Attracting Dynamics of Exponential Maps

Dierk Schleicher

Abstract

We give a complete classification of hyperbolic components in the space of iterated maps $z \mapsto \lambda \exp(z)$, and we describe a preferred parametrization of those components. This leads to a complete classification of all exponential maps with attracting dynamics.

Contents

1 Introduction

This paper is part of the program to describe the dynamics of exponential maps $z \mapsto \lambda e^z$ and the structure of parameter space, in the spirit of the well-developed body of knowledge about polynomial dynamics. The polynomial theory was pioneered by Douady and Hubbard [DH] who systematically investigated the Mandelbrot set as the simplest non-trivial example of a holomorphic parameter space. Since then, there has been a lot of further work in this field.

The description of the exponential parameter space was begun in the 1980's by Baker and Rippon [BR], by Eremenko and Lyubich [EL1, EL2, EL3], and by Devaney, Goldberg and Hubbard [DGH]. These papers discuss certain fundamental properties of hyperbolic components and of bifurcations (in the case of [EL1, EL2, EL3] as an example of a study of finite type entire maps), but a description of the global structure of parameter space

Stony Brook IMS Preprint #2000/4
April 2000

was in terms of pictures and conjectures. In this paper, we give a complete description of hyperbolic components of the exponential parameter space. This was part of Chapter III of the author's habilitation thesis [S1] (of May, 1999) which developed a description of the exponential parameter space in analogy to Douady and Hubbard's Orsay Notes [DH] about the Mandelbrot set.

Our object is to classify hyperbolic components in the λ parameter plane, where λ ranges over $\mathbb{C} \setminus \{0\}$. It is known from the papers cited above that there is a unique hyperbolic component which is bounded, having period 1, and a unique hyperbolic component which is bounded to the right but unbounded to the left and in the imaginary direction, having period 2. All other hyperbolic components have period 3 or more, and are unbounded to the right. Every hyperbolic component is simply connected, except that the period one component is punctured at 0.

Here is our main result; it is illustrated in Figure 1.

Theorem 1.1 (Classification of Hyperbolic Components)
For every period $n \geq 3$, there are countably many hyperbolic components in the space of exponential maps parametrized by λ. Each of them is characterized by a sequence

$$s_1, s_2, \ldots, s_{n-1}$$

(its "intermediate external addresses"), where $s_1 = 0$, $s_2, \ldots, s_{n-2} \in \mathbb{Z}$ and $s_{n-1} \in (\mathbb{Z} + \frac{1}{2})$. Conversely, every such sequence is realized by a unique hyperbolic component of period n. These hyperbolic components have a natural vertical order in which they stretch out to $+\infty$ along bounded imaginary parts, and this order is the same as the lexicographic order of the sequences $s_1, s_2, \ldots, s_{n-1}$. In particular, between any pair of consecutive hyperbolic components of period n, there are infinitely many hyperbolic components of period $n+1$, ordered like $\mathbb{Z}$.

The numbers $s_1, s_2, \ldots, s_{n-1}$ characterizing any hyperbolic component of period n have a dynamic meaning as follows. Let λ be any parameter in the given period n hyperbolic component, and let

$$U_1 \overset{\approx}{\to} U_2 \overset{\approx}{\to} \ldots \overset{\approx}{\to} U_n$$

be the unique cycle of periodic Fatou components for E_λ, where $0 \in U_1$, where $\lambda \in U_2$, and where U_n contains a left half plane. Then for $1 \leq k < n-1$ the points in U_k are asymptotic to the line

$$L(s_k) := \{z \in \mathbb{C} \colon \mathrm{Im}(z) = 2\pi s_k - c\}$$

as $\mathrm{Re}(z) \to \infty$; here c is the imaginary part of $\log(\lambda)$, choosing the branch with $|c| < \pi$. For $k = n-1$ the points of U_{n-1} form a neighborhood of the line $L(S_{n-1})$ for $\mathrm{Re}(z)$ sufficiently large. Thus s_k specifies precisely which branch of E_λ^{-1} carries U_{k+1} to U_k.

Similarly, in the λ parameter plane, if $n > 3$ then the points in the hyperbolic component are asymptotic to the line

$$\mathrm{Im}(\lambda) = 2\pi s_2 \quad \text{as} \quad \mathrm{Re}(\lambda) \to +\infty \ ,$$

while for $n = 3$ they form a neighborhood of this line near $+\infty$.

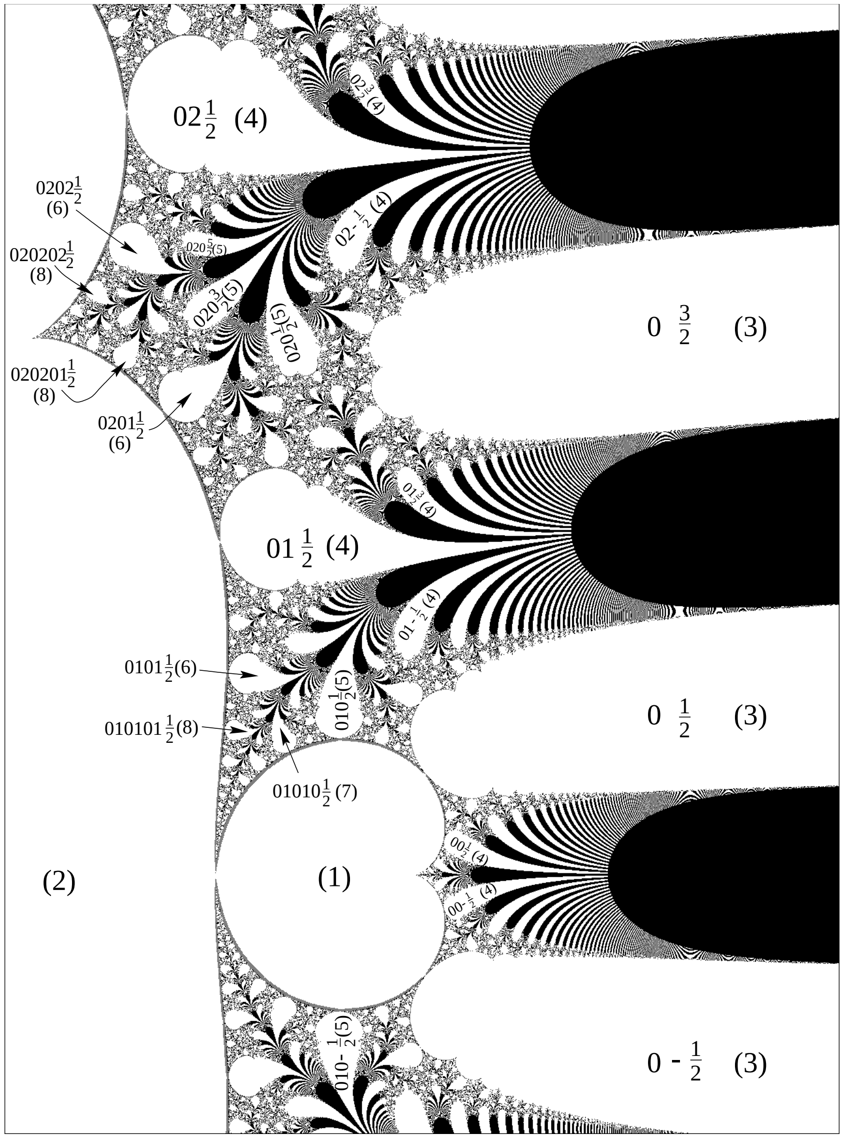

Figure 1: *The space of parameters λ for exponential maps λ, with hyperbolic components indicated in white. Various hyperbolic components are labeled by their intermediate external addresses, or briefly by their periods (in parentheses). The picture has kindly been contributed by Jack Milnor: for every pixel, an approximate test is performed whether or not the corresponding map E_λ has an attracting orbit (with λ at the center of the pixel); in addition, the boundaries of hyperbolic components have been emphasized in order to show their shapes more clearly.*

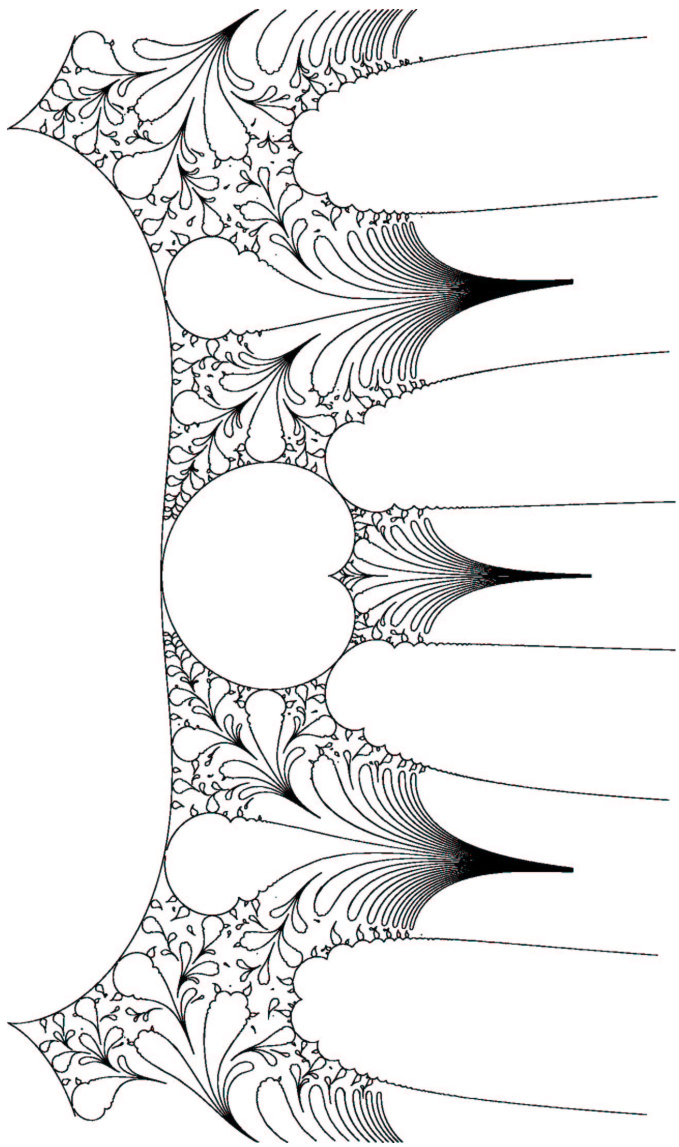

Figure 2: *The same space as in Figure 1, drawn differently by a special purpose program of Günter Rottenfußer: this program traces out the boundaries of hyperbolic components, which is possible with arbitrary precision for any given hyperbolic component. Unlike for the Mandelbrot set, it makes no sense in the exponantial case to test whether the singular orbit "escapes to* ∞*"; instead, in pixel images it is usually tested whether the singular orbit survives some fixed number* N *of iterations without producing numbers too large to store. This is quite different from the existence of an attracting orbit for the given value of* λ*, and logically independent. This picture confirms that pixel test pictures like in Figure 1 are approximately correct.*

Furthermore, if H_1 and H_2 are hyperbolic components of any periods greater than 2, then H_1 lies above H_2 if and only if its symbol sequence is greater, using lexicographic ordering.

In Section 2, we review necessary properties about exponential maps and state results from earlier papers. In particular, we introduce dynamic rays. Then, in Section 3, we give a combinatorial coding to every hyperbolic component in terms of "intermediate external addresses", and we show that each intermediate external address is realized by at least one hyperbolic component. The converse needs a deeper understanding of exponential dynamics and in particular the interplay between attracting Fatou components and dynamic rays. In Section 4, we investigate the combinatorial properties periodic dynamic rays landing at a common point and show that whenever at least three rays land together, the first return map of the landing point permutes them cyclically. Section 5 proves the existence of "dynamic roots" of periodic Fatou components for attracting cycles: these are boundary points which are fixed under the first return map of the Fatou component such that at least two periodic dynamic rays land at this point. Every periodic Fatou component (of period $n \geq 2$) has a unique dynamic root, which helps to break the symmetry of Fatou components and makes a unique coding possible. Using these tools, we can prove uniqueness of the hyperbolic component associated to any intermediate external address (Section 6), and we get a preferred parametrization within each hyperbolic component. Finally, in Section 7, we outline some further related results of [S1] which will be published separately.

Some notation. We write our exponential maps as $z \mapsto E_\lambda(z) := \lambda e^z = \exp(z + \kappa)$ with $\lambda = \exp(\kappa)$, where $\lambda \in \mathbb{C} \setminus \{0\}$ and $\kappa \in \mathbb{C}$. We will often need $F(t) = e^t - 1$, in particular for $t \in \mathbb{R}$. Let $\mathbb{D} := \{z \in \mathbb{C} : |z| < 1\}$ and $\mathbb{D}^* := \mathbb{D} \setminus \{0\}$. We write that a curve or sequence in $\mathbb{C}$ converges to $+\infty$ or to $-\infty$ to indicate that the real parts converge to $\pm\infty$, while the imaginary parts are bounded.

Acknowledgments. I would like to thank Misha Lyubich and Jack Milnor for many helpful and inspiring discussions, and to the Institute for Mathematical Sciences at Stony Brook for continued support and encouragement. I am also grateful for several helpful comments from the audience at various seminars where these results were presented, in particular in Stony Brook (spring 1999 and spring 2000) and at Boston University (spring 1999).

2 Exponential Dynamics

In this section, we will review known properties of exponential dynamics. Points z with $E_\lambda^{\circ k}(z) \to \infty$ are known as *escaping points*; they are completely classified [SZ2]: if the singular value itself does not escape, then the escaping points are on disjoint curves called *dynamic rays* (or *hairs*) labeled by *external addresses* $\underline{s} = s_1 s_2 s_3 \ldots$, which are infinite sequences over $\mathbb{Z}$ (if the singular value does escape, then there is a well-understood exception). The dynamic ray at external address $\underline{s}$ is an injective curve $g_{\underline{s}}\colon\ ]t_{\underline{s}}, \infty[\ \to \mathbb{C}$ (or $g_{\underline{s}}\colon\ [t_{\underline{s}}, \infty[\ \to \mathbb{C}$) with $\mathrm{Re}(g_{\underline{s}}(t)) \to +\infty$ as $t \to \infty$, while $\mathrm{Im}(g_{\underline{s}}(t))$ is bounded. Every point on a dynamic ray is an escaping point, and every escaping point is on such a ray. We have the dynamic relation

$$E_\lambda(g_{\underline{s}}(t)) = g_{\sigma(\underline{s})}(F(t))$$

where σ is the shift map on external addresses, dropping the first entry. The quantity $t_{\underline{s}} \geq 0$ depends on $\underline{s}$ in a well-understood way. The meaning of the external address of a ray is the following: the set $E_\lambda^{-1}(\mathbb{R}^-)$ is a countable union of vertical lines, spaced at distance $2\pi i\mathbb{Z}$, and $\mathbb{C} \setminus E_\lambda^{-1}(\mathbb{R}^-)$ are horizontal strips, labeled by $\mathbb{Z}$ so that the strip with label 0 contains the singular value 0 (perhaps on its boundary). Then at least for sufficiently large $t > t_{\underline{s}}$, the external address $\underline{s}$ of $g_{\underline{s}}$ is the sequence $s_1 s_2 s_3 \dots$ of strips visited by the orbit of $g_{\underline{s}}(t)$. Not all possible sequences are allowed; the set of allowed sequences is completely understood, and it contains all bounded sequences. In this paper, we need only rays $g_{\underline{s}}$ for which $\underline{s}$ is bounded; in this case, $t_{\underline{s}} = 0$. We say that such a ray *lands* at a point $w \in \mathbb{C}$ if $\lim_{t \searrow 0} g_{\underline{s}}(t)$ exists and is equal to w.

If an exponential map has an attracting periodic point, then the singular value is in a periodic Fatou component which we call the *characteristic Fatou component*. All periodic points, except the unique attracting one, are repelling. We will need a construction and results from [SZ1, Section 4.3]: let $n \geq 2$ be the period of the attracting orbit, let $U_1, U_2, \dots, U_n = U_0$ be the cycle of periodic Fatou components, labeled cyclically modulo n so that U_1 is the characteristic Fatou component, and let $a_1, a_2, \dots, a_n$ be the attracting periodic orbit labeled so that $a_k \in U_k$ for all k. Let V_{n+1} be a closed neighborhood of a_1 corresponding to a disk in linearizing coordinates, large enough so as to contain the singular value in its interior. Let

$$V := \bigcup_{k=0}^{n} \left\{ z \in U_k : E_\lambda^{\circ(n+1-k)}(z) \in V_{n+1} \right\} .$$

Since $E_\lambda^{-1}(V_{n+1}) \subset U_n$ contains a left half plane, $V \cap U_k$ (for $k = n-1, n-2, \dots, 1$) contains a band towards $+\infty$, and $V \cap U_0$ contains infinitely many bands towards $+\infty$, spaced equally at integer translates of $2\pi i$. The construction assures that $E_\lambda(V) \subset V$ and that all $V \cap U_k$ are connected and simply connected.

Let $R := \mathbb{C} \setminus (V \cap U_0)$; it consists of countably many connected components which we will call "regions" and denote $R_{\mathtt{j}}$: let $R_{\mathtt{0}}$ be the connected component containing the singular value and $R_{\mathtt{j}} := R_{\mathtt{0}} + 2\pi i \mathtt{j}$, for $j \in \mathbb{Z}$. Then $R = \cup_{\mathtt{j}} R_{\mathtt{j}}$. Any orbit (z_k) within the Julia set then has an associated *itinerary* $\mathtt{u}_1, \mathtt{u}_2, \mathtt{u}_3, \dots$, where $\mathtt{u}_k = \mathtt{j}$ iff $z_k \in R_{\mathtt{j}}$. We should emphasize that this itinerary is different from the external address used for example in the construction of dynamic rays: the external address is constructed using inverse images of the negative real axis, which is dynamically not a natural concept; however, the itineraries as defined here are dynamically natural; compare the discussion in [SZ1, Sections 4 and 5]. (Note the different fonts for external addresses $s_1 s_2 \dots$ and itineraries $\mathtt{u}_1 \mathtt{u}_2 \dots$.)

For exponential dynamics with an attracting periodic orbit, it is shown in [SZ1] that every periodic dynamic ray lands at a repelling periodic point (Theorem 3.1); that every repelling periodic point is the landing point of a finite positive number of periodic rays (Theorem 5.3); and that a periodic ray lands at a periodic point if and only if ray and point have identical itineraries (Proposition 4.5). In particular, different periodic points have different itineraries, while different rays have the same itinerary if and only if they land together.

The set of parameters λ or κ for which there is a (necessarily unique) attracting periodic orbit is known to be open; a connected component where this happens will be called a *hyperbolic component* (this is a slight abuse of notation: hyperbolic dynamics in a strict sense would require a uniformly expanding metric in a neighborhood of the Julia set, but the Julia set is never compact for exponential maps).

The period of the attracting orbit is necessarily constant throughout any hyperbolic component; we will show that (except for low periods) the number of hyperbolic components for any given period is always countably infinite. The following results are known from [EL2, EL3, BR, DGH]: in λ-space, there is a unique hyperbolic component of period 1 which is a bounded neighborhood of the puncture $\lambda = 0$ of parameter space; there is a unique period 2 component which "almost" occupies a left half plane (in the sense that for every $\vartheta \in\,]\pi/2, 3\pi/2[$, there is an $R > 0$ such that for all $r > R$, the parameter $\lambda = r\exp(i\vartheta)$ has an attracting orbit of period 2). All hyperbolic components of period $n \geq 2$ are simply connected. Similarly, in κ-space, all hyperbolic components are simply connected, the unique period 1 component contains a left half plane, and there are countably many period 2 components.

On any hyperbolic component W, there is an associated *multiplier map* $\mu\colon W \to \mathbb{D}^*$, which is a holomorphic covering map. Indeed, W is simply connected and μ is a universal cover (except for the unique period 1 component in λ-space, for which μ is a conformal isomorphism onto $\mathbb{D}^*$).

Lemma 2.1 (Strong Attraction Only at Far Parameters)
For any period n and any $r < 1$, there is an $R > 0$ such that any parameter κ for which there is an attracting orbit of multiplier $|\mu| \leq r$ has $|\kappa| > R$.

For any hyperbolic component W of period $n \geq 2$, there is a unique homotopy class of curves $\gamma(\,[0,\infty[\,) \to W$ with $\lim_{t\to\infty}\gamma(t) = \infty$ such that $\mu(\gamma(t)) \to 0$ as $t \to \infty$ (where the homotopy should fix the endpoint ∞ but need not fix the other endpoint).

Proof. If $a_1, a_2, \ldots, a_n$ is any periodic orbit of period n under $\lambda\exp$, then $(E_\lambda^{\circ n}(a_m))' = \prod_k a_k$. If the orbit is attracting, then there is some $|a_k| < 1$. Hence if $|\kappa| \leq R$, then all $|a_k|$ are bounded above in terms of R and n. Hence there is a $\xi \in \mathbb{R}$ such that all $\mathrm{Re}(a_k) \geq \xi$, but this implies that $|a_{k+1}| \geq \exp(\xi + \mathrm{Re}(\kappa))$. Hence we have a lower bound for all $|a_k|$ and hence for $|\mu|$.

Any curve $\gamma'\colon(\,[0,\infty[\,) \to \mathbb{D}^*$ with $\lim_{t\to 1}\gamma'(t) = 0$ lifts under any branch of the inverse multiplier map to a curve $\gamma\colon(\,[0,\infty[\,) \to \mathbb{C}$ with $\lim_{t\to 1}\gamma(t) = \infty$ by what we just proved. Conversely, any two curves $\gamma_1, \gamma_2\colon(\,[0,\infty[\,) \to W$ which limit at ∞ project under μ to two curves $\gamma_1', \gamma_2'\colon\ [0,\infty[\, \to \mathbb{D}^*$, and these are homotopic; hence γ_1 and γ_2 are also homotopic. □

Remark. It would be conceivable that there are different homotopy types of curves within any hyperbolic component which tend to ∞: some branch of $\mu^{-1}(\,[e^{i\vartheta}/2, e^{i\vartheta}[\,)$ could tend to ∞ as $|\mu| \to 1$. It was conjectured by Eremenko and Lyubich [EL2] that this does not happen; a proof of this conjecture was given in [S1, Section V].

3 Classification of Attracting Components

Hyperbolic components of Multibrot sets have the helpful property that they have a unique "center" in which the dynamics is postcritically finite. If an exponential map has an attracting orbit, it can never be postsingularly finite; the "center has disappeared to $+\infty$". Fair enough, it turns out that hyperbolic components of exponential maps have another feature unknown to the polynomial case: since they stretch out to $+\infty$ like parameter rays, they can be described by a slight generalization of external addresses: we need finite sequences of integers, followed by a symbol $\pm\infty$.

Definition 3.1 (Intermediate External Address)
An intermediate external address *of period* $n \geq 2$ *is a sequence* $s_1 s_2 \dots s_{n-2} s_{n-1}$ *with* $s_k \in \mathbb{Z}$ *for* $k \leq n-2$ *and* $s_{n-1} \in (\mathbb{Z} + \frac{1}{2})$.

The lexicographic order on external addresses (infinite sequences over $\mathbb{Z}$) extends naturally to intermediate external addresses such as $\underline{s} = s_1 s_2 \dots s_{n-1}$. Note that an intermediate external address of period n has length $n-1$ with exactly the last entry a half-integer; it labels hyperbolic components which have an attracting orbit of period n, but it is not periodic itself

As usual, we start with a dynamic consideration.

Definition 3.2 (Attracting Dynamic Ray)
Consider an exponential map E_λ *with an attracting orbit of period* $n \geq 2$ *and let* $\underline{s} = s_1 s_2 \dots s_{n-1}$ *be an intermediate external address of period* n*. We say* E_κ *has an* attracting dynamic ray at external address $\underline{s}$ *if there is a curve* γ: $[0, \infty[\to \mathbb{C}$ *within the characteristic Fatou component such that the following hold:*

- $\gamma(0)$ *is on the attracting periodic orbit;*
- $\lim_{t\to\infty} E_\lambda^{\circ k} \gamma(t) = +\infty$ *for* $k = 0, 1, \dots, n-2$*;*
- $\lim_{t\to\infty} E_\lambda^{\circ(n-1)} \gamma(t) = -\infty$*;*
- *every dynamic ray at an external address* $\underline{s}' < \underline{s}$ *is below* γ*;*
- *every dynamic ray at an external address* $\underline{s}' > \underline{s}$ *is above* γ*.*

REMARK. Since γ is in a Fatou component, it must be disjoint from any dynamic ray. Since γ and any dynamic ray both tend to $+\infty$ at bounded imaginary parts, the ray must be *above* or *below* γ in the following sense: for real ξ sufficiently large, γ cuts the half plane $\{z \in \mathbb{C} : \mathrm{Re}(z) > \xi\}$ into two unbounded parts, one above and one below γ, and any dynamic ray must tend to $+\infty$ within one of these two parts.

Lemma 3.3 (Attracting Dynamics Has External Addresses)
For every attracting exponential dynamics of period $n \geq 2$*, there is an intermediate external address* $\underline{s} = s_1 s_2 \dots s_{n-1}$ *of period* n *such that there is an attracting dynamic ray at this external address. This external address is unique up to changing of the combinatorics (adding a common constant to all* s_k*), and it is the same for any parameter from the same hyperbolic component.*

Remark. We can make the intermediate external address unique by choosing the combinatorics so that $s_1 = 0$.

Proof. There is a unique periodic Fatou component which contains a left half plane. By simple connectivity, it contains a unique homotopy class of curves connecting an attracting periodic point to $-\infty$ eventually within a left half plane (or even eventually along $\mathbb{R}^-$). This homotopy class of curves can be pulled back $n-1$ steps to a homotopy class of curves within the characteristic Fatou component connecting the attracting periodic point to $+\infty$. Choose one such curve γ. This curve avoids dynamic rays, and it is easy to check that the supremum of external addresses of rays below γ has well-defined $n-1$ initial entries in $\mathbb{Z}$ (these rays need not run below the entire Fatou component): the curve γ and its first $n-2$ iterates tend to $+\infty$ (with bounded imaginary parts), so the first $n-1$ entries in the supremum are just the labels of the strips containing the iterates of γ (with respect to inverse images of $\mathbb{R}^-$ used in the construction of external addresses). Similarly, the infimum of external addresses of dynamic rays above γ supplies $n-1$ well-defined first entries which differ from the lower external addresses only in the last entry and only by one. It is easy to confirm that the external address does not depend on the choice of γ or on the parameter chosen from its hyperbolic component. □

We thus have a combinatorial coding for every hyperbolic component (except for the unique component of period 1), and our goal is to show that each coding is realized by exactly one hyperbolic component. This will be done in Theorem 3.5 (existence) and Corollary 6.3 (uniqueness).

First we need a lemma to prove the existence of an attracting orbit.

Lemma 3.4 (Singular Orbit in Horizontal Strip)
Suppose that for some parameter λ there is a real number $h > 3$ such $\operatorname{Re}(\lambda) > h$ and the initial segment $z_1 = 0$, $z_2 = \lambda$, ..., z_n of the singular orbit has the property that $|\operatorname{Im}(z_k)| < h$ for $1 \le k \le n$. Suppose moreover that z_n is the first point on the singular orbit with $\operatorname{Re}(z_n) < 0$. Then the map E_λ has an attracting periodic orbit of exact period n, and the attracting basin contains the left half plane $\operatorname{Re}(z) \le \operatorname{Re}(z_n) + 1$. As $\operatorname{Re}(\lambda) \to \infty$ with fixed height h of the strip, the multiplier tends to 0.

The proof needs a couple of unpleasant calculations, but its idea is very simple: the geometry of the strip containing the singular orbit assures that absolute values of orbit points are dominated by the real parts, and the real parts grow exponentially. Once the orbit reaches a point with negative real part, its absolute value dominates the remaining orbit by far. The contraction coming from the exponential map at this point is far greater than the expansion along the previous orbit, starting at the singular value 0. Therefore, any sufficiently small disk around 0 will map after n steps to a much smaller (almost-) disk close to the origin. In order to map this disk into itself, its size has to be chosen so that it is neither too large (or we would lose control in the estimates) nor too small (or it would not contain the images after n steps). It turns out that it works if we choose the disk so that its image at z_n has radius 1.

Proof. The points $z_2, \ldots, z_{n-1}$ of the orbit are contained within the strip $S := \{z \in \mathbb{C} : |\operatorname{Im}(z)| < h\}$ at positive real parts. We show that they all have real parts greater than h.

Indeed, this is true for $z_2 = \lambda$ by assumption, and for the others it follows by induction using $|\lambda| > h > 1$: $|z_k| = |\lambda| \exp(\text{Re}(z_{k-1})) > he^h$, so $|\text{Re}(z_k)| > h$. We also get $\text{Re}(z_n) < -h$.

Now we show for $n' \leq n$

$$\prod_{k=2}^{n'} (|z_k| + 1) < (|z_{n'}| + 1)^2 \ . \tag{1}$$

Indeed, for $n' = 1$ the empty product on the left equals 1, while $\text{Re}(z_1) > 3$ by assumption. For the inductive step, we only need to prove $(|z_k| + 1)^2 < |z_{k+1}| + 1$. We will use the inequality $(\sqrt{2}x + 1)^2 < 3\exp(x)$ for all $x \geq 2$ and estimate for $\text{Re}(z_k) > 0$ as follows:

$$\begin{aligned} (|z_k| + 1)^2 &< (\sqrt{2}\text{Re}(z_k) + 1)^2 \\ &< 3\exp(\text{Re}(z_k)) < |\lambda||\exp(z_k)| < |z_{k+1}| + 1 \ . \end{aligned}$$

Our next claim is about z_n, the first point with negative real part:

$$e|\lambda| \exp(\text{Re}(z_n)) < (|z_n| + 1)^{-2} \ . \tag{2}$$

Indeed, we have $|z_n| = |\lambda| \exp(\text{Re}(z_{n-1})) > he^h > 3\exp(3) > 60$, thus $|\text{Re}(z_n)| > he^h/\sqrt{2} > 40$ and $|z_n| + 1 < \sqrt{2}|\text{Re}(z_n)|$. Using the inequality $2ex^3 < e^x$ for all $x \geq 8$, it follows

$$e|\lambda|(|z_n| + 1)^2 \leq 2e|\lambda||\text{Re}(z_n)|^2 \leq 2e|\text{Re}(z_n)|^3 < \exp(|\text{Re}(z_n)|) \ .$$

Since the real part of z_n is negative, the claim follows.

Now we can start the actual proof of the lemma. Let D_n be the open disk of radius 1 around the point z_n. Pulling back by the dynamics, we obtain open neighborhoods D_{n-1} around z_{n-1}, ..., D_1 around $z_1 = 0$. These pull-backs are contracting at every step: the derivative of E_λ at a point $z'_k \in D_k$ is equal to $E_\lambda(z'_k)$, and its absolute value is bounded above by $|z_{k+1}| + 1$. The inverse map is thus contracting with contraction factor at most $1/(|z_{k+1}| + 1)$, and the domain D_1 contains a disk around the origin with radius at least $\rho = \prod_{k=2}^{n}(|z_k| + 1)^{-1} > (|z_n| + 1)^{-2}$ by Equation (1) above.

On the other hand, all the points in D_n are contained in the left half plane $\text{Re}(z) \leq \text{Re}(z_n) + 1$. The image points of this half plane have distance to the origin at most $|\lambda| \exp(\text{Re}(z_n) + 1) = e|\lambda| \exp(\text{Re}(z_n))$. All these image points are thus contained in D_1 by Equation (2) above, and the n-th iterate of the dynamics sends D_0 strictly into itself (even injectively). Therefore, there is an attracting orbit of period at most n. The period clearly cannot be smaller than n. If within the same strip with imaginary parts bounded by h, $\text{Re}(\lambda)$ becomes large, the size of the image of D_n within D_1 gets much smaller compared to the size of D_1, and the multiplier tends to 0. □

Now we come to the existence theorem. We restrict to periods $n \geq 3$ because the hyperbolic components of periods 1 and 2 are completely classified: in λ- and κ-space, there is a unique hyperbolic component of period 1 which would be labeled by the empty intermediate external address. For period 2, the unique hyperbolic component in λ-space. It is labeled $\mathbb{Z} + \frac{1}{2}$ and contains a left end of $\mathbb{R}^-$. Hence, in κ-space, for each $n \in \mathbb{Z}$ there is a unique hyperbolic component of period 2 which contains a right end of the horizontal line at imaginary part $2\pi(n + \frac{1}{2})$; it is labeled $n + \frac{1}{2}$.

Theorem 3.5 (Existence of Hyperbolic Components)
For every $n \geq 3$ and every intermediate external address $\underline{s} = s_1 s_2 \ldots s_{n-1}$ of period n, there is a hyperbolic component (in λ- or κ-space) in which every exponential map has an attracting dynamic ray at external address $\underline{s}$ (up to relabeling the combinatorics).

In λ-space, this hyperbolic component contains an analytic curve tending to $+\infty$ with imaginary parts converging to $2\pi s_2$ such that along this curve the multipliers of the attracting orbit tend to 0. In κ-space, the imaginary parts converge to 0.

PROOF. We may relabel the combinatorics by adding $-s_1$ to all entries, so that we may assume $s_1 = 0$. Let $s^{\pm}_{n-1} := s_{n-1} \pm \frac{1}{2} \in \mathbb{Z}$ and define two periodic external addresses of period $n-1$ via $s^- := \overline{s_1 s_2 \ldots s^-_{n-1}}$ and $s^+ := \overline{s_1 s_2 \ldots s^+_{n-1}}$. Let $A := 1 + \max_k\{|s_k|\}$ and $x := 1$.

In [SZ2, Proposition 3.4], the existence of dynamic ray ends $g_{\underline{s}^+}$ and $g_{\underline{s}^-}$ was shown for any parameter $\kappa \in \mathbb{C}$ with $|\mathrm{Im}(\kappa)| \leq \pi$ and for potentials $t \geq x + 2\log(|\kappa| + 3)$, together with the bound

$$g_{\underline{s}^{\pm}}(t) = t - \kappa + 2\pi i s_1 + r_{\underline{s}^{\pm}}(t) \quad \text{with} \quad |r_{\underline{s}^{\pm}}(t)| < 2e^{-t}(|\kappa| + 3 + 2\pi AC')\ ,$$

where $C' < 2.5$ is a universal constant. The same statement with the same bound holds also for all $\sigma^k(\underline{s}^{\pm})$ (replacing s_1 by the appropriate entry, of course). In particular, if we let $t^* := x + 2\log(|\kappa| + 3 + 2\pi AC')$, then

$$|g_{\underline{s}^{\pm}}(t) - (t - \kappa + 2\pi i s_1)| < e^{-(t-t^*)}\ .$$

After $n-2$ iterations, these two rays map to

$$E_\lambda^{\circ(n-2)}(g_{\underline{s}^{\pm}}(\,]t^*, \infty[\,)) = g_{\sigma^{n-2}(\underline{s}^{\pm})}(\,]F^{\circ(n-2)}(t^*), \infty[\,)$$

with

$$g_{\sigma^{n-2}(\underline{s}^{\pm})}(t) = t - \kappa + 2\pi i s^{\pm}_{n-1} + r^{\pm}$$

for $t \geq t^*$ with $|r^{\pm}(t)| < e^{-(t-t^*)}$. Define a curve

$$\gamma'_\kappa\colon\ [F^{\circ(n-2)}(t^*), \infty[\ \to \mathbb{C} \text{ via } \gamma'_\kappa(t) = t - \kappa + 2\pi i s_{n-1}\ ;$$

it has the property that $E_\lambda(\gamma'_\kappa) \subset \mathbb{R}^-$. The construction assures that the two ray ends $g_{\sigma^{n-2}(\underline{s}^{\pm})}(\,]F^{\circ(n-2)}(t^*), \infty[\,)$ are above respectively below γ'_κ (asymptotically by $i\pi$), and all three curves are disjoint. Moreover, any dynamic ray $g_{\underline{s}'}$ with $\underline{s}' < s_{n-1}$ (that is, $s'_1 < s_{n-1}$) is eventually below γ', and if $\underline{s}' > s_{n-1}$ (that is, $s'_1 > s_{n-1}$), then the ray $g_{\underline{s}'}$ is eventually above γ'.

Pulling back along the same branch of $E_\lambda^{\circ(-(n-2))}$, it follows that there is a curve $\gamma_\kappa\colon\ [t^*, \infty[$ "between" $g_{\underline{s}^{\pm}}\ [t^*, \infty[$ with $E_\lambda^{\circ(n-2)}(\gamma_\kappa(t)) = \gamma'_\kappa(F^{\circ(n-2)}(t))$, where the "between" is defined only for large t. More precisely,

$$\text{if } t \geq t^* + \delta, \text{ then } |g_{\underline{s}^+}(t) - \gamma_\kappa(t)| + |g_{\underline{s}^-}(t) - \gamma_\kappa(t)| < \varepsilon \tag{3}$$

because we have $\mathrm{Re}(E_\lambda(g_{\underline{s}^{\pm}}(t^*+\delta))) = \mathrm{Re}(g_{\sigma(\underline{s}^{\pm})}(F(t^*+\delta))) = F(t^*+\delta) - \mathrm{Re}(\kappa) + r^{\pm} > \delta$, and further iterates move much farther to the right, so $\left|(E_\lambda^{\circ(n-2)})'(g_{\underline{s}^{\pm}}(t^* + \delta))\right| \gg 1$; the ε-bound now follows from the Koebe distortion theorem and our estimate $|\gamma'_\kappa(t) - g_{\sigma^{n-2}(\underline{s}^{\pm})}(t)| < \pi + 1$.

The curve γ_κ clearly satisfies the second and third conditions for attracting dynamic rays; the last two are asymptotically satisfied in the sense that for every $\underline{s}'$, the ray $g_{\underline{s}'}$ is above or below γ_κ (as needed) for sufficiently large t depending on $\underline{s}'$. The first condition requires an attracting orbit, which not every parameter κ has.

For any $R \geq 0$, let $I_R := \ [\kappa_R^-, \kappa_R^+]$ with $\kappa_R^\pm := R + 2\pi i s_1 \pm i\pi$. Since $s_1 = 0$, we have $|\kappa| \leq R + \pi$ for all $\kappa \in I_R$. The ray ends $g_{\underline{s}^\pm}$ exist for potentials $t \geq t_R := x + 2\log(R + \pi + 3 + 2\pi AC') > t^*$. Fix $\varepsilon := 1/2$ and an appropriate δ from (3). We can then be sure that for all $t \geq t_R + \delta$, we have $\mathrm{Im}(\gamma_{\kappa_R^-}(t)) < 0$ and $\mathrm{Im}(\gamma_{\kappa_R^+}(t)) > 0$, while for all $\kappa \in I_R$, we have

$$\begin{aligned} \mathrm{Re}(g_{\underline{s}^\pm}(t_R + \delta)) &< (t_R + \delta) - R + 1/2 \quad \text{and} \\ \mathrm{Re}(\gamma_\kappa(t_R + \delta)) &< (t_R + \delta) - R + 1 \ . \end{aligned}$$

Now fix R large enough so that $t_R + \delta - R + 1 < 0$; this is possible since $t_R = O(\log R)$. As κ moves from κ_R^- to κ_R^+, there must be an intermediate value κ^* where $\gamma_{\kappa^*}(t) = 0$ for some $t > t_R + \delta$. The point of this construction is, of course, that κ^* has an attracting periodic orbit of period n with the required properties, at least when R is sufficiently large.

Indeed, with $\lambda^* := \exp(\kappa^*)$, the first n postsingular points 0, $E_{\lambda^*}(0)$, …, $E_{\lambda^*}^{\circ(n-1)}(0)$ are in the strip $|\mathrm{Im}(z)| \leq 2\pi A + \pi + 1$: the $2\pi A$ comes from the estimates on $2\pi|s_k|$; the π is the bound on $\mathrm{Im}(\kappa)$, and the final 1 is the error bound for the rays. The postsingular orbit $E_{\lambda^*}(0) = \lambda^*, \ldots E_{\lambda^*}^{\circ(n-2)}(0)$ has positive real parts, while $\mathrm{Re}(E_{\lambda^*}^{\circ(n-1)}(0)) \ll 0$. Now if R is large enough, then Lemma 3.4 shows that there is an attracting orbit of exact period n for κ^*, and $E_{\lambda^*}^{\circ(n-1)}(\gamma_{\kappa^*}(\ [0,\infty[\ ))$ is in the attracting basin. Hence $\gamma_{\kappa^*}(\ [0,\infty[\ )$ is also in the attracting basin with $\gamma_{\kappa^*}(0) = 0$, so it is in the characteristic Fatou component. By Lemma 3.4, the multiplier tends to 0 as R gets large.

Connect γ_{κ^*} to the attracting periodic point within the characteristic Fatou component and call this resulting curve $\gamma\colon (\ [0,\infty[\ ) \to \mathbb{C}$. It clearly satisfies the first three conditions for attracting dynamic rays. We argued above that the last two conditions were satisfied at least for large t. But since the curve γ is in the characteristic Fatou component, it is disjoint from all dynamic rays, and it is indeed an attracting dynamic ray at external address $s_1 s_2 \ldots s_{n-1}$.

The parameter κ^* sits in a hyperbolic component with the required combinatorics, and by Lemma 3.3 all parameters within this same component have attracting dynamic rays with the same property.

To show the statement about the curve within the hyperbolic component, start with $\gamma_{\kappa^*}(t) = 0 = g_{\underline{s}^\pm}(t) + o(1) = t - \kappa + o(1)$. If $n \geq 4$, then $E_{\lambda^*}(\gamma_{\kappa^*}(t)) = \lambda^* = g_{\sigma(\underline{s}^\pm)}(t) + o(1) = F(t) - \kappa^* + 2\pi i s_2 + o(1)$. In our construction, $|\mathrm{Im}(\kappa^*)| \leq \pi$, and $\mathrm{Re}(F(t)) \gg \mathrm{Re}(\kappa^*)$ for R large. Since $n \geq 4$, then the next E_{λ^*}-image must again have large positive real parts, so $\mathrm{Im}(\kappa^*)$ must tend to 0 and $\mathrm{Im}(\lambda^*) \to 2\pi i s_2$. If $n = 3$, then $E_{\lambda^*}(\gamma_{\kappa^*}(t)) = \lambda^* = g_{\sigma(\underline{s}^-)}(t) + i\pi + o(1) = F(t) - \kappa^* + 2\pi i s_2 + o(1)$ with $s_2 \in (\mathbb{Z} + \frac{1}{2})$. This time, the next image of λ^* must have large negative real parts, so again $\mathrm{Im}(\kappa^*) \to 0$ and $\mathrm{Im}(\lambda^*) \to 2\pi s_2$. In all cases, since the multiplier tends to 0 as $R \to \infty$, the existence of the curve follows from Lemma 2.1. □

We know from Lemma 2.1 that every hyperbolic component has a preferred homotopy class of curves stretching out to ∞ (that there are no other homotopy classes of such curves is shown in [S1, Section V]). These preferred homotopy classes of curves give a natural vertical order to hyperbolic components, much as the order for dynamic or parameter rays: we say that *some hyperbolic component is above another hyperbolic component* if the corresponding homotopy classes of curves have the appropriate vertical order.

Corollary 3.6 (Relative Position of Hyperbolic Component)
The vertical order of hyperbolic components is the same as the lexicographic order of their intermediate external addresses.

Proof. This follows from the previous proof as follows: if $\underline{s}' > \underline{s}''$ are two intermediate external addresses, then there is a periodic external address $\underline{s}$ between them: $\underline{s}' > \underline{s} > \underline{s}''$. In the construction in the proof of Theorem 3.5, the attracting dynamic ray for $\underline{s}'$ is always above $g_{\underline{s}}$, and the attracting dynamic ray for $\underline{s}''$ is always below $g_{\underline{s}}$. □

Remark. This vertical order can also be expressed in terms of parameter rays [S1, S2]: with the notation of the proof of Corollary 3.6, there is a parameter ray in the complex plane associated to external address $\underline{s}$, and the hyperbolic components for $\underline{s}'$ and $\underline{s}''$ are above respectively below this ray.

As this paper was being submitted, a manuscript by Devaney, Fagella and Jarque [DFJ] was released which contains the same sufficient condition for the existence of hyperbolic components as in our Theorem 3.5.

4 Characteristic Rays and Permutations

In this section, we investigate periodic points at which at least two periodic dynamic rays land, and show that the first return map of the periodic points permutes its rays transitively. This property is well known from quadratic polynomials; it depends on the fact that there is a single singularity, not on the degree of the map.

Definition 4.1 (Essential Orbit, Characteristic Point & Rays)
A periodic orbit will be called essential *if at least two dynamic rays land at each of its points. Suppose that a point z on an essential orbit is the landing point of two dynamic rays which separate the singular value from all the other points on the orbit of z; then the point z will be called the* characteristic point *of its orbit. The* characteristic rays *of the orbit will be the two dynamic rays landing at the characteristic point which separate the singular value from all the other rays landing at the same orbit.*

The following result describes the combinatorics of dynamic rays landing together. The statement is the same as for polynomials, but the usual proof (using "widths of sectors") does not apply without modification. Still, essential ideas are borrowed from Milnor [M1].

Lemma 4.2 (Permutation of Dynamic Rays)
Every essential periodic orbit has exactly one characteristic point and exactly two characteristic rays at this point. If more than two dynamic rays land at any periodic point, then the first return map of the periodic point permutes these rays cyclically.

PROOF. Let $z_1, z_2, \ldots, z_n = z_0$ be a periodic orbit of period n, labeled in the order of the dynamics, and let $r \geq 2$ be the number of dynamic rays at each of these points. This number is constant along the orbit. The r rays landing at any point z_k cut the complex plane into r connected components which will be called the "sectors" at z_k.

For any z_k, consider any sector which does not contain a left half plane. Let m' be the position of the first difference in the external addresses of the two rays bounding the sector (with $m' = 1$ if the first entries are different). We define the *singular index* of the sector to be $m := m' + 1$. For the sector which does contain a left half plane, we let the singular index be $m := 1$. Clearly, any sector with index $m \geq 2$ maps homeomorphically onto a sector with index $m - 1$ (for $m = 2$, it follows from the fact that the sector must contain a horizontal line segment which stretches infinitely to the right and maps onto an infinite segment of $\mathbb{R}^-$). It follows that for any sector not containing a left half plane, the index is one greater than the number of iterations it takes for this sector to map over a left half plane.

If the index of a sector equals 1 so that the sector contains a left half plane, then the sector will not map forward homeomorphically, but the image sector at the image point will contain the singular value. If the image sector also contains a left half plane, then the index remains 1; otherwise, the image point separates the singular value from a left half plane, and the index is strictly greater than 1. Each z_k has exactly one sector with index 1 (the unique sector containing a left half plane).

Every sector of every point z_k is periodic (as a local sector near z_k defined by the dynamic rays bounding it, not as a subset of $\mathbb{C}$), and so is the sequence of the indices. The index sequence of any sector must of course contain at least one index 1, and it cannot be the constant sequence: if all the entries of one cycle of sectors were equal to 1, then all the other cycles of sectors could never have index 1, a contradiction.

Therefore, every sequence of indices contains one or several positions at which it increases from 1 to a greater value. Associated to these positions are image sectors which contain the singular value but no left half plane. These sectors are ordered by inclusion. Fix the smallest such sector; it is bounded by two dynamic rays landing at one point z_k. By relabeling cyclically, let z_1 be this landing point. This is the characteristic point of the orbit and the two rays bounding the smallest sector are the characteristic rays. This shows the first statement.

Let $\alpha_1 > \alpha_2 > \ldots > \alpha_r$ be the indices of the sectors at z_1; no two of them can be equal because otherwise the corresponding sectors would map forward homeomorphically until they contained a left half plane at the same time. Of course, $\alpha_r = 1$ is the sector containing a left half plane.

Consider any cycle C of sectors and let α be the largest index within its period. Since indices are always decreasing unless they are equal to 1, the index α must occur for a sector containing the singular value but not a left half plane. Let z_k be the periodic point at which this sector is based. If $z_k = z_1$, then the sector with index α is the sector at z_1 containing the singular value. Hence all sectors for which the maximum is realized at z_1 are on the same orbit. This is true even if the sequence of indices contains several maxima and one of them is realized at z_1.

If, however, $z_k \neq z_1$, then $\alpha \leq \alpha_{r-1}$ because the point z_1 is in the sector with index α, and so are all the sectors at z_1 with indices $\alpha_1 > \ldots > \alpha_{r-1}$. The cycle C of sectors must map through z_1, but the only sector at z_1 it can map through is the sector containing a left half plane.

It follows that there are at most two cycles of sectors: their representatives at z_1 must include either the sector containing the singular value or the sector containing a left half plane, or both. Suppose that not all sectors are on the same orbit. Then the sector at z_1 containing a left half plane is fixed under the first return map of z_1 and has period n, and all the other $r-1$ sectors at z_1 are on the same orbit, so they are permuted transitively by the first return map of the dynamics and have period $(r-1)n$. But all sectors must have equal periods because all dynamic rays have equal periods, and this is possible only if $r=2$. □

5 Dynamic Roots

For an understanding of the dynamics, the most important rays are those which land together. We will now show that such are associated to attracting Fatou components.

Theorem 5.1 (Two Rays at Boundary Fixed Point)
Every periodic Fatou component with attracting dynamics of period $n \geq 2$ has a unique point on its boundary which is fixed by the first return map of the component and which is the landing point of at least two periodic dynamic rays.

Proof. First we prove uniqueness. Let U_1 be the characteristic Fatou component and let $z_1 \in \partial U_1$ be a boundary point which is periodic of period dividing n such that at least two dynamic rays land at z_1. We will need the two *characteristic rays* of this orbit (cf. Definition 4.1 and Lemma 4.2): these are two dynamic rays which land at a common point on the orbit of z_1 and separate the singular value from the remaining periodic orbit of z_1. Since z_1 cannot be separated from the singular value by such a ray pair, it must itself be the landing point of these two characteristic rays. These two rays and their landing point z_1 cut $\mathbb{C}$ into two open parts; let V be the one containing the singular value.

Now suppose that there is another periodic point $z_1' \in \partial U_1$ which is fixed by the first return map of U_1 and which is the landing point of two periodic dynamic rays. We have $z_1' \subset V$ because V contains $\overline{U}_1 - \{z_1\}$. The two characteristic dynamic rays landing at z_1' are then contained in V as well. But by symmetry between z_1 and z_1', it also follows that the two characteristic rays landing at z_1' bound an open sector V' which contains z_1 and all its rays, and this is a contradiction. This proves the claim about uniqueness.

For existence, we need the partition constructed in [SZ1, Section 4.3] and reviewed in Section 2. Let $\mathtt{u}_1, \mathtt{u}_2, \ldots, \mathtt{u}_{n-1}$ be the first $n-1$ entries of the itinerary of the singular value (such that the singular value 0 is in the region labeled $\mathtt{u}_1$, etc). By definition of the labels, $\mathtt{u}_1 = \mathtt{0}$. The n-th label is not defined.

For a dynamic ray $g_{\underline{s}}$ with bounded external address $\underline{s}$, it may happen that there are two curves within U_1 which connect the singular value to $+\infty$ such that they separate $g_{\underline{s}}$ from U_0. In this case, we will say that the ray $g_{\underline{s}}$ is *surrounded by* U_1, and rays $g_{\underline{s}'}$ with bounded external addresses $\underline{s}'$ sufficiently close to $\underline{s}$ will also be surrounded by U_1, so this is an open property in the sequence space $\mathcal{S}$. The Fatou component U_1 contains infinitely many non-homotopic curves connecting the singular value to $+\infty$, namely pull-backs of curves connecting $E_\lambda^{\circ(n-1)}(0)$ to ∞ within the Fatou component U_n containing a left half plane. Hence any ray which is not surrounded by U_1 is either *above* or *below* U_1 in the sense that the ray approaches $+\infty$ above or below all such curves within U_1. Similarly, we will say that rays are above or below $U_2, \dots, U_{n-1}$ (but not U_0).

Let $\underline{s}$ be the infimum of all bounded external addresses for which the dynamic ray $g_{\underline{s}}$ is above U_1 with respect to the lexicographic order. The first entry is finite because U_1 provides a lower bound for the rays, and U_1 itself is bounded above by the partition in U_0.

CLAIM. *The sequence $\underline{s}$ is a periodic external address of period n (without symbols $\pm\infty$), and its itinerary starts* $\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}$.

PROOF OF CLAIM. The first entry must be $\mathtt{u}_1$ because there are dynamic rays above U_1 in the same region as U_1.

Mapping forward one step, the external address $\sigma(\underline{s})$ will be the infimum of all bounded $\underline{s}'$ such that the dynamic ray $g_{\underline{s}'}$ is above U_2, and it follows that the second entry in the itinerary of $\underline{s}$ is $\mathtt{u}_2$. This argument can be repeated for the k-th entry of the itinerary and the external address $\sigma^{k-1}(\underline{s})$ as long as the Fatou component U_{k-1} does not surround U_{n-1}: in that case, U_{k-1} and everything it surrounds will map homeomorphically into one of the strips in the complement of U_0 with real parts bounded below and imaginary parts bounded above and below, and the vertical orders of dynamic rays in this region will be preserved.

However, if U_{k-1} does surround U_{n-1}, then the complement of U_{k-1} is "turned inside out": the region surrounded by U_{k-1} maps over a left half plane, while a sufficiently far left half plane outside of U_{k-1} map into U_1, and U_k must surround U_1. Then the external address $\sigma^{k-1}(\underline{s})$ will be the infimum of all bounded $\underline{s}'$ such that the dynamic ray $g_{\underline{s}'}$ is below U_1 but not separated from U_1 by a curve in U_k (compare Figure 3): roughly speaking, the vertical order within the region containing U_{k-1} is from bottom to top (1) the lower region boundary (part of U_0), then (2) curves in U_{k-1} to $+\infty$, (3) the component U_{n-1}, (4) more curves in U_{k-1} to $+\infty$, (5) dynamic rays at external addresses near $\sigma^{k-2}(\underline{s})$, (6) the upper region boundary (part of U_0 again). The region boundaries map to U_1, so a large circle in $\mathbb{C}$ starting and ending at U_1 will meet: (1) U_1, (2) curves in U_k to $+\infty$, (3) a left half plane, (4) curves in U_k to $+\infty$, (5) dynamic rays at external addresses near $\sigma^{k-1}(\underline{s})$, (6) U_1 again. Since all of U_k must be contained within a single region of the partition, the component U_k will surround U_1 and the dynamic rays at external addresses near $\sigma^{k-1}(\underline{s})$. We also see that the k-th entry in the itinerary of $g_{\underline{s}}$ is $\mathtt{u}_k$.

These arguments can be repeated if another $U_{k'-1}$ surrounds U_{n-1} (for $k' > k$). In the special case that this happens in such a way that U_{n-1} is not separated from the dynamic rays near $\sigma^{k-1}(\underline{s})$ by curves in U_k, then the situation is restored to the initial configuration so that $\sigma^{k-1}(\underline{s})$ is the infimum of external addresses of dynamic rays above U_k. In any case, the first $n-1$ entries in the sequence $\underline{s}$ are finite, and the first $n-1$ entries in the itinerary of $g_{\underline{s}}$ are equal to $\mathtt{u}_1 \dots \mathtt{u}_{n-1}$.

Once we arrive at step n, it may or may not have happened that an earlier component U_{k-1} had surrounded U_{n-1}, for some $k = 3, 4, \dots, n-1$. If this never happened, then $\sigma^{n-2}(\underline{s})$ is the infimum of external addresses of all dynamic rays above U_{n-1}, and rays at external addresses $\underline{s}'$ near $\sigma^{n-2}(\underline{s})$ are not separated from U_{n-1} by the region boundary in U_0. In the

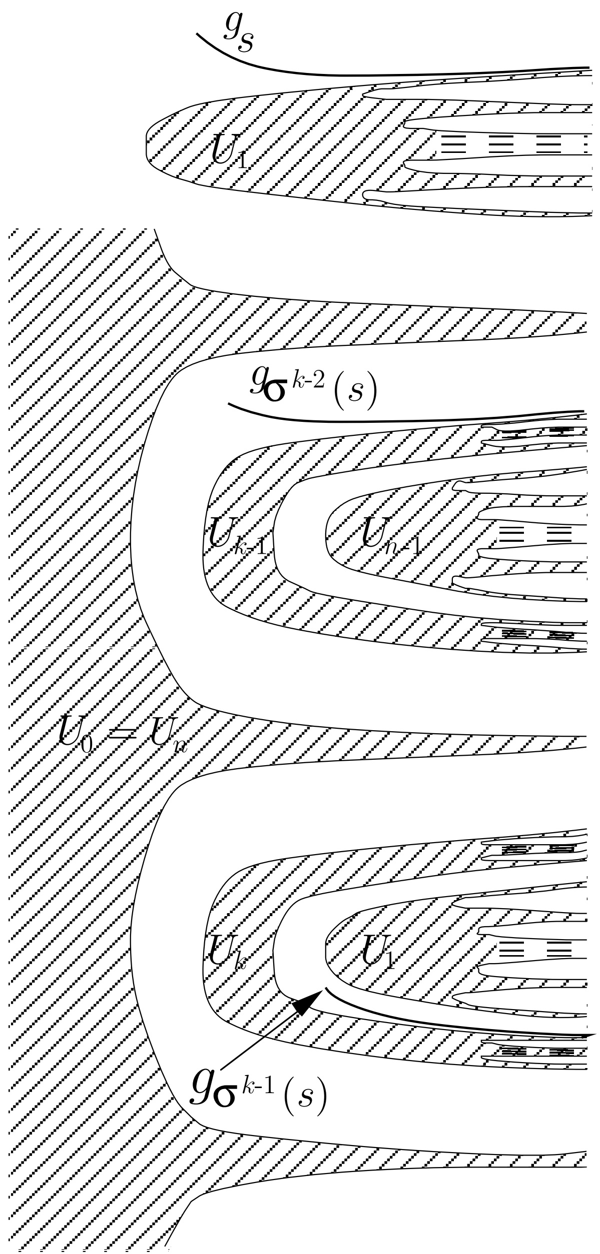

Figure 3: *Illustration of the construction of the dynamic root. Top: the dynamic ray* $g_{\underline{s}}$ *is just above the characteristic Fatou component. Middle: the situation after* $k-2$ *iterations, where the Fatou component* U_{k-1} *surrounds* U_{n-1}*. Bottom: one step later.*

next step, dynamic rays at external addresses $\underline{s}''$ near $\sigma^{n-1}(\underline{s})$ are below U_1 and in the same region as U_1, so the n-th entry in the itinerary of $\underline{s}$ is $\mathtt{u}_1$ like the first entry, and $\sigma^{n-1}(\underline{s})$ is the infimum of all external addresses of rays below U_1 in the same region as U_1. Mapping forward one last step, we see that $\sigma^n(\underline{s})$ is the infimum of external addresses of all dynamic rays above U_1, so $\sigma^n(\underline{s}) = \underline{s}$. It follows that $\underline{s}$ is periodic of period n, all its entries are finite, the first n entries of its itinerary are equal to $\mathtt{u}_1\mathtt{u}_2\ldots\mathtt{u}_{n-1}\mathtt{u}_1$, and they repeat periodically. This proves the CLAIM.

If it does happen that a component U_k surrounds U_{n-1} for some $k \leq n-2$, then let k be the last index for which this happens. We saw that U_{k+1} surrounds U_1, and $\sigma^k(\underline{s})$ is the infimum of external addresses of rays below U_1 but not separated from U_1 by U_{k+1}. Mapping $n-2-k$ steps further, it follows that U_{n-1} surrounds U_{n-1-k}, and $\sigma^{n-2}(\underline{s})$ is the infimum of external addresses of rays below U_{n-1-k} but not separated from U_{n-1-k} by U_{n-1}. In the next step, $\sigma^{n-1}(\underline{s})$ is the infimum of external addresses of rays below U_{n-k} but not separated from U_{n-k} by $U_n = U_0$ or the region boundary. Hence the n-th entry in the itinerary of $g_{\underline{s}}$ is $\mathtt{u}_{n-k}$ with $n-k \geq 2$. One more step forward, $\sigma^n(\underline{s})$ is the infimum of external addresses of rays below U_{n-k+1} but not separated from U_{n-k+1} by U_1. In other words, $\sigma^n(\underline{s})$ is the infimum of external addresses of rays above U_1. Again we have $\sigma^n(\underline{s}) = \underline{s}$, so $\underline{s}$ is periodic of period n, all its entries are finite, the first n entries of its itinerary are $\mathtt{u}_1\mathtt{u}_2\ldots\mathtt{u}_{n-1}\mathtt{u}_{n-k}$, and they repeat periodically.

We can argue similarly for the supremum $\underline{s}'$ of external addresses of all dynamic rays below U_1. Exactly as above, it follows that this is a periodic external address of period n, and the ray $g_{\underline{s}'}$ has the same itinerary as $g_{\underline{s}}$. By the results from [SZ1] mentioned in Section 2, the two dynamic rays $g_{\underline{s}}$ and $g_{\underline{s}'}$ land at a common repelling periodic point z, say. In order to prove that $z \in \partial U_1$, let ℓ be the hyperbolic distance of z to ∂U_1 in the hyperbolic domain consisting of $\mathbb{C}$ with the closure of the singular orbit removed. Assume that $\ell > 0$. The hyperbolic distance between the unique periodic inverse image of z and $U_0 = U_n$ is then less than ℓ. We take $n-1$ further pull-back steps along the periodic orbit of z; since the itinerary of z in those steps is the same as that of the singular orbit, the branches for the pull-back of z are those mapping U_n to $U_{n-1}, U_{n-2}, \ldots, U_1$, and hyperbolic distances are decreased in every step. After n steps, z is mapped back to itself and its hyperbolic distance to U_1 cannot have decreased. This contradiction shows that $\ell = 0$ and $z \in U_1$.

This proves the statement for the periodic Fatou component U_1, and for the others it follows easily. □

REMARK. The same statement holds also for parabolic dynamics; the proof requires only the same modifications as in [SZ1, Section 4.3].

Definition 5.2 (Dynamic Root)
In any exponential dynamics with attracting orbit of period $n \geq 2$, the unique point of the characteristic Fatou component which is fixed under the first return map of the component and which is the landing point of at least two dynamic rays (as described in Theorem 5.1) will be called the dynamic root *of the characteristic Fatou component.*

Lemma 5.3 (Rays at Dynamic Root)
In attracting exponential dynamics, the two characteristic rays of the dynamic root of the characteristic Fatou component separate this Fatou component from all other periodic Fatou components.

Proof. Let U_1 be the characteristic Fatou component and let z_1 be its dynamic root. The characteristic dynamic rays at z_1 separate the singular value and thus U_1 from all other points on the orbit of z_1. Any periodic Fatou component U_i has a point z_i on its boundary. If $z_i \neq z_1$, then U_i is separated from the singular value by the characteristic ray pair. Let the periods of the attracting orbit and of z_1 be n and k, respectively. By uniqueness of the dynamic root (Theorem 5.1), each U_i has a unique z_i on its boundary, so the number of different periodic Fatou components with z_1 on its boundary is exactly n/k and the first return map of z_1 must permute these k components cyclically. Hence the gaps between cyclically adjacent periodic Fatou components at z_1 are also permuted cyclically, and at least one of the must contain a periodic dynamic ray landing at z_1; hence all gaps do, and all periodic Fatou components at z_1 are separated by periodic dynamic rays landing at z_1. (Conversely, it follows from Lemma 4.2 that all the rays landing at z_1 are separated by periodic Fatou components provided at least two periodic Fatou components have z_1 as their common dynamic root.) □

6 Uniqueness of the Classification

In this section, we will prove that the hyperbolic component associated to any intermediate external address is unique. Using the dynamic roots of periodic Fatou components, we set up a quasiconformal map between any two candidate dynamics and show that it can be promoted to a conformal conjugacy.

Lemma 6.1 (Preferred Curve to Boundary Fixed Point)
For an exponential map E_λ with an attracting periodic orbit of period n, let U_1 be the characteristic Fatou component. If the multiplier of the attracting orbit is real and positive, then there is a preferred invariant curve within U_1 which contains the entire singular orbit of $E_\lambda^{\circ n}$ and which lands at a well-defined point $w \in \partial U_1$ with $E_\lambda^{\circ n}(w) = w$.

Proof. Let U_1 be the characteristic Fatou component, let n be its period and let a_1 be the attracting periodic point within U_1. Then there is a unique closed neighborhood D of a_1 which corresponds to a round disk in linearizing coordinates of a_1 and which contains the singular value 0 on its boundary. Let $\gamma_D \subset D$ be the curve corresponding to a diameter in linearizing coordinates such that $0 \in \gamma_D$. Then the first return map $E_\lambda^{\circ n}$ of U_1 sends D into itself and γ_D into itself.

The singular value 0 cuts γ_D into two radii; let γ_+ be the one which ends at 0 and γ_- the other one (then $\gamma_D = \gamma_+ \cup \gamma_- \cup \{0\}$). There is a unique curve $\gamma \subset U_1$ which extends γ_- and which satisfies $E_\lambda^{\circ n}(\gamma) = \gamma$: such a curve can be constructed by an infinite sequence of pull-backs, starting at γ^- and always choosing the branch which extends γ^-. Then $E_\lambda^{\circ n}\colon \gamma \to \gamma$ is a homeomorphism.

The curve γ can easily be parametrized as $\gamma:\mathbb{R}\to U_1$ so that $E_\lambda^{\circ n}(\gamma(t))=\gamma(t+1)$ (this is far from unique). We have $\lim_{t\to+\infty}\gamma(t)=a_1$. We want to show that $\lim_{t\to-\infty}\gamma(t)$ exists in ∂U_1. We will use a modification of the known standard proofs for landing of external dynamic rays of polynomials. Let

$$U' := U_1 \setminus \overline{\bigcup_{k\geq 0} E_\lambda^{\circ kn}(0)}\ .$$

Then $E_\lambda^{\circ n}: U'\to U'$ is a covering map, hence a local isometry with respect to the unique normalized hyperbolic metric of U'. Let $(w_k)\subset\gamma$ be a sequence of points such that $E_\lambda^{\circ n}(w_k)=w_{k-1}$ for $k\geq 1$ (e.g. $w_k:=\gamma(-k)$). Then the hyperbolic length of the segment of γ between w_k and w_{k-1} is the same for all k. By continuity, it follows that there is an $s>0$ such that the hyperbolic distance in U' between any $z\in\gamma$ and $E_\lambda^{\circ n}(z)$ is at most s. But since $\gamma(t)\to\partial U_1$ as $t\to-\infty$ (points $\gamma(t)$ for large negative t need longer and longer to iterate near a_1), and the density of the hyperbolic metric tends to ∞ near ∂U_1, it follows that $|\gamma(t)-E_\lambda^{\circ n}(\gamma(t))|\to 0$ as $t\to-\infty$. Therefore, any limit point of γ is a fixed point of $E_\lambda^{\circ n}$. Since the limit set is connected and the set of fixed points is discrete, it follows that γ lands indeed at a well-defined boundary point of U_1 which is fixed under $E_\lambda^{\circ n}$. The curve of the claim is $\gamma\cup\{0\}\cup\gamma^+$. □

REMARK. In fact, it is not difficult to show that γ is a hyperbolic isometry of U_1 [S1]: this is easier if the first return dynamics is conjugated to the map $M\circ\exp:\mathbb{H}^-\to\mathbb{H}^-$, where $\mathbb{H}^-$ is the left half plane, $\exp:\mathbb{H}^-\to\mathbb{D}^*$ is a universal cover and $M:\mathbb{D}\to\mathbb{H}^-$ is any conformal isomorphism.

Theorem 6.2 (Conformal Conjugation)
Suppose that two exponential maps have attracting orbits of equal period $n\geq 2$ with positive real multiplier which both have attracting dynamic rays at the same intermediate external address $s_1s_2\ldots s_{n-1}$. Suppose in addition that for both maps the preferred invariant curve within the characteristic Fatou component lands at the dynamic root. Then both maps are identical.

PROOF. Let E_λ and $E_{\lambda'}$ be two exponential maps satisfying the assumptions of the theorem. Let $a_1,a_2,\ldots,a_n=a_0$ be the attracting orbit of E_λ with a_1 in the characteristic Fatou component U_1, let $U_2,\ldots,U_n=U_0$ be the other periodic Fatou components labeled cyclically, and let V_{n+1} be a closed round disk with respect to linearizing coordinates of a_1, large enough so as to contain 0 in its interior. For $k=n,n-1,\ldots,2,1,0$, let $V_k\subset U_k$ be the domain $E_\lambda^{\circ(k-(n+1))}(V_{n+1})\cap U_k$. Then $V_1\supset V_{n+1}$ and $V_0\supset V_n$. For $k=0,1,2,\ldots,n-1$, let $\tilde V_k$ consist of all translates of V_k under $2\pi i\mathbb{Z}$ (where $\tilde V_0=V_0$ is introduced only for notational convenience). Denote the corresponding sets for $E_{\lambda'}$ as U'_k, V'_k, and $\tilde V'_k$, where the size of V'_{n+1} is chosen so that the dynamics on it is conformally conjugate to the dynamics on V_{n+1}, respecting the singular value.

We will construct a quasiconformal homeomorphism $\varphi:\mathbb{C}\to\mathbb{C}$ from the dynamic plane of E_λ to the dynamic plane of $E_{\lambda'}$ which will eventually turn into a conformal conjugation. Since the multipliers at the attracting fixed points are the same, we can clearly define $\varphi\colon V_k \to V'_k$ so that it respects the dynamics on V_k in the sense that for $z \in V_k$, we have

$$\varphi(E_\lambda(z)) = E_{\lambda'}(\varphi(z)) \ . \tag{4}$$

It becomes unique on V_1 by the requirement that $\varphi(0) = 0$, and in view of (4) it is unique on all V_k. Extend the definition to $\varphi\colon \tilde{V}_k \to \tilde{V}'_k$ by translation, which is again unique and respects (4).

All the connected components of V_k (for $1 \le k \le n-1$) extend to $+\infty$ in a unique homotopy class. The natural vertical order of these connected components at large real parts is the same for E_λ and $E_{\lambda'}$ because each connected component is immediately above a dynamic ray at known external address.

Clearly, by compactness, $|\varphi(z) - z|$ is bounded on ∂V_{n+1}. But by translation symmetry, $|\varphi(z) - z|$ is then bounded on ∂V_n. This argument can be repeated along the finitely many pull-back steps, and it shows that there is a global constant $\delta \ge 0$ such that $|\varphi(z) - z| \le \delta$ on all ∂V_k.

The next step in the construction is to extend φ to a quasiconformal homeomorphism from $\mathbb{C}$ to itself. We first do that at real parts greater than R, where R is a large positive constant to be chosen depending on the two exponential parameters and $s_1 s_2 \dots s_{n-1}$. To the right of real part R, the gaps between the connected components of the $\tilde{V}_k$ are finite in number, up to symmetry translation by $2\pi i\mathbb{Z}$. Let G by any such gap (bounded to the left by the vertical line at real part R and above and below by two adjacent connected components of the $\tilde{V}_k$). In general, the vertical width of this gap will decrease exponentially as the real parts increase. Map G forward under E_λ until the vertical width becomes positive or one of the two bounding components maps over a left half plane. This happens after a finite number $m < n$ of steps which is entirely encoded in $s_1 s_2 \dots s_{n-1}$ and the indices k of the two components. In particular, it is the same for G and its counterpart G' for $E_{\lambda'}$. There exists a quasiconformal homeomorphism $\varphi'\colon E_\lambda^{\circ m}(G) \to E_{\lambda'}^{\circ m}(G')$ which extends the known map φ on the $E_\lambda^{\circ m}$-image of the upper and lower boundary of G because the vertical width of $E_\lambda^{\circ m}(G)$ and $E_{\lambda'}^{\circ m}(G')$ are bounded below, and $|\varphi(z) - z| < \delta$ for all z on the vertical boundaries. Now φ can be defined on G as the natural pull-back of φ' by the dynamics: we set $\varphi := E_{\lambda'}^{\circ(-m)} \circ \varphi \circ E_\lambda^{\circ m}$ on G, of course choosing the branch of $E_{\lambda'}^{\circ(-m)}$ which lands in G'.

Doing this extension on finitely many gaps and translating by $2\pi i\mathbb{Z}$, we can define φ for all real parts greater than R. The map φ is already defined on V_0, and $\mathbb{C} \setminus V_0$ is the $2\pi i\mathbb{Z}$-image of a single connected piece. Each of these pieces contains finitely many domains on which φ is already defined (one in each $\tilde{V}_k$ for $k = 1, 2, \dots, n-1$), and the vertical order of these domains is the same for E_λ and $E_{\lambda'}$. All of the domains on which φ is still undefined are translates of each other under $2\pi i\mathbb{Z}$, and each of them is compact in $\mathbb{C}$. Hence φ does extend to a quasiconformal map from $\mathbb{C}$ to itself, and its distortion is bounded away from 1. Note that we do not claim that the extension of φ away from the $\tilde{V}_k$ respects the dynamics.

The next step is to promote φ to a conformal conjugation. Let $\varphi_1 := \varphi$. The construction assures that there is another quasiconformal homeomorphism $\varphi_2\colon \mathbb{C} \to \mathbb{C}$ with $\varphi_1 \circ E_\lambda = E_{\lambda'} \circ \varphi_2$ and which coincides with φ_1 on $\cup_k \tilde{V}_k$: we set $\varphi_2 := E_{\lambda'}^{-1} \circ \varphi_1 \circ E_\lambda$ and only have to choose the right branch of $E_{\lambda'}^{-1}$. On V_0 we can do that by stipulating that φ_2 coincides with φ_1 (this is just (4)); the other branches are then forced by continuity since each connected component of $\mathbb{C}\setminus V_0$ is simply connected. In order to see that φ_2 coincides with φ_1 on $\tilde{V}_k$, all we need to verify is that it maps the unique periodic component V_k onto the unique periodic component V'_k. It is here that the assumption about the preferred invariant curves comes in. For E_λ, this invariant curve is a diameter of V_{n+1} with respect to linearizing coordinates: it starts at the singular value 0 and traverses V_{n+1}, leaving V_{n+1} on its way to the dynamic root of U_1; let γ be this invariant curve, and let γ' be the counterpart for $E_{\lambda'}$.

The restriction $\gamma\cap V_{n+1}$ is a diameter and extends to a curve in $V_1\supset V_{n+1}$ starting at $+\infty$ and ending at some $v_1\in\partial V_1$, running through 0; we write $V_1\cap\gamma = \,]+\infty,0[\,\cup\,[0,v_1]$ as the union of two subcurves. Now $E_\lambda^{-1}(\,]\infty,0[\,)$ are countably many curves in V_0 connecting $-\infty$ to $+\infty$, and $E_\lambda^{-1}(\,]0,v_1[\,)$ are curves connecting $-\infty$ to the countably many connected components of $\mathbb{C}\setminus V_0$. Exactly one of these curves runs through the attracting periodic point in V_0, and this curve singles out a preferred connected component of $\mathbb{C}\setminus V_0$. Since we have constructed φ so that $\varphi(V_{n+1})=V'_{n+1}$ and $\varphi(0)=0$ with $0\in(\partial V_{n+1}\cap\gamma)$ and $0\in(\partial V'_{n+1}\cap\gamma')$, we have $\varphi(\gamma)=\gamma'$. Hence $\varphi=\varphi_1$ maps the preferred connected component of $\mathbb{C}\setminus V_0$ to the preferred connected component of $\mathbb{C}\setminus V'_0$: it is always the one containing the dynamic root of V_0 resp. V'_0. The map φ_2 must have the same property. But this preferred connected component of $\mathbb{C}\setminus V_0$ is the connected component containing the singular value: the dynamic root of U_{n-1} separates U_{n-1} from a left half plane, so the dynamic root of U_n separates U_n from the image of a left half plane, which is a neighborhood of the singular value. Thus φ_2 maps V_1 onto V'_1 and not onto one of its translates. For the other V_k this follows now because the connected components of $\mathbb{C}\setminus V_0$ any V_k sits in is encoded in the attracting dynamic ray and thus in $s_1 s_2\dots s_{n-1}$. Therefore, φ_1 and φ_2 coincide on all V_k.

The same step can be continued, yielding a sequence of quasiconformal homeomorphisms $\varphi_j\colon\mathbb{C}\to\mathbb{C}$. All these coincide on $\cup_k\tilde{V}_k$ and all are quasiconformal with the same bound on the dilatation. Every φ_j is conformal on $E_\lambda^{-(j-1)}(\cup_k\tilde{V}_k)$. Now $\bigcup_j E_\lambda^{-(j-1)}(\cup_k\tilde{V}_k)$ fills up the entire Fatou set, while the Julia set has measure zero by [EL1, EL3]. Hence the support of the bounded quasiconformal dilatation of the φ_j converges to zero, so the φ_j converge uniformly to a conformal isomorphism of $\mathbb{C}$, up to postcomposition with an automorphism. However, the point 0 is fixed, and so is the vertical translation symmetry by $2\pi i$. Hence the φ_j converge uniformly to the identity, and since $\varphi_j\circ E_\lambda = E_{\lambda'}\circ\varphi_{j+1}$, it follows that $E_\lambda = E_{\lambda'}$. This is what we claimed. □

Corollary 6.3 (Uniqueness of Classification)
For any intermediate external address $s_1 s_2\dots s_{n-1}$ with $s_1=0$, there is a unique hyperbolic component of period n in λ-space at this address. In κ-space, this is true for any $s_1\in\mathbb{Z}$ (resp. $s_1\in(\mathbb{Z}+\frac{1}{2})$ if $n=2$).

Proof. Any hyperbolic component contains a parameter for which the attracting orbit has multiplier $+1/2$, say (this is clear) and for which the preferred invariant curve lands at the dynamic root (this can be achieved by a quasiconformal deformation, which connects the initial dynamics to the desired one by a path in parameter space which must run within the same hyperbolic component). The external address $s_1 s_2\dots s_{n-1}$ is an invariant of the entire component and unchanged in this procedure. If there were two hyperbolic components of period n at the same intermediate external address, then they would contain an identical parameter by Theorem 6.2, and this is a contradiction. This finishes the claim in λ-space. Since κ-space is a universal cover, we have classified all hyperbolic components of period $n \geq 3$ with $|\mathrm{Im}(\kappa)| < \pi$, and the others are obtained simply by translation of κ by $2\pi i n$ (with $n \in \mathbb{Z}$) and adding $-n$ to all entries in the intermediate external address. $\square$

Corollary 6.4 (Preferred Coordinates in Hyperbolic Component)
Any hyperbolic component of period $n \geq 2$ can uniquely be parametrized by the multiplier μ of the attracting orbit, together with the label of the sector above or below the central line of the component. The label of the sector is naturally an element of $\mathbb{Z} \setminus \{0\}$.

Proof. Let W be the given hyperbolic component. We know that the multiplier map $\mu: W \to \mathbb{D}^*$ is a universal cover, so every $\kappa \in W$ is uniquely specified by μ, together with the data on which sheet of the cover it is on. The sectors of W are bounded by $\mu^{-1}(\ ]0,1[\ )$, and exactly one sector boundary is such that the preferred invariant curve lands at the dynamic root. This property is preserved along the entire branch of $\mu^{-1}(\ ]0,1[\ )$ by continuity, and uniqueness of the sector boundary is Theorem 6.2.

We label the sectors "to the right" of this preferred sector boundary by the positive integers and the sectors "to the left" by the negative integers (omitting the integer 0 as a sector number: this integer corresponds to the exterior of the component; compare [LS, Section 12]). The choice of "right" and "left" are determined so that the "right sector" near the preferred boundary has multipliers $r\exp(2\pi i\vartheta)$ with ϑ small positive. $\square$

External addresses for internal rays. An *internal ray at angle* $\vartheta \in \mathbb{S}^1$ of a hyperbolic component W is any branch of $\mu^{-1}(\ ]0, e^{2\pi i\vartheta}[\ )$, and there are countably many rays with the same angle. In fact, by Corollary 6.4, an internal ray is canonically labeled by a unique real number ϑ so that the fractional part of ϑ is the angle and the integer part distinguishes the sectors: all we need to require is that the internal ray for $\vartheta = 0$ is the ray between sectors -1 and 1; then if $\vartheta > 0$ (resp. $\vartheta < 0$) is an integer, the corresponding internal ray is between sectors ϑ and $\vartheta + 1$ (resp. between sectors ϑ and $\vartheta - 1$). Theorem 7.1 extends this parametrization to ∂W. This way, the external address $s_1, s_2, \ldots, s_{n-1}$ of the hyperbolic component W (with $s_1, \ldots, s_{n-2} \in \mathbb{Z}$ and $s_{n-1} \in (\mathbb{Z} + \frac{1}{2})$) extends naturally to an external address $s_1, s_2, \ldots, s_{n-1}, \vartheta$ of internal rays. Note that if $\vartheta > \vartheta'$, then the internal ϑ-ray approaches $+\infty$ *below* the ϑ'-ray: this reflects the fact that within hyperbolic components of period n, the n-th image of the dynamics is "upside down" in a left half plane. The space of parameter rays [S2] and internal rays of hyperbolic components is totally ordered by the vertical order of the horizontal approach to $+\infty$, and this vertical order is the same as the lexicographic order of the corresponding external addresses of parameter rays and intermediate external addresses for internal rays, with the extra rule that the order for the real value ϑ after the half-integer is reversed.

7 Further Results

In this section, we give a more complete description of hyperbolic components of the exponential family by stating relevant further results from [S1]; these will be published separately.

Theorem 7.1 (Boundary of Hyperbolic Components)
For any hyperbolic component W of period $n \geq 2$, there is a canonical homeomorphism $h_W: \mathbb{R} \to \partial W$, which is uniquely defined by the following conditions

- *for any $\vartheta \in \mathbb{R}$, the parameter $h_W(\vartheta)$ has an indifferent orbit with multiplier* $\exp(2\pi i h_W(\vartheta))$*;*
- *the curve h_W is smooth and analytic for $\vartheta \in (\mathbb{R} \setminus \mathbb{Z})$ and has cusps for $\vartheta \in (\mathbb{Z} \setminus \{0\})$;*
- *for every branch of μ^{-1}, the image of $]0, e^{2\pi it}[$ lands (as $|\mu| \to 1$) at a boundary point in ∂W which equals $h_W(\vartheta + n)$ for some $n \in \mathbb{Z}$;*
- *there is a unique branch of $\mu^{-1}(\,]0,1[\,)$ for which the preferred invariant curve of the characteristic Fatou component lands at the dynamic root; for this branch, the image of $]0,1[$ lands (as $|\mu| \to 1$) at the point $h_W(0)$.*

In particular, the boundary of any hyperbolic component is connected, and the component contains a unique homotopy class of curves connecting an arbitrary point to $+\infty$. Two hyperbolic components can have a common boundary point only if one period strictly divides the other.

All these statements are true in λ-space and in κ-space alike.

The unique boundary point $h_W(0)$ is called the *root* of W.

In [EL2], Eremenko and Lyubich stated three conjectures about the exponential parameter space. Two of these are: density of hyperbolic dynamics in the space of exponential maps and connectedness of the boundary of any hyperbolic component. To state the third, we say that two hyperbolic components are *adjacent* if they have a common boundary point, and they are *on the same tree* if they can be connected by a finite sequence of adjacent hyperbolic components. The conjecture then says that there are infinitely many such trees. The second and third conjectures are now proved (the second one being contained in Theorem 7.1); the first is still an open problem.

There is an interesting interplay between hyperbolic components and *parameter rays*: these are curves in parameter space for which the singular value escapes on given dynamic rays. They are uniquely labeled by external addresses, like dynamic rays. From the point of view of hyperbolic components, the most interesting parameter rays are those with periodic external addresses.

Theorem 7.2 (Landing of Periodic Parameter Rays)
Every parameter ray of some period n lands at a parabolic parameter, which is on the boundary of some hyperbolic component of period n. Every parabolic parameter κ is the landing point of one or two periodic parameter rays. If the parabolic orbit has period k and the period of the Fatou components is n, then

- *if $k < n$, then κ is on the boundary of exactly one hyperbolic component W of period k and W' of period n, exactly two periodic parameter rays land at κ, both of period n, and $\kappa = h_{W'}(0) = h_W(\vartheta)$ with $\vartheta \in ((\mathbb{Z}/k) \setminus \mathbb{Z})$;*

- *if $k = n$, then κ is on the boundary of a unique hyperbolic component W of period n, $\kappa = h_W(n)$ for the homeomorphism h_W from Theorem 7.1 with $n \in \mathbb{Z}$, and κ is the landing point of two parameter rays iff $n = 0$.*

Our combinatorial coding of hyperbolic components so far has been in terms of intermediate external addresses; this is a concept which has no direct counterpart for the Mandelbrot set. We can also code the component by the external addresses of the two parameter rays landing at its root (which are the same as the external addresses of the two characteristic rays landing at the dynamic root of the characteristic Fatou component, for any parameter within the component): this says something about how components are nested, but in order to relate the external addresses of a component W to those of a component directly attached to it, more information is needed (the external addresses of the rays landing at the boundary points $h_W(n)$ with $n \in \mathbb{Z} \setminus \{0\}$). A different coding can be given in terms of kneading sequences.

Definition 7.3 (Kneading Sequence of Hyperbolic Components)
For any hyperbolic component of period $n \geq 2$, pick any of its parameters and let $\overline{\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}\mathtt{u}_n}$ be the common itinerary of the two characteristic rays landing at the dynamic root of the characteristic Fatou component (where $\mathtt{u}_n \in \{\mathtt{u}_1, \dots, \mathtt{u}_{n-1}\}$ as in Theorem 5.1). Then we say that $\overline{\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}\mathtt{u}_n}$ is the kneading sequence *of the hyperbolic component.*

The hyperbolic component consists of countably many sectors, labeled by non-zero integers as in Corollary 6.4. Then to the sector labeled $\mathtt{m}$, we associate the periodic kneading sequence $\overline{\mathtt{u}_1\mathtt{u}_2\dots\mathtt{u}_{n-1}(\mathtt{u}_n+\mathtt{m})}$.

This way, all sectors have kneading sequence which are different from each other and from the kneading sequence of the component. It turns out that for any hyperbolic component bifurcating from sector $\mathtt{m}$, the kneading sequence starts with the first n entries of the kneading sequence of the sector, in analogy to the finite degree unicritical case [LS, Section 12] (a related result has been announced independently by Fagella and Jarque as work in progress). In fact, these kneading sequence can naturally be converted into *internal addresses* as in [LS], and they contain a lot of information about the exponential parameter space.

References

[BR] I. Noël Baker and Phil J. Rippon: *Iteration of exponential functions.* Ann. Acad. Sci. Fenn., Series A.I. Math. **9** (1984), 49–77.

[DFJ] Robert Devaney, Núria Fagella, Xavier Jarque: *Hyperbolic components of the complex exponential family.* Manuscript (2000).

[DGH] Clara Bodelón, Robert Devaney, Lisa Goldberg, Michael Hayes, John Hubbard and Gareth Roberts: *Hairs for the complex exponential family.* International Journal of Bifurcation and Chaos **9** 8 (1999), 1517–1534.

[DH] Adrien Douady and John Hubbard: *Etude dynamique des polynômes complexes.* Prépublications mathématiques d'Orsay **2** (1984) and **4** (1985).

[EL1] Alexandre Eremenko and Mikhail Lyubich: *Итерации целых функции (Iterates of Entire Functions).* Preprint, physico-technical institute of low-temperatures Kharkov **6** (1984).

[EL2] Alexandre Eremenko and Mikhail Lyubich: *Структурная устоичивость в некоторых семействах целых функции (Structural stability in some families of entire functions)*. Preprint, physico-technical institute of low-temperatures Kharkov **29** (1984).

[EL3] Alexandre Eremenko and Mikhail Lyubich: *Dynamical properties of some classes of entire functions*. Annales de l'Institut Fourier, Grenoble **42** 4 (1992), 989–1020.

[LS] Eike Lau and Dierk Schleicher: *Internal addresses in the Mandelbrot set and irreducibility of polynomials*. Preprint, Institute of Mathematical Sciences at Stony Brook **19** (1994).

[M1] John Milnor: *Periodic orbits, external rays and the Mandelbrot set; an expository account. Astérisque* **261** (2000), 277–333.

[S1] Dierk Schleicher: *On the dynamics of iterated exponential maps*. Habilitationsschrift, Technische Universität München (1999).

[S2] Dierk Schleicher: *Exponential maps with escaping singular orbit*. In preparation.

[SZ1] Dierk Schleicher and Johannes Zimmer: *Dynamic rays for exponential maps*. Preprint, Institute for Mathematical Sciences at Stony Brook **9** (1999).

[SZ2] Dierk Schleicher and Johannes Zimmer: *Escaping points for exponential maps*. Manuscript (2000).

Dierk Schleicher, Institute for Mathematical Sciences, SUNY SB, Stony Brook, NY 11794-3660, USA, dierk@math.sunysb.edu.