FUNDAMENTA
MATHEMATICAE
206 (2009)

Intertwined internal rays in Julia sets of rational maps

by

Robert L. Devaney (Boston, MA)

Abstract. We show how the well-known concept of external rays in polynomial dynamics may be extended throughout the Julia set of certain rational maps. These new types of rays, which we call internal rays, meet the Julia set in a Cantor set of points, and each of these rays crosses infinitely many other internal rays at many points. We then use this construction to show that there are infinitely many disjoint copies of the Mandelbrot set in the parameter planes for these maps.

External rays in both the dynamical and parameter planes provide an extremely important tool in complex dynamics. External rays arise as follows. For complex polynomials, there is always a basin of attraction of the fixed point at ∞. If the polynomial has degree n, it is well known that, in a neighborhood of ∞, the polynomial is analytically conjugate to the map $z \mapsto z^n$ also defined in a neighborhood of infinity (the so-called Böttcher coordinate). Now z^n maps straight rays of the form $t\exp(2\pi i\theta)$ for $t > 1$ to straight rays of the form $t^n \exp(2\pi i n\theta)$. So, under the conjugacy, the images of these straight rays in the basin of ∞ are also preserved by the polynomial. These images are the external rays for the polynomial. How these rays "land" on the Julia set often completely determines the structure of the Julia set for this map.

In this paper, we shall concentrate on rational maps, not polynomials. Specifically, we consider the family of maps

$$F_\lambda(z) = z^n + \frac{\lambda}{z^n}$$

where $\lambda \in \mathbb{C}$ and $n > 2$. We are interested in this family for two reasons. First, like polynomials, these maps always have a superattracting fixed point at ∞, so we have external rays for these maps. Second, like the family $z^n + c$, $n > 1$, there is essentially only one critical orbit for each member of this family, so the natural parameter plane is the λ-plane.

2010 *Mathematics Subject Classification*: Primary 37F10; Secondary 37F45.
Key words and phrases: internal ray, Julia set, Mandelbrot set, Sierpiński curve.

DOI: 10.4064/fm206-1-9

Our goal in this paper is to introduce a different type of ray into the dynamical planes of these maps, namely, *internal rays.* Basically, internal rays in the dynamical plane are curves that connect the endpoints of external rays to the origin (the only pole) by winding in a specific way through the Julia set. Unlike the external rays, however, the internal rays always cross other internal rays, usually at multiple points; hence they are "intertwined." In this paper, we will describe a special Cantor set of such internal rays. These rays are important in that how they wind through the Julia sets of these maps is the same no matter what the parameter λ is; for other internal rays, how they wind through the Julia set does depend upon the parameter. Moreover, how these special internal rays cross each other subdivides the Julia sets into regions where the structure is then easy to understand. As an example, we shall use the structure of the internal rays in the dynamical planes for these maps to prove the following result:

THEOREM. *In the parameter plane for $z^n + \lambda/z^n$, $n > 2$, there are countably many distinct copies of the Mandelbrot set, i.e., Mandelbrot sets that are not subsets of larger Mandelbrot sets. In particular, there are $(n-2)\cdot(2n-2)^{k-1}$ distinct baby Mandelbrot sets with base period k in the parameter plane for F_λ.*

In a subsequent paper [4] we shall describe an analogous set of internal rays in the parameter planes for these maps.

It is a pleasure to acknowledge the fundamental work of Michał Misiurewicz in dynamical systems over the years. Indeed, it was his pathbreaking paper on the dynamics of the complex exponential [10] that originally got me interested in holomorphic dynamics.

1. Preliminaries. Let

$$F_\lambda(z) = z^n + \lambda/z^n$$

where $n > 2$ and $\lambda \in \mathbb{C}$. The *Julia set* of F_λ, denoted by $J(F_\lambda)$, is defined to be the set of points at which the family of iterates of F_λ fails to be a normal family in the sense of Montel. There are many other equivalent definitions of the Julia set. For example, the Julia set is the closure of the set of repelling periodic points of F_λ, and, in our special case, it is also the boundary of the set of points whose orbits escape to ∞. As a consequence, the Julia set is the set of points on which F_λ behaves chaotically, since arbitrarily close to any point in $J(F_\lambda)$ there is both a repelling periodic point and a point whose orbit escapes to ∞. The complement of the Julia set is called the *Fatou set.*

One checks easily that F_λ has $2n$ free critical points given by $c_\lambda = \lambda^{1/2n}$. The origin and ∞ are also critical points but they are not free since ∞ is fixed and 0 is mapped directly to ∞. Despite the large number of free critical points for these maps, there are only two critical values: n of the critical points are mapped to $2\sqrt{\lambda}$ and the other n are mapped to $-2\sqrt{\lambda}$. In fact, there is really only one critical orbit, since, when n is even, both of the critical values are then mapped to the same point, while, when n is odd, the entire orbits of the critical values are symmetric under $z \mapsto -z$. Thus this family of maps, like the quadratic polynomial family, is a natural one-parameter family of maps. The parameter plane (the λ-plane) is then a record of the behavior of the free critical orbit, just as in the case of the Mandelbrot set. There are also $2n$ prepoles at the points $p_\lambda = (-\lambda)^{1/2n}$, so $F_\lambda(p_\lambda) = 0$.

Let C_λ be the circle given by $|z| = |\lambda|^{1/2n}$. Note that C_λ contains all of the critical points and prepoles of F_λ. A straightforward computation shows that F_λ maps C_λ $2n$-to-one onto the straight line connecting the two critical values. We call C_λ the *critical circle* and its image the *critical line*. One may also check that any other circle centered at the origin is mapped n-to-one onto an ellipse whose foci are $\pm v_\lambda$.

There are several symmetries for these maps. Let ν be the primitive $2n$th root of unity. Then we have $F_\lambda(\nu z) = \nu^n F_\lambda(z)$, so it follows easily that $J(F_\lambda)$ is invariant under $z \mapsto \nu z$. Also, let $H_\lambda(z)$ be one of the n involutions given by $H_\lambda(z) = \lambda^{1/n}/z$. Then we have $F_\lambda(H_\lambda(z)) = F_\lambda(z)$, so the Julia set is also preserved by each of these involutions.

The parameter plane for F_λ also possesses several symmetries. First of all, we have $\overline{F_\lambda(z)} = F_{\overline{\lambda}}(\overline{z})$ so that F_λ and $F_{\overline{\lambda}}$ are conjugate via the map $z \mapsto \overline{z}$. The parameter plane is also symmetric under $z \mapsto \omega z$ where ω is an $(n-1)$st root of unity, though the symmetrically located maps are not always conjugate (though their second iterates are conjugate). See [6] for details.

When $|z|$ is large, the term λ/z^n in the formula for F_λ is negligible, so $F_\lambda(z) \approx z^n$ near ∞. Consequently, the point at ∞ is a superattracting fixed point for F_λ, so we have an immediate basin of attraction B_λ at ∞. Since F_λ has a pole of order n at 0, there is an open neighborhood of 0 that is mapped n-to-one onto a neighborhood of ∞ in B_λ. If the immediate basin of ∞ is disjoint from this neighborhood around the origin, then there is an open set about 0 that is mapped n-to-one onto B_λ. This set is called the *trap door*, since any orbit that eventually enters B_λ must do so by passing through the trap door. We denote the trap door by T_λ.

As in the case of the quadratic polynomials $z^2 + c$, the orbits of the free critical points may tend to ∞. However, unlike the quadratic case, there are three distinct ways these critical orbits escape, and these lead to three different types of Julia sets for these maps. The following theorem was proved in [6].

Theorem (The escape trichotomy).

1. *If* v_λ *lies in* B_λ, *then* $J(F_\lambda)$ *is a Cantor set.*
2. *If* v_λ *lies in* $T_\lambda \neq B_\lambda$, *then* $J(F_\lambda)$ *is a Cantor set of concentric simple closed curves surrounding the origin.*
3. *In all other cases,* $J(F_\lambda)$ *is a connected set. In particular, if* $F_\lambda^j(v_\lambda) \in T_\lambda \neq B_\lambda$ *for some* $j \geq 1$, *then* $J(F_\lambda)$ *is a Sierpiński curve.*

We remark that case 2 of this theorem does not occur when $n = 2$; this is one of the reasons why we restrict n to be larger than 2. See Figure 1 for a picture of the parameter planes for the cases $n = 3$ and $n = 4$. In each case, the exterior region contains the parameters for which the Julia set is a Cantor set; this is the *Cantor set locus*. The small, central disk is the region where the Julia set is a Cantor set of circles; this is the *McMullen domain* as it was McMullen who first discovered this type of Julia set (see [8]). The complement of these two regions is the *connectedness locus*. The "holes" in this region are the sets where the Julia set is a Sierpiński curve; we call these regions *Sierpiński holes*. See [1], [2], and [5] for other ways that Sierpiński curve Julia sets arise in this family.

Note that there are two large Mandelbrot sets along the real axis in the parameter plane for $n = 3$ and three when $n = 4$. In general, there are $n - 1$ symmetrically located Mandelbrot sets in the parameter plane for $z^n + \lambda/z^n$. These are the so-called principal Mandelbrot sets; their existence was shown in [3]. In this paper we shall prove the existence of infinitely many other Mandelbrot sets some of which are visible in Figure 1.

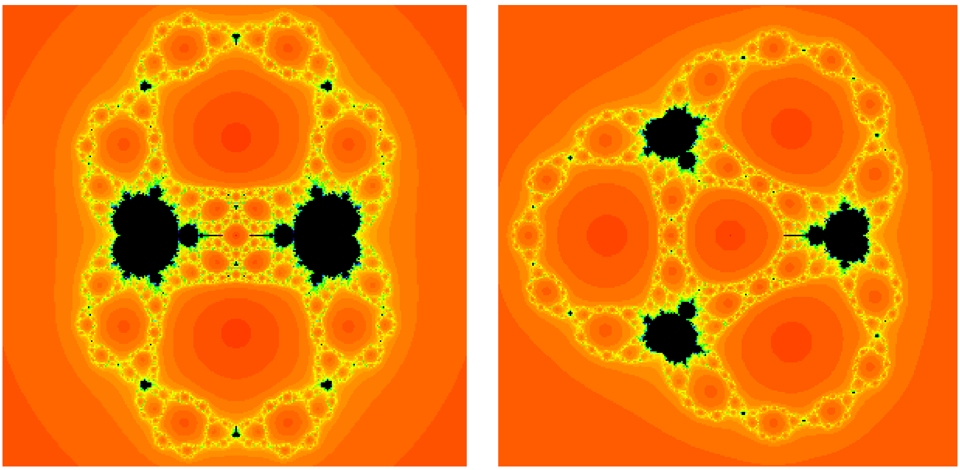

Fig. 1. The parameter planes when $n = 3$ and $n = 4$

Finally, since ∞ is superattracting, there is a Böttcher coordinate in a neighborhood of ∞. This is an analytic map ϕ_λ that conjugates F_λ to $z \mapsto z^n$

near ∞. It is known that

$$\phi_\lambda(z) = \lim_{k\to\infty} (F_\lambda^k(z))^{n^{-k}}.$$

Furthermore, ϕ_λ is unique up to multiplication by an $(n-1)$st root of unity and we may choose $\phi'_\lambda(\infty) = 1$. See [9], [11].

2. Internal rays in the dynamical plane. In this section, for simplicity, we restrict attention to the family of maps

$$F_\lambda(z) = z^3 + \lambda/z^3.$$

At the end of this paper we sketch the straightforward modifications needed to extend these results to the case where $n > 3$. Because of the $(n-1)$-fold symmetry in the parameter plane noted earlier, we may restrict attention to parameters λ that lie in the upper half-plane which we denote by $\mathcal{H}$. Since F_λ and $F_{-\lambda}$ are conjugate by the map $z \mapsto -z$, the construction below works equally well for parameters for which $\pi < \operatorname{Arg}\lambda < 2\pi$. The case where $\lambda \in \mathbb{R}$ is somewhat special but straightforward; this case will be discussed later.

Let $c_0 = c_0(\lambda) = \lambda^{1/6}$ denote the critical point that lies on the real axis when $\lambda \in \mathbb{R}^+$ and that varies analytically as λ ranges through $\mathcal{H}$. Let $c_j = c_j(\lambda) = c_0 \exp(2\pi i j/6)$, so the c_j are arranged in counterclockwise order around the origin as j increases. Three of the critical points (namely, c_0, c_2, and c_{-2}) are mapped to $v_\lambda = 2\sqrt{\lambda}$ which lies in the region $0 < \operatorname{Arg} v_\lambda < \pi/2$; the other three critical points are mapped to $-v_\lambda$, which lies in the negative of this region.

The straight line connecting the origin to ∞ and passing through c_j is called a *critical point ray* and is denoted by ζ_j. For $j = 0, 1, 2, 3$, let S_j be the closed sector bounded by the critcal point rays ζ_j and ζ_{j+1}. Define S_{-1} to be $-S_1$ and S_{-2} to be $-S_2$. So these sectors are arranged in counterclockwise order about the origin as $S_0, S_1, S_2, S_3, S_{-1}, S_{-2}$. The reason for this somewhat peculiar ordering will become clear when we define the internal rays for these maps. S_j is called a *prepole sector* since there is a unique prepole at the "center" of each S_j. One checks easily that the image of ζ_j under F_λ is a straight ray connecting one of the two critical values to ∞; we call this ray a *critical value ray.* Note that, since λ does not lie in $\mathbb{R}^+$, the image of each critical point ray lies in the interior of either S_0 or S_3 for each λ. As a consequence, it follows that F_λ maps each of S_1, S_2, S_{-1}, and S_{-2} univalently over a region that contains the union of these four sectors. The internal rays that we consider in this paper will always lie in one of these four sectors.

Let γ_λ denote the circle of radius $r = r(\lambda)$ centered at the origin and lying in B_λ. We choose r so that $r(\lambda)$ depends smoothly on λ and $F_\lambda(\gamma_\lambda)$

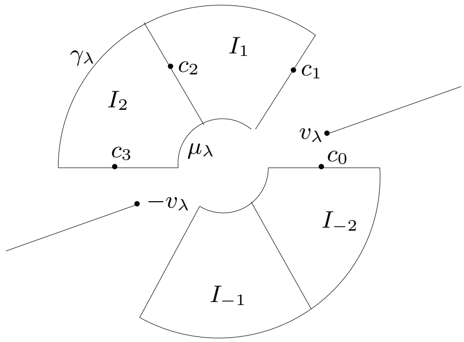

Fig. 2. The region $\mathcal{I}$

lies strictly outside γ_λ. Let μ_λ be the circle of radius $|\lambda|^{1/3}/r$; that is, $\mu_\lambda = H_\lambda(\gamma_\lambda)$. So $F_\lambda(\gamma_\lambda) = F_\lambda(\mu_\lambda)$ and F_λ also maps μ_λ strictly outside γ_λ.

Now consider the closed region in each S_j bounded on the outside by γ_λ and on the inside by μ_λ. Denote this region by $I_j = I_j(\lambda)$. Let $\mathcal{I} = \mathcal{I}(\lambda)$ denote the union of the four regions I_1, I_2, I_{-1}, and I_{-2}. The two regions I_0 and I_3 are not contained in $\mathcal{I}$. See Figure 2. Let Λ_λ denote the set of points whose orbits remain in $\mathcal{I}$ for all iterations. Then we have:

Proposition. *The set Λ_λ is homeomorphic to a Cantor set for each $\lambda \in \mathcal{H}$. Also, $F_\lambda|\Lambda_\lambda$ is conjugate to the one-sided shift on the four symbols $\pm 1, \pm 2$. The sets Λ_λ vary analytically with $\lambda \in \mathcal{H}$.*

Proof. For each $\lambda \in \mathcal{H}$, F_λ maps the boundary curves γ_λ and μ_λ strictly outside γ_λ, and hence outside $\mathcal{I}$. Also, F_λ maps the two critical point ray boundaries of each of the I_j in $\mathcal{I}$ to the two critical value rays, both of which lie in the interior of S_0 or S_3 for each $\lambda \in \mathcal{H}$ and hence also outside $\mathcal{I}$. Therefore it follows that F_λ maps each I_j in $\mathcal{I}$ univalently onto a region that completely covers each of the other I_k in $\mathcal{I}$. Standard arguments from complex dynamics then give that Λ_λ is a Cantor set with $F_\lambda|\Lambda_\lambda$ conjugate to the one-sided shift map on the four symbols $\pm 1, \pm 2$. Since the I_j vary analytically with λ, we have that the points in Λ_λ also vary analytically with λ. ∎

Remark. We emphasize that the set Λ_λ is only a subset of $J(F_\lambda)$; indeed, J itself may be connected.

Let Σ denote the space of one-sided sequences of the four symbols $\pm 1, \pm 2$, and let $\sigma : \Sigma \to \Sigma$ be the one-sided shift map. By the preceding proposition, each point in Λ_λ corresponds a unique itinerary $s = (s_0, s_1, s_2, \ldots) \in \Sigma$.

We shall next be concerned with a special subset Γ of Σ. The subset Γ consists of all sequences in Σ satisfying:

1. the digits -1 and 2 can only be followed by either 1 or 2;
2. the digits 1 and -2 can only be followed by either -1 or -2.

Clearly, Γ is a shift-invariant subset of Σ. Indeed, the shift map on Γ can be thought of as the subshift of finite type on the four symbols $\pm 1, \pm 2$ that obeys the above rules.

We can think of Γ in another way. Let τ denote the angle-tripling map on the unit circle, i.e., $\tau(\theta) = 3\theta \bmod 1$ where θ is defined mod 1. Consider the four arcs in the circle given by $J_1 = [1/4, 1/3], J_2 = [5/12, 1/2], J_{-1} = [3/4, 5/6]$, and $J_{-2} = [11/12, 1]$. Then it follows immediately that τ expands J_1 and J_{-2} over both J_{-1} and J_{-2}, while τ expands J_2 and J_{-1} over J_1 and J_2. Thus the set of points on the unit circle whose orbits under τ remain for all iterations in these four arcs is homeomorphic to the set Γ and, moreover, the map τ on this set is conjugate to the shift map on Γ. With a slight abuse of notation, we denote the set of all of these angles on the unit circle by Γ as well. We similarly denote the itinerary of the angle θ under τ by $s(\theta) = (s_0, s_1, s_2, \ldots) \in \Gamma$.

As a remark, we could equally well have defined Γ to be the set of angles that remain in the two arcs $[1/4, 1/2]$ and $[3/4, 1]$ under angle tripling mod 1. We prefer the above definition since it matches with the subshift definition given earlier.

As mentioned earlier, since ∞ is superattracting, we have a Böttcher coordinate $\phi_\lambda : B_\lambda \to \mathbb{C} - \overline{\mathbb{D}}$. This is an analytic map that conjugates F_λ in a neighborhood of ∞ to the map $z \mapsto z^3$ also near ∞. We may choose ϕ_λ so that $\phi'_\lambda(\infty) = 1$. It is known that, when λ lies in the connectedness locus or in the McMullen domain, ϕ_λ may be extended to a map that takes B_λ univalently onto the exterior of the closed unit disk in the Riemann sphere (see [9], [11]). When λ lies in the Cantor set locus, the domain of ϕ_λ is smaller due to the presence of critical points in B_λ.

Given an angle θ mod 1, the *external ray of angle* θ, denoted by $\xi_\lambda^\theta(t)$, is the image of the straight ray $t \mapsto t\exp(2\pi i\theta)$, $t \in (1, \infty]$, under the inverse map ϕ_λ^{-1}. We set $\xi_\lambda^\theta(\infty) = \infty$ for each $\theta \in \Gamma$. When λ lies in the connectedness locus or the McMullen domain, each $\xi_\lambda^\theta(t)$ is defined for all $t > 1$, and all of these external rays accumulate on the boundary of B_λ, though they may not limit on a unique landing point in ∂B_λ if ∂B_λ is not locally connected. However, the following proposition shows that we always do have a unique landing point whenever the angle θ lies in Γ; in particular, this also occurs in case λ lies in the Cantor set locus.

PROPOSITION. *Let $\theta \in \Gamma$ and $\lambda \in \mathcal{H}$. Then the external ray ξ_λ^θ may be extended so that it lands on (i.e., tends to a limit as $t \to 1$ at) the unique point in Λ_λ whose itinerary is the same as that of θ under the angle-tripling map τ.*

Proof. Recall that each of the prepole sectors S_1, S_2, S_{-1}, and S_{-2} is mapped univalently over the union of all four of these sectors by F_λ and that the exterior of the circle γ_λ of radius $r = r(\lambda)$ lies in B_λ. Given $\lambda \in \mathcal{H}$, let $N_r = \{z \mid |z| \geq r(\lambda)\}$ and let $N_r^j = N_r \cap S_j$. Since $\phi'_\lambda(\infty) = 1$, we may further assume that $r(\lambda)$ is chosen large enough so that, for each $\theta \in \Gamma$, the portion of the external ray ξ_λ^θ lying in N_r is precisely the set of points in N_r whose orbits have the exact same itinerary through the sets N_r^j as the angle θ has under τ, namely $s(\theta)$.

Now we can extend this concept to the entire collection of sectors S_j with $j = \pm 1, \pm 2$ by pulling back the portions of the external rays that lie in N_r by F_λ^{-1}. More precisely, take the portion of the external ray with itinerary $(s_1, s_2, s_3, \ldots)$ that lies in $N_r^{s_1}$ and pull it back under F_λ^{-1} into S_{s_0}. This extends the portion of ξ_λ^θ lying in $N_r^{s_0}$ further back into S_{s_0}. Continuing in this fashion with the portion of the ray with itinerary $(s_n, s_{n+1}, s_{n+2}, \ldots)$ in $N_r^{s_n}$ shows that the entire external ray $\xi_\lambda^\theta(t)$ for all $t > 1$ is defined and lies in S_{s_0}.

Now suppose that $\xi_\lambda^\theta(t)$ accumulates on some point $z \in \mathbb{C}$ as $t \to 1$. Then z cannot lie in B_λ. This follows immediately when λ does not lie in the Cantor set locus since ϕ_λ is defined and univalent on all of B_λ. If λ does lie in the Cantor set locus and z lies in B_λ, then we may iterate F_λ forward enough times so that $F_\lambda^k(z)$ lies in the region where the Böttcher coordinate is defined and analytic. But then the corresponding external ray would accumulate on a point in B_λ as $t \to \infty$ and that cannot happen. Thus $\xi_\lambda^\theta(t)$ must accumulate on a point in $J(F_\lambda)$. But this accumulation point must then have the same itinerary through the S_j as that of θ. However, we know that there is a unique such point in the Julia set, namely the point in Λ_λ with this itinerary. Thus each external ray ξ_λ^θ lands at the unique point in ∂B_λ whose itinerary is the same as that of θ whenever $\theta \in \Gamma$. ■

Given an itinerary $s = s(\theta) = (s_0, s_1, s_2, \ldots)$, we define the *negative* of s to be the itinerary $-s(\theta) = (-s_0, -s_1, -s_2, \ldots)$ and the *shift* of s to be $\sigma(s) = (s_1, s_2, s_3, \ldots)$. Also, let $\hat{\theta} = \theta + 1/2 \bmod 1$ and $\hat{\tau}(\theta) = \tau(\theta) + 1/2 \bmod 1$. Consequently, $s(\hat{\tau}(\theta)) = -\sigma(s(\theta))$. For example, if $\theta = 1/3$, then $s(1/3) = (1, \overline{-2})$ and $\hat{\tau}(1/3) = 1/2$. So we have $s(\hat{\tau}(1/3)) = s(1/2) = (\overline{2}) = -\sigma(s(1/3))$ as required.

Let $\mathbb{S}$ denote the union of the four sectors $S_{\pm 1}$ and $S_{\pm 2}$. Given a point z whose entire orbit lies in $\mathbb{S}$, we may define its itinerary as above to be $s(z) = (s_0, s_1, s_2, \ldots)$ where $s_j = k$ implies that $F_\lambda^j(z) \in S_k$. Note that, even though a pair of sectors overlap along the critical point rays, points on these rays are mapped immediately outside of $\mathbb{S}$. Consequently, the itinerary $s(z)$ is always well-defined when the entire orbit of z lies in $\mathbb{S}$.

Definition. Let $\theta \in \Gamma$ and suppose that $s(\theta) = (s_0, s_1, s_2, \ldots)$. The *full ray of angle* θ, denoted by ω_λ^θ, consists of all points in $\mathbb{S}$ whose itinerary is of the form $(s_0, \pm s_1, \pm s_2, \pm s_3, \ldots)$ together with the origin and ∞. That is, ω_λ^θ

consists of all points in S_{s_0} whose itinerary is the same as that of θ, except that we may change the sign of any of the s_j with $j > 0$.

Note that in this definition we always choose θ so that $s(\theta) \in \Gamma$. However, the points on the full ray with angle θ have itineraries that lie in the larger set Σ, not just Γ. Also, note that the external ray ξ_λ^θ is part of the full ray ω_λ^θ since all points on the external ray have itinerary $s(\theta)$. This implies that the set of points in Λ_λ whose itinerary corresponds to an angle in Γ lies in ∂B_λ when λ lies in the connectedness locus or the McMullen domain. This is the dynamical significance of the set Γ. We also denote the pair of full rays $\omega_\lambda^\theta \cup \omega_\lambda^{\hat\theta}$ by Ω_λ^θ (or, equally well, $\Omega_\lambda^{\hat\theta}$). Finally, this defintion explains why we choose the peculiar ordering of the S_j when j is negative; we only have to change the signs of the s_j to determine all points in the full ray.

DEFINITION. The *internal ray of angle* θ, denoted ν_λ^θ, is the complement of the external ray ξ_λ^θ in the full ray ω_λ^θ.

THEOREM. *Let* $\lambda \in \mathcal{H}$ *and* $\theta \in \Gamma$. *Then the full ray* ω_λ^θ *has the following properties:*

1. ω_λ^θ *meets* $J(F_\lambda)$ *in a Cantor set of points.*
2. ω_λ^θ *is mapped univalently onto the pair of full rays* $\Omega_\lambda^{\tau(\theta)}$.
3. ω_λ^θ *is a continuous curve extending from* 0 *to* ∞, *i.e., there is a continuous, one-to-one map taking the closed half-line* $[0,\infty]$ *onto* ω_λ^θ.

Proof. For part 1, consider the portion of ω_λ^θ that lies inside the bounded region $I_{s_0} \subset S_{s_0}$ and contains no points whose orbits tend to ∞, i.e., the entire orbit lies in $\mathcal{I}$. This portion of ω_λ^θ corresponds exactly to the set of points in Λ_λ whose itinerary in Σ is a sequence of the form $(s_0, \pm s_1, \pm s_2, \pm s_3, \ldots)$. As shown earlier, there is exactly one such point for each given sequence, and these points lie in $J(F_\lambda)$. It follows immediately that the points with these itineraries form a Cantor set in I_{s_0}.

For part 2, we see that ω_λ^θ can be divided into two disjoint subsets, those points with itinerary $(s_0, +s_1, \pm s_2, \pm s_3, \pm s_4, \ldots)$ and those with itinerary $(s_0, -s_1, \pm s_2, \pm s_3, \pm s_4, \ldots)$. These portions of ω_λ^θ are mapped one-to-one by F_λ onto $\omega_\lambda^{\tau(\theta)}$ and $\omega_\lambda^{\hat\tau(\theta)}$ respectively since $F_\lambda | S_{s_0}$ is univalent.

Finally, for part 3, we first note that there is a preimage of the external ray $\xi_\lambda^{\hat\tau(\theta)}$ that lies in S_{s_0} and connects a point in $J(F_\lambda)$ to the origin. This arc then lies in ω_λ^θ since its itinerary is $(s_0, -s_1, -s_2, -s_3 \ldots)$. If λ lies in the connectedness locus or the McMullen domain, then this arc lies in the trap door.

Now consider the pair of full rays ω_λ^θ and $\omega_\lambda^{\hat\theta}$. We first claim that this set of points is a closed, connected set. To see this, note that the set of points whose itinerary begins with either s_0 or $-s_0$ is a closed connected set in $\overline{\mathbb{C}}$

(here as elsewhere we include both the origin and ∞ in this set). Indeed, this set is just the two closed sectors S_{s_0} and S_{-s_0}. In the Riemann sphere, this set is a pair of closed "disks" that meet at two points, namely at 0 and ∞. Recall that each of these sectors is mapped univalently over $S_{s_1} \cup S_{-s_1}$. We therefore have a preimage of $S_{s_1} \cup S_{-s_1}$ in each of S_{s_0} and S_{-s_0}, and each of these preimages connects 0 to ∞. So the set of points whose itinerary begins $\pm s_0, \pm s_1$ is then a string of four closed disks that are contained in $S_{s_0} \cup S_{-s_0}$. Each of these disks meet exactly two of the other disks, one at either 0 or ∞ and the other at one of the two preimages of 0 that lie in the previous set. Then the preimage of this set is a string of eight closed disks, each of which meets exactly two others at 0, ∞, or their first or second preimages, and this string of disks is again contained in the previous set. Continuing inductively, we see that the set of points with this set of itineraries is a nested intersection of closed, connected sets in $\overline{\mathbb{C}}$, each of which is a string of 2^n closed disks, and each of these disks meets exactly two others at different points on its boundary. Hence the intersection is also a closed, connected set.

Now any point in ω_λ^θ that is not in $J(F_\lambda)$ must lie in the Fatou set. If such a point does not eventually tend to ∞, its orbit must be bounded. But then this point would correspond to some point in Λ_λ which then must lie in $J(F_\lambda)$. Therefore, the orbits of all other points in ω_λ^θ must eventually tend to ∞. This implies that they must, at some iteration, say the nth, land on an external ray in $\mathbb{S}$ whose itinerary is of the form $(\pm s_n, \pm s_{n+1}, \pm s_{n+2}, \ldots)$. But there are only two choices for such an external ray. To see this, recall that the itinerary of any external ray whose orbit lies in $\mathbb{S}$ corresponds to that of an angle in Γ. We know two possible such external angles, namely $\tau^n(\theta)$ and $\hat{\tau}^n(\theta)$, i.e., the itineraries $(s_n, s_{n+1}, s_{n+2}, \ldots)$ and $(-s_n, -s_{n+1}, -s_{n+2}, \ldots)$. Any other possible itinerary involving these digits would then change some digit following s_n or $-s_n$ in this sequence. Say s_{n+k} is changed where $k > 0$. Then, to keep this sequence in ω_λ^θ, we can only change s_{n+k} to $-s_{n+k}$. But, if s_{n+k} is allowed to follow s_{n+k-1}, then the rules governing the itineraries in Γ say that $-s_{n+k}$ cannot follow s_{n+k-1}.

As a consequence, ω_λ^θ consists of the external ray ξ_λ^θ, the preimage of $\xi_\lambda^{\hat{\tau}(\theta)}$ connecting to the origin, the Cantor set of points described earlier, and all points with the required itinerary that eventually land on the external rays with angle $\tau^n(\theta)$ or $\hat{\tau}^n(\theta)$, $n > 1$, together with 0 and ∞. But, for each $n > 1$, there are only a finite number of such preimages, in fact, exactly 2^{n-1} such arcs. Each of these curves then connects to a single point in the Cantor set portion of ω_λ^θ since we know that the original external rays of angle $\pm\theta$ have this property. The other endpoint of these curves is one of the 2^{n-2} nth preimages of ∞ in ω_λ^θ. Since, as we showed above, the union of all of the points in ω_λ^θ is closed and connected, it then follows that the entire set must be a closed curve that extends from 0 to ∞. Here the preimages of the two external rays with angle $\pm\tau^n(\theta)$ fill in the "gaps" in the Cantor set, since there are a pair of preimages of the external rays that connect to each preimage of 0 and also to a pair of points in $J(F_\lambda)$. This completes the proof. ∎

As a consequence of this result, we can identify any point on the internal ray ν_λ^θ uniquely as follows.

Theorem (Internal ray specification). *Let $z \in \nu_\lambda^\theta$. Then we may identify z uniquely by:*

1. *If $z \in \nu_\lambda^\theta \cap \Lambda_\lambda$, then we simply specify its itinerary in Σ.*
2. *If z lies in the complement of this set in ν_λ^θ, then there is a first iterate, say the kth where $k \geq 1$, at which the orbit of this point reaches B_λ. Then this point lies on the external ray with angle either $\tau^k(\theta)$ or $\hat{\tau}^k(\theta)$, so on some point of the form either $\xi_\lambda^{\tau^k(\theta)}(t)$ or $\xi_\lambda^{\hat{\tau}^k(\theta)}(t)$. So we may specify this point uniquely by giving θ, t, and which of the two external angles $\tau^k(\theta)$ or $\hat{\tau}^k(\theta)$ the orbit of z lands on.*

Remark. As mentioned above, the case where λ lies in the lower half-plane is entirely analogous, as the parameter plane is symmetric with respect to complex conjugation. The case where $\lambda \in \mathbb{R}$ is somewhat different. For example, if $\lambda \in \mathbb{R}^+$, the critical values have orbits that remain in $\mathbb{R}$. If this orbit is bounded, we can then extend the definition of the internal rays to include those in the sectors S_0 and S_3. The difference arises in that certain external rays now separate into a pair of internal rays when they enter the Julia set. For example, the external ray along $\mathbb{R}^+$ now joins up with the internal ray with itinerary $\overline{0}$ in S_0 and also with the internal ray $\overline{2}$ in S_2. Despite these differences, most of the above construction then proceeds in similar fashion, though the symbolic dynamics is now more complicated. We leave the details of this special case to the reader.

3. Intertwined internal rays. In this section, we show how the internal rays whose angle θ lies in Γ are always "intertwined." By intertwined, we mean that each of these internal rays crosses infinitely many other such internal rays. More precisely, we shall prove:

Theorem (Intertwined internal rays). *Let $\theta \in \Gamma$ and suppose the first digit of $s(\theta)$ is s_0. Then the internal ray of angle θ crosses every other internal ray whose angle lies in Γ and whose itinerary begins with s_0 at some interior point of S_{s_0}.*

Proof. For definiteness, we first consider the external rays that lie in the sector S_2. These external rays are mapped univalently onto the set of all external rays of angle $\tau(\theta) \in \Gamma$ lying in $S_1 \cup S_2$. Moreover, the external rays in S_2 are arranged in the exact same order in S_2 as θ is ordered in Γ. That is, if $\theta_1 < \theta_2$, then $\xi_\lambda^{\theta_1}$ lies above $\xi_\lambda^{\theta_2}$ in the sector S_2.

Now each external ray ξ_λ^θ in S_2 connects to a unique internal ray ν_λ^θ that also lies in S_2, and each of these internal rays contains an arc in the Fatou set that extends from the origin to a first point in $J(F_\lambda)$. Call this arc μ_λ^θ. Then μ_λ^θ is the preimage in S_2 of the external ray $\xi_\lambda^{\hat{\tau}(\theta)}$. The collection of arcs μ_λ^θ in S_2 is just the collection of all preimages of external rays that lie in $S_{-1} \cup S_{-2}$ (with angles in Γ).

We claim that the arcs μ_λ^θ are arranged around the origin in S_0 in the exact opposite order of the ξ_λ^θ as θ increases. This occurs since, close to the origin, $F_\lambda \approx \lambda/z^3$, so F_λ reverses the ordering of rays near the origin as they are mapped to the external rays near ∞. That is, if $\theta_1 < \theta_2$, then $\mu_\lambda^{\theta_2}$ lies above $\mu_\lambda^{\theta_1}$ in the natural ordering of these rays around the origin in S_2. Since ξ_λ^θ must connect to μ_λ^θ for each θ, it then follows that any two internal rays of different angles must cross at some point in S_2 that is not the origin nor ∞.

For rays lying in the other sectors similar arguments work as well. ∎

As a specific example of intertwining, let $\theta = 1/2$. Then $s(1/2) = (\overline{2})$, and $\xi_\lambda^{1/2}$ is the fixed external ray lying in S_2. Similarly, when $\theta = 0$, ξ_λ^0 lies in S_{-2} and has itinerary $(\overline{-2})$. So $\omega_\lambda^{1/2}$ is mapped over the pair of full rays $\omega_\lambda^{1/2}$ and ω_λ^0, which we have named $\Omega_\lambda^{1/2}$ (or Ω_λ^0). Similarly, $\theta = 1/4$ has itinerary $s(1/4) = (\overline{1, -1})$ so $\xi_\lambda^{1/4}$ lies in S_1 and is mapped to $\xi_\lambda^{3/4}$ in S_{-1}, while $\xi_\lambda^{3/4}$ is mapped back to $\xi_\lambda^{1/4}$. So $\omega_\lambda^{1/4}$ lies in S_1 and is mapped over $\Omega_\lambda^{1/4}$, while $\omega_\lambda^{3/4}$ lies in S_{-1} and is also mapped onto the same pair of full rays. Note that $\Omega_\lambda^{1/2}$ is a simple closed curve in the Riemann sphere that passes through the origin and ∞ and lies in $S_2 \cup S_{-2}$, while $\Omega_\lambda^{1/4}$ is a similar simple closed curve that lies in $S_1 \cup S_{-1}$. Hence these curves only meet at 0 and ∞. Also note that $F_\lambda(\Omega_\lambda^{1/2}) = \Omega_\lambda^{1/2}$ and $F_\lambda(\Omega_\lambda^{1/4}) = \Omega_\lambda^{1/4}$; indeed, these are the only pairs of full rays that are mapped over themselves by F_λ.

Now consider the preimages of $\Omega_\lambda^{1/2}$ and $\Omega_\lambda^{1/4}$ in $\mathbb{S}$. One computes that $F_\lambda(\Omega_\lambda^{1/3}) = \Omega_\lambda^0 = \Omega_\lambda^{1/2}$ and $F_\lambda(\Omega_\lambda^{5/12}) = \Omega_\lambda^{1/4}$. Each of the two curves $\Omega_\lambda^{1/3}$ and $\Omega_\lambda^{5/12}$ must pass through the two preimages of 0 in the pair of sectors in which they reside, so $\Omega_\lambda^{5/12}$ meets $\Omega_\lambda^{1/2}$ in a total of four points, the origin, ∞, and the two preimages of 0. Similarly, $\Omega_\lambda^{1/3}$ meets $\Omega_\lambda^{1/4}$ in four points. See Figure 3. In similar fashion, one computes that the preimage of $\Omega_\lambda^{5/12}$ in $S_2 \cup S_{-2}$, i.e., $\Omega_\lambda^{17/36}$, meets $\Omega_\lambda^{1/2}$ in eight points, the previous four points plus four of the second preimages of 0 lying in this region. Similarly, the preimages of $\omega_\lambda^{5/12}$ in $S_1 \cup S_{-1}$, namely $\Omega_\lambda^{11/36}$, meets $\Omega_\lambda^{1/3}$ in eight points.

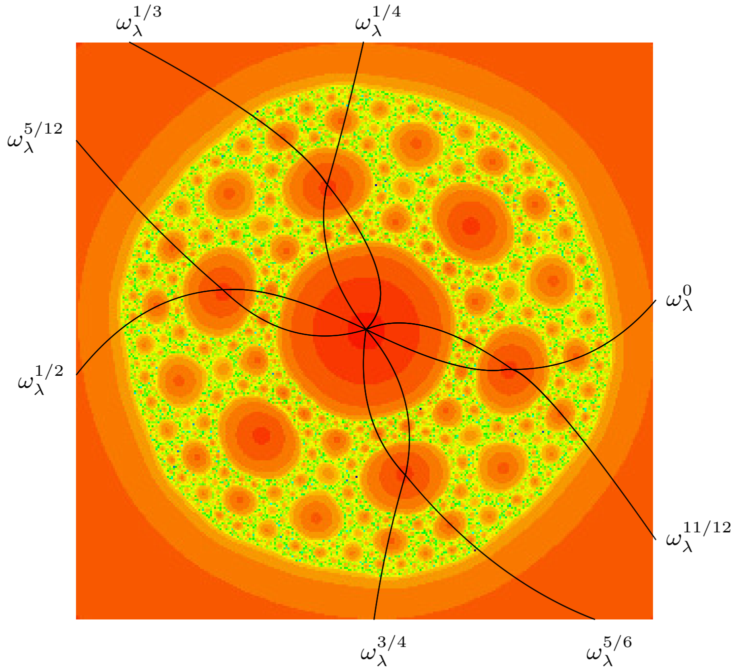

Fig. 3. Some of the rays in the dynamical plane for a fixed $\lambda \in \mathcal{H}$

Continuing in this fashion, we see that any full ray that is a jth preimage of $\Omega_\lambda^{1/2}$ or $\omega_\lambda^{1/4}$ meets these curves in exactly 2^j points.

In particular, note that every internal ray whose itinerary begins with the digit s_0 must pass through the preimage of 0 that lies in the sector S_{s_0}. Indeed, infinitely many internal rays whose itinerary is allowable cross at each kth preimage of 0.

4. A model for the internal rays. In this section we construct a piecewise linear model that exhibits the structure of the internal rays described in the previous two sections. The model for this collection of rays will be drawn in the Sierpiński carpet. This is not intended to say that the Julia sets through which we have drawn the internal rays are always Sierpiński curves; they sometimes are but they need not be such sets. We use the carpet only to facilitate the drawing of the model for these internal rays.

Recall that the Sierpiński carpet consists of infinitely many smaller copies of itself. At stage one, we have the entire carpet. At stage 2, we may break the carpet up into eight copies of itself, each of which is one-third the size of the original. At stage 3, there are 8^2 copies of the carpet that are $1/3^2$ the size of the original, and so forth. We will draw three types of internal straight rays in certain of these self-similar portions of the carpet: horizontal rays, diagonal rays, and connecting rays. By a horizontal ray, we always mean a straight line passing left to right through the center of a Sierpiński carpet, and by a diagonal ray, we always mean a straight line connecting the upper left corner of a carpet to the lower right corner.

We proceed inductively. At stage one, we draw a horizontal and diagonal ray through the entire carpet. The horizontal ray corresponds to the internal rays $\nu_\lambda^0 \cup \nu_\lambda^{1/2}$ while the diagonal ray corresponds to $\nu_\lambda^{1/4} \cup \nu_\lambda^{3/4}$. Note that these rays meet only at the center of the carpet.

Note that these two rays cross through four of the eight copies of the carpet at stage 2. In each of these four smaller carpets, we either have a smaller diagonal or a smaller horizontal ray. So, at this stage, we add a new ray to this carpet of the opposite type; that is, if there is a diagonal ray in this stage 2 carpet, we add a horizontal ray. So we have added four portions of rays so far. Now what we do is join up these portions of the rays with connecting rays which are straight lines passing through the center of the original carpet and joining the endpoints of a pair of rays just constructed. That is, we join up the two new horizontal rays with a connecting ray and the two new diagonal rays as well. See Figure 4. This produces two new rays that correspond to the internal rays $\nu_\lambda^{1/3}$, $\nu_\lambda^{5/12}$, $\nu_\lambda^{5/6}$, and $\nu_\lambda^{11/12}$.

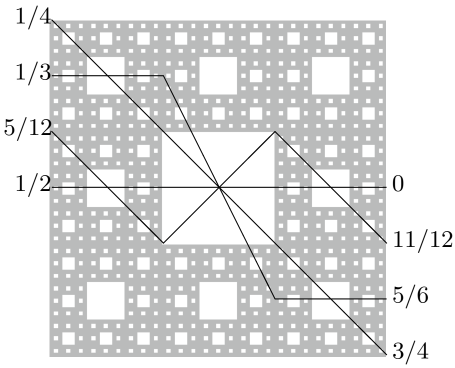

Fig. 4. The first two stages of the construction of the piecewise linear model for the intertwined rays. The rays that correspond to $\nu_\lambda^{p/q}$ in Figure 3 are indicated by p/q.

Now we proceed inductively. At stage 3, our previously constructed (in stage 1 and stage 2) diagonal and horizontal rays (not the connecting rays) cross through exactly 4^2 copies of the carpet at stage 3, so we again adjoin a ray of the opposite type in each of these carpets. Then we connect up these rays as follows. First, inside the smaller stage 2 carpets, we pass a connecting line through the center of this smaller carpet and join the ends of symmetrically located rays. Then, in the original carpet, we connect up the endpoints with connecting rays passing through the center of the largest

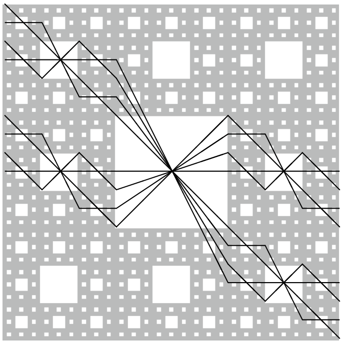

Fig. 5. The first three stages of the construction of the piecewise linear model for the intertwined rays.

carpet. See Figure 5. Taking the limit of this process yields an uncountable collection of curves in the carpet. Given a vertical line through the carpet, this collection of curves meets this line in a Cantor set of points provided that the vertical line does not pass through the center of any of the complementary squares. On the other hand, this collection of curves meets the vertical lines that pass through the centers of the complementary squares in only finitely many points, indeed, in exactly 2^{k-1} points if the complementary square is at the center of a stage k carpet.

One checks easily that, using the mapping properties of F_λ on the various sectors S_j, this construction corresponds exactly to that of the internal rays in the dynamical plane.

5. Baby Mandelbrot sets. Our goal in this section is to prove the existence of infinitely many disjoint copies of the Mandelbrot set in the parameter plane for $z^3 + \lambda/z^3$. We continue to assume that λ lies in $\mathcal{H}$, the upper half-plane, at least initially. Let $\mathcal{R}_\lambda$ denote the collection of all full rays whose angle θ lies in Γ. The set $\mathcal{R}_\lambda$ is a closed subset of the Riemann sphere whose complement consists of a countable collection of open disks (recall that 0 and ∞ lie on all of the full rays). One of these complementary disks contains the critical value v_λ that lies in the upper half-plane; call this disk U_0. Note that U_0 contains the entire sector S_0 as well as portions of both S_1 and S_{-2}. Another complementary disk is $U_3 = -U_0$. This disk contains the entire sector S_3 and also pieces of S_2 and S_{-1}. Note that U_0 contains the critical points c_0 and c_1 that lie on the boundary of S_0, while U_3 contains the critical points c_3 and c_{-2} lying on the boundary of S_3. See Figure 6.

There are two other open sets in the complement of $\mathcal{R}_\lambda$ that contain critical points. Let U_2 be the open disk containing c_2, and U_{-1} the open disk containing c_{-1}. Then U_2 is the region that separates the full rays in S_2 from those in S_1, while U_{-1} separates the full rays in S_{-1} from those in S_{-2}.

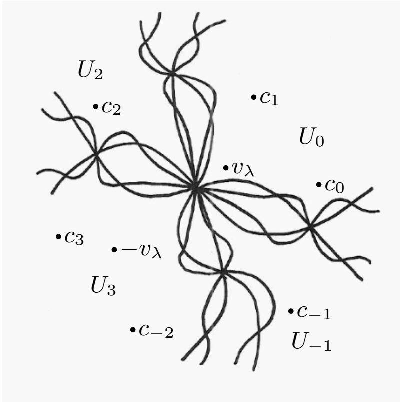

Fig. 6. Some of the rays in $\mathcal{R}_\lambda$ and the U_j.

Let $\mathcal{X}_\lambda$ denote the complement of the two disks U_0 and U_3. Then $\mathcal{X}_\lambda$ is a closed subset of the Riemann sphere that contains $\mathcal{R}_\lambda$, U_2, U_{-1}, and countably many other open disks.

We next discuss how F_λ acts on the complementary domains in $\mathcal{R}_\lambda$. First, $\mathcal{R}_\lambda$ is F_λ-invariant and F_λ maps this set of full rays four-to-one onto itself. It follows that the open disk U_2 is mapped two-to-one onto U_0 since $F_\lambda(c_2) = v_\lambda$. Similarly, U_{-1} is mapped two-to-one onto U_3. These are the only open disks in $\mathcal{X}_\lambda$ that are mapped outside $\mathcal{X}_\lambda$.

There must be two other preimages of $\mathcal{X}_\lambda$, and so one of these preimages lies in U_0; call this preimage $\mathcal{X}_\lambda^0$. The other preimage lies in U_3 and is called $\mathcal{X}_\lambda^3 = -\mathcal{X}_\lambda^0$. Note that each of these sets is mapped one-to-one onto $\mathcal{X}_\lambda$. Thus we have all of the preimages of $\mathcal{X}_\lambda$ and all we need are the four remaining preimages of U_0 and U_3. These necessarily lie in U_0 and U_3.

Since v_λ lies in U_0, it follows that there is an open set in $U_0 - \mathcal{X}_\lambda^0$ that contains c_0 and is mapped two-to-one onto U_0 and another open set that contains c_1 and is mapped two-to-one onto U_3. Similarly, there are a pair of open sets in $U_3 - \mathcal{X}_\lambda^3$ each of which contains a critical point and is mapped two-to-one onto either U_0 or U_3. To summarize all of this, we have shown:

Proposition.

1. *F_λ maps U_2 two-to-one onto U_0.*
2. *F_λ maps U_{-1} two-to-one onto U_3.*
3. *F_λ maps U_0 (and U_3) two-to-one onto $U_0 \cup U_3$ and one-to-one onto $\mathcal{X}_\lambda$.*
4. *F_λ maps $\mathcal{X}_\lambda$ two-to-one onto both U_3 and U_0 and four-to-one onto itself.*

In order to prove the existence of baby Mandelbrot sets for this family, we first recall the theory of polynomial-like maps. Suppose $V' \subset V$ are a pair of bounded, open, simply connected subsets of $\mathbb{C}$ with V' relatively compact in V. A map $G : V' \to V$ is called a *polynomial-like* map of degree two if G is holomorphic and proper of degree two. Hence such a map has a unique critical point $c \in V'$. The filled Julia set of G is defined in the natural manner as the set of points whose orbits never leave the subset V' under iteration of G. By the results in [7], it is known that G is topologically conjugate to some quadratic polynomial in a neighborhood of the polynomial's filled Julia set in $\mathbb{C}$, hence the name polynomial-like.

Now suppose that we have a family of polynomial-like maps $G_\mu : V'_\mu \to V_\mu$ depending on a parameter $\mu \in \mathbb{C}$ and satisfying:

1. The parameter μ lies in an open set in $\mathbb{C}$ that contains a closed disk W, and the boundaries of V'_μ and V_μ vary analytically as μ varies.
2. The map $(\mu, z) \mapsto G_\mu(z)$ depends holomorphically on both μ and z.
3. Each $G_\mu : V'_\mu \to V_\mu$ is polynomial-like of degree two.

Then we may consider the set of parameters in W for which the orbit of the critical point, c_μ, does not escape from V'_μ and so the corresponding filled Julia set is connected. Suppose that for each μ in the boundary of W we have that $G_\mu(c_\mu)$ lies in $V_\mu - V'_\mu$ and that, moreover, $G_\mu(c_\mu) - c_\mu$ winds once around 0 as μ winds once around the boundary of W. Then, in this case, it is also shown in [7] that the set of μ-values for which the orbit of c_μ does not escape from V'_μ is homeomorphic to the Mandelbrot set and that the polynomial to which G_μ corresponds under this homeomorphism is conjugate to G_μ on some neighborhood of its Julia set. This result thus gives a criterion for proving the existence of small copies of a Mandelbrot set inside $\mathcal{H}$.

As a warm-up to how we will proceed to prove the existence of baby Mandelbrot sets, consider the following scenario. Recall that the region $\mathcal{X}^0_\lambda$ is mapped univalently over $\mathcal{X}_\lambda$. Hence there is an open disk $V'_\lambda \subset \mathcal{X}^0_\lambda$ that is mapped univalently onto the region U_2. But then U_2 is mapped two-to-one onto $U_0 \supset V'_\lambda$. So it appears that F^2_λ is a polynomial-like map taking V'_λ onto the larger set U_0. In addition, let the subset of the parameter plane we are considering be the entire region $\mathcal{H}$. Then the critical orbit for $F^2_\lambda|V'_\lambda$ contains c_2 and lands, at the second iteration, on v_λ. As λ winds once around $\mathcal{H}$, the critical value then winds once around the sector given by $0 \le \operatorname{Arg} z \le \pi/2$.

There are two problems with this argument. First, the disk V'_λ actually extends all the way to ∞, so V'_λ is not properly contained in $F^2_\lambda(V'_\lambda)$. Second, as we wind around the boundary of $\mathcal{H}$, the critical value is no longer defined for all λ.

To remedy these defects, we first restrict attention to the subset $\mathcal{S}$ of $\mathcal{H}$ given by $\mathcal{S} = \{\lambda \in \mathbb{C} \mid |\lambda| < 2 \text{ and } \operatorname{Im} \lambda > 0\}$. Secondly, we constrain the open disks U_j to lie in a region that is bounded away from both 0 and ∞. We accomplish this as follows. Choose a level curve β_0 of the Green's function lying in B_λ and surrounding $J(F_\lambda)$. Then F_λ maps β_0 strictly outside itself.

Let β_1 be the curve in T_λ that is mapped onto $F_\lambda(\beta_0)$. Let A_λ denote the annulus bounded by β_0 and β_1. So we see that, if $\lambda \in \mathcal{S}$, F_λ maps both boundary curves of A_λ strictly outside the β_0. Furthermore, all of the critical points of F_λ lie in A_λ.

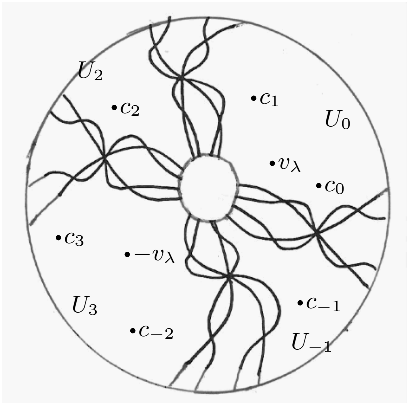

Fig. 7. The regions $\mathcal{R}_\lambda$ and the U_j in A_λ

We now consider the restriction of the sets $\mathcal{R}_\lambda$, $\mathcal{X}_\lambda$, and the U_j to the annulus A_λ. With a slight abuse of notation, we continue to denote these sets by $\mathcal{R}_\lambda$, $\mathcal{X}_\lambda$, and U_j. Since F_λ takes $\mathcal{X}_\lambda$ four-to-one over itself (and both critical points in $\mathcal{X}_\lambda$ are mapped outside of $\mathcal{X}_\lambda$), it follows that there are four preimages of U_2 in $\mathcal{X}_\lambda$. There are then 16 open disks lying in $\mathcal{X}_\lambda$ whose first images under F_λ also lie in $\mathcal{X}_\lambda$ and whose image under F_λ^2 is U_2. Continuing, there are 4^{k-2} open disks in $\mathcal{X}_\lambda$ whose images under F_λ^j, $j = 0, \dots, k-2$, lie in $\mathcal{X}_\lambda$ and such that F_λ^{k-2} maps each of these sets univalently onto U_2. Hence F_λ^{k-1} maps each of these sets in two-to-one fashion onto a set that contains U_0 and so the image of these sets completely contains the closure of the set $\mathcal{X}_\lambda^0 \subset U_0$. Moreover, the critical value for F_λ^{k-1} on these sets is v_λ.

Now F_λ takes $\mathcal{X}_\lambda^0$ univalently over a region containing $\mathcal{X}_\lambda$. So there are 4^{k-2} open disks in $\mathcal{X}_\lambda^0$ each of which is mapped univalently onto one of the 4^{k-2} open disks in $\mathcal{X}_\lambda$ that are mapped by F_λ^{k-2} over U_2. Hence F_λ^k maps each of these open disks in two-to-one fashion onto a disk that properly contains each of them. Thus F_λ^k is a polynomial-like map of degree two on each of these subsets. If we then let λ travel around the boundary of $\mathcal{S}$, then the critical value of F_λ^k on each of these sets, namely v_λ, travels once around the quarter of a disk bounded by $|z| = 2\sqrt{2}$, $0 \le \operatorname{Arg}\lambda \le \pi/2$, together with the intervals in the positive real and imaginary axes connecting the origin and $2\sqrt{2}$. So the critical value winds once around each of the preimages in $\mathcal{X}_\lambda^0$. This shows that there are 4^{k-2} Mandelbrot sets with base period k

(i.e., the period of the cycle corresponding to parameters drawn from the main cardioid of these Mandelbrot sets is k).

REMARKS. 1. Technically, we should not move the parameter around the boundary of $\mathcal{S}$ since, when we pass through the origin, the critical value $2\sqrt{\lambda}$ no longer varies analytically. However, if we take a slight diversion along a semi-circle surrounding the origin and lying in the McMullen domain, the previous argument then works as well. This proves the existence of infinitely many copies of the Mandelbrot set in $\mathcal{S}$.

2. For parameters drawn from these Mandelbrot sets, there are actually two attracting cycles with base period k. This follows since $F_\lambda(-z) = -F_\lambda(z)$. In the case above, there is only one preimage of U_2 lying in $\mathcal{X}_\lambda^0$ and none in $\mathcal{X}_\lambda^3$. So the negatives of these preimages give another polynomial-like map with exactly one preimage in $\mathcal{X}_\lambda^3$ and none in $\mathcal{X}_\lambda^0$.

We call the Mandelbrot sets constructed above *Type 1* Mandelbrot sets, for there is another way Mandelbrot sets arise in the parameter plane, and parameters from these Mandelbrot sets have somewhat different dynamical behavior. Consider the 4^{k-2} preimages of U_2 described above and lying in $\mathcal{X}$. There are the same number of preimages of these sets in $\mathcal{X}_\lambda^3$, not $\mathcal{X}_\lambda^0$ as was considered above. But then each of these sets is mapped by F_λ^k two-to-one over the exact same region containing U_0 above. So, for definiteness, call one of these preimages V. Then F_λ^k maps V two-to-one over a region that properly contains the set $-V$, not V. Now consider the map $-F_\lambda^k$ on V. This map is now a polynomial-like map of degree two on V, and, as above, in the parameter plane corresponding to the map $-F_\lambda^k$ there is again a copy of the Mandelbrot set. But the dynamics of the second iterate of $-F_\lambda^k$ and the map of F_λ^{2k} are the same, so there is also a copy of a Mandelbrot set for our family corresponding to the same parameters. The difference here is that maps drawn from the main cardioid of this set now have an attracting cycle of period $2k$. Also, there are two critical points lying in the immediate basin of this cycle, not 1, and the cycle is symmetric under $z \mapsto -z$. We say that these Mandelbrot sets also have base period k. This produces an additional collection of 4^{k-2} baby Mandelbrot sets. We say that these Mandelbrot sets also have base period k, and we call them *Type 2* baby Mandelbrot sets. We get a similar number of Type 1 and 2 Mandelbrot sets in the lower half of the parameter plane by the complex conjugation symmetry, so this gives a total of 4^{k-1} Type 1 and Type 2 baby Mandelbrot sets in the parameter plane. Therefore we have proved:

THEOREM. *Let $k \geq 2$. For the family of maps $F_\lambda(z) = z^3 + \lambda/z^3$, there are $2 \cdot 4^{k-2}$ Type 1 baby Mandelbrot sets with base period k in the parameter plane for F_λ and the same number of Type 2 baby Mandelbrot sets with base period k.*

6. Final remarks. In this paper we have concentrated on the family of maps $z^3 + \lambda/z^3$. But all of the above results (except for the piecewise linear model) go through in essentially the same manner for the families $z^n + \lambda/z^n$ with $n > 3$. The difference here is that we now work with $2n$ sectors given by $S_0, S_1, \ldots, S_n$ together with $S_{-j} = -S_j$ for $j = 1, \ldots, n-1$. The parameter plane now has $(n-1)$-fold symmetry, so we restrict attention to parameters that satisfy $0 < \operatorname{Arg}\lambda < 2\pi/(n-1)$. Then, as before, the two critical values $\pm 2\sqrt{\lambda}$ lie in S_0 and S_n. We then get a Cantor set of points Λ_λ lying in the Julia set on which F_λ is conjugate to the shift map on the $2n-2$ symbols $\pm 1, \ldots, \pm(n-1)$. Let Σ now denote this sequence space.

Then we can define the relevant subset $\Gamma \subset \Sigma$ in two ways. First, Γ consists of all points on the unit circle whose angles mod 1 always lie in the pair of arcs given by $[1/(n+1), 1/2]$ and $[n/(n+1), 1]$. As a subshift of Σ, Γ can also be defined as all sequences of symbols $\pm 1, \pm 2, \ldots \pm (n-1)$ where

1. odd positive integers are only followed by negative integers;
2. even positive integers are only followed by positive integers;
3. odd negative integers are only followed by positive integers when n is odd, but, if n is even, odd negative integers are only followed by negative integers;
4. even negative integers are only followed by negative integers when n is odd, but, if n is even, even negative integers are only followed by positive integers.

Then the definition of the full and internal rays goes through exactly as above.

We can also use the same techniques to produce copies of baby Mandelbrot sets as in the previous section. One difference here is that, when n is even, we no longer have the $z \mapsto -z$ symmetry that creates Type 1 and 2 Mandelbrot sets of base period k; instead, since all of the critical points map to the same point at iteration two, the above techniques always produce attracting cycles with period k, never $2k$. In any event, a similar count yields the existence of $(n-2)(2n-2)^{k-1}$ baby Mandelbrot sets of base period k.

References

[1] R. L. Devaney, *A myriad of Sierpiński curve Julia sets*, in: Difference Equations, Special Functions and Orthogonal Polynomials, World Sci., 2007, 131–148.

[2] —, *The McMullen domain: Satellite Mandelbrot sets and Sierpiński holes*, Conformal Geom. Dynam. 11 (2007), 164–190.

[3] —, *Baby Mandelbrot sets adorned with halos in families of rational maps*, in: Complex Dynamics; Twenty-Five Years After the Appearance of the Mandelbrot Set, Contemp. Math. 396, Amer. Math. Soc., 2006, 37–50.

[4] R. L. Devaney, *Intertwined internal parameter plane rays for rational maps*, to appear.
[5] —, *Cantor necklaces and structurally unstable Sierpiński curve Julia sets for rational maps*, Qual. Theory Dynam. Systems 5 (2004), 337–359.
[6] R. L. Devaney, D. M. Look and D. Uminsky, *The escape trichotomy for singularly perturbed rational maps*, Indiana Univ. Math. J. 54 (2005), 1621–1634.
[7] A. Douady and J. Hubbard, *On the dynamics of polynomial-like mappings*, Ann. Sci. École Norm. Sup. (4) 18 (1985), 287–343.
[8] C. McMullen, *Automorphisms of rational maps*, in: Holomorphic Functions and Moduli, Vol. 1, Math. Sci. Res. Inst. Publ. 10, Springer, New York, 1988, 31–60.
[9] J. Milnor, *Dynamics in One Complex Variable*, Vieweg, 1999.
[10] M. Misiurewicz, *On iterates of* e^z, Ergodic Theory Dynam. Systems 1 (1981), 103–106.
[11] C. Petersen and G. Ryd, *Convergence of rational rays in parameter spaces*, in: The Mandelbrot Set: Theme and Variations, London Math. Soc. Lecture Note Ser. 274, Cambridge Univ. Press, 2000, 161–172.

Robert L. Devaney
Department of Mathematics
Boston University
Boston, MA 02215, U.S.A.
E-mail: bob@bu.edu

Received 22 December 2008;
in revised form 7 July 2009