

Astérisque

JOHN MILNOR

Periodic orbits, external rays and the Mandelbrot set: an expository account

Astérisque, tome 261 (2000), p. 277-333

<http://www.numdam.org/item?id=AST_2000__261__277_0>

© Société mathématique de France, 2000, tous droits réservés.

L'accès aux archives de la collection « Astérisque » (<http://smf4.emath.fr/Publications/Asterisque/>) implique l'accord avec les conditions générales d'utilisation (<http://www.numdam.org/conditions>). Toute utilisation commerciale ou impression systématique est constitutive d'une infraction pénale. Toute copie ou impression de ce fichier doit contenir la présente mention de copyright.

NUMDAM

Article numérisé dans le cadre du programme
Numérisation de documents anciens mathématiques
<http://www.numdam.org/>

PERIODIC ORBITS, EXTERNAL RAYS AND THE MANDELBROT SET: AN EXPOSITORY ACCOUNT

by

John Milnor

Dedicated to Adrien Douady on the occasion of his sixtieth birthday

Abstract. — A presentation of some fundamental results from the Douady-Hubbard theory of the Mandelbrot set, based on the idea of “orbit portrait”: the pattern of external rays landing on a periodic orbit for a quadratic polynomial map.

1. Introduction




FIGURE 1. Julia set for $z \mapsto z^2 + (\frac{1}{4}e^{2\pi i/3} - 1)$ showing the six rays landing on a period two parabolic orbit. The associated orbit portrait has characteristic arc $\mathcal{I} = (22/63, 25/63)$ and valence $v = 3$ rays per orbit point.

A key point in Douady and Hubbard’s study of the Mandelbrot set M is the theorem that every parabolic point $c \neq 1/4$ in M is the landing point for exactly two external rays with angles which are periodic under doubling. (See [DH2]. By

1991 Mathematics Subject Classification. — 30D05.

Key words and phrases. — Mandelbrot set, periodic orbit.

definition, a parameter point is *parabolic* if and only if the corresponding quadratic map has a periodic orbit with some root of unity as multiplier.) This note will try to provide a proof of this result and some of its consequences which relies as much as possible on elementary combinatorics, rather than on more difficult analysis. It was inspired by §2 of the recent thesis of Schleicher [S1], which contains very substantial simplifications of the Douady-Hubbard proofs with a much more compact argument, and is highly recommended. (See also [S2], [LS].) The proofs given here are rather different from those of Schleicher, and are based on a combinatorial study of the angles of external rays for the Julia set which land on periodic orbits. (Compare [A], [GM].) As in [DH1], the basic idea is to find properties of M by a careful study of the dynamics for parameter values outside of M . The results in this paper are mostly well known; there is a particularly strong overlap with [DH2]. The only claim to originality is in emphasis, and the organization of the proofs. (Similar methods can be used for higher degree polynomials with only one critical point. Compare [S3], [E], and see [PR] for a different approach. For a theory of polynomial maps which may have many critical points, see [K].)

We will assume some familiarity with the classical Fatou-Julia theory, as described for example in [Be], [CG], [St], or [M2].

Standard Definitions. — (Compare Appendix A.) Let $K = K(f_c)$ be the *filled Julia set*, that is the union of all bounded orbits, for the quadratic map

$$f(z) = f_c(z) = z^2 + c.$$

Here both the parameter c and the dynamic variable z range over the complex numbers. The *Mandelbrot set* M can be defined as the compact subset of the *parameter plane* (or c -plane) consisting of all complex numbers c for which $K(f_c)$ is connected. We can also identify the complex number c with one particular point in the *dynamic plane* (or z -plane), namely the *critical value* $f_c(0) = c$ for the map f_c . The parameter c belongs to M if and only if the orbit $f_c : 0 \mapsto c \mapsto c^2 + c \mapsto \dots$ is bounded, or in other words if and only if $0, c \in K(f_c)$. Associated with each of the compact sets $K = K(f_c)$ in the dynamic plane there is a *potential function* or *Green's function* $G^K : \mathbb{C} \rightarrow [0, \infty)$ which vanishes precisely on K , is harmonic off K , and is asymptotic to $\log |z|$ near infinity. The family of *external rays* of K can be described as the orthogonal trajectories of the level curves $G^K = \text{constant}$. Each such ray which extends to infinity can be specified by its angle at infinity $t \in \mathbb{R}/\mathbb{Z}$, and will be denoted by \mathcal{R}_t^K . Here c may be either in or outside of the Mandelbrot set. Similarly, we can consider the potential function G^M and the external rays \mathcal{R}_t^M associated with the Mandelbrot set. We will use the term *dynamic ray* (or briefly *K-ray*) for an external ray of the filled Julia set, and *parameter ray* (or briefly *M-ray*) for an external ray of the Mandelbrot set. (Compare [S1], [S2].)




FIGURE 2. Schematic diagram illustrating the orbit portrait (1).

Definition. — Let $\mathcal{O} = \{z_1, \dots, z_p\}$ be a periodic orbit for f . Suppose that there is some rational angle $t \in \mathbb{Q}/\mathbb{Z}$ so that the dynamic ray $\mathcal{R}_t^{K(f)}$ lands at a point of \mathcal{O} . Then for each $z_i \in \mathcal{O}$ the collection A_i consisting of all angles of dynamic rays which land at the point z_i is a finite and non-vacuous subset of \mathbb{Q}/\mathbb{Z} . The collection $\{A_1, \dots, A_p\}$ will be called the *orbit portrait* $\mathcal{P} = \mathcal{P}(\mathcal{O})$. As an example, Figure 1 shows a quadratic Julia set having a parabolic orbit with portrait

$$\mathcal{P} = \{ \{22/63, 25/63, 37/63\}, \{11/63, 44/63, 50/63\} \}. \quad (1)$$

It is often convenient to represent such a portrait by a schematic diagram, as shown in Figure 2. (For details, and an abstract characterization of orbit portraits, see §2.)

The number of elements in each A_i (or in other words the number of K -rays which land on each orbit point) will be called the *valence* v . Let us assume that $v \geq 2$. Then the v rays landing at z cut the dynamic plane up into v open regions which will be called the *sectors* based at the orbit point $z \in \mathcal{O}$. The *angular width* of a sector S will mean the length of the open arc I_S consisting of all angles $t \in \mathbb{R}/\mathbb{Z}$ with $\mathcal{R}_t^K \subset S$. (We use the word ‘arc’ to emphasize that we will identify \mathbb{R}/\mathbb{Z} with the ‘circle at infinity’ surrounding the plane of complex numbers.) Thus the sum of the angular widths of the v distinct sectors based at an orbit point z is always equal to $+1$. The following result will be proved in 2.11.

Theorem 1.1 (The Critical Value Sector S_1). — *Let \mathcal{O} be an orbit of period $p \geq 1$ for $f = f_c$. If there are $v \geq 2$ dynamic rays landing at each point of \mathcal{O} , then there is one and only one sector S_1 based at some point $z_1 \in \mathcal{O}$ which contains the critical value $c = f(0)$, and whose closure contains no point other than z_1 of the orbit \mathcal{O} . This*

critical value sector S_1 can be characterized, among all of the pv sectors based at the various points of \mathcal{O} , as the unique sector of smallest angular width.

It should be emphasized that this description is correct whether the filled Julia set K is connected or not.

Our main theorem can be stated as follows. Suppose that there exists some polynomial f_{c_0} which admits an orbit \mathcal{O} with portrait \mathcal{P} , again having valence $v \geq 2$. Let $0 < t_- < t_+ < 1$ be the angles of the two dynamic rays $\mathcal{R}_{t_{\pm}}^K$ which bound the critical value sector S_1 for f_{c_0} .

Theorem 1.2 (The Wake $W_{\mathcal{P}}$). — *The two corresponding parameter rays $\mathcal{R}_{t_{\pm}}^M$ land at a single point $\mathbf{r}_{\mathcal{P}}$ of the parameter plane. These rays, together with their landing point, cut the plane into two open subsets $W_{\mathcal{P}}$ and $\mathbb{C} \setminus \overline{W}_{\mathcal{P}}$ with the following property: A quadratic map f_c has a repelling orbit with portrait \mathcal{P} if and only if $c \in W_{\mathcal{P}}$, and has a parabolic orbit with portrait \mathcal{P} if and only if $c = \mathbf{r}_{\mathcal{P}}$.*




FIGURE 3. The boundary of the Mandelbrot set, showing the wake $W_{\mathcal{P}}$ and the root point $\mathbf{r}_{\mathcal{P}} = \frac{1}{4}e^{2\pi i/3} - 1$ associated with the orbit portrait of Figure 1, with characteristic arc $\mathcal{I}_{\mathcal{P}} = (22/63, 25/63)$.

In fact this will follow by combining the assertions 3.1, 4.4, 4.8, and 5.4 below.

Definitions. — This open set $W_{\mathcal{P}}$ will be called the \mathcal{P} -wake in parameter space (compare Atela [A]), and $\mathbf{r}_{\mathcal{P}}$ will be called the *root point* of this wake. The intersection $M_{\mathcal{P}} = M \cap \overline{W}_{\mathcal{P}}$ will be called the \mathcal{P} -limb of the Mandelbrot set. The open arc $I_{S_1} = (t_-, t_+)$ consisting of all angles of dynamic rays \mathcal{R}_t^K which are contained in the interior of S_1 , or all angles of parameter rays \mathcal{R}_t^M which are contained in $W_{\mathcal{P}}$, will be called the *characteristic arc* $\mathcal{I} = \mathcal{I}_{\mathcal{P}}$ for the orbit portrait \mathcal{P} . (Compare 2.6.)

In general, the orbit portraits with valence $v = 1$ are of little interest to us. These portraits certainly exist. For example, for the base map $f_0(z) = z^2$ which lies outside of every wake, *every* orbit portrait has valence $v = 1$. As we follow a path in parameter space which crosses into the wake $W_{\mathcal{P}}$ through its root point, either one orbit with a portrait of valence one degenerates to form an orbit of lower period with portrait \mathcal{P} , or else two different orbits with portraits of valence one fuse together to form an orbit with portrait \mathcal{P} . (If we cross into $W_{\mathcal{P}}$ through a parameter ray $\mathcal{R}_{t_{\pm}}^M$, the picture is similar except that the landing point of the dynamic ray $\mathcal{R}_{t_{\pm}}^K$ jumps discontinuously. If t_+ and t_- belong to the same cycle under angle doubling, then the landing points of both of these dynamics rays jump discontinuously.)

However, there is one exceptional portrait of valence one: The *zero portrait* $\mathcal{P} = \{\{0\}\}$ will play an important role. It is not difficult to check that the dynamic ray \mathcal{R}_0^K of angle 0 for f_c lands at a well defined fixed point if and only if the parameter value c lies in the complement of the parameter ray $\mathcal{R}_0^M = \mathcal{R}_1^M = (1/4, \infty)$. Furthermore, this fixed point necessarily has portrait $\{\{0\}\}$. Thus the wake, consisting of all $c \in \mathbb{C}$ for which f_c has a repelling fixed point with portrait $\{\{0\}\}$, is just the complementary region $\mathbb{C} \setminus [1/4, \infty)$. The characteristic arc $\mathcal{I}_{\{\{0\}\}}$ for this portrait, consisting of all angles t such that $\mathcal{R}_t^K \subset W_{\{\{0\}\}}$, is the open interval $(0, 1)$, and the root point $\mathbf{r}_{\{\{0\}\}}$, the unique parameter value c such that f_c has a parabolic fixed point with portrait $\{\{0\}\}$, is the landing point $c = 1/4$ for the zero parameter ray.




FIGURE 4. Boundaries of the wakes of ray period four or less.

Definition. — It will be convenient to say that a portrait \mathcal{P} is *non-trivial* if it either has valence $v \geq 2$ or is equal to this zero portrait.

Remark. — An alternative characterization would be the following. An orbit portrait $\{A_1, \dots, A_p\}$ is non-trivial if and only if it is *maximal*, in the sense that there is no orbit portrait $\{A'_1, \dots, A'_q\}$ with $A'_1 \supsetneq A_1$. This statement follows easily from 1.5 and 2.7 below. Still another characterization would be that \mathcal{P} is non-trivial if and only if it is the portrait of some parabolic orbit. (See 5.4.)

Corollary 1.3 (Orbit Forcing). — *If \mathcal{P} and \mathcal{Q} are two distinct non-trivial orbit portraits, then the boundaries $\partial W_{\mathcal{P}}$ and $\partial W_{\mathcal{Q}}$ of the corresponding wakes are disjoint subsets of \mathbb{C} . Hence the closures $\overline{W}_{\mathcal{P}}$ and $\overline{W}_{\mathcal{Q}}$ are either disjoint or strictly nested. In particular, if $\mathcal{I}_{\mathcal{P}} \subset \mathcal{I}_{\mathcal{Q}}$ with $\mathcal{P} \neq \mathcal{Q}$, then it follows that $\overline{W}_{\mathcal{P}} \subset W_{\mathcal{Q}}$.*

Thus whenever $\overline{\mathcal{I}}_{\mathcal{P}} \subset \mathcal{I}_{\mathcal{Q}}$, the existence of a repelling or parabolic orbit with portrait \mathcal{P} forces the existence of a repelling orbit with portrait \mathcal{Q} . We will write briefly $\mathcal{P} \Rightarrow \mathcal{Q}$. On the other hand, if $\mathcal{I}_{\mathcal{P}} \cap \mathcal{I}_{\mathcal{Q}} = \emptyset$ then no f_c can have both an orbit with portrait \mathcal{P} and an orbit with portrait \mathcal{Q} .

See Figure 5 for a schematic description of orbit forcing relations for orbits with ray period 4 or less, corresponding to the collection of wakes illustrated in Figure 4. (Evidently this diagram, as well as analogous diagrams in which higher periods are included, has a tree structure, with no loops.)

Proof of 1.3, assuming 1.2. — First note that $W_{\mathcal{P}}$ and $W_{\mathcal{Q}}$ cannot have a boundary ray in common. For the landing point of such a common ray would have to have one parabolic orbit with portrait \mathcal{P} and one parabolic orbit with portrait \mathcal{Q} . But a quadratic map, having only one critical point, cannot have two distinct parabolic orbits. In fact this argument shows that $\partial W_{\mathcal{P}} \cap \partial W_{\mathcal{Q}} = \emptyset$. Note that the parameter point $c = 0$ (corresponding to the map $f_0(z) = z^2$) does not belong to any wake $W_{\mathcal{P}}$ with $\mathcal{P} \neq \{\{0\}\}$. Since rays cannot cross each other, it follows easily that either

$$\overline{W}_{\mathcal{P}} \subset W_{\mathcal{Q}}, \quad \text{or} \quad \overline{W}_{\mathcal{Q}} \subset W_{\mathcal{P}}, \quad \text{or} \quad \overline{W}_{\mathcal{P}} \cap \overline{W}_{\mathcal{Q}} = \emptyset,$$

as required. □

For further discussion and a more direct proof, see §7.

To fill out the picture, we also need the following two statements. To any orbit portrait $\mathcal{P} = \{A_1, \dots, A_p\}$ we associate not only its *orbit period* p but also its *ray period* rp , that is the period of the angles $t \in A_i$ under doubling modulo one. In many cases, rp is a proper multiple of p . (Compare Figure 1.) Suppose in particular that $c \in M$ is a parabolic parameter value, that is suppose that f_c has a periodic orbit where the multiplier is an r -th root of unity, $r \geq 1$. Then one can show that the ray period for the associated portrait is equal to the product rp . (See for example [GM].) This is also the period of the Fatou component containing the critical point.




FIGURE 5. Forcing tree for the non-trivial orbit portraits of ray period $n \leq 4$. Each disk in this figure contains a schematic diagram of the corresponding orbit portrait, with the first n forward images of the critical value sector labeled. (Compare Figure 4; and compare the “disked-tree model” for the Mandelbrot set in Douady [D5].)

This ray period rp is the most important parameter associated with a parabolic point c or with a wake W_p .

It follows from 1.2 that every non-trivial portrait which occurs at all must occur as the portrait of some uniquely determined parabolic orbit. The converse statement will be proved in 4.8:

Theorem 1.4 (Parabolic Portraits are Non-Trivial). — *If c is any parabolic point in M , then the portrait $\mathcal{P} = \mathcal{P}(\mathcal{O})$ of its parabolic orbit is a non-trivial portrait. That is, if we exclude the special case $c = 1/4$, then at least two K -rays must land on each parabolic orbit point.*

It then follows immediately from 1.2 that the parabolic parameter point c must be equal to the root point r_p of an associated wake. It also follows from 1.2 that the angles of the M -rays which bound a wake W_p are always periodic under doubling. In §5 we use a simple counting argument to prove the converse statement. (This imitates Schleicher, who uses a similar counting argument in a different way.)

Theorem 1.5 (Every Periodic Angle Occurs). — *If $t \neq 0$ in \mathbb{R}/\mathbb{Z} is periodic under doubling, then \mathcal{R}_t^M is one of the two boundary rays of some (necessarily unique) wake.*

Further consequences of these ideas will be developed in §6 which shows that each wake contains a uniquely associated hyperbolic component, §8 which describes how each wake contains an associated small copy of the Mandelbrot set, and §9 which shows that each limb is connected even if its root point is removed. There are two appendices giving further supporting details.

Acknowledgement. — I want to thank M. Lyubich and D. Schleicher for their ideas, which play a basic role in this presentation. I am particularly grateful to Schleicher, to S. Zakeri, and to Tan Lei for their extremely helpful criticism of the manuscript. Also, I want to thank both the Gabriella and Paul Rosenbaum Foundation and the National Science Foundation (Grant DMS-9505833) for their support of mathematical activities at Stony Brook.

2. Orbit Portraits

This section will begin the proofs by describing the basic properties of orbit portraits. We will need the following. Let $f(z) = z^2 + c$ with filled Julia set K .

Lemma 2.1 (Mapping of Rays). — *If a dynamic ray \mathcal{R}_t^K lands at a point $z \in \partial K$, then the image ray $f(\mathcal{R}_t^K) = \mathcal{R}_{2t}^K$ lands at the image point $f(z)$. Furthermore, if three or more rays $\mathcal{R}_{t_1}^K, \mathcal{R}_{t_2}^K, \dots, \mathcal{R}_{t_k}^K$ land at $z \neq 0$, then the cyclic order of the angles t_i around the circle \mathbb{R}/\mathbb{Z} is the same as the cyclic order of the doubled angles $2t_i \pmod{\mathbb{Z}}$ around \mathbb{R}/\mathbb{Z} .*

Proof. — Since each $\mathcal{R}_{t_j}^K$ is assumed to be a smooth ray, it cannot pass through any precritical point. Hence $\mathcal{R}_{2t_j}^K$ also cannot pass through a precritical point, and must be a smooth ray landing at $f(z)$. Now suppose that we are given three or more rays with angles $0 \leq t_1 < t_2 < \dots < t_k < 1$, all landing at z . These rays, together with their landing point, cut the plane up into sectors S_1, \dots, S_k , where each S_i is bounded by $\mathcal{R}_{t_i}^K$ and $\mathcal{R}_{t_{i+1}}^K$ (with subscripts modulo k). The cyclic ordering of these various rays can be measured within an arbitrarily small neighborhood of the landing point z , since any transverse arc which crosses $\mathcal{R}_{t_i}^K$ in the positive direction must pass from S_{i-1} to S_i . Since f maps a neighborhood of z to a neighborhood of $f(z)$ by an orientation preserving diffeomorphism, it follows that the image rays must have the same cyclic order. \square

Now let us impose the following.

Standing Hypothesis 2.2. — $\mathcal{O} = \{z_1, \dots, z_p\}$ is a periodic orbit for a quadratic map $f_c(z) = z^2 + c$, with orbit points numbered so that $f(z_j) = z_{j+1}$, taking subscripts modulo p . Furthermore there is at least one rational angle $t \in \mathbb{Q}/\mathbb{Z}$ so that the dynamic ray \mathcal{R}_t^K associated with f lands at some point of this orbit \mathcal{O} .

If c belongs to the Mandelbrot set M , or in other words if the filled Julia set K is connected, then this condition will be satisfied if and only if the orbit \mathcal{O} is either repelling or parabolic. (Compare [Hu], [M3].) On the other hand, for $c \notin M$, all periodic orbit are repelling, but the condition may fail to be satisfied either because the rotation number is irrational (compare [GM, Figure 16]), or because the K -rays which 'should' land on \mathcal{O} bounce off precritical points en route ([GM, Figure 14]).

As in §1, let $A_j \subset \mathbb{R}/\mathbb{Z}$ be the set of all angles of K -rays which land on the point $z_j \in \mathcal{O}$.

Lemma 2.3 (Properties of Orbit Portraits). — *If this Standing Hypothesis 2.2 is satisfied, then:*

- (1) *Each A_j is a finite subset of \mathbb{Q}/\mathbb{Z} .*
- (2) *For each j modulo p , the doubling map $t \mapsto 2t \pmod{\mathbb{Z}}$ carries A_j bijectively onto A_{j+1} preserving cyclic order around the circle,*
- (3) *All of the angles in $A_1 \cup \dots \cup A_p$ are periodic under doubling, with a common period rp , and*
- (4) *the sets A_1, \dots, A_p are pairwise unlinked; that is, for each $i \neq j$ the sets A_i and A_j are contained in disjoint sub-intervals of \mathbb{R}/\mathbb{Z} .*




FIGURE 6. Julia set for $z \mapsto z^2 - 7/4$, showing the six K -rays landing on a period three parabolic orbit. Each number (j) in parentheses is close to the orbit point z_j (and also to $f^{oj}(0)$).

As in §1, the collection $\mathcal{P} = \{A_1, \dots, A_p\}$ is called the *orbit portrait* for the orbit \mathcal{O} . As examples, Figure 6 shows an orbit of period and ray period three, with portrait

$$\mathcal{P} = \{ \{3/7, 4/7\}, \{6/7, 1/7\}, \{5/7, 2/7\} \},$$

Figure 7 shows a period three orbit with ray period six, and with portrait

$$\mathcal{P} = \{ \{4/9, 5/9\}, \{8/9, 1/9\}, \{7/9, 2/9\} \},$$




FIGURE 7. Julia set for $z \mapsto z^2 - 1.77$, showing the six K -rays landing on a period three orbit. In contrast to Figure 6, these six rays are permuted cyclically by the map.

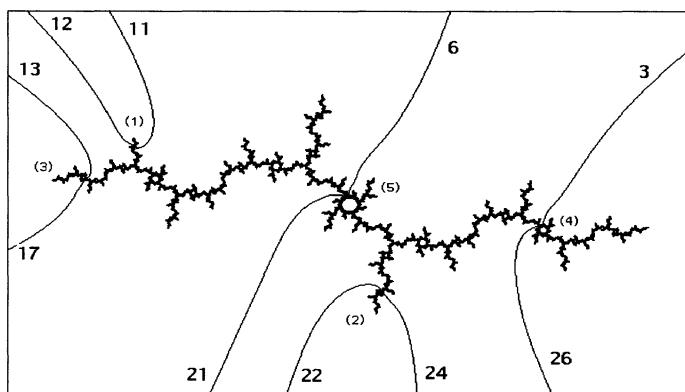


FIGURE 8. Julia set $J(f_c)$ for $c = -1.2564 + .3803i$, showing the ten rays landing on a period 5 orbit. Here the angles are in units of $1/31$.

while Figure 8 shows an orbit of period and ray period five, with portrait

$$\mathcal{P} = \left\{ \left\{ \frac{11}{31}, \frac{12}{31} \right\}, \left\{ \frac{22}{31}, \frac{24}{31} \right\}, \left\{ \frac{13}{31}, \frac{17}{31} \right\}, \left\{ \frac{26}{31}, \frac{3}{31} \right\}, \left\{ \frac{21}{31}, \frac{6}{31} \right\} \right\}.$$




FIGURE 9. Schematic diagrams associated with the orbit portraits of Figures 6, 7, 8. The angles are in units of $1/7$, $1/9$ and $1/31$ respectively.

Proof of 2.3. — Since some A_i contains a rational number modulo \mathbb{Z} , it follows from 2.1 that some A_j contains an angle t_0 which is periodic under doubling. Let the period be $n \geq 1$, so that $2^n t_0 \equiv t_0 \pmod{\mathbb{Z}}$. Applying 2.1 n times, we see that the mapping $\eta(t) \equiv 2^n t \pmod{\mathbb{Z}}$ maps the set $A_j \subset \mathbb{R}/\mathbb{Z}$ injectively into itself, preserving cyclic order and fixing t_0 . In fact we will show that every element of A_j is fixed by η . For otherwise, if $t \in A_j$ were not fixed, then choosing suitable representatives modulo \mathbb{Z} we would have for example $t_0 = \eta(t_0) < t < \eta(t) < t_0 + 1$. Since η preserves cyclic order, it would then follow inductively that

$$t_0 < t < \eta(t) < \eta^{\circ 2}(t) < \eta^{\circ 3}(t) < \cdots < t_0 + 1.$$

Hence the successive images of t would converge to a fixed point of η . But this is impossible since every fixed point of η is repelling. Thus η fixes every point of A_j . But the fixed points of η are precisely the rational numbers of the form $i/(2^n - 1)$, so it follows that A_j is a finite set of rational numbers. It follows easily that all of the A_k are pointwise fixed by η . This proves (1), (2) and (3) of 2.3; and (4) is clearly true since rays cannot cross each other. \square

It is often convenient to compactify the complex numbers by adding a circle of points $e^{2\pi i t} \infty$ at infinity, canonically parametrized by $t \in \mathbb{R}/\mathbb{Z}$. Within the resulting closed topological disk \mathbb{D} , we can form a *diagram* \mathcal{D} illustrating the orbit portrait \mathcal{P} by drawing all of the K -rays joining the circle at infinity to \mathcal{O} . These various rays are disjoint, except that each $z \in \mathcal{O}$ is a common endpoint for exactly v of these rays.

Note that this diagram \mathcal{D} deforms continuously, preserving its topology, as we move the parameter point c , provided that the periodic orbit \mathcal{O} remains repelling, and provided that the associated K -rays do not run into precritical points. (Compare [GM, Appendix B].)

In fact, given \mathcal{P} , we can construct a diagram homeomorphic to \mathcal{D} as follows. Start with the unit circle, and mark all of the points $e(t) = e^{2\pi i t}$ corresponding to angles t in the union $\mathbf{A}_{\mathcal{P}} = A_1 \cup \cdots \cup A_p$. Now for each A_i , let \hat{z}_i be the center of gravity of the corresponding points $e(t)$, and join each of these points to \hat{z}_i by a straight line segment. It follows easily from Condition (4) that these line segments will not cross each other. (In practice, in drawing such diagrams, we will not usually use straight lines and centers of gravity, but rather use some topologically equivalent picture, fixing the boundary circle, which is easier to see. Compare Figures 2, 5, 9.)

It will be convenient to temporarily introduce the term *formal orbit portrait* for a collection $\mathcal{P} = \{A_1, \dots, A_p\}$ of subsets of \mathbb{R}/\mathbb{Z} which satisfies the four conditions of 2.3, whether or not it is actually associated with some periodic orbit. In fact we will prove the following.

Theorem 2.4 (Characterization of Orbit Portraits). — *If \mathcal{P} is any formal orbit portrait, then there exists a quadratic polynomial f and an orbit \mathcal{O} for f which realizes this portrait \mathcal{P} .*

This will follow from Lemma 2.9 below. To begin the proof, let us study the way in which the angle doubling map acts on a formal orbit portrait. As in §1, the number of angles in each A_j will be called the *valence* v for the formal portrait \mathcal{P} . It is easy to see that any formal portrait of valence $v = 1$ can be realized by an appropriate orbit for the map $f(z) = z^2$. Hence it suffices to study the case $v \geq 2$. For each $A_j \in \mathcal{P}$ the v connected components of the complement $\mathbb{R}/\mathbb{Z} \setminus A_j$ are connected open arcs with total length $+1$. These will be called the *complementary arcs* for A_j .

Lemma 2.5 (The Critical Arcs). — *For each A_j in the formal orbit portrait \mathcal{P} , all but one of the complementary arcs is carried diffeomorphically by the angle doubling map onto a complementary arc for A_{j+1} . However, the remaining complementary arc for A_j has length greater than $1/2$. Its image under the doubling map covers one particular complementary arc for A_{j+1} twice, and every other complementary arc for A_{j+1} just once.*

Definition. — This longest complementary arc will be called the *critical arc* for A_j . The arc which it covers twice under doubling will be called the *critical value arc* for A_{j+1} . (This language will be justified in 2.9 below.)

Proof of 2.5. — If $I \subset \mathbb{R}/\mathbb{Z}$ is a complementary arc for A_j of length less than $1/2$, then clearly the doubling map carries I bijectively onto an arc $2I$ of twice the length, bounded by two points of A_{j+1} . This image arc cannot contain any other point of A_{j+1} , since the doubling map from A_j to A_{j+1} preserves cyclic order. It follows easily that these image arcs cannot overlap. Since we cannot fit v arcs of total length $+2$ into the circle without overlap, and since there cannot be any complementary arc of length exactly $1/2$, it follows that there must be exactly one “critical” complementary arc for A_j which has length greater than $1/2$. Suppose that it has length $(1 + \varepsilon_j)/2$. Then the $v - 1$ non-critical arcs for A_j have total length $(1 - \varepsilon_j)/2$, and their images under doubling form $v - 1$ complementary arcs for A_{j+1} with total length $1 - \varepsilon_j$. Since the doubling map is exactly two-to-one, it follows easily that it maps the critical arc for A_j onto the entire circle, doubly covering one “critical value arc” for A_{j+1} which has length ε_j , and covering every other complementary arc for A_{j+1} just once. \square

Lemma 2.6 (The Characteristic Arc for \mathcal{P}). — *Among the complementary arcs for the various $A_j \in \mathcal{P}$, there exists a unique arc $\mathcal{I}_{\mathcal{P}}$ of shortest length. This shortest arc is a critical value arc for its A_j , and is contained in all of the other critical value arcs.*

Definition. — This shortest complementary arc $\mathcal{I}_{\mathcal{P}}$ will be called the *characteristic arc* for \mathcal{P} . (Compare 2.11.)

Proof of 2.6. — There certainly exists at least one complementary arc $\mathcal{I}_{\mathcal{P}}$ of minimal length ℓ among all of the complementary arcs for all of the $A_j \in \mathcal{P}$. This $\mathcal{I}_{\mathcal{P}}$ must be a critical value arc, since otherwise it would have the form $2J$ where J is some complementary arc of length $\ell/2$. Suppose then that $\mathcal{I}_{\mathcal{P}}$ is the critical value arc for

A_{j+1} , doubly covered by the critical arc I_c for A_j . Since $\mathcal{I}_{\mathcal{P}}$ is minimal, it follows from 2.3(4) that this open arc $\mathcal{I}_{\mathcal{P}}$ cannot contain any point of the union $\mathbf{A}_{\mathcal{P}} = A_1 \cup \cdots \cup A_p$. Hence its preimage under doubling also cannot contain any point of $\mathbf{A}_{\mathcal{P}}$. This preimage consists of two arcs I' and $I'' = I' + 1/2$, each of length $\ell/2$. Note that both of these arcs are contained in I_c . In fact the arc I_c of length $(1 + \ell)/2$ is covered by these two open arcs of length $\ell/2$ lying at either end, together with the closed arc $I_c \setminus (I' \cup I'')$ of length $(1 - \ell)/2$ in the middle.

Now consider any $A_k \in \mathcal{P}$ with $k \neq j$. It follows from the unlinking property 2.3(4) that the entire set A_k must be contained either in the arc $(\mathbb{R}/\mathbb{Z}) \setminus I_c$ of length $(1 - \ell)/2$, or in I_c and hence in the arc $I_c \setminus (I' \cup I'')$ which also has length $(1 - \ell)/2$. In either case, it follows that the union of all non-critical arcs for A_k is contained in this same arc of length $(1 - \ell)/2$, and hence that the image of this union under doubling is contained in the arc

$$2((\mathbb{R}/\mathbb{Z}) \setminus I_c) = 2(I_c \setminus (I' \cup I'')) = (\mathbb{R}/\mathbb{Z}) \setminus \mathcal{I}_{\mathcal{P}}$$

of length $1 - \ell$. Therefore, the critical value arc for A_{k+1} contains the complementary arc $\mathcal{I}_{\mathcal{P}}$, as required. It follows that this minimal arc $\mathcal{I}_{\mathcal{P}}$ is unique. For if there were an $\mathcal{I}'_{\mathcal{P}}$ of the same length, then this argument would show that each of these two must contain the other, which is impossible. \square

Remark. — This characteristic arc never contains the angle zero. In fact let I_c be the critical arc whose image under doubling covers $\mathcal{I}_{\mathcal{P}}$ twice. If $0 \in I_{\mathcal{P}}$, then it is not hard to see that one endpoint of I_c must lie in $\mathcal{I}_{\mathcal{P}}$ and the other endpoint must lie outside, in $1/2 + \mathcal{I}_{\mathcal{P}}$. But this is impossible by 2.3(4) and the minimality of $\mathcal{I}_{\mathcal{P}}$.

Recall that the union $\mathbf{A}_{\mathcal{P}} = A_1 \cup \cdots \cup A_p$ contains pv elements, each of which has period rp under doubling. Hence this union splits up into

$$\frac{pv}{rp} = \frac{v}{r}$$

distinct cycles under doubling. If \mathcal{P} is the portrait of a periodic orbit \mathcal{O} , then the ratio v/r can be described as the *number of cycles of K -rays* which land on the orbit \mathcal{O} . As examples, we have $v = r = 3$ for Figure 1 and $v = r = 2$ for Figure 7 so that there is only one cycle under doubling, but $v = 2$ and $r = 1$ for Figures 6 and 8 so that there are two distinct cycles. In fact we next show that there are at most two cycles in all cases.

Lemma 2.7 (Primitive versus Satellite). — *Any formal orbit portrait of valence $v > r$ must have $v = 2$ and $r = 1$. It follows that there are just two possibilities:*

Primitive Case. If $r = 1$, so that every ray which lands on the period p orbit is mapped to itself by $f^{\circ p}$, then at most two rays land on each orbit point.

Satellite Case. If $r > 1$, then $v = r$ so that exactly r rays land on each orbit point, and all of these rays belong to a single cyclic orbit under angle doubling.

This terminology will be justified in §6. (Compare Figure 12.)

Proof of 2.7. — Suppose that $v > r$ and $v \geq 3$. Let $\mathcal{I}_{\mathcal{P}}$ be the characteristic arc. We suppose that $\mathcal{I}_{\mathcal{P}}$ is the critical value arc in the complement of A_1 . Let I_- the complementary arc for A_1 which is just to the left of $\mathcal{I}_{\mathcal{P}}$ and let I_+ be the complementary arc just to the right of $\mathcal{I}_{\mathcal{P}}$. To fix our ideas, suppose that I_- has length $\ell(I_-) \geq \ell(I_+)$. Since I_+ is not the critical value arc for A_1 , we see, arguing as in 2.6, that it must be the image under iterated doubling of the critical value arc I' for some A_j . That is, we have $I_+ = 2^m I'$ for some $m \geq 1$. Hence $\ell(I') < \ell(I_+)$.

The hypothesis that $v > r$ implies that the two endpoints of $\mathcal{I}_{\mathcal{P}}$ belong to different cycles under doubling. Thus the left endpoints of I' and $\mathcal{I}_{\mathcal{P}}$ belong to distinct cycles, hence $I' \neq \mathcal{I}_{\mathcal{P}}$. Therefore, by 2.6, I' strictly contains $\mathcal{I}_{\mathcal{P}}$. This arc I' cannot strictly contain the neighboring arc I_+ , since it is shorter than I_+ . Hence it must have an endpoint in I_+ , and therefore, by 2.3(4), it must have both endpoints in I_+ . But this implies that I' contains I_- , which is impossible since $\ell(I') < \ell(I_+) \leq \ell(I_-)$. Thus, if $v > r$ it follows that $v \leq 2$, hence $r = 1$ and $v = 2$, as asserted. \square

Lemma 2.8 (Two Rays determine \mathcal{P}). — *Let $\mathcal{P} = \{A_1, \dots, A_p\}$ be a formal orbit portrait of valence $v \geq 2$, and let $\mathcal{I}_{\mathcal{P}} = (t_-, t_+)$ be its characteristic arc, as described above. Then a quadratic polynomial f_c has an orbit with portrait \mathcal{P} if and only if the two K -rays with angles t_- and t_+ for the filled Julia set of f_c land at a common point.*

Proof. — If f_c has an orbit with portrait \mathcal{P} , this is true by definition. Conversely, if these rays land at a common point z_1 , then the orbit of z_1 is certainly periodic. Let \mathcal{P}' be the portrait for this actual orbit. We will denote its period by p' , its valence by v' , and so on. Note that the ray period rp is equal to $r'p'$, the common period of the angles t_- and t_+ under doubling.

Primitive Case. — Suppose that $r = 1$ so that $v/r = 2$, and so that each of these angles t_{\pm} has period exactly p under doubling. If $p' < p$ hence $r' > 1$, then it would follow from 2.7 applied to the portrait \mathcal{P}' that t_- and t_+ must belong to the same cycle under doubling, contradicting the hypothesis that $v/r = 2$.

Satellite Case. — If $r > 1$ hence $v = r$, then t_- and t_+ do belong to the same cycle under doubling, say $2^k t_- \equiv t_+ \pmod{\mathbb{Z}}$. Clearly it follows that $r' > 1$ hence $v' = r'$. Furthermore, it follows easily that multiplication by 2^k acts transitively on A_1 , and hence that all of the rays \mathcal{R}_t^K with $t \in A_1$ land at the same point z_1 . In other words $A_1 \subset A'_1$. This implies that $r \leq r'$ hence $p \geq p'$. If p were strictly greater than p' , then it would follow that $A_{1+p'}$ is also contained in A'_1 . But the two sets A_1 and $A_{1+p'}$ are unlinked in \mathbb{R}/\mathbb{Z} . Hence there is no way that multiplication by 2^p can act non-trivially on $A_1 \cup A_{1+p'}$ carrying each of these two sets into itself and preserving cyclic order on their union. This contradiction implies that $A_1 = A'_1$ and $p = p'$, and hence that $\mathcal{P} = \mathcal{P}'$, as required. \square

Now let c be some parameter value outside the Mandelbrot set. Then, following Douady and Hubbard, the point c , either in the dynamic plane or in the parameter plane, lies on a unique external ray, with the same well defined angle $t(c) \in \mathbb{R}/\mathbb{Z}$ in either case. (Compare Appendix A.)

Lemma 2.9 (Outside the Mandelbrot Set). — *Let $\mathcal{P} = \{A_1, \dots, A_p\}$ be a formal orbit portrait with characteristic arc $\mathcal{I}_{\mathcal{P}}$, and let c be a parameter value outside of the Mandelbrot set. Then the map $f_c(z) = z^2 + c$ admits a periodic orbit with portrait \mathcal{P} if and only if the external angle $t(c)$ belongs to this open arc $\mathcal{I}_{\mathcal{P}}$.*

Proof. — The two dynamic rays $\mathcal{R}_{t(c)/2}^K$ and $\mathcal{R}_{(1+t(c))/2}^K$ meet at the critical point 0, and together cut the dynamic plane into two halves. Furthermore, every point of the Julia set $\partial K = K$ is uniquely determined by its symbol sequence with respect to this partition. Correspondingly, the two diametrically opposite points $t(c)/2$ and $(1 + t(c))/2$ on the circle \mathbb{R}/\mathbb{Z} cut the circle into two semicircles, and almost every point $t \in \mathbb{R}/\mathbb{Z}$ has a well defined symbol sequence with respect to this partition under the doubling map. Two rays \mathcal{R}_t^K and \mathcal{R}_u^K land at a common point of K if and only if the external angles t and u have the same symbol sequence.

First suppose that the angle $t(c)$ lies in the characteristic arc $\mathcal{I}_{\mathcal{P}}$. Then, with notation as in the proof of 2.6, the two points $t(c)/2$ and $(1 + t(c))/2$ lie in the two components I' and I'' of the preimage of $\mathcal{I}_{\mathcal{P}}$. For every $A_j \in \mathcal{P}$, all of the points of A_j lie in a single component of $\mathbb{R}/\mathbb{Z} \setminus (I' \cup I'')$. Hence the rays \mathcal{R}_t^K with $t \in A_j$ land at a common point $z_j \in K$. It follows from 2.8 that these points lie in an orbit with portrait \mathcal{P} , as required.

On the other hand, if $t(c)$ lies outside of $\overline{\mathcal{I}_{\mathcal{P}}}$, then it is easy to check that the two endpoints of $\mathcal{I}_{\mathcal{P}}$ are separated by the points $t(c)/2$ and $(1 + t(c))/2$. Hence these two endpoints, both belonging to $A_1 \in \mathcal{P}$, land at different points of K . Hence f_c has no orbit with portrait \mathcal{P} .

Finally, in the limiting case where $t(c)$ is precisely equal to one of the two endpoints t_{\pm} of $\mathcal{I}_{\mathcal{P}}$, since these angles are periodic under doubling, it follows that the ray $\mathcal{R}_{t_{\pm}}^K$ passes through a precritical point, and hence does not have any well defined landing point in K . This completes the proof of 2.9. \square

Evidently the Realization Theorem 2.4 is an immediate corollary. Since we have proved 2.4, we can now forget about the distinction between “formal” orbit portraits and portraits which are actually realized. We can describe further properties of portraits and their associated diagrams as follows.

Definition 2.10. — Suppose that we start with any periodic orbit \mathcal{O} with valence $v \geq 2$ and period $p \geq 1$, and fix some point $z_i \in \mathcal{O}$. As in §1, the v rays landing at z_i cut the dynamic plane \mathbb{C} up into v open subsets which we call the *sectors* based at z_i . Evidently there is a one-to-one correspondence between sectors based at z_i and complementary arcs for the corresponding set of angles $A_i \subset \mathbb{R}/\mathbb{Z}$, characterized by

the property that \mathcal{R}_t^K is contained in the open sector S if and only if t is contained in the corresponding complementary arc. By definition, the *angular size* $\alpha(S) > 0$ of a sector is the length of the corresponding complementary arc, which we can think of as its “boundary at infinity”. It follows that $\sum_S \alpha(S) = 1$, where the sum extends over the v sectors based at some fixed $z_i \in \mathcal{O}$.

Remark. — The angular size of a sector has nothing to do with the angle between the rays at their common landing point, which is often not even defined.

Altogether there are pv rays landing at the various points of the orbit \mathcal{O} . Together these rays cut the plane up into $pv - p + 1$ connected components. The closures of these components will be called the pieces of the *preliminary puzzle* associated with the diagram \mathcal{D} or the associated portrait \mathcal{P} . Note that every closed sector \bar{S} can be expressed as a union of preliminary puzzle pieces, and that every preliminary puzzle piece is equal to the intersection of the closed sectors containing it. This construction will be modified and developed further in Sections 7 and 8.

For every point z_i of the orbit, note that just one of the v sectors based at z_i contains the critical point 0. We will call this the *critical sector* at z_i , while the others will be called the *non-critical sectors* at z_i . Another noteworthy sector at z_i (not necessarily distinct from the critical sector) is the *critical value sector*, which contains $f(0) = c$.

Lemma 2.11 (Properties of Sectors). — *The diagram $\mathcal{D} \subset \mathbb{C}$ associated with any orbit \mathcal{O} of valence $v \geq 2$ has the following properties:*

- (a) *For each $z_i \in \mathcal{O}$, the critical sector at z_i has angular size strictly greater than $1/2$. It follows that the $v - 1$ non-critical sectors at z_i have total angular size less than $1/2$.*
- (b) *The map f carries a small neighborhood of z_i diffeomorphically onto a small neighborhood of $z_{i+1} = f(z_i)$, carrying each sector based at z_i locally onto a sector based at z_{i+1} , and preserving the cyclic order of these sectors around their base point. The critical sector at z_i always maps locally, near z_i , onto the critical value sector based at z_{i+1} .*
- (c) *Globally, each non-critical sector S at z_i is mapped homeomorphically by f onto a sector $f(S)$ based at z_{i+1} , with angular size given by $\alpha(f(S)) = 2\alpha(S)$. However, the critical sector at z_i maps so as to cover the entire plane, covering the critical value sector at z_{i+1} twice with a ramification point at $0 \mapsto c$, and covering every other sector just once.*
- (d) *Among all of the pv sectors based at the various points of \mathcal{O} , there is a unique sector of smallest angular size, corresponding to the characteristic arc \mathcal{I}_P . This smallest sector contains the critical value, and does not contain any other sector.*

(As usual, the index i is to be construed as an integer modulo p .) The proof, based on 2.6 and the fact that f is exactly two-to-one except at its critical point, is straightforward and will be left to the reader. Evidently Theorem 1.1 follows. \square

Now let us take a closer look at the dynamics of the diagram \mathcal{D} or of the associated portrait \mathcal{P} . The iterated map $f^{\circ p}$ fixes each point $z_i \in \mathcal{O}$, permuting the various rays which land on z_i but preserving their cyclic order. Equivalently, the p -fold iterate of the doubling map carries each finite set $A_i \subset \mathbb{Q}/\mathbb{Z}$ onto itself by a bijection which preserves the cyclic order. For any fixed $i \bmod p$, we can number the angles in A_i as $0 \leq t^{(1)} < t^{(2)} < \dots < t^{(v)} < 1$. It then follows that

$$2^p t^{(j)} \equiv t^{(j+k)} \pmod{\mathbb{Z}},$$

taking superscripts modulo v , where k is some fixed residue class modulo v .

Definition 2.12. — The ratio $k/v \pmod{\mathbb{Z}}$ is called the combinatorial *rotation number* of our orbit portrait. It is easy to check that this rotation number does not depend on the choice of orbit point z_i . Let d be the greatest common divisor of v and k . Then we can express the rotation number as a fraction q/r in lowest terms, where $k = qd$ and $v = rd$. (In the special case of rotation number zero, we take $q = 0$ and $r = 1$.)

In all cases, note that the denominator $r \geq 1$ is equal to the period of the angles $t^{(j)} \in A_i$ under the mapping $t \mapsto 2^p t \pmod{\mathbb{Z}}$ from A_i to itself. It follows easily that the period of $t^{(j)}$ under angle doubling is equal to the product rp . Thus this definition of r as the denominator of the rotation number is compatible with our earlier notation rp for the ray period.

Notation Summary. — Since we have been accumulating quite a bit of notation, here is a brief summary:

Orbit period p : the number of distinct element in our orbit \mathcal{O} ,

Ray period rp : the period of each angle $t \in A_1 \cup \dots \cup A_p$ under doubling.

Rotation number q/r : describes the action of multiplication by 2^p on each set A_i .

Valence v : number of angles in each A_i , for a total of pv angles altogether.

Cycle number v/r : the number of disjoint cycles of size rp in the union $A_1 \cup \dots \cup A_p$.

According to 2.7, this cycle number is always equal to 1 for a satellite portrait, and is at most 2 in all cases. Thus, in the case $v \geq 2$ there are just two possibilities as follows:

Primitive Case. — The rotation number is zero. There are $v = 2$ rays landing at each orbit point, for a total of $2p$ rays. These split up into two cycles of p rays each under doubling.

Satellite Case. — The rotation number is $q/r \not\equiv 0$. There are $v = r$ rays landing at each orbit point, for a total of $pv = rp$ rays altogether. These rp rays are permuted cyclically under angle doubling, so that the number of cycles is $v/r = 1$.

As examples, Figures 6, 8 illustrate primitive portraits with rotation number zero, while Figures 1, 7 show satellite portraits with rotation number $1/3$ and $1/2$. We will see in §6 that primitive portraits correspond to primitive hyperbolic components in the Mandelbrot set, that is, to those with a cusp point.

3. Parameter Rays

This section will prove the following preliminary version of Theorem 1.2.

Let \mathcal{P} be any orbit portrait of valence $v \geq 2$, and let $\mathcal{I}_{\mathcal{P}} = (t_-, t_+)$ be its characteristic arc, where $0 < t_- < t_+ < 1$. If the quadratic polynomial $f_c = z^2 + c$ has an orbit \mathcal{O} with portrait \mathcal{P} , recall that the two dynamic rays $\mathcal{R}_{t_-}^K$ and $\mathcal{R}_{t_+}^K$ for f_c land at a common orbit point, and together bound a sector S_1 which has minimal angular size among all of the sectors based at points of the orbit \mathcal{O} . This S_1 can also be characterized as the smallest of these sectors which contains the critical value c . (Compare Lemmas 2.6, 2.9, 2.11.)

Theorem 3.1 (Parameter Rays and the Wake). — *The two parameter rays $\mathcal{R}_{t_-}^M$ and $\mathcal{R}_{t_+}^M$ with these same angles land at a common parabolic point in the Mandelbrot set. Furthermore, these two rays, together with their common landing point, cut the parameter plane into two open subsets $W_{\mathcal{P}}$ and $\mathbb{C} \setminus \overline{W_{\mathcal{P}}}$ with the following property: The quadratic map f_c has a repelling orbit with portrait \mathcal{P} if and only if $c \in W_{\mathcal{P}}$.*

Proof. — Let $\mathbf{A}_{\mathcal{P}} = A_1 \cup \dots \cup A_p$ be the set of all angles for the orbit portrait \mathcal{P} , and let $n = rp$ be the common period of these angles under doubling. The set $F_n \subset M$ of possibly exceptional parameter values will consist of those c for which $f_c^{o_n}$ has a fixed point of multiplier $+1$. Since $F_n \subset \mathbb{C}$ is an algebraic variety and is not the entire complex plane, it is necessarily a finite set. As noted in [GM], if c belongs to the Mandelbrot set but $c \notin F_n$, then the various dynamic rays $\mathcal{R}_t^{K(f_c)}$ with $t \in \mathbf{A}_{\mathcal{P}}$ all land on repelling periodic points, and the pattern of which of these rays land at a common point remains stable under perturbation of c throughout some open neighborhood within parameter space.

Now suppose that c lies outside of the Mandelbrot set. Then c , considered as a point in parameter space, belongs to some uniquely defined parameter ray $\mathcal{R}_{t(c)}^M$, and considered as a point in the dynamic plane for f_c , belongs to the dynamic ray $\mathcal{R}_{t(c)}^K$ with this same angle. In this case, a dynamic ray \mathcal{R}_t^K for f_c has a well defined landing point in $K = K(f_c)$ if and only if the forward orbit $\{2t, 4t, 8t, \dots\}$ under doubling does not contain this critical value angle $t(c)$. Since the angles in $\mathbf{A}_{\mathcal{P}}$ are periodic, it follows that all of the dynamic rays \mathcal{R}_t^K with $t \in \mathbf{A}_{\mathcal{P}}$ have well defined landing points in K if and only if $t(c) \notin \mathbf{A}_{\mathcal{P}}$.

Let $t \in \mathbf{A}_{\mathcal{P}}$ and let $c_0 \in M$ be any accumulation point for the parameter ray \mathcal{R}_t^M . Since every neighborhood of c_0 contains parameter values $c \in \mathcal{R}_t^M$ for which

the dynamic ray $\mathcal{R}_t^{K(f_c)}$ does not land, it follows that c_0 must belong to F_n . Thus every accumulation point for \mathcal{R}_t^M belongs to the finite set F_n . Since the set of all accumulation points of a ray is connected, this proves that \mathcal{R}_t^M must actually land at a single point of F_n .

These parameter rays \mathcal{R}_t^M with $t \in \mathbf{A}_{\mathcal{P}}$, together with the points of F_n , cut the complex parameter plane up into finitely many open sets U_i , and the pattern of which of the corresponding dynamic rays \mathcal{R}_t^K with $t \in \mathbf{A}_{\mathcal{P}}$ land at a common periodic point remains fixed as c varies through any U_i . Since every U_i is unbounded, it follows from Lemma 2.9 that for $c \in U_i$ the map f_c has an orbit with portrait \mathcal{P} if and only if U_i is that open set which contains the points in $\mathbb{C} \setminus M$ with external angle $t(c)$ in (t_-, t_+) . It follows that the two rays $\mathcal{R}_{t_-}^M$ and $\mathcal{R}_{t_+}^M$ must land at a common point of F_n , so as to separate the parameter plane. For otherwise, if they had different landing points, the connected set U_i containing points with external angles in (t_-, t_+) would also contain points with other external angles, which is impossible.

Define the *root point* $\mathbf{r}_{\mathcal{P}} \in M$ to be this common landing point, and define the *wake* $W_{\mathcal{P}}$ to be that connected component of $\mathbb{C} \setminus (\mathcal{R}_{t_-}^M \cup \mathcal{R}_{t_+}^M \cup \mathbf{r}_{\mathcal{P}})$ which does not contain 0. For $c \in W_{\mathcal{P}} \setminus F_n$, it follows from the discussion above that f_c does have a repelling orbit with portrait \mathcal{P} , while for $c \in \mathbb{C} \setminus (W_{\mathcal{P}} \cup F_n)$ it follows that f_c does not have any repelling orbit with portrait \mathcal{P} . Thus, to complete the proof of 3.1, we need only consider those f_c with c in the finite set F_n .

First suppose that some point $c_0 \in F_n \setminus W_{\mathcal{P}}$ had a repelling orbit with portrait \mathcal{P} . Then any nearby parameter value would have a nearby repelling orbit with the same landing pattern for rays with angles in $\mathbf{A}_{\mathcal{P}}$. A priori it might seem possible that some extra ray, perhaps one landing on a parabolic orbit for f_{c_0} , might land on this same repelling orbit after perturbation. (Compare [GM, Fig. 12].) However, this is ruled out by 2.8. Hence all nearby parameter values must belong to $W_{\mathcal{P}}$, which is impossible.

Now consider parameter points $c \in W_{\mathcal{P}}$. I am indebted to Tan Lei for pointing out the following very elegant argument due to Peter Haïssinsky which replaces my own more complicated reasoning. As noted above, for every $c \in W_{\mathcal{P}}$ the rays $\mathcal{R}_{t_-}^{K(f_c)}$ and $\mathcal{R}_{t_+}^{K(f_c)}$ land at well defined periodic points of the Julia set $J(f_c)$. Let $z = z(m, t_{\pm}, c)$ be the unique point on the ray $\mathcal{R}_{t_{\pm}}^{K(f_c)}$ which has potential $G(z) = 1/m$. The functions $c \mapsto z(m, t_{\pm}, c)$ are evidently holomorphic; in fact $z(m, t, c) = \phi_c^{-1}(\exp(2\pi it + 1/m))$, where the function

$$\phi_c : \mathbb{C} \setminus K \xrightarrow{\cong} \mathbb{C} \setminus \overline{\mathbb{D}} \quad \text{can be defined locally as} \quad \phi_c(z) = \lim_{k \rightarrow \infty} \sqrt[k]{f_c^{\circ k}(z)}$$

(choosing appropriate branches of the iterated square root) and hence is holomorphic as a function of both variables. These maps $c \mapsto z(m, t_-, c)$ or $c \mapsto z(m, t_+, c)$ form a normal family throughout $W_{\mathcal{P}}$, since they miss the three points 0, c and ∞ . Choosing a convergent subsequence and passing to the limit as $m \rightarrow \infty$, we see that the landing

points also depend holomorphically on c . Since the landing points for $\mathcal{R}_{t-}^{K(f_c)}$ and $\mathcal{R}_{t+}^{K(f_c)}$ coincide for $c \in W_{\mathcal{P}} \setminus F_n$, it follows by continuity that they coincide for all $c \in W_{\mathcal{P}}$. It follows also that the multiplier $\lambda(c)$ of this common landing point depends holomorphically on the parameter $c \in W_{\mathcal{P}}$, with $|\lambda(c)| \geq 1$ since the landing point must be repelling or parabolic. But the absolute value of a non-zero holomorphic function cannot have a local minimum unless it is constant, so all of these landing points must actually be repelling. Together with 2.8, this completes the proof of 3.1. \square

We will deal with parabolic orbits with portrait \mathcal{P} in the next two sections.

4. Near Parabolic Maps

Let \hat{c} be a parabolic point in parameter space. This section will study the dynamic behavior of the quadratic map f_c for c in a neighborhood of \hat{c} . (Compare [DH2, §14(CH)], [Sh2].)

Let \mathcal{O} be the parabolic orbit for $f_{\hat{c}}$ with period $p \geq 1$ and with representative point \hat{z} . Then the multiplier $\hat{\lambda} = (f_{\hat{c}}^{\circ p})'(\hat{z})$ is a primitive r -th root of unity for some $r \geq 1$. Let \mathcal{P} be the associated orbit portrait, with ray period $rp \geq p$. We will first prove the following.

Theorem 4.1 (Deformation Preserving the Orbit Portrait). — *There exists a smooth path in parameter space ending at the parabolic point \hat{c} and consisting of parameter values c with the following property: The associated map f_c has both a repelling orbit of period p and an attracting orbit of period rp . Furthermore, this repelling orbit has portrait \mathcal{P} , and lies on the boundary of the immediate basin for the attracting orbit. As c tends to \hat{c} , these two orbits both converge towards the original parabolic orbit \mathcal{O} .*




FIGURE 10. (Courtesy of S. Zakeri). The left sketch shows a parabolic fixed point with $r = 3$, the middle shows the modified version with an attracting orbit of period 3, and the right shows a modified version with an attracting fixed point. Here the arrows indicate the action of $f^{\circ 3}$.

(Compare Figure 10, middle.) The proof will depend on the following.

Lemma 4.2 (Convenient Coordinates). — *For any complex number λ close to $\widehat{\lambda}$ there exists at least one parameter value c close to \widehat{c} and point z_λ close to \widehat{z} so that z_λ is a periodic point for the map f_c with period p and with multiplier λ . Furthermore there is a local holomorphic change of coordinate $z = \phi_\lambda(w)$ with $z_\lambda = \phi_\lambda(0)$ so that the map $F = F_\lambda = \phi_\lambda^{-1} \circ f_c^{\circ p} \circ \phi_\lambda$ takes the form*

$$F(w) = \lambda w + R(\lambda, w)$$

for w near zero, and so that its r -th iterate takes the form

$$F^{\circ r}(w) = \phi_\lambda^{-1} \circ f_c^{\circ rp} \circ \phi_\lambda = \lambda^r w (1 + w^r + R'(\lambda, w)), \quad (2)$$

where the remainder terms R and R' satisfy $|R|, |R'| \leq \text{constant} |w|^{r+1}$ uniformly for λ in some neighborhood of $\widehat{\lambda}$ and for w in some neighborhood of zero.

(In 4.5, we will sharpen this statement by showing that the phrase “at least one” in 4.2 can be replaced by “exactly one”.)

Proof of 4.2 in the Primitive Case. — First suppose that \mathcal{P} is a primitive portrait, so that the multiplier $(f_{\widehat{c}}^{\circ p})'(\widehat{z})$ is equal to $+1$ for $\widehat{z} \in \mathcal{O}$, with $r = 1$. In this case, \widehat{z} is a fixed point of multiplicity two for the iterate $f_{\widehat{c}}^{\circ p}$, and splits into two nearby fixed points under perturbation. (It cannot have a higher multiplicity, since a fixed point of multiplicity $\mu > 2$ would have $\mu - 1 \geq 2$ attracting Leau-Fatou petals, each with at least one critical point in its basin, which is impossible for a quadratic map.) As c traverses a small loop around \widehat{c} , these two fixed points a priori may be (and in practice always will be) interchanged. However, if we loop twice around \widehat{c} , then each of these fixed points must return to its original position. Thus, if we introduce a new parameter u by the equation $c = \widehat{c} + u^2$, then we can choose these fixed points as holomorphic functions, $z_\iota = z_\iota(u)$ for $\iota = 1, 2$, with $z_1(0) = z_2(0) = \widehat{z}$. Evidently the u -plane is a two-fold branched cover of the c -parameter plane. Let $\lambda_\iota(u) = (f_c^{\circ p})'(z_\iota(u))$ be the multiplier for the orbit of z_ι , and note that $\lambda_1(0) = \lambda_2(0) = 1$. Since the holomorphic function $u \mapsto \lambda_1(u)$ cannot be constant, it takes on all values close to $+1$ as u varies through a neighborhood of 0.

Expanding the function $f_c^{\circ p}$ as a power series about its fixed point z_1 , we obtain

$$f_c^{\circ p}(z_1(u) + h) - z_1(u) = \lambda_1(u)h + a(u)h^2 + (\text{higher terms in } h) \quad (3)$$

for h and u close to zero, where $c = \widehat{c} + u^2$. Here the coefficient $a(u)$ is also a holomorphic function of u , with $a(0) \neq 0$ since the fixed point multiplicity is two. It follows that $a(u) \neq 0$ for u sufficiently small. Denoting the expression (3) by $g_u(h)$, and replacing the variable $h = z - z_1$ by $w = \alpha_u h$ where $\alpha_u = a(u)/\lambda_1(u)$, we see easily that the function

$$F_u(w) = \alpha_u g_u(w/\alpha_u)$$

has the required form (2). □

Proof of 4.2 in the Satellite Case. — We now suppose that $\hat{\lambda}$ is a primitive r -th root of unity, with $r > 1$. Then we can solve for the period p point $z = z(c)$ as a holomorphic function of c for c in some neighborhood of \hat{c} , with $z(\hat{c}) = \hat{z}$. Hence the multiplier $\lambda(c) = (f_c^{\circ p})'(z(c))$ will also be a holomorphic function of c , taking the value $\hat{\lambda} \in \sqrt[r]{1}$ when $c = \hat{c}$. Similarly $\lambda(c)^r$ is a holomorphic function, taking the value $\lambda(\hat{c})^r = 1$ when $c = \hat{c}$. This function $\lambda(c)^r$ clearly cannot be constant, so it takes all values close to $+1$ as c varies through a neighborhood of \hat{c} .

We will construct a sequence of holomorphic changes of variable which conjugate the map $z \mapsto f_c^{\circ p}(z)$ in a neighborhood of $z = z(c)$ to maps $h \mapsto g_{c,k}(h)$ in a neighborhood of $h = 0$, where $1 \leq k \leq r$, so that

$$g_{c,k}(h) = \lambda(c)h(1 + a_k(c)h^k + (\text{higher terms in } h))$$

for some constant $a_k(c)$. Here c can be any point in some neighborhood of \hat{c} . To begin the construction, let

$$g_{c,1}(h) = f^{\circ p}(z(c) + h) - z(c).$$

This certainly has the required properties. Now inductively set

$$g_{c,k+1}(h) = \phi^{-1} \circ g_{c,k} \circ \phi(h) \quad \text{where} \quad \phi(h) = h + bh^{k+1}$$

for $1 \leq k < r$. We claim that the constant $b = b(c)$ can be uniquely chosen so that $g_{c,k+1}$ will have the required form. In fact a brief computation shows that

$$g_{c,k+1}(h) = \lambda h(1 + (a + b - \lambda^k b)h^k + (\text{higher terms})).$$

But $\lambda^k \neq 1$ since λ is close to $\hat{\lambda}$, which is a primitive r -th root of unity with $1 \leq k < r$. Hence there is a unique choice of b so that $a + b - \lambda^k b = 0$, as required.

In particular, pushing this argument as far as possible, we can take $k = r$ and replace $f_c^{\circ p}$ near $z = z(c)$ by $g_{c,r}(h) = \lambda h(1 + ah^r + \dots)$ near $h = 0$. Hence we can replace $f_c^{\circ rp}$ near $z(c)$ by

$$g_{c,r}^{\circ r}(h) = \lambda^r h(1 + a' h^r + (\text{higher terms})),$$

where computation shows that $a' = (1 + \lambda^r + \lambda^{2r} + \dots + \lambda^{(r-1)r})a$. Here the coefficient a' of h^r must be non-zero when $\lambda = \hat{\lambda}$, and hence for λ close to $\hat{\lambda}$. For otherwise, the Leau-Fatou flowers around the points of the parabolic orbit would give rise to more than one periodic cycle of attracting petals for f_c . This is impossible, since each such cycle must contain a critical point, and a quadratic polynomial has only one critical point. Finally, after a scale change, replacing $g_{c,r+1}(h)$ by $F_c(w) = \alpha_c g_{c,r+1}^{\circ r}(w/\alpha_c)$ for suitably chosen α_c , we obtain simply

$$F_c^{\circ r}(w) = \lambda^r w(1 + w^r + (\text{higher terms in } w)),$$

as required. □

Proof of 4.1. — First note that we can choose a smooth path in parameter space so that the multiplier λ^r of Lemma 4.2 is real and belongs to some interval $(1, 1 + \eta)$. This follows easily from the fact that λ is a non-constant holomorphic function of c in the case $r > 1$, or of $u = \sqrt{c - \hat{c}}$ in the case $r = 1$. Note that the map $F^{\circ r}$ of 4.2 satisfies

$$|F^{\circ r}(w)| = \lambda^r \cdot |w| \cdot (1 + \operatorname{Re}(w^r) + (\text{higher terms})) \quad (4)$$

and

$$\arg(F^{\circ r}(w)) = \arg(w) + \operatorname{Im}(w^r) + (\text{higher terms}) \quad (5)$$

whenever λ^r is real and positive; and note also that $F^{\circ r}$ has a locally defined holomorphic inverse of the form

$$F^{-r}(w) = w (1 - w^r/\lambda^{2r} + (\text{higher terms}))/\lambda^r,$$

which satisfies

$$|F^{-r}(w)| = |w| (1 - \operatorname{Re}(w^r)/\lambda^{2r} + (\text{higher terms}))/\lambda^r \quad (4')$$

and

$$\arg(F^{-r}(w)) = \arg(w) - \operatorname{Im}(w^r)/\lambda^{2r} + (\text{higher terms}). \quad (5')$$




FIGURE 11. A repelling petal V_ε and attracting petal V'_ε for the map $F(w) \approx w + w^2$ (illustrating the primitive case, before perturbation).

As a representative repelling petal for $F^{\circ r}$ let us choose a small wedge shaped region V_ε described in polar coordinates by setting $w = \rho e^{2\pi i t}$ with $0 \leq \rho \leq \varepsilon$ and $|t| < 1/(8r)$. (Compare Figure 11 for the case $r = 1$.) If $\lambda^r \geq 1$ with λ^r sufficiently close to 1, it follows easily from (4') and (5') that V_ε maps into itself under F^{-r} , with all orbits converging towards the boundary fixed point at $w = 0$. If a dynamic ray for $f_{\hat{c}}$ lands at \hat{z} , then it must land through one of the r repelling petals, for example through the image of V_ε in the z -plane. For c sufficiently close to \hat{c} , this image must still contain a full segment, from some point z to $f_c^{\circ rp}(z)$, of the perturbed ray, hence this perturbed ray must still land at the repelling point which corresponds to $w = 0$.

Note that no new rays land at this point, after perturbation. There are only finitely many rays which have period p . But every dynamic ray of period p for $f_{\hat{c}}$ with angle

not in the set $\mathbf{A}_{\mathcal{P}}$ of angles for \mathcal{P} must land on some disjoint repelling point, and this condition will be preserved under perturbation. Thus the perturbed orbit, for $\lambda^r > 1$, still has portrait \mathcal{P} .

As an attracting petal for $F^{\circ r}$ we can choose the set $V'_\varepsilon = e^{\pi i/r} V_\varepsilon$ consisting of all $w = \rho e^{2\pi i t}$ with $0 \leq \rho \leq \varepsilon$ and $\frac{3}{8r} \leq t \leq \frac{5}{8r}$. If $\lambda^r > 1$ with λ^r close to 1, then using (4) and (5) we can check that $F^{\circ r}$ maps V'_ε into itself. However, the origin is a repelling point, so orbits cannot converge to it. In fact, if K is the compact set obtained from V'_ε by removing a very small neighborhood of the origin, then $F^{\circ r}$ maps K into its own interior. It follows easily that all orbits in $V'_\varepsilon \setminus \{0\}$ converge to an interior fixed point. This must be a strictly attracting point, and must correspond to an attracting orbit of period rp for the map f_c . \square

Corollary 4.3 (Parabolic Points as Root Points). — *If $f_{\hat{c}}$ has a parabolic orbit whose portrait \mathcal{P} is non-trivial, then \hat{c} must be equal to the root point $\mathbf{r}_{\mathcal{P}}$ of the \mathcal{P} -wake.*

Note. — The hypothesis that \mathcal{P} is non-trivial is actually redundant. (See 4.8.) It will be shown in 5.4 that every parabolic point is the root point of only one wake, so that the root point of the \mathcal{P} -wake always has portrait equal to \mathcal{P} .

Proof of 4.3. — Since $f_{\hat{c}}$ has a parabolic orbit with portrait \mathcal{P} , it certainly cannot have a repelling orbit with portrait \mathcal{P} . Hence it cannot be inside the \mathcal{P} -wake by 3.1. On the other hand, by 4.1 it must belong to the boundary of the \mathcal{P} -wake. By construction, the root point $\mathbf{r}_{\mathcal{P}}$ is the only boundary point of $W_{\mathcal{P}}$ which belongs to the Mandelbrot set. \square

Here is a complementary statement to 4.1, in the case $r > 1$.

Lemma 4.4 (A Deformation Breaking the Portrait). — *Under the hypothesis of 4.1, there also exists a smooth path of parameter values c , converging to \hat{c} , so that each f_c has an attracting orbit of period p , and a repelling orbit of period rp which lies on the boundary of its immediate basin. Furthermore, the dynamic rays with angles in $\mathbf{A}_{\mathcal{P}} = A_1 \cup \cdots \cup A_p$ all land on this repelling orbit.*

(Compare Figure 10, right.) For such values of c (still assuming that $r > 1$), it follows that there is *no* periodic orbit with portrait \mathcal{P} . Together with 4.1, this gives an alternative proof that \hat{c} is on the boundary of the \mathcal{P} -wake.

Proof. — The proof of 4.4 is completely analogous to the proof of 4.1, and will be left to the reader: One simply deforms so that $\lambda^r < 1$, instead of $\lambda^r > 1$. \square

The following assertion helps to make the statement of 4.2 more precise.

Lemma 4.5 (Local Uniqueness). — *Under the hypothesis of 4.2, there exist unique single valued functions $c = c(\lambda)$ and $z = z(\lambda)$, defined and holomorphic for λ in a neighborhood of $\hat{\lambda}$, so that $z(\lambda)$ is a periodic point of period p and multiplier λ for the*

map $f_{c(\lambda)}$, with $\widehat{c} = c(\widehat{\lambda})$ and $\widehat{z} = z(\widehat{\lambda})$. This function $c(\lambda)$ is univalent in the satellite case, but has a simple critical point at $\widehat{\lambda}$ in the primitive case.

The implications of this lemma for the geometry of the Mandelbrot set will be described in 6.1 and 6.2.

Proof of 4.5. — First consider the satellite case, with $\widehat{\lambda} \neq 1$. Then clearly the period p orbit and its multiplier $\lambda(c)$ depend smoothly on c throughout some neighborhood of \widehat{c} . We will show that the derivative $d\lambda/dc$ is non-zero at \widehat{c} . For otherwise, we could write

$$\lambda^r(c) = 1 + a(c - \widehat{c})^k + (\text{higher terms})$$

with $k \geq 2$. Hence we could vary c from \widehat{c} in two or more different directions so that $\lambda^r > 1$ and in two or more intermediate directions so that $\lambda^r < 1$. The former points would be within the \mathcal{P} -wake and the later points would be outside it; but this configuration is impossible by 3.1. Thus $d\lambda/dc \neq 0$, and it follows by the Inverse Function Theorem that the inverse mapping $\lambda \mapsto c(\lambda)$ is well defined and holomorphic throughout a neighborhood of $\widehat{\lambda}$, as required.

In the primitive case, the situation is different, but the proof is similar. In this case, setting $c = \widehat{c} + u^2$, we must express the multiplier λ_1 for one of the two nearby period p points as a holomorphic function of u , and show that the derivative $d\lambda_1/du$ is non-zero at $u = 0$. Otherwise, if the derivative $d\lambda_1(u)/du$ were equal to zero for $u = 0$, then we could write

$$\lambda_1(u) = 1 + a u^k + (\text{higher terms})$$

for some $k \geq 2$. It would follow that we could vary u from 0 in two or more different directions so that $\lambda_1 > 1$ and in two or more separating directions so that $\lambda_2 > 1$. All of these points would be within the \mathcal{P} -wake, but the rays landing on the periodic point z_1 would have to jump discontinuously so as to land on z_2 as we pass from $\lambda_1 > 1$ to $\lambda_2 > 1$, and such points of discontinuity must be outside the \mathcal{P} -wake. Even allowing for the fact that the u -plane is a two-fold covering of the c -plane, such a configuration is incompatible with 3.1. Therefore, λ_1 and u must determine each other holomorphically in a neighborhood of $\widehat{\lambda} \leftrightarrow 0$. In particular, it follows that the parameter value $c = \widehat{c} + u^2$ can be expressed as a holomorphic function of λ_1 , with a simple critical point at $\lambda_1 = \widehat{\lambda}$. \square

To conclude this section, we will prove that the portrait of a parabolic periodic point is always non-trivial. We will use a somewhat simplified form of the Hubbard tree construction to show that every parabolic orbit with ray period $rp \geq 2$ must have portrait with valence $v \geq 2$. First some general remarks about locally connected subsets of the plane.

Lemma 4.6 (A Canonical Retraction). — *Let $K \subset \mathbb{C}$ be compact, connected, locally connected, and full, and let U be a connected component of the interior of K . Then*

the closure \overline{U} is homeomorphic to the closed unit disk, and there is a unique retraction ρ_U from \mathbb{C} onto \overline{U} which carries each external ray, and also each connected component of the complement $K \setminus \overline{U}$, to a single point of the circle ∂U . There are at least two distinct external rays landing at a point $z_0 \in \partial U$ if and only if $K \setminus \{z_0\}$ is disconnected, or if and only if there is some connected component X of $K \setminus \overline{U}$ with $\rho_U(X) = \{z_0\}$.

Proof. — (Compare [D5].) The statement that \overline{U} is a disk follows easily from well known results of Carathéodory. Furthermore, according to Carathéodory, there is a unique retraction from \mathbb{C} onto K which maps each external ray to its landing point. Composing this with the retraction $K \rightarrow \overline{U}$ which maps each component X of $K \setminus \overline{U}$ to the unique intersection point $z_0 \in \overline{X} \cap \overline{U}$, we obtain the required retraction ρ_U .

For any such X , note that there must be at least one maximal open interval of angles t such that the ray \mathcal{R}_t^K lands in X . The endpoints of such a maximal interval are the angles for the required pair of rays landing on z_0 . Conversely, if there were two rays landing on z_0 but no component X attached in between, then there would be an entire open interval of angles t so that \mathcal{R}_t^K lands at z_0 . But this is impossible by a classical theorem of F. and M. Riesz. (See for example [M2, App. A].) \square

In particular, let $K = K(f)$ be the filled Julia set for a hyperbolic quadratic polynomial. (We are actually interested in the parabolic case, but will work first with the hyperbolic case, since that will suffice for our purposes, and since it is much easier to prove local connectivity in the hyperbolic case.)

Lemma 4.7 (The Dynamic Root Point). — *Suppose that $f = f_c$ has an attracting orbit of period $n \geq 2$. Let K be its filled Julia set, and let U_0 and $U_1 \subset K$ be the Fatou components containing the critical point 0 and the critical value c respectively. Then the canonical retraction $\rho_{U_1} : \mathbb{C} \rightarrow \overline{U}_1$ carries the component U_0 to the unique point $\mathbf{r}_c \in \partial U_1$ which is fixed by $f^{\circ n}$. Hence at least two dynamic rays land at this point.*

(See for example Figures 1, 6.) Following Schleicher, I will call \mathbf{r}_c the *dynamic root point* for the Fatou component U_1 .

Proof. — Let $U_0 \rightarrow U_1 \xrightarrow{\sim} U_2 \xrightarrow{\sim} \cdots \xrightarrow{\sim} U_n = U_0$ be the Fatou components containing the critical orbit. Then $f^{\circ n}$ maps each circle ∂U_j onto itself by an expanding map of degree two. Hence there is a canonical homeomorphism $a_j : \partial U_j \rightarrow \mathbb{R}/\mathbb{Z}$ which conjugates $f^{\circ n}$ to the angle doubling map on the standard circle. For each $z \in \mathbb{C} \setminus U_j$, the image $a_j(\rho_{U_j}(z))$ will be called the *internal angle* of the point z with respect to U_j . The map f from ∂U_j to ∂U_{j+1} preserves the internal angles of boundary points for $0 < j < n$, but doubles them for the case $j = 0$ of the critical component.

Define the *t-wake* $L_t(U_j)$ to be the set of all $z \in \mathbb{C} \setminus U_j$ with $a_j(\rho_{U_j}(z)) = t \in \mathbb{R}/\mathbb{Z}$. These wakes are pairwise disjoint sets with union equal to $\mathbb{C} \setminus U_j$. In general f maps to *t-wake* of U_j homeomorphically onto the *t-wake* of U_{j+1} for $0 < j < n$, and onto the *2t-wake* of U_{j+1} when $j = 0$. However, there is one exceptional value of t for each

U_j with $0 < j < n$. Namely, if the wake $L_t(U_j)$ contains the critical component U_0 then it certainly cannot map homeomorphically, and its image may be much larger than $L_t(U_{j+1})$.

Let $\mathcal{A}_j \subset \mathbb{R}/\mathbb{Z}$ be the finite set consisting of all angles $t \in \mathbb{R}/\mathbb{Z}$ such that the wake $L_t(U_j)$ contains one of the components U_k (where necessarily $j \neq k$). Then it follows that $\mathcal{A}_1 \subset \mathcal{A}_2 \subset \cdots \subset \mathcal{A}_n$ and $2\mathcal{A}_n = \mathcal{A}_0$. On the other hand, since K is full, the various U_j must be connected together in a tree-like arrangement (the *Hubbard tree*). There cannot be any cycles. Hence at least one of the \mathcal{A}_i must consist of a single angle. It follows easily that $\mathcal{A}_1 = \{0\}$, and the conclusion follows. \square

Corollary 4.8 (Parabolic Orbit Portraits are Non-Trivial). — *If c is any parabolic point of the Mandelbrot set other than $c = 1/4$, and if \mathcal{O} is the parabolic orbit for f_c , then at least two dynamic rays land on each point of \mathcal{O} .*

(This is just a restatement of Theorem 1.4 of §1.)

Proof. — In the satellite case this is trivially true, while in the primitive case it follows from 4.7, using 4.1 to pass from the parabolic to the hyperbolic case. This completes the proof of Theorem 1.4. \square

5. The Period n Curve in (Parameter \times Dynamic) Space

It is convenient to define a sequence of numbers $\nu_2(n)$ inductively by the formula

$$2^k = \sum_{n|k} \nu_2(n), \quad \text{or} \quad \nu_2(k) = \sum_{n|k} \mu(k/n) 2^n,$$

to be summed over all divisors $n \geq 1$ of k , where $\mu(k/n) \in \{\pm 1, 0\}$ is the Möbius function. In fact we will be mainly interested in the quotients $\nu_2(n)/2$ and $\nu_2(n)/n$. The first few values are

n	1	2	3	4	5	6	7	8	9	10
$\nu_2(n)/2$	1	1	3	6	15	27	63	120	252	495
$\nu_2(n)/n$	2	1	2	3	6	9	18	30	56	99

Define the *period n curve* $\text{Per}_n \subset \mathbb{C}^2$ to be the locus of zeros of the polynomial $Q_n(c, z)$ which is defined by the formula

$$f_c^{\circ k}(z) - z = \prod_{n|k} Q_n(c, z), \quad \text{or} \quad Q_k(c, z) = \prod_{n|k} (f_c^{\circ k}(z) - z)^{\mu(k/n)},$$

taking the product over all divisors n of k . For example,

$$Q_1(c, z) = z^2 + c - z, \quad Q_2(c, z) = \frac{(z^2 + c)^2 + c - z}{z^2 + c - z} = z^2 + z + c + 1.$$

Note that each point $(c, z) \in \text{Per}_n$ determines a periodic orbit

$$z = z_0 \mapsto z_1 \mapsto \cdots \mapsto z_n = z_0$$

for the map f_c . Let $\lambda_n = \lambda_n(c, z) = \partial f_c^{\circ n}(z)/\partial z = 2^n z_1 \cdots z_n$. For a generic choice of c , this orbit has period exactly n , and λ_n is the multiplier. However, if z is a parabolic periodic point for f_c with ray period $n = rp > p$, then (c, z) belongs both to Per_n with $\lambda_n = 1$, and to Per_p with $\lambda_p \in \sqrt[p]{1}$. (In fact, the two curves Per_n and Per_p intersect transversally at (c, z) .)

Remarks. — Compare [M4] for a somewhat analogous discussion for cubic polynomials. The fact that Q_n is really a polynomial can be verified by expressing $f_c^{\circ k}(z) - z$ as a product of irreducible polynomials, and checking that each of these irreducible factors has a well defined period n dividing k . The factors are all distinct since $\partial(f_c^{\circ j}(z) - z)/\partial z \neq 0$ at every zero of this polynomial when $|c|$ is large. It is shown in [Bou], and also in [S1], [LS], that the algebraic curve Per_n (or the polynomial Q_n) is actually irreducible; however, we will not make any use of that fact.

Lemma 5.1 (Properties of the Period n Curve). — *This algebraic curve $\text{Per}_n \subset \mathbb{C}^2$ is non-singular. The projection $(c, z) \mapsto c$ is a proper map of degree $\nu_2(n)$ from Per_n to the parameter plane, while the projection $(c, z) \mapsto z$ is a proper map of degree $\nu_2(n)/2$ to the dynamic plane. Finally, the function $(c, z) \mapsto \lambda_n(c, z)$ is a proper map of degree $n\nu_2(n)/2$ to the λ_n -plane.*

Note that the cyclic group of order n , which we will denote by \mathbb{Z}_n , acts on Per_n , a generator carrying (c, z) to $(c, f_c(z))$.

Lemma 5.2 (Properties of $\text{Per}_n/\mathbb{Z}_n$). — *The quotient $\text{Per}_n/\mathbb{Z}_n$ is a smooth algebraic curve consisting of all pairs (c, \mathcal{O}) where \mathcal{O} is a periodic orbit for f_c which is either non-parabolic of period n , or parabolic with attracting petals of period n . At any point where $\lambda_n \neq 1$, the coordinate c can be used as local uniformizing parameter, while in a neighborhood of a point with $\lambda_n = 1$, the multiplier $\lambda_n = \lambda_n(c, z)$ serves as a local uniformizing parameter for this curve. The projection maps $(c, \mathcal{O}) \mapsto c$ and $(c, \mathcal{O}) \mapsto \lambda_n$ are proper, with degrees $\nu_2(n)/n$ and $\nu_2(n)/2$ respectively.*

The proof that Per_n and $\text{Per}_n/\mathbb{Z}_n$ are non-singular will be divided into three cases, as follows.

Generic Case. — First consider a point $(\hat{c}, \hat{z}) \in \text{Per}_n$ with $\lambda_n(\hat{c}, \hat{z}) \neq 1$. Then, by the Implicit Function Theorem, we can solve the equation $f_c^{\circ n}(z) = z$ locally for z as a smooth function of c . It follows that both of the curves Per_n and $\text{Per}_n/\mathbb{Z}_n$ are locally smooth, with c as local uniformizing parameter.

Primitive Parabolic Case. — Now consider a point $(\hat{c}, \hat{z}) \in \text{Per}_n$ with $\lambda_n(\hat{c}, \hat{z}) = 1$, where \hat{z} has period exactly n under $f_{\hat{c}}$. According to the proof of 4.5, if we set $c = \hat{c} + u^2$, then both z and $\lambda_n = \lambda_n(c, z)$ can be expressed locally as smooth functions of u with $d\lambda_n/du \neq 0$. It follows that both Per_n and $\text{Per}_n/\mathbb{Z}_n$ are locally smooth at this point, and that we can use either u or λ_n as local uniformizing parameter.

(Similarly $dz/du \neq 0$, so we could use z as local uniformizing parameter for Per_n . However dc/du is zero when $u = 0$, so c cannot be used as local parameter.)

Satellite Parabolic Case. — Again suppose that $\lambda_n(\widehat{c}, \widehat{z}) = 1$, but now assume that the period p of \widehat{z} is strictly less than the ray period $n = rp$. For c near \widehat{c} , let $z = z(c)$ be the equation of the unique period p point near \widehat{z} . Using the change of variable $w = \alpha(z - z(c)) + (\text{higher terms})$ of 4.2, the map $f_c^{\circ n}$ corresponds to

$$w \mapsto F^{\circ r}(w) = \lambda^r w (1 + w^r + (\text{higher terms})), \quad (6)$$

where $\lambda = \lambda(w)$ is the multiplier of this period p orbit. The equation for a fixed point is $w = \lambda^r w (1 + w^r + (\text{higher terms}))$. Dividing by w (since we want the fixed point with $w \neq 0$ or with $z \neq z(c)$), this becomes

$$1 = \lambda^r (1 + w^r + (\text{higher terms})) \quad \text{or} \quad \lambda^r = 1 - w^r + (\text{higher terms}).$$

Thus we can express λ as a holomorphic function of w , with a critical point at $w = 0$. Therefore, by 4.5, we can also express c as a holomorphic function of w . Since w is defined as a holomorphic function of z and c with $\partial w / \partial z \neq 0$, it follows that Per_n is locally smooth with local uniformizing parameter z or w .

Now note that there is a unique local change of coordinate $w \mapsto \phi(w)$ with $\phi'(0) = 1$ so that $\lambda^r = 1 - \phi(w)^r$. Since the expression $\phi(w)^r$ is invariant under the \mathbb{Z}_n action of Per_n , it follows easily that this action can be described by the formula $\phi(w) \mapsto \widehat{\lambda} \phi(w)$. It follows that $\phi(w)^r = 1 - \lambda^r$ is a local uniformizing parameter for the quotient curve $\text{Per}_n / \mathbb{Z}_n$. Therefore, either λ or c can also be taken as local uniformizing parameter. In particular, it follows that the multiplier λ_n of the period $n = rp$ orbit can be expressed as a smooth function of the multiplier $\lambda = \lambda_p$ of the period p orbit. Note that

$$d\lambda_n / d(\lambda^r) = -r \quad (7)$$

at the parabolic point. (Compare [CM, (4.3)].) This can be verified by direct computation from (6), or by using the holomorphic fixed point formula [M2] for the function $f_c^{\circ n}$ to show that the expression

$$\frac{r}{1 - \lambda_n} + \frac{1}{1 - \lambda^r}$$

depends smoothly on the parameter c throughout some neighborhood of the parabolic point. Therefore λ_n can also be used as local uniformizing parameter for $\text{Per}_n / \mathbb{Z}_n$.

The degrees of the various projection maps can easily be computed algebraically, by counting solutions to the appropriate polynomial equations. Here is a more geometric argument, which also provides a quite explicit description of the ends of the curve Per_n , and hence proves that these mappings are proper. Let us consider the limiting case as $|c| \rightarrow \infty$. Setting $c = -v^2$ with $|v| > 2$, let $\pm\Delta$ be the open disk of radius 1 centered at $\pm v$. It is not difficult to check that both Δ and $-\Delta$ map holomorphically onto a disk $f(\Delta)$ which contains $\overline{\Delta} \cup (-\overline{\Delta})$. The (filled) Julia set K can then be

described explicitly as follows. Given an arbitrary sequence of signs $\varepsilon_0, \varepsilon_1, \dots$, there is one and only one orbit $z_0 \mapsto z_1 \mapsto \dots$ in K with $z_j \in \varepsilon_j \Delta$ for every $j \geq 0$. This is proved using the Poincaré metric for the inverse maps $f(\Delta) \rightarrow \pm \Delta \subset f(\Delta)$. In particular, the number of solutions of period n is equal to the number of sign sequences of period n , which is easily seen to be $\nu_2(n)$. Thus the degree of the projection to the c -plane is $\nu_2(n)$. It follows also that the product $z_1 \cdots z_n = \lambda/2^n$ is given asymptotically by

$$\lambda/2^n \sim \pm z^n \sim \pm v^n = \pm(-c)^{n/2} \quad \text{as} \quad |v| \rightarrow \infty.$$

Thus the degree of the projection to the λ -plane is n times the degree of the projection to the z -plane, and is $n/2$ times the degree of the projection to the c -plane. \square

Thus we have a diagram of smooth algebraic curves and proper holomorphic maps with degrees as indicated:

$$\begin{array}{ccc} \text{Per}_n & \xrightarrow{n} & \text{Per}_n / \mathbb{Z}_n & \xrightarrow{\nu_2(n)/2} & \lambda_n\text{-plane} \\ \downarrow \nu_2(n)/2 & & \downarrow \nu_2(n)/n & & \\ z\text{-plane} & & c\text{-plane} & & \end{array}$$

For a generic choice of c , it follows that the map f_c has exactly $\nu_2(n)/n$ periodic orbits of period n , while for generic choice of λ_n there are exactly $\nu_2(n)/2$ pairs (c, \mathcal{O}) consisting of a parameter value c and a period n orbit of multiplier λ_n for the map f_c . The discussion shows that the correspondence $(c, \mathcal{O}) \mapsto (c, \lambda_n)$ yields a smooth immersion of $\text{Per}_n / \mathbb{Z}_n$ into \mathbb{C}^2 . (Caution: Presumably some f_c may have two different period n orbits with the same multiplier, so this immersion may have self-intersections.)

Corollary 5.3 (Counting Parabolic Points). — *The number of parabolic points in the Mandelbrot set with ray period $rp = n$ is equal to $\nu_2(n)/2$.*

Proof. — This is the same as the number of points in the pre-image of $+1$ under the projection $(c, \mathcal{O}) \mapsto \lambda_n(c, \mathcal{O})$ from $\text{Per}_n / \mathbb{Z}_n$ to the λ_n -plane. According to 5.2, the degree of this projection is $\nu_2(n)/2$, and $+1$ is a regular value. The conclusion follows. \square

We are now ready to prove the main results, as stated in §1.

Corollary 5.4. — *There are exactly two parameter rays whose angles are periodic under doubling landing at each parabolic point $\hat{c} \neq 1/4$. Hence distinct wakes have distinct root points; and for each non-trivial portrait \mathcal{P} , the root point of the \mathcal{P} -wake has a parabolic orbit with portrait \mathcal{P} .*

(For angles which are not periodic, compare 9.4.)

Corollary 5.5. — *Every parameter ray \mathcal{R}_t^M whose angle has period $n \geq 2$ under doubling forms one of the two boundary rays for one and only one wake $W_{\mathcal{P}}$, where \mathcal{P} is some portrait with ray period n .*

Proof of 5.4 and 5.5. — According to 5.3, the number of parabolic points \hat{c} with ray period $n \geq 2$ is equal to $\nu_2(n)/2$, and according to Theorem 1.4 each such point is the landing point of at least two rays, which necessarily have ray period n . Thus altogether there are at least $\nu_2(n)$ distinct rays of period n . On the other hand, since the map $t \mapsto 2^n t \pmod{\mathbb{Z}}$ has $2^n - 1$ fixed points, it follows inductively that the number of angles with period exactly $n \geq 2$ is precisely equal to $\nu_2(n)$. Thus there cannot be more than two rays landing at any such point \hat{c} . It follows that \hat{c} is the root point of at most one wake. For if \hat{c} were the root point of two different wakes, then (even if they shared a boundary ray) it would be the landing point for at least three different parameter rays. Using 4.3, it now follows that each such \hat{c} is the root point $\mathbf{r}_{\mathcal{P}}$ for exactly one wake $W_{\mathcal{P}}$, and furthermore that each $f_{\mathbf{r}_{\mathcal{P}}}$ has a parabolic orbit with portrait \mathcal{P} .

Here we have assumed that $n \geq 2$. However, for $n = 1$ there is clearly just one parameter ray $\mathcal{R}_0^M = (1/4, \infty)$ which is fixed under doubling, and its landing point $\hat{c} = 1/4$ is the unique parabolic point with ray period $n = 1$. This completes the proof of 5.4 and 5.5. Clearly Theorems 1.2 and 1.5, as stated in §1, follow immediately. \square

To conclude this section, here is a more explicit description of the first few period n curves:

Period 1. — The curve $\text{Per}_1 = \text{Per}_1/\mathbb{Z}_1 \cong \mathbb{C}$ can be identified with the λ_1 -plane. It is a 2-fold branched cover of the c -plane, ramified at the root point $r_{\{\{0\}\}} = 1/4$, and can be described by the equations $z = \lambda_1/2$, $c = z - z^2$. Note that the unit disk $|\lambda_1| < 1$ in the λ_1 -plane maps homeomorphically onto the region bounded by the cardioid in the c -plane.

Period 2. — The quotient $\text{Per}_2/\mathbb{Z}_2 \cong \mathbb{C}$ can be identified either with the λ_2 -plane or with the c -plane, where $\lambda_2 = 4(1+c)$. The curve $\text{Per}_2 \cong \mathbb{C}$ is a 2-fold branched cover with coordinate z , branched at the point $\lambda_2 = 1$ which corresponds to the period 2 root point $c = r_{\mathcal{P}} = -3/4$ with portrait $\mathcal{P} = \{\{1/3, 2/3\}\}$. It is described by the equation $z^2 + z + (c+1) = 0$, with \mathbb{Z}_2 -action $z \leftrightarrow f_c(z) = -z - 1$.

Period 3. — (See [GF].) The quotient $\text{Per}_3/\mathbb{Z}_3 \cong \mathbb{C}$ can be identified with a 2-fold branched cover of the c -plane, branched at the root point $r_{\mathcal{P}} = -7/4$ of the real period 3 component, where $\mathcal{P} = \{\{3/7, 4/7\}, \{6/7, 1/7\}, \{5/7, 2/7\}\}$. If we choose a parameter u on this quotient by setting $c = -(u^2 + 7)/4$, then computation shows that the multiplier is given by the cubic expression $\lambda_3 = u^3 - u^2 + 7u + 1$. The curve Per_3 itself is conformally isomorphic to a thrice punctured Riemann sphere. It can be

described as a 3-fold cyclic branched cover of this u -plane, branched with ramification index 3 at the two points $u = (1 \pm \sqrt{-27})/2$ where $\lambda_3 = 1$.

6. Hyperbolic Components

By definition, a *hyperbolic component* H of *period* n in the Mandelbrot set is a connected component of the open set consisting of all parameter values c such that f_c has a (necessarily unique) attracting orbit of period n . We will first study the geometry of a hyperbolic component near a parabolic boundary point.

Lemma 6.1 (Geometry near a Satellite Boundary Point). — *Let \hat{c} be a parabolic point with orbit portrait \mathcal{P} having ray period $rp > p$. Then \hat{c} lies on the boundary of exactly two hyperbolic components. One of these has period rp and lies inside the \mathcal{P} -wake, while the other has period p and lies outside the \mathcal{P} -wake. Locally the boundaries of these components are smooth curves which meet tangentially at \hat{c} .*




FIGURE 12. Detail of the Mandelbrot boundary, showing the rays landing at the root points of a primitive period 4 component and a satellite period 12 component.

Proof. — According to 4.1, \hat{c} lies on the boundary of a hyperbolic component H_{rp} of period rp which lies inside the \mathcal{P} -wake, while according to 4.4 it lies on the boundary of a component H_p of period p which lies outside the \mathcal{P} -wake. Let \mathcal{O}_{rp} and \mathcal{O}_p be the associated periodic orbits, with multipliers λ_{rp} and λ_p . According to 4.5, the

multiplier λ_p can be used as a local uniformizing parameter for the c -plane near \hat{c} . Therefore the boundary ∂H_p , with equation $|\lambda_p| = 1$, is locally smooth. Similarly, it follows from equation (7) of §5, that we can take λ_{rp} as local uniformizing parameter, so the locus $|\lambda_{rp}| = 1$ is also locally smooth. These two boundary curves are necessarily tangent to each other since the two hyperbolic components cannot overlap, or by direct computation from (7).

To see that there are no other components with \hat{c} as boundary point, first note that all periodic orbits for the map $f_{\hat{c}}$, other than its designated parabolic orbit, must be strictly repelling. For any orbit with multiplier $|\lambda| \leq 1$ must either attract the critical orbit (in the attracting or parabolic case) or at least be in the ω -limit set of the critical orbit (in the Cremer case), or have Fatou component boundary in this ω -limit set (in the Siegel disk case). Since the unique critical orbit converges to the parabolic orbit, all other periodic orbits must be repelling.

Now choose some large integer N . If we choose c sufficiently close to \hat{c} , then all repelling periodic orbits of period $\leq N$ for $f_{\hat{c}}$ will deform to repelling periodic orbits of the same period for f_c . Thus any non-repelling orbit of period $\leq N$ for f_c must be one of the two orbits \mathcal{O}_p and \mathcal{O}_{rp} which arise from perturbation of the parabolic orbit. In other words, any hyperbolic component H' of period $\leq N$ which intersects some small neighborhood of \hat{c} must be either H_p or H_{rp} . In particular, any hyperbolic component which has \hat{c} as boundary point must coincide with either H_p or H_{rp} . \square

By definition, the component H_{rp} is a *satellite* of H_p , attached at the parabolic point \hat{c} . (It follows from (7) that $|d\lambda_{rp}/d\lambda_p| = r^2$ at \hat{c} , so to a first approximation the component H_p is r^2 times as big as its satellite H_{rp} . Compare [CM].)

Lemma 6.2 (Geometry near a Primitive Boundary Point). — *If the portrait \mathcal{P} of the parabolic point \hat{c} has ray period $rp = p$, then \hat{c} lies on the boundary of just one hyperbolic component H , which has period p and lies inside the \mathcal{P} -wake. The boundary of H near \hat{c} is a smooth curve, except for a cusp at the point \hat{c} itself.*

Proof. — As in the proof of 4.2, we set $c = \hat{c} + u^2$ and find a period p point $z(u)$ with multiplier $\lambda(u)$ which depends smoothly on u , with $d\lambda/du \neq 0$. Hence the locus $|\lambda(u)| = 1$ is a smooth curve in the u -plane, while its image in the c -plane has a cusp at $c = \hat{c}$. The rest of the argument is completely analogous to the proof of 6.1. \square

Lemma 6.3 (The Root Point of a Hyperbolic Component). — *Every parabolic point of ray period $n = rp$ is on the boundary of one and only one hyperbolic component of period n . Conversely, every hyperbolic component of period n has one and only one parabolic point of ray period n on its boundary. In this way, we obtain a canonical one-to-one correspondence between parabolic points and hyperbolic components in parameter space.*

Proof. — The first statement follows immediately from 6.1 and 6.2. Conversely, if H is a hyperbolic component of period n , then we can map H holomorphically into the open unit disk \mathbf{D} by sending each $c \in H$ to the multiplier of the unique attracting orbit for f_c . In order to extend to the closure \overline{H} , it is convenient to lift to the curve $\text{Per}_n/\mathbb{Z}_n$, using the proper holomorphic map $(c, \mathcal{O}) \mapsto c$ of §5. Evidently H lifts biholomorphically to an open set $H^\natural \subset \text{Per}_n/\mathbb{Z}_n$, which then maps holomorphically to the λ_n -plane under the projection $(c, \mathcal{O}) \mapsto \lambda_n(c, \mathcal{O})$. (Here H^\natural is a connected component of the set of (c, \mathcal{O}) such that \mathcal{O} is an attracting period n orbit for f_c .) Since the projection to the λ_n -plane is open and proper, it follows easily that the closure \overline{H}^\natural maps *onto* the closed disk $\overline{\mathbf{D}}$. In particular, there exists a point $(\widehat{c}, \widehat{\mathcal{O}})$ of \overline{H}^\natural with $\lambda_n(\widehat{c}, \widehat{\mathcal{O}}) = +1$. Evidently this \widehat{c} is a parabolic boundary point of H with ray period dividing n , and it follows from 6.1 and 6.2 that it must have ray period precisely n .

According to 4.7, for each $c \in H$ there is a unique repelling orbit of lowest period on the boundary of the immediate basin for the attracting orbit of f_c . Furthermore, according to 4.1, the portrait $\mathcal{P} = \mathcal{P}_H$ for this orbit is the same as the portrait for the parabolic orbit of $f_{\widehat{c}}$. Since there is only one parabolic point with specified portrait by Theorem 1.2, this proves that there can only one such point $\widehat{c} \in \partial H$. \square

Definition. — This distinguished parabolic point on the boundary ∂H of a hyperbolic component is called the *root point* of the hyperbolic component H . We know from 1.2 and 1.4 that the parabolic points of ray period n can be indexed by the non-trivial orbit portraits of ray period n . Hence the hyperbolic components of period n can also be indexed by non-trivial portraits of ray period n . We will write $H = H_{\mathcal{P}}$ (or $\mathcal{P} = \mathcal{P}_H$) if H is the hyperbolic component with root point $r_{\mathcal{P}}$. We will say that H is a *primitive component* or a *satellite component* according as the associated portrait is primitive or satellite.

Remark 6.4. — Of course there are many other parabolic points in ∂H . For each root of unity $\mu = e^{2\pi i q/s} \neq 1$ a similar argument shows that there is at least one point $(\widehat{c}_\mu, \mathcal{O}_\mu) \in \partial H^\natural$ with $\lambda_n(\widehat{c}_\mu, \mathcal{O}_\mu) = \mu$. In fact the following theorem implies that \widehat{c}_μ is unique. This \widehat{c}_μ is the root point for a hyperbolic component H' of period $sn > n$, with associated orbit portrait \mathcal{P}' of period n and rotation number q/s . By definition, \mathcal{P}' is the (q/s) -satellite of \mathcal{P} , and H' is the (q/s) -satellite of H .

We next prove the following basic result of Douady and Hubbard. Again let H be a hyperbolic component of period n and let $H^\natural \subset \text{Per}_n/\mathbb{Z}_n$ be the set of pairs (c, \mathcal{O}) with $c \in H$, where \mathcal{O} is the attracting orbit for f_c .

Theorem 6.5 (Uniformization of Hyperbolic Components). — *The closure \overline{H} is homeomorphic to the closed unit disk $\overline{\mathbf{D}}$. In fact there is a canonical homeomorphism*

$$\overline{\mathbb{D}} \cong \overline{H}^\natural \rightarrow \overline{H}$$

which carries each point $\lambda \neq 1$ in \mathbb{D} to the unique point $c \in \overline{H}$ such that f_c has a period n orbit of multiplier λ . This homeomorphism extends holomorphically over a neighborhood of \mathbb{D} , with just one critical point $1 \in \mathbb{D}$ mapping to the root point $\hat{c} \in H$ in the primitive case, and with no critical points in the satellite case. The closures of the various hyperbolic components are pairwise disjoint, except for the tangential contact between a component and its satellite as described in 6.1.

Proof. — Recall that $\lambda_n : \text{Per}_n/\mathbb{Z}_n \rightarrow \mathbb{C}$ is a proper holomorphic map of degree $\nu_2(n)/2$. We will first show that there are no critical values of λ_n within the closed unit disk \mathbb{D} . This will imply that the inverse image $\lambda_n^{-1}(\mathbb{D})$ is the disjoint union of $\nu_2(n)/2$ disjoint sets \overline{H}^\natural , each of which maps diffeomorphically onto \mathbb{D} . First note that there are no critical values of λ_n on the boundary circle $\partial\mathbb{D}$. In the case of a root of unity $\mu \in \partial\mathbb{D}$, every (c, \mathcal{O}) with $\lambda_n(c, \mathcal{O}) = \mu$ must be parabolic, and it follows from 6.1 and 6.2 that the derivative of λ_n at (c, \mathcal{O}) is non-zero. Consider then a point $(\hat{c}, \mathcal{O}) \in \partial H^\natural$ such that $\lambda_n(\hat{c}, \mathcal{O})$ is not a root of unity. According to 5.2, we can use c as local uniformizing parameter throughout a neighborhood of (\hat{c}, \mathcal{O}) . If this were a critical point of λ_n , then it would follow that we could find two different line segments emerging from \hat{c} which map into \mathbb{D} , separated by two line segments which map outside of \mathbb{D} . In other words, one of the following two possibilities would have to occur.

Case 1. — There are two different hyperbolic components with \hat{c} as non-root boundary point. Each of these components must have a root point, and be contained in its associated wake. But these two components cannot be separated by any rational parameter ray, hence each one must be contained in the wake of the other, which is impossible.

Case 2. — The single hyperbolic component H must approach \hat{c} from two different directions, separated by two directions which lie outside of H . In other words. There must be a simple closed loop $L \subset \overline{H}$ which encloses points lying outside of \overline{H} . Now the collection of iterates $f_c^{\circ k}(0)$ must be uniformly bounded for $c \in L$, and hence also for all c in the region bounded by L . Thus this entire region must lie within the interior of the Mandelbrot set, which is impossible since this region contains parabolic points.

Thus both cases are impossible, and λ_n must be locally injective near the boundary of H^\natural . It follows easily that H^\natural maps onto \mathbb{D} by a proper map of some degree $d \geq 1$, and similarly that the boundary ∂H^\natural wraps around the boundary circle $\partial\mathbb{D}$ exactly d times. Now a counting argument shows that this degree is $+1$. In fact the number of H or H^\natural of period n is equal to $\nu_2(n)/2$ by 6.3 and 5.3. Since the degree of the map λ_n on $\text{Per}_n/\mathbb{Z}_n$ is also $\nu_2(n)/2$ by 5.2, it follows that each H^\natural must map with degree $d = 1$. Therefore λ_n maps each \overline{H}^\natural biholomorphically onto \mathbb{D} .

Next consider the projection $(c, \mathcal{O}) \mapsto c$ from the compact set \overline{H}^h onto \overline{H} . This is one-to-one, and hence a homeomorphism, by a theorem of Douady and Hubbard which asserts that a polynomial of degree d can have at most $d-1$ non-repelling cycles. (Compare [Sh1]. Alternatively, it follows from the classical Fatou-Julia theory that a polynomial with one critical point can have at most one attracting cycle. If two distinct points of ∂H^h mapped to a single point of ∂H , then, as in Case 2 above, a path between these points in H^h would map to a loop in \overline{H} which could enclose no boundary points of H , leading to a contradiction.)

According to 5.2, the parameter c can be used as local uniformizing parameter for $\text{Per}_n/\mathbb{Z}_n$ unless $\lambda_n = 1$. Hence the only possible critical value for the projection $\overline{H}^h \rightarrow \overline{H}$ is the root point. In fact, by 6.1 and 6.2, the root point is actually a critical value if and only if H is a primitive component.

Finally suppose that two different hyperbolic components have a common boundary point. If this boundary point is parabolic, then one of these components must be a satellite of the other by 6.1 and 6.2. If the point were non-parabolic, then the argument of Case 1 above would yield a contradiction. This completes the proof of 6.5. \square

7. Orbit Forcing

Recall that an orbit portrait is *non-trivial* if either it has valence $v \geq 2$, or it is the zero portrait $\{\{0\}\}$. The following statement follows easily from 1.3. However, it seems of interest to give a direct and more constructive proof; and the methods used will be useful in the next section.

Lemma 7.1 (Orbit Forcing). — *Let \mathcal{P} and \mathcal{Q} be distinct non-trivial orbit portraits. If their characteristic arcs satisfy $I(\mathcal{P}) \subset I(\mathcal{Q})$, then every f_c with a (repelling or parabolic) orbit of portrait \mathcal{P} must also have a repelling orbit of portrait \mathcal{Q} .*

Compare Figure 5, and see 1.3 and for further discussion. The proof of 7.1 begins as follows.

Puzzle Pieces. — Recall from 2.10 that the pv rays landing on a periodic orbit for $f = f_c$ separate the dynamic plane into $pv - p + 1$ connected components, the closures of which are called the (unbounded) *preliminary puzzle pieces* associated with the given orbit portrait. (As in [K], we work with puzzle pieces which are closed but not compact. The associated *bounded* pieces can be obtained by intersecting each unbounded puzzle piece with the compact region enclosed by some fixed equipotential curve.)

Most of these preliminary puzzle pieces Π have the *Markov property* that f maps Π homeomorphically onto some union of preliminary puzzle pieces. However, the puzzle piece containing the critical point is exceptional: Its image under f covers the critical value puzzle piece twice, and also covers some further puzzle pieces once. To obtain

a modified puzzle with more convenient properties, we will subdivide this exceptional piece into two connected sub-pieces.

Let Π_1 be the preliminary puzzle piece containing the critical value. Then $\partial\Pi_1$ consists of the two rays whose angles bound the characteristic arc for \mathcal{P} , together with their common landing point, say z_1 . The pre-image $\Pi_0 = f^{-1}(\Pi_1)$ is bounded by two rays landing at the point $z_0 = f^{-1}(z_1) \cap \mathcal{O}$, together with two rays landing at the symmetric point $-z_0$. Note that Π_0 is a connected set containing the critical point, and that the map f from Π_0 onto Π_1 is exactly two-to-one, except at the critical point 0, which maps to c .

The pv rays landing on \mathcal{O} , together with these two additional rays landing on $-z_0$, cut the complex plane up into $pv - p + 2$ closed subsets which we will call the pieces of the corrected puzzle associated with \mathcal{P} . These will be numbered as $\Pi_0, \Pi_1, \dots, \Pi_{pv-p+1}$, with Π_0 and Π_1 as above. The central piece Π_0 will be called the *critical puzzle piece*, and Π_1 will be called the *critical value puzzle piece*. This corrected puzzle satisfies the following.

Modified Markov Property. — *The puzzle piece Π_0 maps onto Π_1 by a 2-fold branched covering, while every other puzzle piece maps homeomorphically onto a finite union of puzzle pieces.*

We can represent the allowed transitions by a *Markov matrix* M_{ij} , where

$$M_{ij} = \begin{cases} 1 & \text{if } \Pi_i \text{ maps homeomorphically, with } f(\Pi_i) \supset \Pi_j \\ 0 & \text{if } f(\Pi_i) \text{ and } \Pi_j \text{ have no interior points in common,} \end{cases}$$

and where $M_{01} = 2$ since Π_0 double covers Π_1 . Since f is quadratic, note that the sum of entries in any column is equal to 2. Equivalently, this same data can be represented by a *Markov graph*, with one vertex for each puzzle piece, and with M_{ij} arrows from the i -th vertex to the j -th.

As an example, for the puzzle shown in Figure 13, we obtain the Markov graph of Figure 14, or the following Markov matrix

$$[M_{ij}] = \begin{bmatrix} 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{bmatrix} \quad (8)$$

To illustrate the idea of the proof of 7.1, let us show that any f having an orbit with this portrait \mathcal{P} must also have a repelling orbit with portrait $\mathcal{Q} = \{\{1/7, 2/7, 4/7\}\}$. (Compare the top implication in Figure 5.) Inspecting the next to last row of the matrix (8), we see that $f(\Pi_4) = \Pi_0 \cup \Pi_4 \cup \Pi_5$. Therefore, there is a branch g of f^{-1} which maps the interior of Π_4 holomorphically onto some proper subset of itself.




FIGURE 13. Julia set with a parabolic orbit of period four with characteristic arc $I(\mathcal{P}) = (3/15, 4/15)$, showing the six corrected puzzle pieces; and a corresponding schematic diagram. (For the corresponding preliminary puzzle, see the top of Figure 5.)




FIGURE 14. Markov graph associated with the matrix (8), with one vertex for each puzzle piece. Since f is quadratic, there are two arrows pointing to each vertex.

This mapping g must strictly decrease the Poincaré metric for the interior of Π_4 . On the other hand, it is easy to check that the $1/7$, $2/7$ and $4/7$ rays are all contained in the interior of Π_4 . Hence their landing points, call them w_1 , w_2 and w_3 , are also contained in Π_4 , necessarily in the interior, since the points of $K \cap \partial\Pi_4$ have period four. Now

$$g : w_1 \mapsto w_3 \mapsto w_2 \mapsto w_1,$$

and all positive distances are strictly decreased. Thus if the distance from w_i to w_j were greater than zero, then applying g three times we would obtain a contradiction. This proves that $w_1 = w_2 = w_3$, as required. This fixed point must be repelling, since g clearly cannot be an isometry.

A similar argument proves the following statement. Suppose that $f = f_c$ has an orbit \mathcal{O} with some given portrait \mathcal{P} . By a *Markov cycle* for \mathcal{P} we will mean an infinite sequence of non-critical puzzle pieces $\Pi_{i_1}, \Pi_{i_2} \dots$ which is periodic, $i_j = i_{j+m}$ with period $m \geq 1$, and which satisfies $f(\Pi_{i_\alpha}) \supset \Pi_{i_{\alpha+1}}$, so that $M_{i_\alpha i_{\alpha+1}} = 1$, for every α modulo m .

Lemma 7.2 (Realizing Markov Cycles). — *Given such a Markov cycle, there is one and only one periodic orbit $z_1 \mapsto \dots \mapsto z_m$ for f_c with period dividing m so that each z_α belongs to Π_{i_α} , and this orbit is necessarily repelling unless it coincides with the given orbit \mathcal{O} (which may be parabolic). In particular, for any angle t which is periodic under doubling, if the dynamic ray with angle $2^\alpha t$ lies in Π_{i_α} for all integers α , then this ray must land at the point z_α .*

(Note that the period of t may well be some multiple of m , as in the example just discussed.)

Proof. — There is a unique branch of f_c^{-1} which carries the interior of $\Pi_{i_{\alpha+1}}$ holomorphically onto a subset of Π_{i_α} . Let g_{i_α} be the composition of these m maps, in the appropriate reversed order so as to carry the interior of Π_{i_α} into itself.

A similar construction applies to the associated external angles. Let $J_i \subset \mathbb{R}/\mathbb{Z}$ be the set of all angles of dynamic rays which are contained in Π_i . Thus each J_i is a finite union of closed arcs, and together the J_i cover \mathbb{R}/\mathbb{Z} without overlap. Now there is a unique branch of the 2-valued map $t \mapsto t/2$ which carries $J_{i_{\alpha+1}}$ into J_{i_α} with derivative $1/2$ everywhere. Taking an m -fold composition, we map each J_{i_α} into itself with derivative $1/2^m$. This composition may well permute the various connected components of J_{i_α} . However, some iterate must carry some component of J_{i_α} into itself, and hence have a unique fixed point t in that component. The landing point of the corresponding dynamic ray will be a periodic point $z_\alpha \in \Pi_{i_\alpha}$.

Case 1. — If this landing point belongs to the interior of Π_{i_α} , then it is fixed by some iterate of our map g_{i_α} . This map g_{i_α} cannot be an isometry, hence it must contract the Poincaré metric. Therefore every orbit under g_{i_α} must converge towards z_α . Thus z_α is an attracting fixed point for g_{i_α} , and hence is a repelling periodic point for f .

Case 2. — If the landing point belongs to the boundary of Π_{i_α} then it must belong to $\mathcal{O} \cup \{-z_0\}$, and hence to the original orbit \mathcal{O} since $-z_0$ is not periodic. Evidently this case will occur only when the angle t belongs to the union $A_1 \cup \dots \cup A_p$ of angles in the given portrait \mathcal{P} . \square

Note. — It is essential for this argument that our given Markov cycle $\{\Pi_{i_\alpha}\}$ does not involve the critical puzzle piece Π_0 . In fact, as an immediate corollary we get the following statement:

Corollary 7.3 (Non-Repelling Cycles). — *Any non-repelling periodic orbit for f must intersect the critical puzzle piece Π_0 as well as the critical value puzzle piece Π_1 .*

Proof of 7.1. — If $I(\mathcal{P}) \subset I(\mathcal{Q})$, then it follows from Lemma 2.9 that there exists a map \hat{f} having both an orbit with portrait \mathcal{P} and an orbit with portrait \mathcal{Q} . The latter orbit determines a Markov cycle in the puzzle associated with \mathcal{P} . (The condition $I(\mathcal{P}) \subset I(\mathcal{Q})$ guarantees that this cycle avoids the critical puzzle piece.) Now for any map f with an orbit of portrait \mathcal{P} , we can use this Markov cycle, together with 7.2, to construct the required periodic orbit and to guarantee that the rays associated with the portrait \mathcal{Q} land on it, as required. \square

In fact an argument similar to the proof of 7.2 proves a much sharper statement. Let \mathcal{O} be a repelling periodic orbit with non-trivial portrait \mathcal{P} .

Lemma 7.4 (Orbits Bounded Away From Zero). — *Given an infinite sequence of non-critical puzzle pieces $\{\Pi_{i_k}\}$ for $k \geq 0$ with $f(\Pi_{i_k}) \supset \Pi_{i_{k+1}}$, there is one and only one point $w_0 \in K(f)$ so that the orbit $w_0 \mapsto w_1 \mapsto \cdots$ satisfies $w_k \in \Pi_{i_k}$ for every $k \geq 0$. It follows that the action of f on the compact set $K_{\mathcal{P}}$ consisting of all $w_0 \in K(f)$ such that the forward orbit $\{w_k\}$ never hits the interior of Π_0 is topologically conjugate to the one-sided subshift of finite type, associated to the matrix $[M_{ij}]$ with 0-th row and column deleted. In particular, the topology of $K_{\mathcal{P}}$ depends only on \mathcal{P} , and not on the particular choice of f within the \mathcal{P} -wake.*

Proof Outline. — First replace each puzzle piece Π_i by a slightly thickened puzzle piece, as described in [M3]. (Compare §8, Figure 18.) The interior of this thickened piece is an open neighborhood $N_i \supset \Pi_i$, with the property that $f(N_i) \supset \overline{N_j}$ whenever $f(\Pi_i) \supset \Pi_j$. It then follows that there is a branch of f^{-1} which maps N_j into N_i , carrying $K \cap \Pi_j$ into $K \cap \Pi_i$, and reducing distances by at least some fixed ratio $r < 1$ throughout the compact set $K \cap \Pi_j$. Further details are straightforward. \square

Presumably this statement remains true for a parabolic orbit, although the present proof does not work in the parabolic case. (Compare [Ha].)

8. Renormalization

One remarkable property of the Mandelbrot boundary is that it is densely filled with small copies of itself. (See Figures 11, 14 for a magnified picture of one such small copy.) This section will provide a rough outline, without proofs, of the Douady-Hubbard theory of renormalization, or the inverse operation of tuning, which provides a dynamical explanation for these small copies. It is based on [D4] as well as [DH3], [D3]. (Compare [D1], [M1]. For the Yoccoz interpretation of this construction, see [Hu], [M3], [Mc], [Ly]. For a more general form of renormalization, see [Mc], [RS].)




FIGURE 15. Detail near the period 4 hyperbolic component H_P of Figure 12, where $P = P(1/5, 4/15)$, showing the first eight of the parameter sectors which must be pruned away from M to leave the small Mandelbrot set consisting of P -renormalizable parameter values.

To begin the construction, consider any orbit portrait \mathcal{P} of ray period $n \geq 2$ and valence $v \geq 2$. Let c be a parameter value in $W_P \cup \{r_P\}$, so that $f = f_c$ has a periodic orbit \mathcal{O} with portrait \mathcal{P} , and let $S = S(f)$ be the critical value sector for this orbit (so that \bar{S} is the critical value puzzle piece). To a first approximation, we could try to say that f is “ \mathcal{P} -renormalizable” if the orbit of c under $f^{\circ n}$ is completely contained in S . In fact this is a necessary and sufficient condition whenever the map $f^{\circ n-1}|_S$ is univalent. However, in examples such as that of Figures 1, 2 one needs a slightly sharper condition.

Let $\mathcal{I}_P = (t_-, t_+)$ be the characteristic arc for this portrait, so that ∂S consists of the dynamic rays of angle t_- and t_+ together with their common landing point z_1 , and let $\ell = t_+ - t_-$ be the length of this arc.

Lemma 8.1 (A (Nearly) Quadratic-Like Map). — *The dynamic rays of angle $t'_1 = t_- + \ell/2^n$ and $t'_2 = t_+ - \ell/2^n$ land at a common point $z' \neq z_1$ in $S \cap f^{-n}(z_1)$. Let $S' \subset S$ be the region bounded by ∂S together with these two rays and their common landing*




FIGURE 16. The n -fold pull-back of the critical sector S along the orbit \mathcal{O} , illustrated schematically for the orbit diagram of period $n = 4$ which has characteristic arc $(1/5, 4/15)$. Compare Figures 5 (top), 12, 16.

point. Then the map $f^{\circ n}$ carries S' onto S by a proper map of degree two, with critical value equal to the critical value $f(0) = c$.

This region S' can be described as the n -fold “pull-back” of S along the orbit \mathcal{O} . (Compare Figure 16, which also shows the first three forward images of S' .)

Proof of 8.1. — First suppose that $c \in W_{\mathcal{P}}$ is outside the Mandelbrot set. Then, following Appendix A, we can bisect the complex plane by the two rays leading from infinity to the critical point. (Compare the proof of 2.9.) In order to check that the two rays of angle t'_1 and t'_2 have a common landing point, we need only show that they have the same symbol sequence with respect to the resulting partition. In other words, we must show, for every $k \geq 0$, that the $2^k t'_1$ and $2^k t'_2$ rays lie on the same side of the bisecting critical ray pair. For $k \geq n$ this is clear since $2^n t'_1 \equiv t_+$ and $2^n t'_2 \equiv t_-$ modulo \mathbb{Z} .

Now consider the critical puzzle piece Π_0 of §7. Evidently Π_0 is a neighborhood, of angular radius $\ell/4$, of the bisecting critical ray pair. For $k < n - 1$ the dynamic rays with angle $2^k t_-$ and $2^k t_+$ both lie in the same component of $\mathbb{C} \setminus \Pi_0$. Since $2^k t'_j$ differs from $2^k t_j$ by at most $\ell/4$, it follows that the $2^k t'_1$ and $2^k t'_2$ rays have the same symbol. Finally, for $k = n - 1$, it is not difficult to check that the $2^k t'_1$ and $2^k t'_2$ rays both land at the same point $-z_0 \neq z_0$. This proves that the t'_1 and t'_2 rays land at the same point, different from z_1 , when $c \notin M$. A straightforward continuity argument now proves the same statement for all $c \in W_{\mathcal{P}}$.

Thus we obtain the required region $S' \subset S$. As in §2, it will be convenient to complete the complex plane by adjoining a circle of points at infinity. Note that the boundary of S' within this circled plane \mathbb{C} consists of two arcs of length $\ell/2^n$ at infinity, together with two ray pairs and their common landing points. As we traverse this boundary once in the positive direction, the image under f^n evidently traverses the boundary of S twice in the positive direction. Using the Argument Principle, it follows that the image of S' is contained in S , and covers every point of S twice, as required. Thus $f^n|_{S'}$ must have exactly one critical point, which can only be c . \square

Thus we have an object somewhat like a quadratic-like map, as studied in [DH3]. Note however that S' is not compactly contained in S .

Definition. — We will say that f is \mathcal{P} -renormalizable if $f(0) = c$ is contained in the closure $\overline{S'}$, and furthermore the entire forward orbit of c under the map f^n is contained in $\overline{S'}$. If this condition is satisfied, and the orbit of c is also bounded so that $c \in M$, then we will say that c belongs to the “small copy” $\mathcal{P} * M$ of the Mandelbrot set which is associated with \mathcal{P} . (This terminology will be justified in 8.2. If the orbit is unbounded, then we may say that c belongs to a \mathcal{P} -renormalizable external ray.)




FIGURE 17. Julia set for the center point of the period 12 satellite component of Figure 12 (the point c of Figure 15), and a detail near the critical value c , showing the first eight of the sectors of the dynamic plane which must be pruned away to leave the small Julia set associated with \mathcal{P} -renormalization, with \mathcal{P} as in Figures 12, 15. (Here the right hand figure has been magnified by a factor of 75.) This can be described as the Julia set of Figure 13 (left) tuned by a “Douady rabbit” Julia set.

Closely associated is the “small filled Julia set” $K' = K(f^n|_{S'})$ consisting of all $z \in \overline{S'}$ such that the entire forward orbit of z under f^n is bounded and contained in $\overline{S'}$. (Compare Figure 17.) Thus the critical value $f(0) = c$ belongs to K' if and only

if f is \mathcal{P} -renormalizable, with $c \in M$. As in the classical Fatou-Julia theory, c belongs to K' if and only if K' is connected.




FIGURE 18. The sector S and the thickened sector T .

In order to tie this construction up with Douady and Hubbard's theory of polynomial-like mappings, we need to thicken the sector S , and then cut it down to a bounded set. (Compare [M3].) We exclude the exceptional special case where c is the root point $r_{\mathcal{P}}$. Thus we will suppose that the periodic point $z_1 \in \partial S$ is repelling. Choose a small disk D_ε about z_1 which is mapped univalently by $f_c^{\circ n}$ and is compactly contained in $f_c^{\circ n}(D_\varepsilon)$. Choose also a very small $\eta > 0$, and consider the dynamic rays with angle $t_- - \eta$ and $t_+ + \eta$. Following these rays until they first meet D_ε , they delineate an open region $T \supset S \cup D_\varepsilon$ in \mathbb{C} . (Compare Figure 18.) Now let T' be the connected component of $f_c^{-n}(T)$ which contains S' . It is not difficult to check that $\overline{T'} \subset T$, and that $f_c^{\circ n}$ carries T' onto T by a proper map of degree two.

To obtain a bounded region, we let U be the intersection of T with the set $\{z \in \mathbb{C} ; G^K(z) < 1\}$, where G^K is the Green's function for $K = K(f_c)$. Similarly, let U' be the intersection T' with $\{z ; G^K(z) < 1/2^n\}$. Then U' is compactly contained in U , and $f_c^{\circ n}$ carries U' onto U by a proper map of degree two. *In other words, $f_c^{\circ n}|_{U'}$ is a quadratic-like map.*

Evidently the forward orbit of a point $z \in U'$ under $f_c^{\circ n}$ is contained in U' if and only if z belongs to the small filled Julia set K' . In particular, for $c \in M$, the map f_c is \mathcal{P} -renormalizable if and only if $c \in K'$, or if and only if K' is connected.

If these conditions are satisfied, then according to [DH3] the map $f_c^{\circ n}$ restricted to a neighborhood of K' is "hybrid equivalent" to some uniquely defined quadratic map $f_{c'}$, with $c' \in M$. Briefly, we will write $c = \mathcal{P} * c'$, or say that c equals \mathcal{P} tuned by c' . Douady and Hubbard show also that this correspondence

$$c' \mapsto \mathcal{P} * c'$$

is a well defined continuous embedding of $M \setminus \{1/4\}$ onto a proper subset of itself. As an example, as c' varies over the hyperbolic component $H_{\{0\}}$ which is bounded by the cardioid, they show that $\mathcal{P} * c'$ varies over the hyperbolic component $H_{\mathcal{P}}$.

It is convenient to supplement this construction, by defining the operation \mathcal{P} , $c' \mapsto \mathcal{P} * c'$ in two further special cases. If c' is the root point $1/4 = r_{\{0\}}$ of M , then we

define

$$\mathcal{P} * (1/4) = r_{\mathcal{P}}$$

to be the root point of the \mathcal{P} -wake. Furthermore, if $\mathcal{P} = \{\{0\}\}$ is the zero orbit portrait, then we define $\{\{0\}\}*$ to be the identity map,

$$\{\{0\}\} * c' = c'$$

for all $c' \in M$. With these definitions, we have the following basic result of Douady and Hubbard.

Theorem 8.2 (Tuning). — *For each non-trivial orbit portrait \mathcal{P} , the correspondence $c \mapsto \mathcal{P} * c$ defines a continuous embedding of the Mandelbrot set M into itself. The image of this embedding is just the “small Mandelbrot set” $\mathcal{P} * M \subset M$ described earlier. Furthermore, there is a unique composition operation $\mathcal{P}, \mathcal{Q} \mapsto \mathcal{P} * \mathcal{Q}$ between non-trivial orbit portraits so that the associative law is valid,*

$$(\mathcal{P} * \mathcal{Q}) * c = \mathcal{P} * (\mathcal{Q} * c)$$

for all \mathcal{P}, \mathcal{Q} and c' . Under this $*$ composition operation, the collection of all non-trivial orbit portraits forms a free (associative but noncommutative) monoid, with the zero orbit portrait as identity element.

The proof is beyond the scope of this note.

We can better understand this construction by introducing a nested sequence of open sets

$$S = S^{(0)} \supset S' = S^{(1)} \supset S^{(2)} \supset \dots$$

in the dynamic plane for f , where $S^{(k+1)}$ is defined inductively as $S^{(k)} \cap f^{-n}(S^{(k)})$ for $k \geq 1$. Thus $S = S^{(0)}$ is bounded by the dynamic rays of angle t_- and t_+ , together with their common landing point z_1 . Similarly, $S^{(1)}$ is bounded by $\partial S^{(0)}$ together with the rays of angle $t_- + \ell/2^n$ and $t_+ - \ell/2^n$, together with their common landing point, which is an n -fold pre-image of z_1 . If $c \in S^{(1)}$, so that $S^{(2)}$ is a 2-fold branched covering of $S^{(1)}$, then $S^{(2)}$ has two further boundary components, namely the rays of angle $t_- + \ell/2^{2n}$ and $t_- + \ell/2^n - \ell/2^{2n}$ and their common landing point, together with the rays of angle $t_+ - \ell/2^n + \ell/2^{2n}$ and $t_+ - \ell/2^{2n}$ and their common landing point, for a total of 4 boundary components. Similarly, if $c \in S^{(2)}$, then $S^{(3)}$ has 8 boundary components, as illustrated in Figure 17.

The angles which are left, after we have cut away the angles in all of these (open) sectors, form a standard middle fraction Cantor set \mathcal{K} , which can be described as follows. Let ϕ be the fraction $1 - 2/2^n$. Start with the closure $[t_-, t_+]$ of the characteristic arc for \mathcal{P} , with length ℓ . First remove the open middle segment of length $\phi \ell$, leaving two arcs of length $\ell/2^n$. Then, from each of these two remaining closed arcs, remove the middle segment of length $\phi \ell/2^n$, leaving four segments of length $\ell/2^{2n}$, and continue inductively. The intersection of all of the sets obtained in this way is the required Cantor set $\mathcal{K} \subset [t_-, t_+]$ of angles. These are precisely the angles of the

dynamic rays which land on the small Julia set $\partial K'$ (at least if we assume that these Julia sets are locally connected).

There is a completely analogous construction in parameter space, as illustrated in Figure 15. As noted earlier, parameter rays of angle t_- and t_+ land on a common point $r_{\mathcal{P}}$, and together form the boundary of the \mathcal{P} -wake. Similarly, the parameter rays of angle $t_- + \ell/2^n$ and $t_+ - \ell/2^n$ must land at a common point. These rays, together with their landing point, cut $W_{\mathcal{P}}$ into two halves. For c in the inner half, with boundary point $r_{\mathcal{P}}$, the critical value of f_c lies in $S' = S^{(1)}$, while for c in the outer half, this is not true. Similarly, for each pair of dynamic rays with a common landing point in ∂K , forming part of the boundary of $S^{(k)}$, there is a pair of parameter rays with the same angles which have a common landing point in ∂M and form part of the boundary of a corresponding region $W_{\mathcal{P}}^{(k)}$ in parameter space. The basic property is that $c \in W_{\mathcal{P}}^{(k)}$ if and only if c belongs to the corresponding region $S^{(k)}$ in the dynamic plane for f_c .

Dynamically, the Cantor set $\mathcal{K} \subset \mathbb{R}/\mathbb{Z}$ can be described as the set of angles in

$$[t_-, t_- + \ell/2^n] \cup [t_+ - \ell/2^n, t_+]$$

such that the entire forward orbit under multiplication by 2^n is contained in this set. Evidently the resulting dynamical system is topologically isomorphic to the one-sided two-shift. Thus each element $t \in \mathcal{K}$ can be coded by an infinite sequence (b_0, b_1, \dots) of bits, where each b_k is zero or one according as $2^{nk}t$ belongs to the left or right subarc. We will write $t = \mathcal{P} * (b_0b_1b_2 \dots)$. Intuitively, we can identify this sequence of bits b_i with the angle $.b_0b_1b_2 \dots = \sum b_k/2^{k+1}$. However, some care is needed since the correspondence $.b_0b_1b_2 \dots \mapsto \mathcal{P} * (b_0b_1 \dots)$ has a jump discontinuity at every dyadic rational angle, i.e., at those angles corresponding to gaps in the Cantor set \mathcal{K} . Thus we must distinguish between the left hand limit $\mathcal{P} * \alpha -$ and the right hand limit $\mathcal{P} * \alpha +$ when α is a dyadic rational.

With this notation, the angles of the bounding rays for the various open sets $S^{(k)}$, or for the corresponding sets $W_{\mathcal{P}}^{(k)}$ in parameter space, are just these left and right hand limits $\mathcal{P} * \alpha \pm$, where α varies over the dyadic rationals; and the composition operation between non-trivial orbit portraits can be described as follows: *If \mathcal{Q} has characteristic arc (t_-, t_+) , then $\mathcal{P} * \mathcal{Q}$ has characteristic arc $(\mathcal{P} * t_-, \mathcal{P} * t_+)$.* For further details, see [D3].

9. Limbs and the Satellite Orbit

Let \mathcal{P} be a non-trivial orbit portrait with period $p \geq 1$ and ray period $rp \geq p$. (Thus \mathcal{P} may be either a primitive or a satellite portrait.) Recall that the *limb* $M_{\mathcal{P}}$ consists of all points which belong both to the Mandelbrot set M and to the closure $\overline{W}_{\mathcal{P}}$ of the \mathcal{P} -wake. By definition, a limb $M_{\mathcal{Q}}$ with $\mathcal{Q} \neq \mathcal{P}$ is a *satellite* of $M_{\mathcal{P}}$ if its

root point \mathbf{r}_Q belongs to the boundary of the associated hyperbolic component H_P . (See 6.4.) We will prove the following two statements. (Compare [Hu], [Sø], [S3].)

Theorem 9.1 (Limb Structure). — *Every point in the limb M_P either belongs to the closure \overline{H}_P of the associated hyperbolic component, or else belongs to some satellite limb M_Q .*

(For a typical example, see Figure 12.) For any parameter value c in the wake W_P , let $\mathcal{O}(c) = \mathcal{O}_P(c)$ be the repelling orbit for f_c which has period p and portrait \mathcal{P} . Clearly this orbit $\mathcal{O}(c)$ varies holomorphically with the parameter value c .

Corollary 9.2 (The Satellite Orbit). — *To any $c \in W_P$ there is associated another orbit $\mathcal{O}^*(c) = \mathcal{O}_P^*(c)$, distinct from $\mathcal{O}(c)$, which has period $n = rp$ and which also varies holomorphically with the parameter value c . As c tends to the root point \mathbf{r}_P , the two orbits $\mathcal{O}(c)$ and $\mathcal{O}^*(c)$ converge towards a common parabolic orbit of portrait \mathcal{P} . (Compare 4.1.) This associated orbit $\mathcal{O}^*(c)$ is attracting if c belongs to the hyperbolic component $H_P \subset W_P$, indifferent for $c \in \partial H_P$, and is repelling for $c \in W_P \setminus \overline{H}_P$, with portrait equal to \mathcal{Q} if c belongs to the satellite wake W_Q .*

As an example, both statements apply to the zero portrait, with $M_{\{0\}}$ equal to the entire Mandelbrot set, with $W_{\{0\}} = \mathbb{C} \setminus (1/4, +\infty)$, and with $H_{\{0\}}$ bounded by the cardioid. In this case, for any $c \in W_{\{0\}}$, the orbit $\mathcal{O}(c)$ consists of the *beta fixed point* $(1 + \sqrt{1 - 4c})/2$ while $\mathcal{O}^*(c)$ consists of the *alpha fixed point* $(1 - \sqrt{1 - 4c})/2$, taking that branch of the square root function with $\sqrt{1} = 1$.

Proof of 9.1. — For each $c \in H_P$ let $\mathcal{O}^*(c)$ be the unique attracting periodic orbit. By the discussion in §6, this orbit extends analytically as we vary c over some neighborhood of the closure \overline{H}_P , provided that we stay within the wake W_P . Furthermore, this orbit becomes strictly repelling as we cross out of \overline{H}_P . Therefore we can choose a neighborhood N of \overline{H}_P which is small enough so that this analytically continued orbit $\mathcal{O}^*(c)$ will be strictly repelling for all $c \in N \cap W_P \setminus \overline{H}_P$. If c also belongs to the Mandelbrot set, so that $c \in N \cap M_P \setminus \overline{H}_P$, it follows that at least one rational dynamic ray lands on the orbit $\mathcal{O}^*(c)$; hence there is an orbit portrait $\mathcal{Q} = \mathcal{Q}(c)$ of period n associated with $\mathcal{O}^*(c)$. Choosing the neighborhood N even smaller if necessary, we will show that the rotation number of $\mathcal{Q}(c)$ is non-zero, and hence that this portrait $\mathcal{Q}(c)$ is non-trivial. In other words, we will prove that c belongs to a limb M_Q which is associated to the orbit $\mathcal{O}^*(c)$.

First consider a point \hat{c} which belongs to the boundary ∂H_P . Then $\mathcal{O}^*(\hat{c})$ is an indifferent periodic orbit, with multiplier on the unit circle. Consider some dynamic ray \mathcal{R}_t^K which has period n , but does not participate in the portrait \mathcal{P} , and hence does not land on the original orbit $\mathcal{O}(\hat{c})$. Such a ray certainly cannot land on $\mathcal{O}^*(\hat{c})$, for that would imply that $\mathcal{O}^*(\hat{c})$ was a repelling or parabolic orbit of rotation number zero. However, for \hat{c} in the boundary of H_P the orbit $\mathcal{O}^*(\hat{c})$ is never repelling, and

is parabolic of rotation number zero only when \hat{c} is the root point of $H_{\mathcal{P}}$, so that $\mathcal{O}^*(\hat{c}) = \mathcal{O}(\hat{c})$. Since we have assumed that the ray \mathcal{R}_t^K does not land on $\mathcal{O}(\hat{c})$, it must land on some repelling or parabolic periodic point which is disjoint from $\mathcal{O}^*(\hat{c})$. In fact it must land on a repelling orbit, since a quadratic map cannot have a parabolic orbit and also a disjoint indifferent orbit. (Compare §6.) Now as we perturb c throughout some neighborhood of \hat{c} it follows that the corresponding ray still lands on a repelling periodic point disjoint from $\mathcal{O}^*(c)$. Since $\mathcal{O}^*(c)$ has period n , but no ray of period n can land on it, this proves that the rotation number of the associated portrait $\mathcal{Q}(c)$ is non-zero, as asserted.

Let X be any connected component of $M_{\mathcal{P}} \setminus \overline{H}_{\mathcal{P}}$. Since the Mandelbrot set is connected, X must have some limit point in $\partial H_{\mathcal{P}}$. Therefore, by the argument above, some point $c \in X$ must belong to a wake $W_{\mathcal{Q}}$ associated with the orbit $\mathcal{O}^*(c)$. Since the portrait \mathcal{Q} has period n , the root point $r_{\mathcal{Q}}$ of its wake must lie on the boundary of some hyperbolic component H' which has period n and is contained in $W_{\mathcal{P}}$. In fact, for suitable choice of c , we claim that H' can only be $H_{\mathcal{P}}$ itself. There are finitely many other components of period n , but these others are all bounded away from $H_{\mathcal{P}}$, while the point $c \in X$ can be chosen arbitrarily close to $\overline{H}_{\mathcal{P}}$. Thus we may assume that $W_{\mathcal{Q}}$ is rooted at a point of $\partial H_{\mathcal{P}}$, and hence is a satellite wake. Since the connected set X cannot cross the boundary of $W_{\mathcal{Q}}$, it follows that X is completely contained within $W_{\mathcal{Q}}$, which completes the proof of 9.1. \square

Proof of 9.2. — As in the argument above, the orbit $\mathcal{O}^*(c)$ is well defined for c in some neighborhood of $W_{\mathcal{P}} \cap \overline{H}_{\mathcal{P}}$, and we can try to extend analytically throughout the simply connected region $W_{\mathcal{P}}$. There is a potential obstruction if we ever reach a point in $W_{\mathcal{P}}$ where the multiplier λ_n of this analytically extended orbit is equal to $+1$. However, this can never happen. In fact such a point would have to belong to the Mandelbrot set, and hence to some satellite limb $M_{\mathcal{Q}}$. But we can extend analytically throughout the associated wake $W_{\mathcal{Q}}$, taking $\mathcal{O}^*(c)$ to be the repelling orbit $\mathcal{O}_{\mathcal{Q}}(c)$ for every $c \in W_{\mathcal{Q}}$. Thus there is no obstruction. It follows similarly that the analytically extended orbit must be repelling everywhere in $W_{\mathcal{P}} \setminus \overline{H}_{\mathcal{P}}$. For if it became non-repelling at some point c , then again c would have to belong to some satellite limb $M_{\mathcal{Q}}$, but $\mathcal{O}^*(c)$ is repelling throughout the wake $W_{\mathcal{Q}}$. \square

Corollary 9.3 (Limb Connectedness). — *Each limb $M_{\mathcal{P}} = M \cap \overline{W}_{\mathcal{P}}$ is connected, even if we remove its root point $r_{\mathcal{P}}$.*

Proof. — The entire Mandelbrot set is connected by [DH1]. It follows that each $M_{\mathcal{P}}$ is connected. For if some limb $M_{\mathcal{P}}$ could be expressed as the union of two disjoint non-vacuous compact subsets, then only one of these two could contain the root point $r_{\mathcal{P}}$. The other would be a non-trivial open-and-closed subset of M , which is impossible.

Now consider the open subset $M_{\mathcal{P}} \setminus \{r_{\mathcal{P}}\}$. This is a union of the connected set $\overline{H}_{\mathcal{P}} \setminus \{r_{\mathcal{P}}\}$, together with the various satellite limbs M_Q , where each M_Q has root point r_Q belonging to $\overline{H}_{\mathcal{P}} \setminus \{r_{\mathcal{P}}\}$. Since each M_Q is connected, the conclusion follows. \square

Remark 9.4. — It follows easily that every satellite root point separates the Mandelbrot set into exactly two connected components, and hence that exactly two parameter rays land at every such point. For a proof of the corresponding statement for a primitive root (other than $1/4$) see [Ta] or [S3].

Appendix A

Totally Disconnected Julia Sets and the Mandelbrot set

This appendix will be a brief review of well known material. For any parameter value c , let $K = K(f_c)$ be the filled Julia set for the map $f_c(z) = z^2 + c$, and let

$$G(z) = G^K(z) = \lim_{n \rightarrow \infty} \frac{1}{2^n} \log |f_c^n(z)|$$

be the canonical potential function or *Green's function*, which vanishes only on K , and satisfies $G(f_c(z)) = 2G(z)$. The level sets $\{z; G(z) = G_0\}$ are called *equipotential curves* for K , and the orthogonal trajectories which extend to infinity are called the *dynamic rays* \mathcal{R}_t^K , where $t \in \mathbb{R}/\mathbb{Z}$ is the angle at infinity.

Now suppose that K is totally disconnected (and hence coincides with the Julia set $J = \partial K$). Then the value $G(0) = G(c)/2 > 0$ plays a special role. In fact there is a canonical conformal isomorphism ψ_c from the open set $\{z; G(z) > G(0)\}$ to the region $\{w; \log |w| > G(0)\}$. The map $z \mapsto f_c(z)$ on this region is conjugate under ψ_c to the map $w \mapsto w^2$, and the equipotentials and dynamic rays in the z -plane correspond to concentric circles and straight half-lines through the origin respectively in the w -plane. In particular, if we choose a constant $G_0 > G(0)$, then the locus $\{z; G(z) = G_0\}$ is a simple closed curve, canonically parametrized by the angle of the corresponding dynamic ray. In particular, the critical value $c \in \mathbb{C} \setminus K$ has a well defined external angle, which we denote by $t(c) \in \mathbb{R}/\mathbb{Z}$. Thus $\psi_c(c)/|\psi_c(c)| = e^{2\pi i t(c)}$, and c belongs to the dynamic ray $\mathcal{R}_{t(c)} = \mathcal{R}_{t(c)}^K$.

However, for $G_0 = G(0)$ this locus $\{z; G(z) = G(0)\}$ is a figure eight curve. The open set $\{z; G(z) < G(0)\}$ splits as a disjoint union $U_0 \cup U_1$, where the U_b are the regions enclosed by the two lobes of this figure eight. (We can express this splitting in terms of dynamic rays as follows. The ray $\mathcal{R}_{t(c)}^K \subset \mathbb{C} \setminus K$ has two preimage rays under f_c , with angles $t(c)/2$ and $(1 + t(c))/2$ respectively. Each of these joins the critical point 0 to the circle at infinity, and together they cut \mathbb{C} into two open subsets, say $V_0 \supset U_0$ and $V_1 \supset U_1$. If c does not belong to the positive real axis, then we can choose the labels for these open sets so that the zero ray is contained in V_0 , and $c \in V_1$.) We then cut the filled Julia set K into two disjoint compact subsets $K_b = K \cap U_b$. These constitute a *Bernoulli partition*. That is, for any one-sided-infinite sequence of

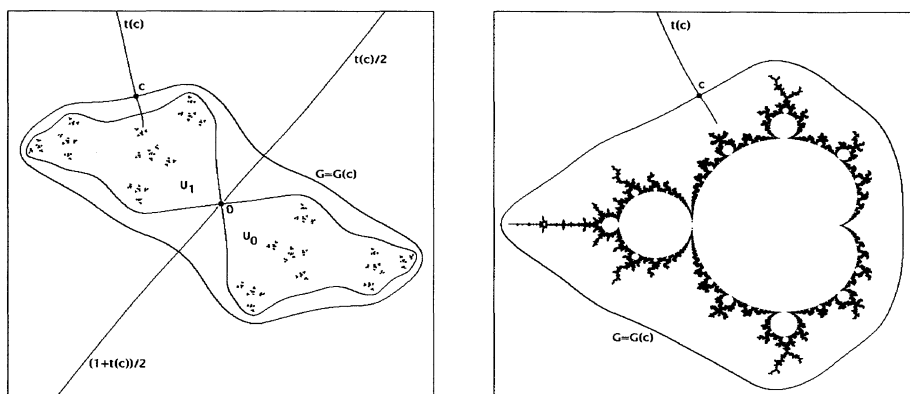


FIGURE 19. Picture in the dynamic plane for a polynomial f_c with $c \notin M$, and a corresponding picture in the parameter plane.

bits $b_0, b_1, \dots \in \{0, 1\}$, there is one and only one point $z \in K$ with $f_c^{o_k}(z) \in K_{b_k}$ for every $k \geq 0$. To prove this statement, let U be the region $\{z; G(z) < G(c)\}$ and let $\phi_b : U \rightarrow U_b$ be the branch of f^{-1} which maps U diffeomorphically onto U_b . Using the Poincaré metric for U , we see that each ϕ_b shrinks distances by a factor bounded away from one, and it follows easily that the diameter of the image

$$\phi_{b_0} \circ \phi_{b_1} \circ \dots \circ \phi_{b_n}(U)$$

shrinks to zero, so that this intersection shrinks to a single point $z \in K$, as $n \rightarrow \infty$. Thus each point of $J = K$ can be uniquely characterized by an infinite sequence of symbols (b_0, b_1, \dots) with $b_j \in \{0, 1\}$. In particular, K is homeomorphic to the infinite cartesian product $\{0, 1\}^{\mathbb{N}}$, where the symbol \mathbb{N} stands for the set $\{0, 1, 2, \dots\}$ of natural numbers. We say that the dynamical system $(K, f_c|_K)$ is a *one sided shift* on two symbols.

Similarly, given any angle $t \in \mathbb{R}/\mathbb{Z}$, if none of the successive images $2^k t \pmod{\mathbb{Z}}$ under doubling is precisely equal to $t(c)/2$ or $(1 + t(c))/2$, then t has an associated symbol sequence, called its $t(c)$ -itinerary, and the ray \mathcal{R}_t^K lands precisely at that point of K which has this symbol sequence. For the special case $t = t(c)$, this symbol sequence characterizes the point $c \in K$, and is called the *kneading sequence* for c or for $t(c)$. (However, if $t(c)$ is periodic, there is some ambiguity since the symbols b_{n-1}, b_{2n-1}, \dots of the kneading sequence are not uniquely defined in the period n case.)

If t is periodic under doubling, then the itinerary is periodic (if uniquely defined), and the ray \mathcal{R}_t^K lands at a periodic point of K . For further discussion, see [LS], as well as Appendix B.

Here we have been thinking of $c = f(0)$ as a point in the dynamic plane (the z -plane), but we can also think of $c \in \mathbb{C} \setminus M$ as a point in the parameter plane (the

c -plane). In fact Douady and Hubbard construct a conformal isomorphism from the complement of M onto the complement of the closed unit disk by mapping $c \in \mathbb{C} \setminus M$ to the point $\psi_c(c) = \exp(G^K(c) + 2\pi it(c)) \in \mathbb{C} \setminus \overline{\mathbf{D}}$. Thus they show that the value of the Green's function on c and the external angle $t(c)$ of c are the same whether c is considered as a point of $\mathbb{C} \setminus K(f_c)$ or as a point of $\mathbb{C} \setminus M$. In particular, the point $c \in \mathbb{C} \setminus M$ lies on the external ray $\mathcal{R}_{t(c)}^M$ for the Mandelbrot set.

Appendix B

Computing Rotation Numbers

This appendix will outline how to actually compute the rotation number q/r of a periodic point for a map f_c with $c \notin M$. Let $\tau = t(c) \in \mathbb{R}/\mathbb{Z}$ be the angle of the external ray which passes through c . We may identify this critical value angle with a number in the interval $0 < \tau \leq 1$. The two preimages of τ under the angle doubling map $m_2 : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ separate the circle \mathbb{R}/\mathbb{Z} into the two open arcs

$$I(0) = I_\tau(0) = \left(\frac{\tau-1}{2}, \frac{\tau}{2} \right) \quad \text{and} \quad I(1) = I_\tau(1) = \left(\frac{\tau}{2}, \frac{\tau+1}{2} \right).$$

(We will write I_τ instead of I whenever we want to emphasize dependence on the critical value angle τ .) For any finite sequence b_0, b_1, \dots, b_k of zeros and ones, let $\bar{I}(b_0, b_1, \dots, b_k)$ be the closure of the open set

$$I(b_0, b_1, \dots, b_k) = I(b_0) \cap m_2^{-1}I(b_1) \cap \dots \cap m_2^{-k}I(b_k)$$

consisting of all $t \in \mathbb{R}/\mathbb{Z}$ with $m_2^{o_i}(t) \in I(b_i)$ for $0 \leq i \leq k$. (Caution: This is not the same as the intersection of the corresponding closures $m_2^{-i}\bar{I}(b_i)$, which may contain additional isolated points.) An easy induction shows that $\bar{I}(b_0, b_1, \dots, b_k)$ is a finite union of closed arcs with total length $1/2^{k+1}$. If $\sigma = (b_0, b_1, \dots)$ is any infinite sequence of zeros and ones, it follows that the intersection

$$\bar{I}(\sigma) = \bigcap_k \bar{I}(b_0, b_1, \dots, b_k)$$

is a compact non-vacuous set of measure zero. For each angle $t \in \mathbb{R}/\mathbb{Z}$ there are two possibilities:

Precritical Case. — If t satisfies $m_2^{o_i}(t) \equiv \tau$ for some $i > 0$, then there will be two distinct infinite symbol sequences with $t \in \bar{I}(b_0, b_1, b_2 \dots)$. In this case, the associated dynamic ray \mathcal{R}_t^K does not land, but rather bounces off some precritical point for the map f_c . (Compare [GM].)

Generic Case. — Otherwise there will be a unique infinite symbol sequence with $t \in \bar{I}(b_0, b_1, \dots)$. The corresponding ray \mathcal{R}_t^K will land at the unique point of the Julia set for f_c which has this same symbol sequence, as described in Appendix A. In particular, if t is periodic under doubling, then \mathcal{R}_t^K must land at a periodic point of the Julia set, possibly with smaller period.

Lemma B.1 (Symbol Sequences and Rotation Numbers). — For any symbol sequence $\sigma = (b_0, b_1, \dots) \in \{0, 1\}^{\mathbb{N}}$ which is periodic of period p , the map $m_2^{\circ p}$ on the compact set $\bar{I}_\tau(\sigma) \subset \mathbb{R}/\mathbb{Z}$ has a well defined rotation number $\text{rot}(b_0, \dots, b_{p-1}; \tau) \in \mathbb{R}/\mathbb{Z}$ which is invariant under cyclic permutation of the bits b_i . This number increases monotonically with τ , and winds $b_0 + \dots + b_{p-1}$ times around the circle as τ increases from 0 to 1.

To see this, we introduce an auxiliary monotone degree one map which is defined on the entire circle and agrees with $m_2^{\circ p}$ on $\bar{I}_\tau(\sigma)$. (Compare [GM].) By definition, a *monotone degree one circle map* $\psi : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ is the reduction modulo \mathbb{Z} of a map $\Psi : \mathbb{R} \rightarrow \mathbb{R}$ which is monotone increasing and satisfies the identity $\Psi(u+1) = \Psi(u) + 1$. Such a Ψ , called a *lift* of ψ , is unique up to addition of an integer constant. The *translation number* of such a map Ψ is defined to be the real number

$$\text{Trans}(\Psi) = \lim_{k \rightarrow \infty} (\Psi^{\circ k}(u) - u)/k.$$

This always exists, and is independent of u . The *rotation number* $\text{rot}(\psi)$ of the associated circle map is now defined to be the image of this real number $\text{Trans}(\Psi)$ under the projection $\mathbb{R} \rightarrow \mathbb{R}/\mathbb{Z}$. This is well defined, since $\text{Trans}(\Psi + 1) = \text{Trans}(\Psi) + 1$. One important property is the identity

$$\text{Trans}(\Psi_1 \circ \Psi_2) = \text{Trans}(\Psi_2 \circ \Psi_1), \quad (9)$$

where Ψ_1 and Ψ_2 are the lifts of two different monotone degree one circle maps. If Ψ_1 is a homeomorphism, this is just invariance under a suitable change of coordinates, and the general case follows by continuity.

Given any $b \in \{0, 1\}$, and given a critical value angle τ , define an auxiliary monotone map $\Phi_{b,\tau}$ by the formula


$$\Phi_{b,\tau}(u) = \begin{cases} \min(2u, \tau) & \text{if } b = 0, \\ \max(2u, \tau) & \text{if } b = 1, \end{cases}$$

for u between $(\tau - 1)/2$ and $(\tau + 1)/2$, extending by the identity $\Phi(u + 1) = \Phi(u) + 1$ for u outside this interval. (See Figure 20.) Note that $\bar{I}(b)$ is just the set of points on the circle where the associated circle map $\phi_{b,\tau}$ is not locally constant, and that $\phi_{b,\tau}(u) \equiv 2u \pmod{\mathbb{Z}}$ whenever $u \in \bar{I}(b)$.

For any symbol sequence σ which is periodic of period p , we set $\Phi_{\sigma,\tau}$ equal to the p -fold composition $\Phi_{b_{p-1},\tau} \circ \dots \circ \Phi_{b_0,\tau}$. (Note that $\bar{I}_\tau(\sigma)$ is just the set of all points $t \in \mathbb{R}/\mathbb{Z}$ such that the orbit of t under the associated circle map $\phi_{\sigma,\tau}$ coincides with the orbit of t under $m_2^{\circ p}$.) This composition is also monotone, with $\Phi(t+1) = \Phi(t) + 1$, and therefore has a well defined translation number, which we denote by

$$\text{Trans}(b_0, \dots, b_{p-1}; \tau) = \text{Trans}(\Phi_{\sigma,\tau}) \in \mathbb{R}.$$

It follows from property (9) that this translation number is invariant under cyclic permutation of the bits b_0, \dots, b_{p-1} . Since each $\Phi_{b,\tau}(u)$ increases monotonically

FIGURE 20. Graphs of $\Phi_{0,\tau}$ and $\Phi_{1,\tau}$ (with $\tau = 0.6$).

with τ , with $\Phi_{b,0}(0) = 0$ and $\Phi_{b,1}(0) = b$, it follows easily that $\text{Trans}(\Phi_{\sigma,\tau})$ depends monotonically on τ , increasing from 0 to $b_0 + \dots + b_{p-1}$ as τ increases from 0 to 1. In other words its image in \mathbb{R}/\mathbb{Z} wraps $b_0 + \dots + b_{p-1}$ times around the circle as τ varies from 0 to 1. By definition, the rotation number $\text{rot}(b_0, \dots, b_{p-1}; \tau)$ of m_2^{op} on the compact set $\bar{I}_\tau(\sigma)$ is equal to the image of the real number $\text{Trans}(\Phi_{\sigma,\tau})$ in the circle \mathbb{R}/\mathbb{Z} . \square

If a map f_c has critical value angle $t(c) = \tau$, then it is not hard to see that $\text{rot}(b_0, \dots, b_{p-1}; \tau)$ coincides with the rotation number as defined in 2.12 for the orbit with periodic symbol sequence $\overline{b_0, \dots, b_{p-1}} = (b_0, \dots, b_{p-1}, b_0, \dots, b_{p-1}, \dots)$, so long as at least one rational ray lands on this orbit. (Compare [GM, Appendix C].)

We will use the notation $\mathcal{S}(q/r)$ for the orbit portrait with orbit period $p = 1$ and rotation number q/r , associated with the q/r -satellite of the main cardioid. (Compare [G].) If \mathcal{P} is an arbitrary orbit portrait, then $\mathcal{P} * \mathcal{S}(q/r)$ can be described as its (q/r) -satellite portrait. (See 6.4, 8.2.)

To any orbit portrait \mathcal{P} with period $p \geq 1$ and ray period $n = rp \geq p$ we can associate a symbol sequence $\sigma = \sigma(\mathcal{P})$ of period p as follows. Choose any $c \notin M$ in the wake $W_{\mathcal{P}}$, and number the points of the f_c -orbit with portrait \mathcal{P} as $z_0 \mapsto z_1 \mapsto \dots$, where z_0 is on the boundary of the critical puzzle piece and z_1 is on the boundary of the critical value puzzle piece. Now let $\sigma(\mathcal{P})$ be the symbol sequence for z_0 , as described in Appendix A. This is independent of the choice of $c \in W_{\mathcal{P}} \setminus M_{\mathcal{P}}$.

There is an associated *satellite symbol sequence* $\sigma^* = \sigma^*(\mathcal{P})$ of period $n = rp$, constructed as follows. (Compare 9.2.) By definition, the k -th bit of σ^* is identical to the k -th bit of σ for $k \not\equiv 0 \pmod{n}$, but is reversed, so that $0 \leftrightarrow 1$, when $k \equiv 0 \pmod{n}$.




FIGURE 21. The translation number as a function of the critical value exterior angle τ for the period 3 point with symbol sequence $\overline{110} = (1, 1, 0, 1, 1, 0, \dots)$.

Lemma B.2 (Satellite Symbol Sequences). — *For every satellite $\mathcal{P} * \mathcal{S}(q'/r')$ of \mathcal{P} , the symbol sequence $\sigma(\mathcal{P} * \mathcal{S}(q'/r'))$ coincides with the satellite sequence $\sigma^*(\mathcal{P})$. The translation number $\text{Trans}(\sigma(\mathcal{P}), \tau)$ is constant for τ in the characteristic arc $\mathcal{I}_{\mathcal{P}}$, while $\text{Trans}(\sigma^*(\mathcal{P}), \tau)$ increases by $+1$ as τ increases through $\mathcal{I}_{\mathcal{P}}$, taking the value $q'/r' \pmod{\mathbb{Z}}$ on the characteristic arc of $\mathcal{P} * \mathcal{S}(q'/r')$.*

Intuitively, if we tune a map in $H_{\mathcal{P}}$ by a map in $H_{\mathcal{S}(q'/r')}$ then we must replace the Fatou component containing the critical point for the first map by a small copy of the filled Julia set for a (q'/r') -rabbit. Here the period p point z_0 for \mathcal{P} corresponds to the β -fixed point of this small rabbit, while the period n point z_0 for $\mathcal{P} * \mathcal{S}(q'/r')$ corresponds to the α fixed point for this rabbit. Perturbing out of the connectedness locus M , these two points will be separated by the ray pair terminating at the critical point. Further details will be omitted. \square

For example, starting with $\sigma(\{\{0\}\}) = \overline{0}$, where the overline indicates infinite repetition, we find that

$$\sigma(\mathcal{S}(q/r)) = \sigma^*(\{\{0\}\}) = \overline{1},$$

while

$$\sigma^*(\mathcal{S}(1/2)) = \overline{01}, \quad \sigma^*(\mathcal{S}(q/3)) = \overline{011}, \quad \sigma^*(\mathcal{S}(q/4)) = \overline{0111}, \quad \dots$$

We can use this discussion to provide a different insight on the counting argument of §5. Since $\text{Trans}(\sigma^*(\mathcal{P}); \tau)$ increases by $+1$ on the characteristic arc $\mathcal{I}_{\mathcal{P}}$, we see that the total number of portraits (or the total number of characteristic arcs) with ray period $rp = n$ is equal to the sum of $b_0 + \dots + b_{n-1}$ taken over all cyclic equivalence classes of symbol sequences of period exactly n . But the number of such symbol sequences, up to cyclic permutation, is $\nu_2(n)/n$, and the average value of $b_0 + \dots + b_{n-1}$ is equal

to $n/2$, since each symbol sequence with sum different from $n/2$ has an opposite with zero and one interchanged. Therefore, this sum is equal to $\nu_2(n)/2$, as in §5.

Examples. — (Compare Figure 4). Here is a list for all cyclic equivalence classes of symbol sequences of period at most four:

Trans(0; τ) is identically zero.

Trans(1; τ) increases from 0 to 1 for $0 \leq \tau \leq 1$, taking the value q/r in the characteristic arc for $\mathcal{S}(q/r)$.

Trans(1, 0; τ) increases from 0 to 1 as τ passes through $(1/3, 2/3)$, the characteristic arc for $\mathcal{S}(1/2)$.

Trans(1, 0, 0; τ) increases from 0 to 1 as τ passes through the characteristic arc $(3/7, 4/7)$ for the period 3 portrait with root point $c = -1.75$.

Trans(1, 1, 0; τ) increases by one in the arc $(1/7, 2/7)$ for $\mathcal{S}(1/3)$, and by one more in the arc $(5/7, 6/7)$ for $\mathcal{S}(2/3)$. (Compare Figure 21.)

Trans(1, 0, 0, 0; τ) increases by one in the arc $(7/15, 8/15)$, corresponding to the leftmost period 4 component on the real axis.

Trans(1, 1, 0, 0; τ) increases by one in the arcs $(1/5, 4/15)$ and $(11/15, 4/5)$ associated with the period 4 components on the $1/3^{\text{rd}}$ and $2/3^{\text{rd}}$ limbs. (Figure 12.)

Trans(1, 1, 1, 0; τ) increases by one in the arcs $(1/15, 2/15)$ and $(13/15, 14/15)$ for $\mathcal{S}(1/4)$ and $\mathcal{S}(3/4)$, and also in the arc $(2/5, 3/5)$ for the portrait $\mathcal{S}(1/2) * \mathcal{S}(1/2)$ with root point -1.25 .

References

- [A] P. Atela, *Bifurcations of dynamic rays in complex polynomials of degree two*, Ergod. Th. Dynam. Sys. **12** (1992) 401–423.
- [Be] A. Beardon, “Iteration of Rational Functions”, Springer 1991.
- [BFH] B. Bielefeld, Y. Fisher and J. Hubbard, *The classification of critically preperiodic polynomials as dynamical systems*, Jour. AMS **5** (1992) 721–762.
- [Bou] T. Bousch, *Sur quelques problèmes de la dynamique holomorphe*, Thèse, Université de Paris-Sud (1992).
- [Boy1] P. Boyland, *An analog of Sharkovski’s theorem for twist maps*, Comtemp. Math. **81** (1988) 119–133.
- [Boy2] P. Boyland, *Topological methods in surface dynamics*, Topology and its Applications **58** (1994) 223–298.
- [Br] B. Branner, *The Mandelbrot set*, pp. 75–105 of “Chaos and Fractals, the Mathematics behind the Computer Graphics”, Proc. Symp. Appl. Math **39**, AMS 1989.
- [CG] L. Carleson and T. Gamelin, “Complex Dynamics”, Springer 1993.
- [CM] P. Cvitanović and J. Myrheim, *Complex universality*, Comm. Math. Phys. **121** (1989) 225–254.
- [D1] A. Douady, *Systèmes dynamiques holomorphes*, Séminaire Bourbaki, 35^e année 1982–83, n° **599**; Astérisque **105-106** (1983) 39–63.
- [D2] A. Douady, *Julia sets and the Mandelbrot set*, pp. 161–173 of “The Beauty of Fractals”, edit. Peitgen and Richter, Springer (1986)

- [D3] A. Douady, *Algorithms for computing angles in the Mandelbrot set*, pp. 155-168 of "Chaotic Dynamics and Fractals", ed. Barnsley & Demko, Acad. Press (1986)
- [D4] A. Douady, *Chirurgie sur les applications holomorphes*, Proc. Int. Cong. Math. Berkeley, A.M.S. (1987) 724-738.
- [D5] A. Douady, *Descriptions of compact sets in \mathbb{C}* , pp. 429-465 of "Topological Methods in Modern Mathematics", Publish or Perish, Houston (1992)
- [DH1] A. Douady and J. Hubbard, *Itération des polynômes quadratiques complexes*, Comptes Rendus Acad. Sci. Paris **294** (1982) 123-126.
- [DH2] A. Douady and J. Hubbard, *Etude dynamique des polynômes complexes*, Publications Math. d'Orsay **84-02** (1984) (première partie) and **85-04** (1985) (deuxième partie).
- [DH3] A. Douady and J. H. Hubbard, *On the dynamics of polynomial-like mappings*, Ann. Sci. Ec. Norm. Sup. (Paris) **18** (1985) 287-343.
- [E] D. Eberlein, *Rational Parameter Rays of Multibrot Sets*, Diploma Thesis, Technische Universität München, in preparation
- [G] L. Goldberg, *Fixed points of polynomial maps, I*, Ann. Scient. Ec. Norm. Sup., 4^e série, **25** (1992) 679-685.
- [GF] D. Giarrusso and Y. Fisher, *A parametrization of the period 3 hyperbolic components of the Mandelbrot set*, Proc. Amer. Math. Soc. **123** (1995) 3731-3737.
- [GM] L. Goldberg and J. Milnor, *Fixed points of polynomial maps, II*, Ann. Scient. Ec. Norm. Sup., 4^e série, **26** (1993) 51-98.
- [Ha] P. Haïssinsky, *Chirurgie parabolique*, Comptes Rendus Acad. Sci. Paris, **327** (1998) 195-198.
- [Hu] J. Hubbard, *Local connectivity of Julia sets and bifurcation loci: three theorems of J.-C. Yoccoz*, pp. 467-511 and 375-378 of "Topological Methods in Modern Mathematics", Publish or Perish, Houston (1992)
- [HS] J. Hubbard and D. Schleicher, *The Spider Algorithm*, "Complex Dynamics: the Mathematics Behind the Mandelbrot and Julia sets", pp. 155-180 of Proc. Symposia in Appl. Math. **49** AMS 1994.
- [K] J. Kiwi, *Rational Rays and Critical Portraits of Complex Polynomials*, thesis, Stony Brook 1997; Stony Brook IMS Preprint 1997#15.
- [LS] E. Lau and D. Schleicher, *Internal Addresses in the Mandelbrot Set and Irreducibility of Polynomials*, Stony Brook IMS Preprint 1994#19.
- [La] P. Lavaurs, *Une description combinatoire de l'involution définie par M sur les rationnels à dénominateur impair*, Comptes Rendus Acad. Sci. Paris (4) **303** (1986) 143-146.
- [Le] G. Levin, *On the complement of the Mandelbrot set*, Israel J. Math. **88** (1984) 189-212.
- [Lo] J. Los, *On the forcing relation for surface homeomorphisms*, preprint, Nice 1994.
- [Ly] M. Lyubich, *Dynamics of quadratic polynomials, I, II*, Acta Math. **178** (1997) 185-297.
- [Mc] C. McMullen, "Complex Dynamics and Renormalization," Annals of Math. Studies 135, Princeton U. Press 1994.
- [M1] J. Milnor, *Self-similarity and hairiness in the Mandelbrot set*, pp. 211-257 of "Computers in Geometry and Topology", edit. Tangora, Lect. Notes Pure Appl. Math. **114**, Dekker 1989 (cf. p. 218 and §5).
- [M2] J. Milnor, *Dynamics in One Complex Variable: Introductory Lectures*, Stony Brook IMS Preprint 1990#5.

- [M3] J. Milnor, *Local Connectivity of Julia Sets: Expository Lectures*, Stony Brook IMS Preprint 1992#11
- [M4] J. Milnor, *On cubic polynomials with periodic critical point*, in preparation.
- [PR] C. L. Petersen and G. Ryd, *Convergence of rational rays in parameter space*, preprint, Roskilde Univ. 1998.
- [P] A. Poirier, *On postcritically finite polynomials*, parts I and II, Stony Brook IMS Preprint 1993#5 and #7.
- [RS] J. Riedl and D. Schleicher, *On the locus of crossed renormalization*, Proceedings 1997 conference on holomorphic dynamics, Kyoto 1998.
- [S1] D. Schleicher, *Internal addresses in the Mandelbrot set and irreducibility of polynomials*, Thesis, Cornell 1994.
- [S2] D. Schleicher, *Rational external rays of the Mandelbrot set*, to appear in *Asterisque*; See also Stony Brook IMS Preprint 1997#13.
- [S3] D. Schleicher, *On fibers and local connectivity of Mandelbrot and Multibrot sets*, Stony Brook IMS Preprint 1998#13.
- [Sh1] M. Shishikura, *On the quasiconformal surgery of rational functions*, Ann. Sci. École Norm. Sup. **20** (1987), 1–29.
- [Sh2] M. Shishikura, *The Hausdorff dimension of the boundary of the Mandelbrot set and Julia sets*. Ann. of Math. **147** (1998), 225–267.
- [Sø] D. Sørensen, *Complex dynamical systems: rays and non-local connectivity*, thesis, Techn. Univ. Denmark 1994.
- [St] N. Steinmetz, “Rational Iteration”, de Gruyter 1993.
- [Ta] Tan Lei, *Local properties of M at parabolic points*, in preparation.
- [Th] W. Thurston, *On the Geometry and Dynamics of Iterated Rational Maps*, Preprint, Princeton University.

J. MILNOR, Institute for Mathematical Sciences, State University of New York, Stony Brook, NY 11794-3660 • E-mail : jack@math.sunysb.edu • Url : <http://www.math.sunysb.edu/~jack>