$ that contains ω. For a subset $X \subset \mathcal{O}(\Omega_f)$ let $N_r(X) = \bigcup_{\omega \in X} N_r(\omega)$.

Now, as there is no purely topological method to distinguish $\widehat{\mathbb{C}} - \mathcal{O}(\Omega_f)$ from $\widehat{\mathbb{C}} - N_r(\mathcal{O}(\Omega_f))$, we can construct an embedded web and a regular extension f_Θ so that the following criteria are met:

- For $z \in N_{1/2}(\mathcal{O}(\Omega_f))$ we have $f_\Theta(z) = f(z)$.
- In any $N_{1/2}$, the web rays correspond to actual internal rays of f. Furthermore, the labeling of web rays agrees with the one derived from the Böttcher coordinates.

Denote by $\mathcal{W}$ this web and by V the corresponding collection of vertices (there is no further need to write this set as $\psi(V)$). Recall that we are taking ψ_0 to be the identity. The construction implies that near critical points the homeomorphism ψ_1 is a rotation in the Böttcher coordinate.

Untwisting external rays. We first try to improve the construction in the basin of attraction of ∞. Notice that for $r \leq 1/2$ the portion of $\mathcal{W}$ found in $N_r(\infty)$ is a union of ray segments pinned at ∞. Fix $\theta \in \Theta^{\cup}$. The web ray $\mathcal{R}_{d\theta} = \psi_0(\mathcal{R}_{d\theta})$ by definition must agree with the actual ray $R_{d\theta}$ near ∞. Therefore, the portion of the web ray $\psi_1(\mathcal{R}_\theta) \cap N_r(\infty)$ constitutes part of a ray $R_{\theta+j/d}$ (this is because all preimages of $R_{d\theta}$ look like that). Furthermore, we can measure the relative twist of $\psi_1(\mathcal{R}_\theta)$ with respect to $\psi_0(\mathcal{R}_\theta)$ in $\partial N_r(\infty)$, which by construction is a number of the form

j/d plus an integer, that is, a fraction k/d. Stating this fact as an equation

$$\text{twist} - \text{twist}/d = \text{difference in twist},$$

we obtain a rational solution $k/(d-1)$ (same k as above).

To confirm this potential twist is a true one, take $0 < s < r$ and consider $N_{r^d}(\infty) - N_{s^d}(\infty)$. We modify ψ_0 inside this cylinder by performing a twist of $-k/(d-1)$ turns. This forces us to change ψ_1 by $-k/d(d-1)$ turns in $N_r(\infty) - N_s(\infty)$ to keep the diagram commutative. There is no obstruction to neither procedure because ψ_0 is the identity in $N_{r^d}(\infty)$ and ψ_1 is a rotation in $N_r(\infty)$ in the Böttcher coordinate. As we have $N_{r^d}(\infty) \subset N_r(\infty)$, outside $N_r(\infty)$ nothing is touched.

Formally, in the complement of $N_r(\infty) \cup V$ the homeomorphisms ψ_0, ψ_1 are not isotopic relative to the boundary $\partial N_r(\infty)$ as they differ by k/d turns. In $N_r(\infty) - N_{s^d}(\infty)$ the modified ψ_0 and ψ_1 find themselves $-k/d$ turns apart. Combining those details, in the bigger $\widehat{\mathbb{C}} - N_{s^d}(\infty) - V$ the new ψ_0, ψ_1 are isotopic relative to the boundary. Thus, the relative difference between the "new" web rays $\psi_0(\mathcal{R}_\theta)$ and $\psi_1(\mathcal{R}_\theta)$ is null when measured in $\partial N_{s^d}(\infty)$. In particular, the successive lifts $\psi_n(\mathcal{W}) \subset f^{-n}(\psi_0(\mathcal{W}))$ (see the end of Section 11) will have no difference in twist.

However, there is still a small pitfall: when you follow the actual stable route of the ray $\mathcal{R}_\theta$ close to ∞, you turn out to stand in direction $\theta - k/(d-1)$ instead of θ. Of course, this means that the rays are not where they are supposed to be: they will only be after conjugation by the global rotation $\psi(z) = \lambda z$, where $\lambda = e^{-2k\pi i/(d-1)}$.

Untwisting periodic preferred internal rays. The next thing to do is to perform the analogous construction in the basin of attraction of finite periodic critical cycles. Suppose $\omega_0 \mapsto \omega_1 \mapsto \cdots \mapsto \omega_n = \omega_0$ is a critical cycle, and let d_i be the local degree of ω_i. We want to prove that each coordinate in this cycle is also twisted by say t_i turns. Let ℓ_i be the preferred internal web edge adjacent to ω_i, the one always assigned argument 0.

We measure the displacement of $\psi_1(\ell_i) \subset f^{-1}(\psi_0(\mathcal{W}))$ relative to $\psi_0(\ell_i)$, its counterpart in $\psi_0(\mathcal{W})$. Let this value be y_i (which can only be a rational number with denominator d_i). If the coordinates are crooked, then the possible twist of $\psi_0(\ell_i)$ is t_i by construction; while when "lifting back" ℓ_{i+1} to recover $\psi_1(\ell_i)$, its possible twist t_{i+1} appears divided by d_i. If we want to untwist as in the previous paragraph, we must solve the system of equations $t_i = t_{i+1}/d_i + y_i$ for rational t_i with denominator $d_0 d_1 \cdots d_{n-1} - 1$. But note that this can be easily achieved provided we rewrite the system with integer coefficients as

$$\begin{array}{rclcrcl}
d_0 d_1 \cdots d_{n-1} & t_0 & = d_1 \cdots d_{n-1} & t_1 & + d_0 d_1 \cdots d_{n-1} & y_0, \\
d_1 \cdots d_{n-1} & t_1 & = d_2 \cdots d_{n-1} & t_2 & + \ d_1 \cdots d_{n-1} & y_1, \\
& & \vdots & & & \\
d_{n-2} d_{n-1} & t_{n-2} & = \quad d_{n-1} & t_{n-1} & + \ d_{n-2} d_{n-1} & y_{n-2}, \\
d_{n-1} & t_{n-1} & = & t_0 & + \quad d_{n-1} & y_{n-1}.
\end{array}$$

(Note that $d_i y_i$ is always an integer.) With the solutions $t_0, \ldots, t_{n-1}$ we proceed to untwist the conjugacy in all neighborhoods of the cycle at the same time.

Untwisting non-periodic Fatou critical components. The last basins that require some adjustment are the ones determined by strictly preperiodic Fatou critical points. Let ω be such a critical point, and let $\omega' = f^{\circ n}(\omega)$ be the first critical point in its forward orbit. We assume that near ω' everything has already been straightened up. In this case the corrective measures are taken from an equation $t_\omega = y_\omega$, so essentially there is nothing left to do.

Proof of the main result. In order to finish the proof of the main result, namely that the postcritically finite polynomial admits the initial marking, we need to show that the corresponding external and internal rays land where they are expected to. This, indeed, is enough to assess the correct marking to the Julia set critical points. In the Fatou case we still have to confirm that the "supporting" rays do actually support the component.

We take one of the marked vertices v_θ and check now if all internal and external rays supposed to land together actually do. To make things conceptually clear, suppose first that the $\sim_l$-class of θ is not a singleton within $\mathcal{J}^{*\cup}$, that is, there is another angle θ' for which the topological rays $\mathcal{R}_\theta$ and $\mathcal{R}_{\theta'}$ intersect at v_θ.

LEMMA 12.2. *Suppose $\theta \sim_l \theta'$ are two different arguments in $\mathcal{J}^{*\cup}$. Then for the polynomial f the two rays R_θ and $R_{\theta'}$ land at the same point.*

Proof. In the starting web $\mathcal{W}$ the union $\mathcal{R}_\theta \cup \mathcal{R}_{\theta'}$ consists of portions of true rays plus a not completely understood compact leftover that amalgamates everything together. However, in the successive lifts $\psi_n(\mathcal{W}) \subset f^{-n}(\psi_0(\mathcal{W}))$, the faithful rays $R_\theta, R_{\theta'}$ become more dominant and at the same time the remnant part appears less significant relative to the orbifold metric. This means, in practice, that the two rays R_θ and $R_{\theta'}$ cluster together. But as we are in the connected and locally connected Julia set case, this behavior is only possible when these two rays share the landing point. ■

As we can see, the reasoning works because two or more things are pulling from different directions and at the same time they grow more and more analytical. The same occurs when the landing point of a supporting ray is trapped between ∞ and the center of the Fatou component it protects.

LEMMA 12.3. *Let $\theta \in \mathcal{F}^{*\cup}$ be a supporting argument, in theory associated to the critical Fatou point ω. Then the actual external ray of argument θ and the associated ω-internal ray land at the same boundary point of $\partial U(\omega)$.*

Proof. The same idea. In the starting web $\mathcal{W}$ the union $\mathcal{R}_\theta \cup \mathcal{I}_\theta^\omega$ can be described as portions of true rays plus a compact piece. However, in the successive lifts $\psi_n(\mathcal{W}) \subset f^{-n}(\psi_0(\mathcal{W}))$ the actual external and internal rays become bigger and at the same time the complement shrinks. As rays from two different basins—in the locally connected case, of course—come arbitrarily close to each other only if they collide, we are done. ∎

We remark that even if this implies that R_θ can be stretched up to ω, we are still not able to conclude that R_θ supports the component $U(\omega)$. This will be settled shortly.

Finally, let us suppose R_θ is not required to support a component nor has a recognized landing partner. This, by construction, implies that the nominal landing point of $\mathcal{R}_\theta$ is a Julia postcritical point, and as such, becomes one of the untouchables for the isotopies ψ_n.

LEMMA 12.4. *If θ is a postcritical Julia argument* (*that is*, $\theta \in \mathcal{O}(d\mathcal{J}^\cup)$), *then the ray R_θ lands at the marked postcritical vertex v_θ.*

Proof. This time the ray $\mathcal{R}_\theta$ is the union of a subset of R_θ and a small arc connecting it to v_θ. The successive lifts can be depicted the same way, that is, they all finish at v_θ. This brings part of the ray R_θ arbitrarily close to the Julia set point v_θ, which in the postcritically finite case is consistent only if R_θ lands there. ∎

Now it only remains to confirm that the periodic rays R_γ associated with Fatou critical sets actually support the component where they land.

LEMMA 12.5. *Suppose $\gamma \in \mathcal{F}_j^*$ is periodic. Then R_γ supports the Fatou component with center $\omega(\mathcal{F}_j^*)$.*

Proof. Choose $\varepsilon > 0$ so that $\gamma + \varepsilon$ is a special argument for γ. Starting all over again, we find that γ and $\gamma + \varepsilon$ belong to the same $\mathcal{F}_j^*$ and as such are expected to support the same "component". This implies for the moment that they both land at the boundary of $U(\omega_j)$, but nothing else. If $\varepsilon > 0$ was chosen small enough so that in the interval $(\gamma, \gamma + \varepsilon)$ there is no room for a periodic argument of the same period as γ, we are through. In fact, any other ray $R_{\gamma'}$ that shares with R_γ its landing point must have the same period as γ. So, the hypothetical $R_{\gamma'}$ cannot be inserted between R_γ and $R_{\gamma+\varepsilon}$, and consequently by definition R_γ supports $U(\omega_j)$. ∎

THEOREM 12.6. *For any degree d critical portrait $\Theta = (\mathcal{F}, \mathcal{J})$ there is a unique monic centered postcritically finite polynomial f with marking (f, Θ).*

Proof. Along this section we have established that the canonical polynomial associated with Θ admits the correct marking.

To settle uniqueness we remark first that the act of delineating external and internal rays—critical and postcritical—and using f itself determines, lawfully, a topological polynomial associated with the data Θ. If we have two different polynomials to work with, we finish up with two different functions that, according to Theorem 9.12, must be isotopic. However, Theorem 3.2 states that each Thurston class contains at most one polynomial up to conjugation. ∎

Acknowledgments. We would like to thank John Milnor for his insight and helpful suggestions. Some of the arguments presented here are in its final formulation due to him. We are also grateful, among others, to Ben Bielefield, John Hubbard and Jan Kiwi for discussions at different moments of the preparation of the manuscript.

References

[BFH] B. Bielefeld, Y. Fisher and J. Hubbard, *The classification of critically preperiodic polynomials as dynamical systems*, J. Amer. Math. Soc. 5 (1992), 721–762.

[CG] L. Carlsen and T. W. Gamelin, *Complex Dynamics*, Springer, 1993.

[D] A. Douady, *Descriptions of compact sets in* $\mathbb{C}$, in: Topological Methods in Modern Mathematics, Publish or Perish, 1993, 429–465.

[DH1] A. Douady et J. Hubbard, *Étude dynamique des polynômes complexes*, *I*, *II*, Publ. Math. Orsay 1984-1985.

[DH2] —, —, *A proof of Thurston's topological characterization of rational maps*, Acta Math. 171 (1993), 263–297.

[F] Y. Fisher, *The classification of critically preperiodic polynomials*, Ph.D. thesis, Cornell Univ., 1989.

[GM] L. Goldberg and J. Milnor, *Fixed points of polynomial maps II: fixed point portraits*, Ann. Sci. École Norm. Sup. (4) 26 (1993), 51–98.

[K] J. Kiwi, $\mathbb{R}$*eal laminations and the topological dynamics of complex polynomials*, Adv. Math. 184 (2004), 207–267.

[Mi] J. Milnor, *Dynamics in One Complex Variable. Introductory Lectures*, Vieweg, 1999.

[Mu] J. Munkres, *Topology*, Prentice-Hall, 2000.

[P1] A. Poirier, *On postcritically finite polynomials*, Ph.D. thesis, SUNY Stony Brook, 1993.

[P2] A. Poirier, *On postcritically finite polynomials, Part one: Critical portraits*, Preprint IMS Stony Brook 5/1993.

Sección Matemáticas
Departamento de Ciencias
Pontificia Universidad Católica del Perú
Apartado 1761
Lima 100, Perú
E-mail: apoirie@pucp.edu.pe

Received 29 October 2007;
in revised form 11 February 2009